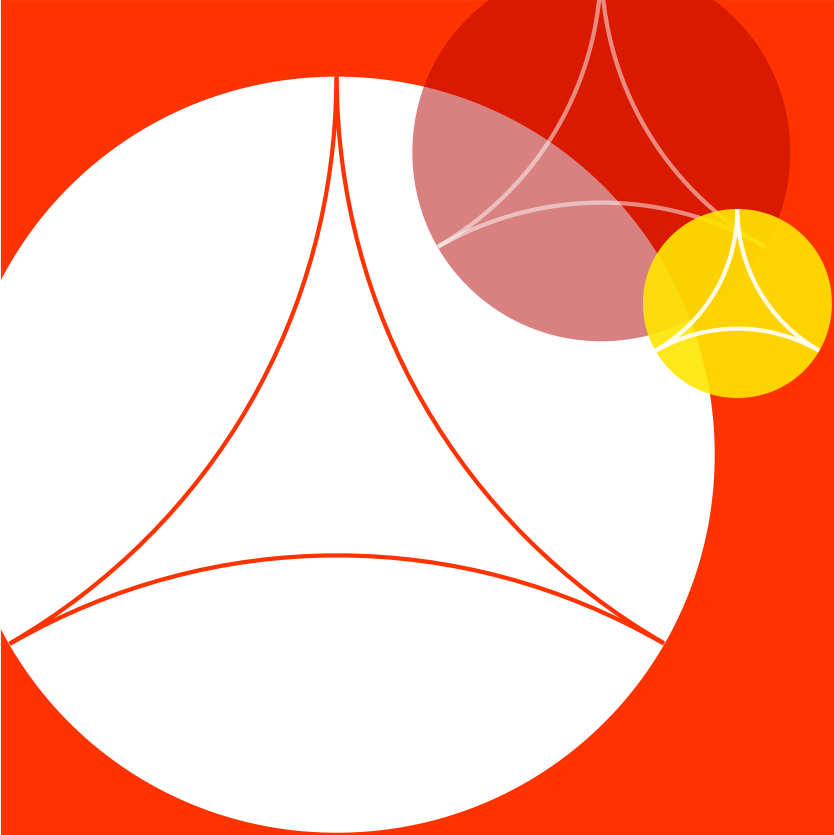

Clara Löh

Geometric group theory, an introduction

January 30, 2015 – 22:12
Version for the course in 2014/15

Please send corrections and suggestions to clara.loeh@mathematik.uni-regensburg.de

Clara Löh
clara.loeh@mathematatik.uni-regensburg.de
http://www.mathematik.uni-regensburg.de/loeh/

Fakultät für Mathematik
Universität Regensburg
93040 Regensburg
Germany

Contents

1 Introduction

What is geometric group theory? Geometric group theory investigates the interaction between algebraic and geometric properties of groups:

- Can groups be viewed as geometric objects and how are geometric and algebraic properties of groups related?
- More generally: On which geometric objects can a given group act in a reasonable way, and how are geometric properties of these geometric objects/actions related to algebraic properties of the group?

How does geometric group theory work? Classically, group-valued invariants are associated with geometric objects, such as, e.g., the isometry group or the fundamental group. It is one of the central insights leading to geometric group theory that this process can be reversed to a certain extent:

1. We associate a geometric object with the group in question; this can be an "artificial" abstract construction or a very concrete model space (such as the Euclidean plane or the hyperbolic plane) or action from classical geometric theories.
2. We take geometric invariants and apply these to the geometric objects obtained by the first step. This allows to translate geometric terms such as geodesics, curvature, volumes, etc. into group theory.
   Usually, in this step, in order to obtain good invariants, one restricts

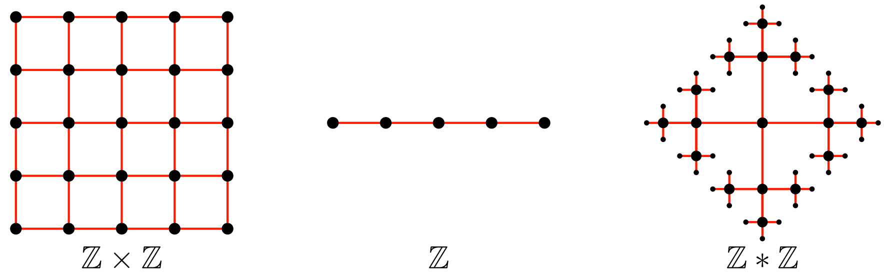

Figure 1.1.: Basic examples of Cayley graphs

attention to finitely generated groups and takes geometric invariants from large scale geometry (as they blur the difference between different finite generating sets of a given group).

3. We compare the behaviour of such geometric invariants of groups with the algebraic behaviour, and we study what can be gained by this symbiosis of geometry and algebra.

A key example of a geometric object associated with a group is the so-called Cayley graph (with respect to a chosen generating set) together with the corresponding word metric. For instance, from the point of view of large scale geometry, the Cayley graph of $\mathbb{Z}$ resembles the geometry of the real line, the Cayley graph of $\mathbb{Z} \times \mathbb{Z}$ resembles the geometry of the Euclidean plane, while the Cayley graph of the free group $\mathbb{Z} * \mathbb{Z}$ on two generators has essential features of the geometry of the hyperbolic plane (Figure 1.1; exact definitions of these concepts are introduced in later chapters).

More generally, in (large scale) geometric group theoretic terms, the universe of (finitely generated) groups roughly unfolds as depicted in Figure 1.2. The boundaries are inhabited by amenable groups and non-positively curved groups respectively – classes of groups that are (at least partially) accessible. However, studying these boundary classes is only the very beginning of understanding the universe of groups; in general, knowledge about these two classes of groups is far from enough to draw conclusions about groups at the inner regions of the universe:

“Hic abundant leones.” [15]
“A statement that holds for all finitely generated groups
has to be either trivial or wrong.” [attributed to M. Gromov]

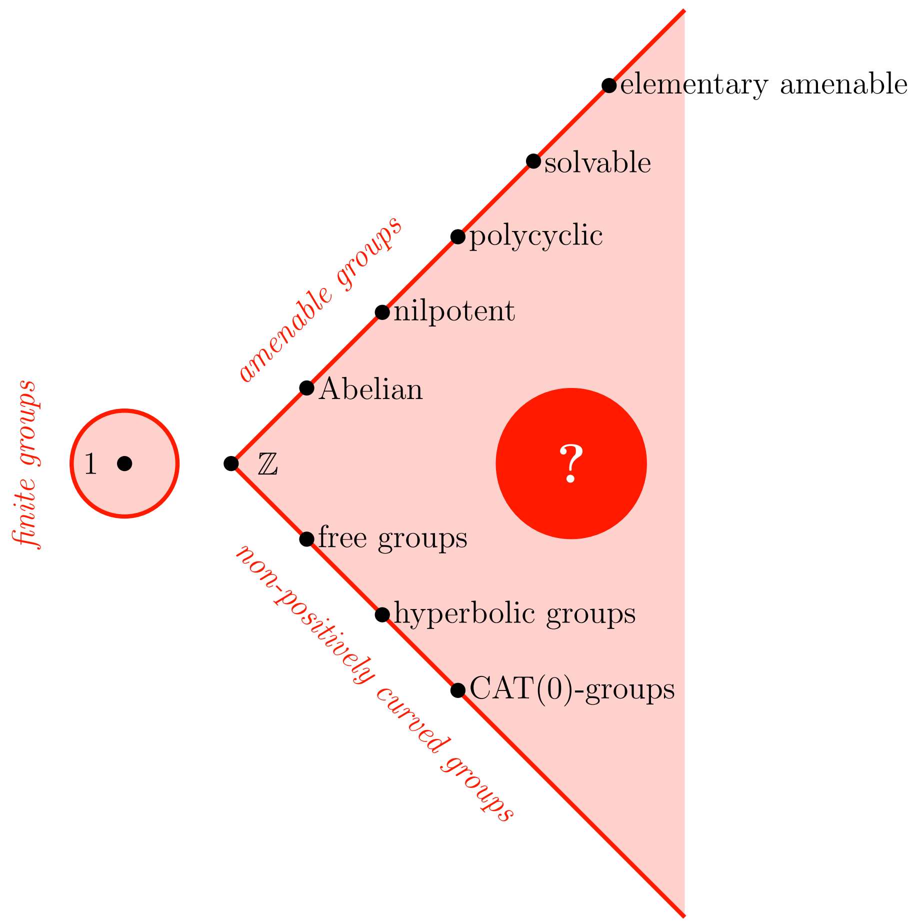

Figure 1.2.: The universe of groups (simplified version of Bridson's universe of groups [15])

Why study geometric group theory? On the one hand, geometric group theory is an interesting theory combining aspects of different fields of mathematics in a cunning way. On the other hand, geometric group theory has numerous applications to problems in classical fields such as group theory and Riemannian geometry.

For example, so-called free groups (an a priori purely algebraic notion) can be characterised geometrically via actions on trees; this leads to an elegant proof of the (purely algebraic!) fact that *subgroups of free groups are free.*

Further applications of geometric group theory to algebra and Riemannian geometry include the following:

- *Recognising that certain matrix groups are free groups;* there is a geometric criterion, the so-called *ping-pong-lemma*, that allows to deduce freeness of a group by looking at a suitable action (not necessarily on a tree).
- *Recognising that certain groups are finitely generated;* this can be done geometrically by exhibiting a good action on a suitable space.
- *Establishing decidability of the word problem for large classes of groups;* for example, Dehn used geometric ideas in his algorithm solving the word problem in certain geometric classes of groups.
- *Recognising that certain groups are virtually nilpotent;* Gromov found a characterisation of finitely generated virtually nilpotent groups in terms of geometric data, more precisely, in terms of the growth function.
- *Proving non-existence of Riemannian metrics satisfying certain curvature conditions on certain smooth manifolds;* this is achieved by translating these curvature conditions into group theory and looking at groups associated with the given smooth manifold (e.g., the fundamental group). Moreover, a similar technique also yields (non-)splitting results for certain non-positively curved spaces.
- *Rigidity results for certain classes of matrix groups and Riemannian manifolds;* here, the key is the study of an appropriate geometry at infinity of the groups involved.
- *Geometric group theory provides a layer of abstraction that helps to understand and generalise classical geometry* – in particular, in the case of negative or non-positive curvature and the corresponding geometry at infinity.

- *The Banach-Tarski paradox (a sphere can be divided into finitely many pieces that in turn can be puzzled together into two spheres congruent to the given one [this relies on the axiom of choice]);* the Banach-Tarski paradox corresponds to certain matrix groups not being "amenable", a notion related to both measure theoretic and geometric properties of groups.
- *A better understanding of many classical groups;* this includes, for instance, mapping class groups of surfaces and outer automorphisms of free groups (and their behaviour similar to certain matrix groups).

Overview of the course. As the main characters in geometric group theory are groups, we will start by reviewing some concepts and examples from group theory, and by introducing constructions that allow to generate interesting groups. Then we will introduce one of the main combinatorial objects in geometric group theory, the so-called Cayley graph, and review basic notions concerning actions of groups. A first taste of the power of geometric group theory will then be presented in the discussion of geometric characterisations of free groups. As next step, we will introduce a metric structure on groups via word metrics on Cayley graphs, and we will study the large scale geometry of groups with respect to this metric structure (in particular, the concept of quasi-isometry). After that, invariants under quasi-isometry will be introduced – this includes, in particular, curvature conditions, the geometry at infinity and growth functions. Finally, we will have a look at the class of amenable groups and their properties.

Literature. The standard resources for geometric group theory are:
- *Topics in Geometric group theory* by de la Harpe [42],
- *Metric spaces of non-positive curvature* by Bridson and Haefliger [17],
- *Trees* by Serre [87].

A short and comprehensible introduction into curvature in classical Riemannian geometry is given in the book *Riemannian manifolds. An introduction to curvature* by Lee [54].

Furthermore, I recommend to look at the overview articles by Bridson on geometric and combinatorial group theory [15, 16]. The original reference for modern large scale geometry of groups is the landmark paper *Hyperbolic groups* [40] by Gromov.

About these lecture notes

These lecture notes are the manuscript of a one-semester course on geometric group theory, taught in 2014/15 at the Universität Regensburg. They evolve(d) out of lecture notes for a corresponding course taught in 2010/11 at the Universität Regensburg. Depending on the background in geometry and topology of the participants, the course might develop in a different direction than its predecessor.

Regensburg, March 2011/October 2014 Clara Löh

2 Generating groups

As the main characters in geometric group theory are groups, we will start by reviewing some concepts and examples from group theory, and by introducing constructions that allow to generate interesting groups. In particular, we will explain how to describe groups in terms of generators and relations, and how to construct groups via semi-direct products and amalgamated products.

2.1 Review of the category of groups

2.1.1 Axiomatic description of groups

For the sake of completeness, we briefly recall the definition of a group; more information on basic properties of groups can be found in any textbook on algebra [52, 82]. The category of groups has groups as objects and group homomorphisms as morphisms.

Definition 2.1.1 (Group)**.** A *group* is a set G together with a binary operation $\cdot \colon G \times G \longrightarrow G$ satisfying the following axioms:

– *Associativity.* For all $g_1, g_2, g_3 \in G$ we have
$$g_1 \cdot (g_2 \cdot g_3) = (g_1 \cdot g_2) \cdot g_3.$$

– *Existence of a neutral element.* There exists a *neutral element* $e \in G$ for "$\cdot$", i.e.,
$$\forall_{g \in G} \;\; e \cdot g = g = g \cdot e.$$
(Notice that the neutral element is uniquely determined by this property.)

– *Existence of inverses.* For every $g \in G$ there exists an *inverse element* $g^{-1} \in G$ with respect to "$\cdot$", i.e.,
$$g \cdot g^{-1} = e = g^{-1} \cdot g.$$

A group G is *Abelian* if composition is commutative, i.e., if $g_1 \cdot g_2 = g_2 \cdot g_1$ holds for all $g_1, g_2 \in G$.

Definition 2.1.2 (Subgroup)**.** Let G be a group with respect to "$\cdot$". A subset $H \subset G$ is a *subgroup* if H is a group with respect to the restriction of "$\cdot$" to $H \times H \subset G \times G$.

Example 2.1.3 (Some (sub)groups)**.**

- *Trivial group*(s) are groups consisting only of a single element e and the composition $(e, e) \longmapsto e$. Clearly, every group contains a trivial group given by the neutral element as subgroup.
- The sets $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ are groups with respect to addition; moreover, $\mathbb{Z}$ is a subgroup of $\mathbb{Q}$, and $\mathbb{Q}$ is a subgroup of $\mathbb{R}$.
- The natural numbers $\mathbb{N}$ do not form a group with respect to addition (e.g., 1 does not have an additive inverse in $\mathbb{N}$); the rational numbers $\mathbb{Q}$ do not form a group with respect to multiplication (0 does not have a multiplicative inverse), but $\mathbb{Q} \setminus \{0\}$ is a group with respect to multiplication.

Now that we have introduced the main objects, we need morphisms to relate different objects to each other. As in other mathematical theories, morphisms should be structure preserving, and we consider two objects to be the same if they have the same structure:

Definition 2.1.4 (Group homomorphism/isomorphism)**.** Let G and H be two groups.

- A map $\varphi\colon G \longrightarrow H$ is a *group homomorphism* if φ is compatible with the composition in G and H respectively, i.e., if
$$\varphi(g_1 \cdot g_2) = \varphi(g_1) \cdot \varphi(g_2)$$
holds for all g_1, $g_2 \in G$. (Notice that every group homomorphism maps the neutral element to the neutral element).
- A group homomorphism $\varphi\colon G \longrightarrow H$ is a *group isomorphism* if there exists a group homomorphism $\psi\colon H \longrightarrow G$ such that $\varphi \circ \psi = \mathrm{id}_H$ and $\psi \circ \varphi = \mathrm{id}_G$. If there exists a group isomorphism between G and H, then G and H are *isomorphic*, and we write $G \cong H$.

Example 2.1.5 (Some group homomorphisms)**.**

- Clearly, any two trivial groups are (canonically) isomorphic. Hence, we speak about "the" trivial group.
- If H is a subgroup of a group G, then the inclusion $H \hookrightarrow G$ is a group homomorphism.
- Let $n \in \mathbb{Z}$. Then
$$\begin{aligned} \mathbb{Z} &\longrightarrow \mathbb{Z} \\ x &\longmapsto n \cdot x \end{aligned}$$

is a group homomorphism; however, addition of $n \neq 0$ is *not* a group homomorphism (e.g., the neutral element is not mapped to the neutral element).
- The exponential map $\exp\colon \mathbb{R} \longrightarrow \mathbb{R}_{>0}$ is a group homomorphism between the additive group $\mathbb{R}$ and the multiplicative group $\mathbb{R}_{>0}$; the exponential map is even an isomorphism (the inverse homomorphism is given by the logarithm).

Definition 2.1.6 (Kernel/image of homomorphisms)**.** Let $\varphi\colon G \longrightarrow H$ be a group homomorphism. Then the subgroup

$$\ker \varphi := \{g \in G \mid \varphi(g) = e\}$$

of G is the *kernel* of φ, and the subgroup

$$\operatorname{im} \varphi := \{\varphi(g) \mid g \in G\}$$

of H is the *image* of φ.

Remark 2.1.7 (Isomorphisms via kernel/image)**.** It is a simple exercise in algebra to verify the following:

1. A group homomorphism is injective if and only if its kernel is the trivial subgroup.
2. A group homomorphism is an isomorphism if and only if it is bijective.
3. In particular: A group homomorphism $\varphi\colon G \longrightarrow H$ is an isomorphism if and only if $\ker \varphi$ is the trivial subgroup and $\operatorname{im} \varphi = H$.

2.1.2 Concrete groups – automorphism groups

The concept, and hence the axiomatisation, of groups developed originally out of the observation that certain collections of "invertible" structure preserving transformations of geometric or algebraic objects fit into the same abstract framework; moreover, it turned out that many interesting properties of the underlying objects are encoded in the group structure of the corresponding automorphism group.

Example 2.1.8 (Symmetric groups)**.** Let X be a set. Then the set S_X of all bijections of type $X \longrightarrow X$ is a group with respect to composition of maps, the so-called *symmetric group over X*. If $n \in \mathbb{N}$, then we abbreviate $S_n := S_{\{1,\dots,n\}}$. In general, the group S_X is *not* Abelian.

This example is generic in the following sense:

Proposition 2.1.9 (Cayley's theorem)**.** *Every group is isomorphic to a subgroup of some symmetric group.*

Proof. Let G be a group. For $g \in G$ we define the map

$$\begin{aligned} f_g \colon G &\longrightarrow G \\ x &\longmapsto g \cdot x. \end{aligned}$$

For all $g, h \in G$ we have $f_g \circ f_h = f_{g \cdot h}$. Therefore, looking at $f_{g^{-1}}$ shows that $f_g \colon G \longrightarrow G$ is a bijection for all $g \in G$. Moreover, it follows that

$$\begin{aligned} f \colon G &\longrightarrow S_G \\ g &\longmapsto f_g \end{aligned}$$

is a group homomorphism, which is easily shown to be injective. So, f induces an isomorphism $G \cong \operatorname{im} f \subset S_G$, as desired. □

Example 2.1.10 (Automorphism groups)**.** Let G be a group. Then the set $\operatorname{Aut}(G)$ of group isomorphisms of type $G \longrightarrow G$ is a group with respect to composition of maps, the *automorphism group of* G. Clearly, $\operatorname{Aut}(G)$ is a subgroup of S_G.

Example 2.1.11 (Isometry groups/Symmetry groups)**.** Let X be a metric space. The set $\operatorname{Isom}(X)$ of all isometries of type $X \longrightarrow X$ forms a group with respect to composition (a subgroup of the symmetric group S_X). For example, in this way the dihedral groups naturally occur as symmetry groups of regular polygons.

Example 2.1.12 (Matrix groups)**.** Let k be a commutative ring with unit, and let V be a k-module. Then the set $\operatorname{Aut}(V)$ of all k-linear isomorphisms $V \longrightarrow V$ forms a group with respect to composition. In particular, the set $\operatorname{GL}(n, k)$ of invertible $n \times n$-matrices over k is a group (with respect to matrix multiplication) for every $n \in \mathbb{N}$. Similarly, also $\operatorname{SL}(n, k)$ is a group.

Example 2.1.13 (Galois groups)**.** Let $K \subset L$ be a Galois extension of fields. Then the set

$$\operatorname{Gal}(L/K) := \bigl\{\sigma \in \operatorname{Aut}(L) \bigm| \sigma|_K = \mathrm{id}_K\bigr\}$$

of field automorphisms of L fixing K is a group with respect to composition, the so-called *Galois group* of the extension L/K.

Example 2.1.14 (Deck transformation groups)**.** Let $\pi\colon X \longrightarrow Y$ be a covering map of topological spaces. Then the set

$$\bigl\{f \in \operatorname{map}(X, X) \bigm| f \text{ is a homeomorphism with } \pi \circ f = \pi\bigr\}$$

of *Deck transformations* forms a group with respect to composition.

In modern language, these examples are all instances of the following general principle: If X is an object in a category C, then the set $\operatorname{Aut}_C(X)$ of C-isomorphisms of type $X \longrightarrow X$ is a group with respect to composition in C.

Exercise 2.1.1 (Groups and categories)**.**

1. Look up the definition of *category*, *isomorphism*, and *automorphism in a category*.
2. Let C be a category and let X be an object in C. Show that the set(!) of automorphisms of X in C forms a group with respect to composition of morphisms in C.
3. What are the (natural) categories in the examples above?
4. Show that, conversely, every group arises as the automorphism group of an object in some category.

Taking automorphism groups of geometric/algebraic objects is only one way to associate meaningful groups to interesting objects. Over time, many group-valued invariants have been developed in all fields of mathematics. For example:

– fundamental groups (in topology, algebraic geometry or operator algebra theory, ...)
– homology groups (in topology, algebra, algebraic geometry, operator algebra theory, ...)
– ...

2.1.3 Normal subgroups and quotients

Sometimes it is convenient to ignore a certain subobject of a given object and to focus on the remaining properties. Formally, this is done by taking quotients. In contrast to the theory of vector spaces, where the quotient of any vector space by any subspace again naturally forms a vector space, we have to be a little bit more careful in the world of groups. Only special subgroups lead to quotient *groups*:

Definition 2.1.15 (Normal subgroup)**.** Let G be a group. A subgroup N of G is *normal* if it is conjugation invariant, i.e., if

$$g \cdot n \cdot g^{-1} \in N$$

holds for all $n \in N$ and all $g \in G$. If N is a normal subgroup of G, then we write $N \lhd G$.

Example 2.1.16 (Some (non-)normal subgroups)**.**

- All subgroups of Abelian groups are normal.
- Let $\tau \in S_3$ be the bijection given by swapping 1 and 2 (i.e., $\tau = (1\ 2)$). Then $\{\mathrm{id}, \tau\}$ is a subgroup of S_3, but it is not a normal subgroup. On the other hand, the subgroup $\{\mathrm{id}, \sigma, \sigma^2\} \subset S_3$ generated by the cycle $\sigma := (1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1)$ is a normal subgroup of S_3.
- Kernels of group homomorphisms are normal in the domain group; conversely, every normal subgroup also is the kernel of a certain group homomorphism (namely of the canonical projection to the quotient (Proposition 2.1.17)).

Proposition 2.1.17 (Quotient group)**.** *Let G be a group, and let N be a subgroup.*

1. *Let $G/N := \{g \cdot N \mid g \in G\}$, where we use the coset notation $g \cdot N := \{g \cdot n \mid n \in N\}$ for $g \in G$. Then the map*

$$\begin{aligned} G/N \times G/N &\longrightarrow G/N \\ (g_1 \cdot N, g_2 \cdot N) &\longmapsto (g_1 \cdot g_2) \cdot N \end{aligned}$$

is well-defined if and only if N is normal in G. If N is normal in G, then G/N is a group with respect to this composition map, the so-called quotient group *of G by N.*

2. *Let N be normal in G. Then the* canonical projection

$$\begin{aligned} \pi \colon G &\longrightarrow G/N \\ g &\longmapsto g \cdot N \end{aligned}$$

is a group homomorphism, and the quotient group G/N together with π has the following universal property: For any group H and any group homomorphism $\varphi\colon G \longrightarrow H$ with $N \subset \ker\varphi$ there is exactly one group homomorphism $\overline{\varphi}\colon G/N \longrightarrow H$ satisfying $\overline{\varphi} \circ \pi = \varphi$:

$$\begin{tikzcd} G \arrow[r, "\varphi"] \arrow[d, "\pi"'] & H \\ G/N \arrow[ur, dashed, "\overline{\varphi}"'] & \end{tikzcd}$$

Proof. Ad 1. Suppose that N is normal in G. In this case, the composition map is well-defined (in the sense that the definition does not depend on the choice of the representatives of cosets) because: Let $g_1, g_2, \overline{g}_1, \overline{g}_2 \in G$ with

$$g_1 \cdot N = \overline{g}_1 \cdot N, \qquad \text{and} \qquad g_2 \cdot N = \overline{g}_2 \cdot N.$$

In particular, there are $n_1, n_2 \in N$ with $\overline{g}_1 = g_1 \cdot n_1$ and $\overline{g}_2 = g_2 \cdot n_2$. Thus we obtain

$$\begin{aligned} (\overline{g}_1 \cdot \overline{g}_2) \cdot N &= (g_1 \cdot n_1 \cdot g_2 \cdot n_2) \cdot N \\ &= (g_1 \cdot g_2 \cdot (g_2^{-1} \cdot n_1 \cdot g_2) \cdot n_2) \cdot N \\ &= (g_1 \cdot g_2) \cdot N; \end{aligned}$$

in the last step we used that N is normal, which implies that $g_2^{-1} \cdot n_1 \cdot g_2 \in N$ and hence $g_2^{-1} \cdot n_1 \cdot g_2 \cdot n_2 \in N$. Therefore, the composition on G/N is well-defined.

That G/N indeed is a group with respect to this composition follows easily from the fact that the group axioms are satisfied in G.

Conversely, suppose that the composition on G/N is well-defined. Then the subgroup N is normal in G because: Let $n \in N$ and let $g \in G$. Then $g \cdot N = (g \cdot n) \cdot N$, and so (by well-definedness)

$$\begin{aligned} N &= (g \cdot g^{-1}) \cdot N \\ &= (g \cdot N) \cdot (g^{-1} \cdot N) \\ &= ((g \cdot n) \cdot N) \cdot (g^{-1} \cdot N) \\ &= (g \cdot n \cdot g^{-1}) \cdot N; \end{aligned}$$

in particular, $g \cdot n \cdot g^{-1} \in N$. Therefore, N is normal in G.

Ad 2. Let H be a group and let $\varphi\colon G \longrightarrow H$ be a group homomorphism with $N \subset \ker \varphi$. It is easy to see that

$$\begin{aligned} \overline{\varphi}\colon G/N &\longrightarrow H \\ g \cdot N &\longmapsto \varphi(g) \end{aligned}$$

is a well-defined group homomorphism, that it satisfies $\overline{\varphi} \circ \pi = \varphi$, and that $\overline{\varphi}$ is the only group homomorphism with this property. □

Example 2.1.18 (Quotient groups)**.**

- Let $n \in \mathbb{Z}$. Then composition in the quotient group $\mathbb{Z}/n\mathbb{Z}$ is nothing but addition modulo n.
- The quotient group $\mathbb{R}/\mathbb{Z}$ is isomorphic to the (multiplicative) circle group $\{z \mid z \in \mathbb{C},\ |z| = 1\} \subset \mathbb{C} \setminus \{0\}$.
- The quotient of S_3 by the subgroup $\{\mathrm{id}, \sigma, \sigma^2\}$ generated by the cycle $\sigma := (1\ 2\ 3)$ is isomorphic to $\mathbb{Z}/2\mathbb{Z}$.

Example 2.1.19 (Outer automorphism groups)**.** Let G be a group. An automorphism $\varphi\colon G \longrightarrow G$ is an *inner automorphism of* G if φ is given by conjugation by an element of G, i.e., if there is an element $g \in G$ such that

$$\forall_{h \in G}\ \ \varphi(h) = g \cdot h \cdot g^{-1}.$$

The subset of $\mathrm{Aut}(G)$ of all inner automorphisms of G is denoted by $\mathrm{Inn}(G)$. Then $\mathrm{Inn}(G)$ is a normal subgroup of $\mathrm{Aut}(G)$ (exercise), and the quotient

$$\mathrm{Out}(G) := \mathrm{Aut}(G)/\operatorname{Inn}(G)$$

is the *outer automorphism group of* G. For example, the outer automorphism groups of finitely generated free groups form an interesting class of groups that has various connections to lattices in Lie groups and mapping class groups [94]. Curiously, one has for all $n \in \mathbb{N}$ that [65]

$$\mathrm{Out}(S_n) \cong \begin{cases} \{e\} & \text{if } n \neq 6 \\ \mathbb{Z}/2 & \text{if } n = 6. \end{cases}$$

2.2 Groups via generators and relations

How can we specify a group? One way is to construct a group as the automorphism group of some object or as a subgroup or quotient thereof. However, when interested in finding groups with certain algebraic features, it might sometimes be difficult to find a corresponding geometric object.

In this section, we will see that there is another – abstract – way to construct groups, namely by generators and relations: We will prove that for every list of elements ("generators") and group theoretic equations ("relations") linking these elements there always exists a group in which these relations hold as non-trivially as possible. (However, in general, it is not possible to decide whether the given wish-list of generators and relations can be realised by a *non-trivial* group.) Technically, generators and relations are formalised by the use of free groups and suitable quotient groups thereof.

2.2.1 Generating sets of groups

We start by reviewing the concept of a generating set of a group; in geometric group theory, one usually is only interested in finitely generated groups (for reasons that will become clear in Chapter 5).

Definition 2.2.1 (Generating set)**.**

- Let G be a group and let $S \subset G$ be a subset. The *subgroup generated by S in G* is the smallest subgroup (with respect to inclusion) of G that contains S; the subgroup generated by S in G is denoted by $\langle S\rangle_G$.
  The set S *generates* G if $\langle S\rangle_G = G$.
- A group is *finitely generated* if it contains a finite subset that generates the group in question.

Remark 2.2.2 (Explicit description of generated subgroups)**.** Let G be a group and let $S \subset G$. Then the subgroup generated by S in G always exists and can be described as follows:

$$\begin{aligned}\langle S\rangle_G &= \bigcap\{H \mid H \subset G \text{ is a subgroup with } S \subset H\}\\ &= \bigl\{s_1^{\varepsilon_1} \cdot \dots \cdot s_n^{\varepsilon_n} \bigm| n \in \mathbb{N},\ s_1, \dots, s_n \in S,\ \varepsilon_1, \dots, \varepsilon_n \in \{-1, +1\}\bigr\}.\end{aligned}$$

Example 2.2.3 (Generating sets)**.**

– If G is a group, then G is a generating set of G.
– The trivial group is generated by the empty set.
– The set $\{1\}$ generates the additive group $\mathbb{Z}$; moreover, also, e.g., $\{2, 3\}$ is a generating set for $\mathbb{Z}$. But $\{2\}$ and $\{3\}$ are no generating sets of $\mathbb{Z}$.
– Let X be a set. Then the symmetric group S_X is finitely generated if and only if X is finite (exercise).

2.2.2 Free groups

Every vector space admits special generating sets: namely those generating sets that are as free as possible (meaning having as few linear algebraic relations between them as possible), i.e., the linearly independent ones. Also, in the setting of group theory, we can formulate what it means to be a free generating set – however, as we will see, most groups do *not* admit free generating sets. This is one of the reasons why group theory is much more complicated than linear algebra.

Definition 2.2.4 (Free groups, universal property)**.** Let S be a set. A group F is *freely generated by S* if F has the following universal property: For any group G and any map $\varphi\colon S \longrightarrow G$ there is a unique group homomorphism $\overline{\varphi}\colon F \longrightarrow G$ extending φ:

$$\begin{tikzcd} S \arrow[r, "\varphi"] \arrow[d, hook] & G \\ F \arrow[ur, dashed, "\overline{\varphi}"'] & \end{tikzcd}$$

A group is *free* if it contains a free generating set.

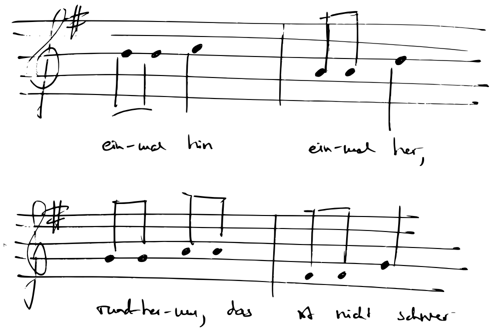

Figure 2.1.: Uniqueness of universal objects; *Brüderchen, komm tanz mit mir*

Example 2.2.5 (Free groups)**.**

- The additive group $\mathbb{Z}$ is freely generated by $\{1\}$. The additive group $\mathbb{Z}$ is *not* freely generated by $\{2, 3\}$ or $\{2\}$ or $\{3\}$; in particular, not every generating set of a group contains a free generating set.
- The trivial group is freely generated by the empty set.
- Not every group is free; for example, the additive groups $\mathbb{Z}/2\mathbb{Z}$ and $\mathbb{Z}^2$ are *not* free (exercise).

The term "universal property" obliges us to prove that objects having this universal property are unique in an appropriate sense; moreover, we will see below (Theorem 2.2.7) that for every set there indeed exists a group freely generated by the given set.

Proposition 2.2.6 (Free groups, uniqueness)**.** *Let S be a set. Then, up to canonical isomorphism, there is at most one group freely generated by S.*

The proof consists of the standard universal-property-yoga (Figure 2.1): Namely, we consider two objects that have the universal property in question. We then proceed as follows:

1. We use the existence part of the universal property to obtain interesting morphisms in both directions.
2. We use the uniqueness part of the universal property to conclude that both compositions of these morphisms have to be the identity

(and hence that both morphisms are isomorphisms).

Proof. Let F and F' be two groups freely generated by S. We denote the inclusion of S into F and F' by φ and φ' respectively.

1. Because F is freely generated by S, the existence part of the universal property of free generation guarantees the existence of a group homomorphism $\overline{\varphi}'\colon F \longrightarrow F'$ such that $\overline{\varphi}' \circ \varphi = \varphi'$. Analogously, there is a group homomorphism $\overline{\varphi}\colon F' \longrightarrow F$ satisfying $\overline{\varphi} \circ \varphi' = \varphi$:

$$\begin{tikzcd} S \arrow[r, "\varphi'"] \arrow[d, hook, "\varphi"'] & F' \\ F \arrow[ur, dashed, "\overline{\varphi}'"'] & \end{tikzcd} \qquad \begin{tikzcd} S \arrow[r, "\varphi"] \arrow[d, hook, "\varphi'"'] & F \\ F' \arrow[ur, dashed, "\overline{\varphi}"'] & \end{tikzcd}$$

2. We now show that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$ and $\overline{\varphi}' \circ \overline{\varphi} = \mathrm{id}_{F'}$ and hence that φ and φ' are isomorphisms: The composition $\overline{\varphi} \circ \overline{\varphi}'\colon F \longrightarrow F$ is a group homomorphism making the diagram

$$\begin{tikzcd} S \arrow[r, "\varphi"] \arrow[d, hook, "\varphi"'] & F \\ F \arrow[ur, dashed, "\overline{\varphi}\circ\overline{\varphi}'"'] & \end{tikzcd}$$

commutative. Moreover, also id_F is a group homomorphism fitting into this diagram. Because F is freely generated by S, the uniqueness part of the universal property thus tells us that these two homomorphisms have to coincide.

These isomorphisms are canonical in the following sense: They induce the identity map on S, and they are (by the uniqueness part of the universal property) the only isomorphisms between F and F' extending the identity on S. $\square$

Theorem 2.2.7 (Free groups, existence)**.** *Let S be a set. Then there exists a group freely generated by S. (By the previous proposition, this group is unique up to isomorphism.)*

Proof. The idea is to construct a group consisting of "words" made up of elements of S and their "inverses" using only the obvious cancellation rules for elements of S and their "inverses." More precisely, we consider the alphabet

$$A := S \cup \overline{S},$$

where $\overline{S} := \{\overline{s} \mid s \in S\}$; i.e., $\overline{S}$ contains an element for every element in S, and $\overline{s}$ will play the role of the inverse of s in the group that we will construct.

– As first step, we define A^* to be the set of all (finite) sequences ("words") over the alphabet A; this includes in particular the empty word ε. On A^* we define a composition $A^* \times A^* \longrightarrow A^*$ by concatenation of words. This composition is associative and ε is the neutral element.
– As second step we define

$$F(S) := A^*/\sim,$$

where $\sim$ is the equivalence relation generated by

$$\begin{aligned} &\forall_{x,y \in A^*} \ \forall_{s \in S} \ \ xs\overline{s}y \sim xy, \\ &\forall_{x,y \in A^*} \ \forall_{s \in S} \ \ x\overline{s}sy \sim xy; \end{aligned}$$

i.e., $\sim$ is the smallest equivalence relation in $A^* \times A^*$ (with respect to inclusion) satisfying the above conditions. We denote the equivalence classes with respect to the equivalence relation $\sim$ by $[\,\cdot\,]$.
It is not difficult to check that concatenation induces a well-defined composition $\cdot : F(S) \times F(S) \longrightarrow F(S)$ via

$$[x] \cdot [y] = [xy]$$

for all $x, y \in A^*$.

The set $F(S)$ together with the composition "$\cdot$" given by concatenation is a group: Clearly, $[\varepsilon]$ is a neutral element for this composition, and associativity of the composition is inherited from the associativity of the composition in A^*. For the existence of inverses we proceed as follows: Inductively (over the length of sequences), we define a map $I \colon A^* \longrightarrow A^*$ by $I(\varepsilon) := \varepsilon$ and

$$\begin{aligned} I(sx) &:= I(x)\overline{s}, \\ I(\overline{s}x) &:= I(x)s \end{aligned}$$

for all $x \in A^*$ and all $s \in S$. An induction shows that $I(I(x)) = x$ and

$$[I(x)] \cdot [x] = [I(x)x] = [\varepsilon]$$

for all $x \in A^*$ (in the last step we use the definition of $\sim$). This shows that inverses exist in $F(S)$.

The group $F(S)$ is freely generated by S: Let $i\colon S \longrightarrow F(S)$ be the map given by sending a letter in $S \subset A^*$ to its equivalence class in $F(S)$; by construction, $F(S)$ is generated by the subset $i(S) \subset F(S)$. As we do not know yet that i is injective, we take a little detour and first show that $F(S)$ has the following property, similar to the universal property of groups freely generated by S: For every group G and every map $\varphi\colon S \longrightarrow G$ there is a unique group homomorphism $\overline{\varphi}\colon F(S) \longrightarrow G$ such that $\overline{\varphi} \circ i = \varphi$. Given φ, we construct a map

$$\varphi^*\colon A^* \longrightarrow G$$

inductively by

$$\begin{aligned} \varepsilon &\longmapsto e, \\ sx &\longmapsto \varphi(s) \cdot \varphi^*(x), \\ \overline{s}x &\longmapsto \bigl(\varphi(s)\bigr)^{-1} \cdot \varphi^*(x) \end{aligned}$$

for all $s \in S$ and all $x \in A^*$. It is easy to see that this definition of φ^* is compatible with the equivalence relation $\sim$ on A^* (because it is compatible with the given generating set of $\sim$) and that $\varphi^*(xy) = \varphi^*(x) \cdot \varphi^*(y)$ for all x, $y \in A^*$; thus, φ^* induces a well-defined map

$$\begin{aligned} \overline{\varphi}\colon F(S) &\longrightarrow G \\ [x] &\longmapsto [\varphi^*(x)], \end{aligned}$$

which is a group homomorphism. By construction $\overline{\varphi} \circ i = \varphi$. Moreover, because $i(S)$ generates $F(S)$ there cannot be another such group homomorphism.

In order to show that $F(S)$ is freely generated by S, it remains to prove that i is injective (and then we identify S with its image under i in $F(S)$): Let s_1, $s_2 \in S$. We consider the map $\varphi\colon S \longrightarrow \mathbb{Z}$ given by $\varphi(s_1) := 1$

and $\varphi(s_2) := -1$. Then the corresponding homomorphism $\overline{\varphi}\colon F(S) \longrightarrow G$ satisfies

$$\overline{\varphi}\big(i(s_1)\big) = \varphi(s_1) = 1 \neq -1 = \varphi(s_2) = \overline{\varphi}\big(i(s_2)\big);$$

in particular, $i(s_1) \neq i(s_2)$. Hence, i is injective. □

Depending on the problem at hand, the declarative description of free groups via the universal property or the constructive description as in the previous proof might be more appropriate than the other. A refined constructive description of free groups in terms of reduced words will be given in the context of group actions on trees (Section 3.3.1).

We conclude by collecting some properties of free generating sets in free groups: First of all, free groups indeed are generated (in the sense of Definition 2.2.1) by any free generating set (Corollary 2.2.8); second, free generating sets are generating sets of minimal size (Proposition 2.2.9); moreover, finitely generated groups can be characterised as the quotients of finitely generated free groups (Corollary 2.2.12).

Corollary 2.2.8. *Let F be a free group, and let S be a free generating set of F. Then S generates F.*

Proof. By construction, the statement holds for the free group $F(S)$ generated by S constructed in the proof of Theorem 2.2.7. In view of the uniqueness result Proposition 2.2.6, we find an isomorphism $F(S) \cong F$ that is the identity on S. Hence, it follows that also the given free group F is generated by S. □

Proposition 2.2.9 (Rank of free groups)**.** *Let F be a free group.*

1. *Let $S \subset F$ be a free generating set of F and let S' be a generating set of F. Then $|S'| \geq |S|$.*
2. *In particular: all free generating sets of F have the same cardinality, called the* rank *of F.*

Proof. The first part can be derived from the universal property of free groups (applied to homomorphisms to $\mathbb{Z}/2\mathbb{Z}$) together with a counting argument (exercise). The second part is a consequence of the first part. □

Definition 2.2.10 (Free group F_n)**.** Let $n \in \mathbb{N}$ and let $S = \{x_1, \dots, x_n\}$, where $x_1, \dots, x_n$ are n distinct elements. Then we write F_n for "the" group freely generated by S, and call F_n the *free group of rank n*.

Caveat 2.2.11. While subspaces of vector spaces cannot have bigger dimension than the ambient space, free groups of rank 2 contain subgroups that are isomorphic to free groups of higher rank, even free subgroups of (countably) infinite rank. Subgroups of this type can easily be constructed via covering theory [64, Chapter VI.8] or via actions on trees (Chapter 4.3).

Corollary 2.2.12. *A group is finitely generated if and only if it is the quotient of a finitely generated free group, i.e., a group G is finitely generated if and only if there exists a finitely generated free group F and a surjective group homomorphism $F \longrightarrow G$.*

Proof. Quotients of finitely generated groups are finitely generated (e.g., the image of a finite generating set is a finite generating set of the quotient).

Conversely, let G be a finitely generated group, say generated by the finite set $S \subset G$. Furthermore, let F be the free group generated by S; by Corollary 2.2.8, the group F is finitely generated. Using the universal property of F we find a group homomorphism $\pi\colon F \longrightarrow G$ that is the identity on S. Because S generates G and because S lies in the image of π, it follows that $\operatorname{im}\pi = G$. □

2.2.3 Generators and relations

Free groups enable us to generate generic groups over a given set; in order to force generators to satisfy a given list of group theoretic equations, we divide out a suitable normal subgroup.

Definition 2.2.13 (Normal generation)**.** Let G be a group and let $S \subset G$ be a subset. The *normal subgroup of G generated by S* is the smallest (with respect to inclusion) normal subgroup of G containing S; it is denoted by $\langle S \rangle^{\triangleleft}_G$.

Remark 2.2.14 (Explicit description of generated normal subgroups)**.** Let G be a group and let $S \subset G$. Then the normal subgroup generated by S in G always exists and can be described as follows:

$$\begin{aligned}\langle S \rangle^{\triangleleft}_G &= \bigcap \{H \mid H \subset G \text{ is a normal subgroup with } S \subset H\} \\ &= \bigl\{ g_1 \cdot s_1^{\varepsilon_1} \cdot g_1^{-1} \cdot \dots \cdot g_n \cdot s_n^{\varepsilon_n} \cdot g_n^{-1} \\ &\qquad \bigm| n \in \mathbb{N}, s_1, \dots, s_n \in S,\ \varepsilon_1, \dots, \varepsilon_n \in \{-1, +1\},\ g_1, \dots, g_n \in G \bigr\}.\end{aligned}$$

Example 2.2.15.

- As all subgroups of Abelian groups are normal, we have $\langle S\rangle^{\triangleleft}_G = \langle S\rangle_G$ for all Abelian groups G and all subsets $S \subset G$.
- We consider the symmetric group S_3 and the permutation $\tau \in S_3$ given by swapping 1 and 2; then $\langle \tau\rangle_{S_3} = \{\mathrm{id}_{\{1,2,3\}}, \tau\}$ and $\langle \tau\rangle^{\triangleleft}_{S_3} = S_3$.

Caveat 2.2.16. If G is a group, and $N \triangleleft G$, then, in general, it is rather difficult to determine what the minimal number of elements of a subset $S \subset G$ is that satisfies $\langle S\rangle^{\triangleleft}_G = N$.

In the following, we use the notation A^* for the set of (possibly empty) words in A; moreover, we abuse notation and denote elements of the free group $F(S)$ over a set S by words in $(S \cup S^{-1})^*$ (even though, strictly speaking, elements of $F(S)$ are equivalence classes of words in $(S \cup S^{-1})^*$).

Definition 2.2.17 (Generators and relations)**.** Let S be a set, and let $R \subset (S \cup S^{-1})^*$ be a subset; let $F(S)$ be the free group generated by S. Then the group

$$\langle S \mid R\rangle := F(S)/\langle R\rangle^{\triangleleft}_{F(S)}$$

is said to be *generated by S with the relations R*; if G is a group with $G \cong \langle S \mid R\rangle$, then $\langle S \mid R\rangle$ is a *presentation of G*.

Relations of the form "$w \cdot w'^{-1}$" are also sometimes denoted as "$w = w'$", because in the generated group, the words w and w' represent the same group element.

The following proposition is a formal way of saying that $\langle S \mid R\rangle$ is a group in which the relations R hold as non-trivially as possible:

Proposition 2.2.18 (Universal property of groups given by generators and relations)**.** *Let S be a set and let $R \subset (S \cup S^{-1})^*$. The group $\langle S \mid R\rangle$ generated by S with relations R together with the canonical map $\pi \colon S \longrightarrow F(S)/\langle R\rangle^{\triangleleft}_{F(S)} = \langle S \mid R\rangle$ has the following universal property: For any group G and any map $\varphi \colon S \longrightarrow G$ with the property that*

$$\varphi^*(r) = e \quad \text{in } G$$

holds for all words $r \in R$ there exists precisely one group homomorphism $\overline{\varphi} \colon \langle S \mid R\rangle \longrightarrow G$ such that $\overline{\varphi} \circ \pi = \varphi$; here, $\varphi^* \colon (S \cup S^{-1})^* \longrightarrow G$ is the canonical extension of φ to words over $S \cup S^{-1}$ (as described in the proof of Theorem 2.2.7). Moreover, $\langle S \mid R\rangle$ (together with π) is determined uniquely (up to canonical isomorphism) by this universal property.

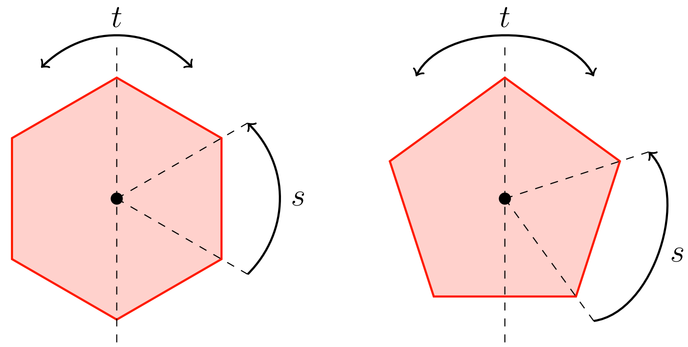

Figure 2.2.: Generators of the dihedral groups D_6 and D_5

Proof. This is a combination of the universal property of free groups (Definition 2.2.4) and of the universal property of quotient groups (Proposition 2.1.17) (exercise). □

Example 2.2.19 (Presentations of groups)**.**

- For all $n \in \mathbb{N}$, we have $\langle x \,|\, x^n \rangle \cong \mathbb{Z}/n$. This can be seen via the universal property or via the explicit construction of $\langle x \,|\, x^n \rangle$.
- We have $\langle x, y \,|\, xyx^{-1}y^{-1} \rangle \cong \mathbb{Z}^2$, as can be seen from the universal property (exercise).

Example 2.2.20 (Dihedral groups)**.** Let $n \in \mathbb{N}_{\geq 3}$ and let $X_n \subset \mathbb{R}^2$ be a regular n-gon (with the metric induced from the Euclidean metric on $\mathbb{R}^2$). Then the isometry group of X_n is a *dihedral group*:

$$\operatorname{Isom}(X_n) \cong \langle s, t \,|\, s^n, t^2, tst^{-1} = s^{-1} \rangle =: D_n;$$

geometrically, s corresponds to a rotation about $2\pi/n$ around the centre of the regular n-gon X_n, and t corresponds to a reflection along a diameter passing through one of the vertices (Figure 2.2). Moreover, one can show that the group $\operatorname{Isom}(X_n)$ contains exactly $2n$ elements, namely, $\mathrm{id}, s, \dots, s^{n-1}, t, t \cdot s, \dots, t \cdot s^{n-1}$. These facts can be established by using the fact that isometries of X_n map (neighbouring) vertices to (neighbouring) vertices and geometric calculations.

Example 2.2.21 (Thompson's group F)**.** *Thompson's group* F is defined as

$$F := \bigl\langle x_0, x_1, \dots \;\bigm|\; \{x_k^{-1} x_n x_k = x_{n+1} \mid k, n \in \mathbb{N}, k < n\} \bigr\rangle.$$

Actually, F admits a presentation by finitely many generators and relations, namely

$$F \cong \big\langle a, b \bigm| [ab^{-1}, a^{-1}ba], [ab^{-1}, a^{-2}ba^2] \big\rangle$$

(exercise). Here, we use the commutator notation "$[x, y] := xyx^{-1}y^{-1}$" both for group elements and for words. Geometrically, the group F can be interpreted in terms of certain PL-homeomorphisms of $[0, 1]$ and in terms of actions on certain binary rooted trees [20].

Thompson's group F has many peculiar properties. For example, the commutator subgroup $[F, F]$, i.e., the subgroup of F generated by the set $\{[g, h] \mid g, h \in F\}$ is an example of an infinite simple group [20]. Moreover, it can be shown that F does not contain subgroups isomorphic to F_2 [20]. However, the question whether F belongs to the class of so-called amenable groups (Chapter 9) is a long-standing open problem in geometric group theory with an interesting history [85]: "False proof of amenability and non-amenability of the R. Thompson group appear about once a year. The interesting thing is that about half of the wrong papers claim amenability and about half claim non-amenability."

Example 2.2.22 (Baumslag-Solitar groups)**.** For $m, n \in \mathbb{N}_{>0}$ the *Baumslag-Solitar group* $\mathrm{BS}(m, n)$ is defined via the presentation

$$\mathrm{BS}(m, n) := \langle a, b \,|\, ba^m b^{-1} = a^n \rangle.$$

For instance, $\mathrm{BS}(1, 1) \cong \mathbb{Z}^2$ (exercise). But the family of Baumslag-Solitar groups also contains many intriguing examples of groups. For example, $\mathrm{BS}(2, 3)$ is a group given by only two generators and a single relation that is *non-Hopfian*, i.e., there exists a surjective group homomorphism $\mathrm{BS}(2, 3) \longrightarrow \mathrm{BS}(2, 3)$ that is *not* an isomorphism [9], namely the homomorphism given by

$$\begin{aligned} \mathrm{BS}(2, 3) &\longmapsto \mathrm{BS}(2, 3) \\ a &\longmapsto a^2 \\ b &\longmapsto b. \end{aligned}$$

However, proving that this homomorphism is *not* injective requires more advanced techniques.

Example 2.2.23 (complicated trivial group)**.** The group

$$G := \langle x, y \mid xyx^{-1} = y^2, yxy^{-1} = x^2 \rangle$$

is trivial because: Let $\overline{x} \in G$ and $\overline{y} \in G$ denote the images of x and y respectively under the canonical projection

$$F(\{x,y\}) \longrightarrow F(\{x,y\}) / \langle \{xyx^{-1}y^{-2}, yxy^{-1}x^{-2}\} \rangle^{\triangleleft}_{F(S)} = G.$$

By definition, in G we obtain

$$\overline{x} = \overline{x} \cdot \overline{y} \cdot \overline{x}^{-1} \cdot \overline{x} \cdot \overline{y}^{-1} = \overline{y}^2 \cdot \overline{x} \cdot \overline{y}^{-1} = \overline{y} \cdot \overline{y} \cdot \overline{x} \cdot \overline{y}^{-1} = \overline{y} \cdot \overline{x}^2,$$

and hence $\overline{x} = \overline{y}^{-1}$. Therefore,

$$\overline{y}^{-2} = \overline{x}^2 = \overline{y} \cdot \overline{x} \cdot \overline{y}^{-1} = \overline{y} \cdot \overline{y}^{-1} \cdot \overline{y}^{-1} = \overline{y}^{-1},$$

and so $\overline{x} = \overline{y}^{-1} = e$. Because $\overline{x}$ and $\overline{y}$ generate G, we conclude that G is trivial.

Caveat 2.2.24 (Word problem)**.** The problem to determine whether a group given by (finitely many) generators and (finitely many) relations is the trivial group is undecidable (in the sense of computability theory); i.e., there is no algorithmic procedure that, given generators and relations, can decide whether the corresponding group is trivial or not [82, Chapter 12].

More generally, the *word problem*, i.e., the problem of deciding for given generators and relations whether a given word in these generators represents the trivial element in the corresponding group, is undecidable. In contrast, we will see in Section **??** that for certain geometric classes of groups the word problem is solvable.

The undecidability of the word problem implies the undecidability of many other problems in pure mathematics. For example, the homeomorphism problem for closed manifolds in dimension at least 4 is undecidable [63], and there are far-reaching consequences for the global shape of moduli spaces [95].

Particularly nice presentations of groups consist of a finite generating set and a finite set of relations:

Definition 2.2.25 (Finitely presented group)**.** A group G is *finitely presented*[1] if there exists a finite generating set S and a finite set $R \subset (S \cup S^{-1})^*$ of relations such that $G \cong \langle S \mid R \rangle$.

However, the examples given above already show that also finitely presented groups can be rather complicated.

Clearly, any finitely presented group is finitely generated. The converse is not true in general:

Example 2.2.26 (A finitely generated group that is not finitely presented)**.** The group

$$\bigl\langle s, t \bigm| \{[t^n s t^{-n}, t^m s t^{-m}] \mid n, m \in \mathbb{Z}\} \bigr\rangle$$

is finitely generated, but not finitely presented [10]. This group is an example of a so-called *lamplighter group* (see also Example 2.3.5).

While it might be difficult to prove that a specific group is not finitely presented (and such proofs usually require some input from algebraic topology), there is a non-constructive argument showing that there are finitely generated groups that are not finitely presented (Corollary 2.2.28):

Theorem 2.2.27 (Uncountably many finitely generated groups)**.** *There exist uncountably many isomorphism classes of groups generated by two elements.*

Before sketching de la Harpe's proof [42, Chapter III.C] of this theorem, we discuss an important consequence:

Corollary 2.2.28. *There exist uncountably many finitely generated groups that are* not *finitely presented.*

Proof. Notice that (up to renaming) there are only countably many finite presentations of groups, and hence that there are only countably many isomorphism types of finitely presented groups. However, there are uncountably many finitely generated groups by Theorem 2.2.27. □

[1]Sometimes the term *finitely presented* is reserved for groups together with a choice of a finite presentation. If only existence of a finite presentation is assumed, then this is sometimes called *finitely presentable.* This is in analogy with the terms *oriented* vs. *orientable* for manifolds.

The proof of Theorem 2.2.27 consists of two steps:

1. We first show that there exists a group G generated by two elements that contains uncountably many different normal subgroups (Proposition 2.2.29).
2. We then show that G even has uncountably many quotient groups that are pairwise non-isomorphic (Proposition 2.2.30).

Proposition 2.2.29 (Uncountably many normal subgroups)**.** *There exists a group generated by two elements that has uncountably many normal subgroups.*

Sketch of proof. The basic idea is as follows: We construct a group G generated by two elements that contains a central subgroup C (i.e., each element of this subgroup is fixed by conjugation with any other group element) isomorphic to the additive group $\bigoplus_{\mathbb{Z}} \mathbb{Z}$. The group C contains uncountably many subgroups (e.g., given by taking subgroups generated by the subsystem of the unit vectors corresponding to different subsets of $\mathbb{Z}$), and all these subgroups of C are normal in G because C is central in G.

To this end we consider the group $G := \langle s, t \mid R \rangle$, where

$$R := \bigl\{ \bigl[[s, t^n s t^{-n}], s\bigr] \bigm| n \in \mathbb{Z} \bigr\} \cup \bigl\{ \bigl[[s, t^n s t^{-n}], t\bigr] \bigm| n \in \mathbb{Z} \bigr\}.$$

Let C be the subgroup of G generated by the set $\{[s, t^n s t^{-n}] \mid n \in \mathbb{Z}\}$. All elements of C are invariant under conjugation with s by the first part of the relations, and they are invariant under conjugation with t by the second part of the relations; thus, C is central in G. Moreover, by carefully inspecting the relations, it can be shown that C is isomorphic to the additive group $\bigoplus_{\mathbb{Z}} \mathbb{Z}$. $\square$

Proposition 2.2.30 (Uncountably many quotients)**.** *For a finitely generated group G the following are equivalent:*

1. *The group G contains uncountably many normal subgroups.*
2. *The group G has uncountably many pairwise non-isomorphic quotients.*

Proof. Clearly, the second statement implies the first one. Conversely, suppose that G has only countably many pairwise non-isomorphic quotients.

If Q is a quotient group of G, then Q is countable (as G is finitely generated); hence, there are only countably many group homomorphisms of type $G \longrightarrow Q$; in particular, there can be only countably many normal subgroups N of G with $G/N \cong Q$. Thus, in total, G can have only countably many different normal subgroups. $\square$

The fact that there exist uncountably many finitely generated groups can be used for non-constructive existence proofs of groups with certain features; a recent example of this type of argument is Austin's proof of the existence of finitely generated groups and Hilbert modules over these groups with irrational von Neumann dimension (thereby answering a question of Atiyah in the negative) [3].

2.3 New groups out of old

In many categories, there are several ways to construct objects out of given components; examples of such constructions are products and sums/push-outs. In the world of groups, these correspond to direct products and free (amalgamated) products.

In the first section, we study products and product-like constructions such as semi-direct products; in the second section, we discuss how groups can be glued together, i.e., free (amalgamated) products.

2.3.1 Products and extensions

The simplest type of group constructions are direct products and their twisted variants, semi-direct products.

Definition 2.3.1 (Direct product)**.** Let I be a set, and let $(G_i)_{i\in I}$ be a family of groups. The *(direct) product group* $\prod_{i\in I} G_i$ of $(G_i)_{i\in I}$ is the group whose underlying set is the cartesian product $\prod_{i\in I}$ and whose composition is given by pointwise composition:

$$\begin{aligned} \prod_{i \in I} G_i \times \prod_{i \in I} G_i &\longrightarrow \prod_{i \in I} G_i \\ \bigl((g_i)_{i \in I}, (h_i)_{i \in I}\bigr) &\longmapsto (g_i \cdot h_i)_{i \in I}. \end{aligned}$$

The direct product of groups has the *universal property* of the category theoretic product in the category of groups, i.e., homomorphisms to the direct product group are in one-to-one correspondence to families of homomorphisms to all factors.

The direct product of two groups is an extension of the second factor by the first one (taking the canonical inclusion and projection as maps):

Definition 2.3.2 (Group extension)**.** Let Q and N be groups. An *extension of Q by N* is an exact sequence

$$1 \longrightarrow N \xrightarrow{\ i\ } G \xrightarrow{\ \pi\ } Q \longrightarrow 1$$

of groups, i.e., i is injective, π is surjective and $\operatorname{im} i = \ker \pi$.

Not every group extension has as extension group the direct product of the kernel and the quotient; for example, we can deform the direct product by introducing a twist on the kernel:

Definition 2.3.3 (Semi-direct product)**.** Let N and Q be two groups, and let $\varphi\colon Q \longrightarrow \operatorname{Aut}(N)$ be a group homomorphism (i.e., Q acts on N via φ). The *semi-direct product of Q by N with respect to φ* is the group $N \rtimes_\varphi Q$ whose underlying set is the cartesian product $N \times Q$ and whose composition is

$$\begin{aligned} (N \rtimes_\varphi Q) \times (N \rtimes_\varphi Q) &\longrightarrow (N \rtimes_\varphi Q) \\ \bigl((n_1, q_1), (n_2, q_2)\bigr) &\longmapsto \bigl(n_1 \cdot \varphi(q_1)(n_2), q_1 \cdot q_2\bigr) \end{aligned}$$

In other words, whenever we want to swap the positions of an element of N with an element of Q, then we have to take the twist φ into account. E.g., if φ is the trivial homomorphism, then the corresponding semi-direct product is nothing but the direct product.

Remark 2.3.4 (Semi-direct products and split extensions)**.** A group extension $1 \longrightarrow N \xrightarrow{i} G \xrightarrow{\pi} Q \longrightarrow 1$ *splits* if there exists a group homomorphism $s\colon Q \longrightarrow G$ such that $\pi \circ s = \mathrm{id}_Q$. If $\varphi\colon Q \longrightarrow \mathrm{Aut}(N)$ is a homomorphism, then

$$1 \longrightarrow N \xrightarrow{i} N \rtimes_\varphi Q \xrightarrow{\pi} Q \longrightarrow 1$$

is a split extension; here, $i\colon N \longrightarrow N \rtimes_\varphi$ is the inclusion of the first component, π is the projection onto the second component, and a split is given by

$$\begin{aligned} Q &\longrightarrow N \rtimes_\varphi Q \\ q &\longmapsto (e, q). \end{aligned}$$

Conversely, in a split extension, the extension group is a semi-direct product of the quotient by the kernel: Let $1 \longrightarrow N \xrightarrow{i} G \xrightarrow{\pi} Q \longrightarrow 1$ be an extension of groups that admits a splitting $s\colon Q \longrightarrow G$. Then

$$\begin{aligned} N \rtimes_\varphi Q &\rightleftarrows G \\ (n, q) &\longmapsto n \cdot s(q) \\ \bigl(g \cdot (s \circ \pi(g))^{-1}, \pi(g)\bigr) &\longleftarrow g \end{aligned}$$

are well-defined mutually inverse group homomorphisms, where

$$\begin{aligned} \varphi\colon Q &\longrightarrow \mathrm{Aut}(N) \\ q &\longmapsto \bigl(n \mapsto s(q) \cdot n \cdot s(q)^{-1}\bigr). \end{aligned}$$

However, there are also group extensions that do *not* split; in particular, not every group extension is a semi-direct product. For example, the extension

$$1 \longrightarrow \mathbb{Z} \xrightarrow{2\cdot} \mathbb{Z} \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow 1$$

does not split because there is no non-trivial homomorphism from the torsion group $\mathbb{Z}/2\mathbb{Z}$ to $\mathbb{Z}$. One way to classify group extensions (with Abelian kernel) is to consider group cohomology [18, 57, Chapter IV,Chapter 1.4.4].

Example 2.3.5 (Semi-direct product groups)**.**

- If N and Q are groups and $\varphi\colon Q \longrightarrow \operatorname{Aut}(N)$ is the trivial homomorphism, then the identity map (on the level of sets) yields an isomorphism $N \rtimes_\varphi Q \cong N \times Q$.
- Let $n \in \mathbb{N}_{\geq 3}$. Then the dihedral group $D_n = \langle s, t \mid s^n, t^2, tst^{-1} = s^{-1}\rangle$ (see Example 2.2.20) is a semi-direct product

$$\begin{aligned} D_n &\rightleftarrows \mathbb{Z}/n\mathbb{Z} \rtimes_\varphi \mathbb{Z}/2\mathbb{Z} \\ s &\longmapsto ([1], 0) \\ t &\longmapsto (0, [1]), \end{aligned}$$

  where $\varphi\colon \mathbb{Z}/2\mathbb{Z} \longrightarrow \operatorname{Aut}\mathbb{Z}/n\mathbb{Z}$ is given by multiplication by -1. Similarly, also the infinite dihedral group $D_\infty = \langle s, t \mid t^2, tst^{-1} = s^{-1}\rangle \cong \operatorname{Isom}(\mathbb{Z})$ can be written as a semi-direct product of $\mathbb{Z}/2\mathbb{Z}$ by $\mathbb{Z}$ with respect to multiplication by -1 (exercise).
- Semi-direct products of the type $\mathbb{Z}^n \rtimes_\varphi \mathbb{Z}$ lead to interesting examples of groups provided the automorphism $\varphi(1) \in \operatorname{GL}(n, \mathbb{Z}) \subset \operatorname{GL}(n, \mathbb{R})$ is chosen suitably, e.g., if $\varphi(1)$ has interesting eigenvalues.
- Let G be a group. Then the *lamplighter group over* G is the semi-direct product group $\bigl(\prod_\mathbb{Z} G\bigr) \rtimes_\varphi \mathbb{Z}$, where $\mathbb{Z}$ acts on the product $\prod_\mathbb{Z} G$ by shifting the factors:

$$\begin{aligned} \varphi\colon \mathbb{Z} &\longrightarrow \operatorname{Aut}\Bigl(\prod_\mathbb{Z} G\Bigr) \\ z &\longmapsto \bigl((g_n)_{n\in\mathbb{Z}} \mapsto (g_{n+z})_{n\in\mathbb{Z}}\bigr) \end{aligned}$$

- More generally, the *wreath product* of two groups G and H is the semi-direct product $\bigl(\prod_H G\bigr) \rtimes_\varphi H$, where φ is the shift action of H on $\prod_H G$. The wreath product of G and H is denoted by $G \wr H$.

2.3.2 Free products and free amalgamated products

We now describe a construction that "glues" two groups along a common subgroup. In the language of category theory, glueing processes are modelled by the universal property of pushouts:

Definition 2.3.6 (Free product with amalgamation, universal property)**.** Let A be a group, and let $\alpha_1\colon A \longrightarrow G_1$ and $\alpha_2\colon A \longrightarrow G_2$ be group homomorphisms. A group G together with homomorphisms $\beta_1\colon G_1 \longrightarrow G$ and $\beta_2\colon G_2 \longrightarrow G$ satisfying $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$ is called an *amalgamated free product of* G_1 *and* G_2 *over* A *(with respect to* α_1 *and* α_2*)* if the following universal property is satisfied: For any group H and any two group homomorphisms $\varphi_1\colon G_1 \longrightarrow H$ and $\varphi_2\colon G_2 \longrightarrow H$ with $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$ there is exactly one homomorphism $\varphi\colon G \longrightarrow H$ of groups with $\varphi \circ \beta_1 = \varphi \circ \beta_2$:

$$\begin{tikzcd}
& G_1 \arrow[dr, "\beta_1"] \arrow[drr, bend left, "\varphi_1"] & & \\
A \arrow[ur, "\alpha_1"] \arrow[dr, "\alpha_2"'] & & G \arrow[r, dashed, "\varphi"] & H \\
& G_2 \arrow[ur, "\beta_2"'] \arrow[urr, bend right, "\varphi_2"'] & &
\end{tikzcd}$$

Such a free product with amalgamation is denoted by $G_1 *_A G_2$ (see Theorem 2.3.9 for existence and uniqueness).

If A is the trivial group, then we write $G_1 * G_2 := G_1 *_A G_2$ and call $G_1 * G_2$ the *free product of* G_1 *and* G_2.

Caveat 2.3.7. Notice that in the situation of the above definition, in general, the free product with amalgamation does depend on the homomorphisms α_1, α_2; however, usually, it is clear implicitly which homomorphisms are meant and so they are omitted from the notation.

Example 2.3.8 (Free (amalgamated) products)**.**

- Free groups can also be viewed as free products of several copies of the additive group $\mathbb{Z}$; e.g., the free group of rank 2 is nothing but $\mathbb{Z} * \mathbb{Z}$ (which can be seen by comparing the respective universal properties and using uniqueness).
- The infinite dihedral group $D_\infty \cong \operatorname{Isom}(\mathbb{Z})$ (Example 2.3.5) is isomorphic to the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$; for instance, reflection at 0 and reflection at $1/2$ provide generators of D_∞ corresponding to the obvious generators of $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$ (exercise).
- The matrix group $\mathrm{SL}(2, \mathbb{Z})$ is isomorphic to the free amalgamated product $\mathbb{Z}/6\mathbb{Z} *_{\mathbb{Z}/2\mathbb{Z}} \mathbb{Z}/4\mathbb{Z}$ [87, Example I.4.2].
- Free amalgamated products occur naturally in topology: By the theorem of Seifert and van Kampen, the fundamental group of a pointed

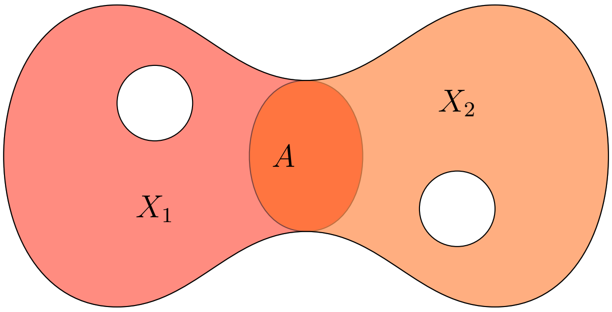

$$\pi_1(X_1 \cup_A X_2) \cong \pi_1(X_1) *_{\pi_1(A)} \pi_1(X_2)$$

Figure 2.3.: The theorem of Seifert and van Kampen, schematically

space glued together of several components is a free amalgamated product of the fundamental groups of the components over the fundamental group of the intersection (the two subspaces and their intersection have to be non-empty and path-connected) [64, Chapter IV] (see Figure 2.3). A quick introduction to fundamental groups is given in Appendix A.1.

Theorem 2.3.9 (Free product with amalgamation, uniqueness and construction)**.** *All free products with amalgamation exist and are unique up to canonical isomorphism.*

Proof. The uniqueness proof is similar to the one that free groups are uniquely determined up to canonical isomorphism by the universal property of free groups (Proposition 2.2.6).

We now prove the existence of free products with amalgamation: Let A be a group and let $\alpha_1 \colon A \longrightarrow G_1$ and $\alpha_2 \colon A \longrightarrow G_2$ be two group homomorphisms. Let

$$G := \bigl\langle \{x_g \mid g \in G_1\} \sqcup \{x_g \mid g \in G_2\} \bigm| \{x_{\alpha_1(a)} {x_{\alpha_2(a)}}^{-1} \mid a \in A\} \cup R_{G_1} \cup R_{G_2} \bigr\rangle,$$

where (for $j \in \{1,2\}$)

$$R_{G_j} := \{x_g x_h {x_k}^{-1} \mid g,h,k \in G_j \text{ with } g \cdot h = k \text{ in } G_j\}.$$

Furthermore, we define for $j \in \{1, 2\}$ group homomorphisms

$$\begin{aligned} \beta_j \colon G_j &\longrightarrow G \\ g &\longmapsto x_g; \end{aligned}$$

the relations R_{G_j} ensure that β_j indeed is compatible with the compositions in G_j and G respectively. Moreover, the relations $\{x_{\alpha_1(a)} x_{\alpha_2(a)}{}^{-1} \mid a \in A\}$ show that $\beta_1 \circ \alpha_1 = \beta_2 \circ \alpha_2$.

The group G (together with the homomorphisms β_1 and β_2) has the universal property of the amalgamated free product of G_1 and G_2 over A because:

Let H be a group and let $\varphi_1 \colon G_1 \longrightarrow H$ and $\varphi_2 \colon G_2 \longrightarrow H$ be homomorphisms with $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$. We define a homomorphism $\varphi \colon G \longrightarrow H$ using the universal property of groups given by generators and relations (Proposition 2.2.18): The map on the set of all words in the generators $\{x_g \mid g \in G\} \sqcup \{x_g \mid g \in G\}$ and their formal inverses induced by the map

$$\begin{aligned} \{x_g \mid g \in G_1\} \sqcup \{x_g \mid g \in G_2\} &\longrightarrow H \\ x_g &\longmapsto \begin{cases} \varphi_1(g) & \text{if } g \in G_1 \\ \varphi_2(g) & \text{if } g \in G_2 \end{cases} \end{aligned}$$

vanishes on the relations in the above presentation of G (it vanishes on R_{G_j} because φ_j is a group homomorphism, and it vanishes on the relations involving A because $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$). Let $\varphi \colon G \longrightarrow H$ be the homomorphism corresponding to this map provided by said universal property.

Furthermore, by construction, $\varphi \circ \beta_1 = \varphi_1$ and $\varphi \circ \beta_2 = \varphi_2$.

As (the image of) $S := \{x_g \mid g \in G_1\} \sqcup \{x_g \mid g \in G_2\}$ generates G and as any homomorphism $\psi \colon G \longrightarrow H$ with $\psi \circ \beta_1 = \varphi_1$ and $\psi \circ \beta_2 = \varphi_2$ has to satisfy "$\psi|_S = \varphi|_S$", we obtain $\psi = \varphi$. In particular, φ is the unique homomorphism of type $G \longrightarrow H$ with $\varphi \circ \beta_1 = \varphi_1$ and $\varphi \circ \beta_2 = \varphi_2$. $\square$

Clearly, the same construction as in the proof above can be applied to any presentation of the summands; this produces more efficient presentations of the amalgamated free product.

Instead of gluing two different groups along subgroups, we can also glue a group to itself along an isomorphism of two of its subgroups:

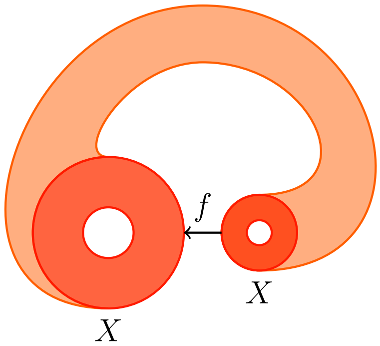

$$\pi_1(\text{mapping torus of } f) \cong \pi_1(X)*_{\pi_1(f)}$$

Figure 2.4.: The fundamental group of a mapping torus, schematically

Definition 2.3.10 (HNN-extension)**.** Let G be a group, let A, $B \subset G$ be two subgroups, and let $\vartheta\colon A \longrightarrow B$ be an isomorphism. Then the *HNN-extension of G with respect to ϑ* is the group

$$G*_\vartheta := \bigl\langle \{x_g \mid g \in G\} \sqcup \{t\} \bigm| \{t^{-1}x_a t = x_{\vartheta(a)} \mid a \in A\} \cup R_G \bigr\rangle,$$

where

$$R_G := \{x_g x_h {x_k}^{-1} \mid g, h, k \in G \text{ with } g \cdot h = k \text{ in } G\}.$$

In other words, using an HNN-extension, we can force two given subgroups to be conjugate; iterating this construction leads to interesting examples of groups [61, 82, Chapter IV, Chapter 12]. HNN-extensions are named after G. Higman, B.H. Neumann, and H. Neumann who were the first to systematically study such groups (Appendix A.2). Topologically, HNN-extensions arise naturally as fundamental groups of mapping tori of maps that are injective on the level of fundamental groups [61, p. 180] (see Figure 2.4).

Remark 2.3.11 (Outlook – ends of groups and actions on trees)**.** The class of (non-trivial) free amalgamated products and of (non-trivial) HNN-extensions plays an important role in geometric group theory; more precisely, they are the key objects in Stallings's classification of groups with infinitely many ends [90] (Theorem 8.2.9), and they are the starting point of Bass-Serre theory [87], which is concerned with nice actions of groups on trees. Moreover, it should be noted that free groups, free products, certain amalgamated free products, and HNN-extensions can be understood in very concrete terms via suitable normal forms (Chapter 3.3.1).

3 Groups → geometry, I: Cayley graphs

One of the central questions of geometric group theory is how groups can be viewed as geometric objects; one way to view a (finitely generated) group as a geometric object is via Cayley graphs:

1. As first step, one associates a combinatorial structure to a group and a given generating set – the corresponding Cayley graph; this step is discussed in this chapter.
2. As second step, one adds a metric structure to Cayley graphs via the so-called word metrics; we will study this step in Chapter 5.

We start by reviewing some basic notation from graph theory (Section 3.1); more information on graph theory can be found in various textbooks [24, 43, 22].

We then introduce Cayley graphs and discuss some basic examples of Cayley graphs (Section 3.2); in particular, we show that free groups can be characterised combinatorially by trees: The Cayley graph of a free group with respect to a free generating set is a tree; conversely, if a group admits a Cayley graph that is a tree, then (under mild additional conditions) the corresponding generating system is a free generating system for the group in question (Section 3.3).

3.1 Review of graph notation

We start by reviewing some basic notation from graph theory; in the following, unless stated explicitly otherwise, we always consider undirected, simple graphs:

Definition 3.1.1 (Graph)**.** A *graph* is a pair $G = (V, E)$ of disjoint sets where E is a set of subsets of V that contain exactly two elements, i.e.,

$$E \subset V^{[2]} := \{e \mid e \subset V,\ |e| = 2\};$$

the elements of V are the *vertices*, the elements of E are the *edges of* G.

In other words, graphs are a different point of view on relations, and normally graphs are used to model relations. Classical graph theory has many applications, mainly in the context of networks of all sorts and in computer science (where graphs are a fundamental basic structure).

Definition 3.1.2 (Adjacent, neighbour)**.** Let (V, E) be a graph.

- We say that two vertices $v, v' \in V$ are *neighbours* or *adjacent* if they are joined by an edge, i.e., if $\{v, v'\} \in E$.
- The number of neighbours of a vertex is the *degree* of this vertex.

Example 3.1.3 (Graphs)**.** Let $V := \{1, 2, 3, 4\}$, and let

$$E := \bigl\{\{1,2\}, \{2,3\}, \{3,1\}\bigr\}.$$

Then the graph $G_1 := (V, E)$ can be illustrated as in Figure 3.1; notice however that differently looking pictures can in fact represent the same graph (a graph is a combinatorial object!). In G_1, the vertices 2 and 3 are neighbours, while 2 and 4 are not.

Similarly, we can consider the following graphs (see Figure 3.1):

$$\begin{aligned}
G_2 &:= \bigl(\{1, \dots, 5\}, \{\{j, k\} \mid j, k \in \{1, \dots, 5\},\ j \neq k\}\bigr),\\
G_3 &:= \bigl(\{1, \dots, 9\}\\
&\qquad , \{\{1,2\}, \{2,3\}, \{3,4\}, \{4,5\}, \{2,6\}, \{3,7\}, \{4,8\}, \{8,9\}\}\bigr).
\end{aligned}$$

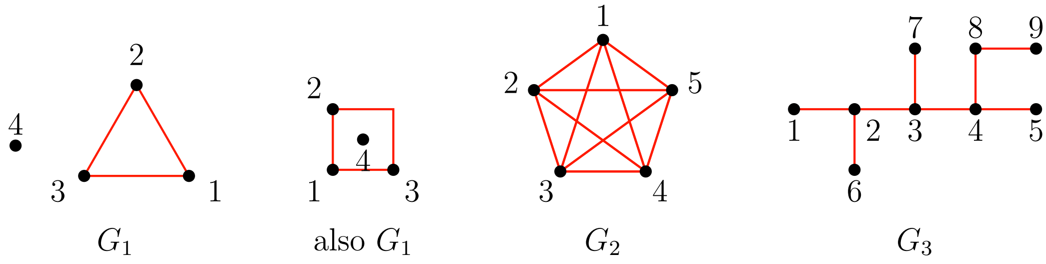

Figure 3.1.: Some graphs

The graph G_2 is a so-called *complete* graph – all vertices are neighbours of each other.

Definition 3.1.4 (Graph isomorphisms)**.** Let $G = (V, E)$ and $G' = (V', E')$ be graphs. The graphs G and G' are *isomorphic*, if there is a *graph isomorphism* between G and G', i.e., a bijection $f\colon V \longrightarrow V'$ such that for all v, $w \in V$ we have $\{v, w\} \in E$ if and only if $\bigl\{f(v), f(w)\bigr\} \in E'$; i.e., isomorphic graphs only differ in the labels of the vertices.[1]

The problem to decide whether two given graphs are isomorphic or not is a difficult problem – in the case of finite graphs, this problem seems to be a problem of high algorithmic complexity, though its exact complexity class is still unknown [49].

In order to work with graphs, we introduce geometric terms for graphs:

Definition 3.1.5 (Paths, cycles)**.** Let $G = (V, E)$ be a graph.

– Let $n \in \mathbb{N} \cup \{\infty\}$. A *path in G of length n* is a sequence $v_0, \dots, v_n$ of different vertices $v_0, \dots, v_n \in V$ with the property that $\{v_j, v_{j+1}\} \in E$ holds for all $j \in \{0, \dots, n-1\}$; if $n < \infty$, then we say that this path *connects the vertices v_0 and v_n*.
– The graph G is called *connected* if any two of its vertices can be connected by a path in G.

[1] Of course, this notion of graph isomorphism can also be obtained as isomorphisms of a category of graphs with suitable morphisms. However, there are several natural choices for such a category; therefore, we prefer the above concrete formulation.

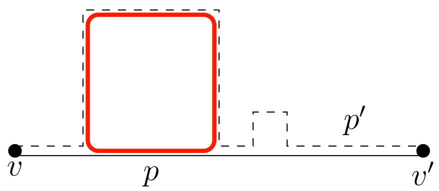

Figure 3.2.: Constructing a cycle (orange) out of two different paths.

- Let $n \in \mathbb{N}_{>2}$. A *cycle in G of length n* is a sequence $v_0, \dots, v_{n-1}$ of different vertices $v_0, \dots, v_{n-1} \in V$ with $\{v_{n-1}, v_0\} \in E$ and moreover $\{v_j, v_{j+1}\} \in E$ for all $j \in \{0, \dots, n-2\}$.

Example 3.1.6. In Example 3.1.3, the graphs G_2 and G_3 are connected, but G_1 is not connected (e.g., in G_1 there is no path connecting the vertex 4 to vertex 1). The sequence $1, 2, 3$ is a path in G_3, but $7, 8, 9$ and $2, 3, 2$ are no paths in G_3. In G_1, the sequence $1, 2, 3$ is a cycle.

Definition 3.1.7 (Tree)**.** A *tree* is a connected graph that does not contain any cycles. A graph that does not contain any cycles is a *forest*; thus, a tree is the same as a connected forest.

Example 3.1.8 (Trees)**.** The graph G_3 from Example 3.1.3 is a tree, while G_1 and G_2 are not.

Proposition 3.1.9 (Characterising trees)**.** *A graph is a tree if and only if for every pair of vertices there exists exactly one path connecting these vertices.*

Proof. Let G be a graph such that every pair of vertices can be connected by exactly one path in G; in particular, G is connected. *Assume* for a contradiction that G contains a cycle $v_0, \dots, v_{n-1}$. Because $n > 2$, the two paths v_0, v_{n-1} and $v_0, \dots, v_{n-1}$ are different, and both connect v_0 with v_{n-1}, which is a contradiction. Hence, G is a tree.

Conversely, let G be a tree; in particular, G is connected, and every two vertices can be connected by a path in G. *Assume* for a contradiction that there exist two vertices v and v' that can be connected by two different paths p and p'. By looking at the first index at which p and p' differ and at the first indices of p and p' respectively where they meet again, we can construct a cycle in G (see Figure 3.2), contradicting the fact that G is a tree. Hence, every two vertices of G can by connected by exactly one path in G. $\square$

3.2 Cayley graphs

Given a generating set of a group, we can organise the combinatorial structure given by the generating set as a graph:

Definition 3.2.1 (Cayley graph)**.** Let G be a group and let $S \subset G$ be a generating set of G. Then the *Cayley graph of G with respect to the generating set S* is the graph $\operatorname{Cay}(G, S)$ whose

- set of vertices is G, and whose
- set of edges is

$$\bigl\{\{g, g \cdot s\} \bigm| g \in G,\ s \in (S \cup S^{-1}) \setminus \{e\}\bigr\}.$$

I.e., two vertices in a Cayley graph are adjacent if and only if they differ by an element of the generating set in question.

Example 3.2.2 (Cayley graphs)**.**

- The Cayley graphs of the additive group $\mathbb{Z}$ with respect to the generating sets $\{1\}$ and $\{2, 3\}$ respectively are illustrated in Figure 3.3. Notice that when looking at these two graphs "from far away" they seem to have the same global structure, namely they look like the real line; in more technical terms, these graphs are quasi-isometric with respect to the corresponding word metrics – a concept that we will study thoroughly in later chapters (Chapter 5 ff.).
- The Cayley graph of the additive group $\mathbb{Z}^2$ with respect to the generating set $\{(1, 0), (0, 1)\}$ looks like the integer lattice in $\mathbb{R}^2$, see Figure 3.4; when viewed from far away, this Cayley graph looks like the Euclidean plane.

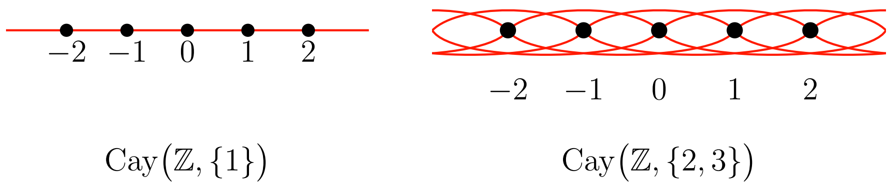

Figure 3.3.: Cayley graphs of the additive group $\mathbb{Z}$

- The Cayley graph of the cyclic group $\mathbb{Z}/6\mathbb{Z}$ with respect to the generating set $\{[1]\}$ looks like a cycle graph (Figure 3.5).
- We now consider the symmetric group S_3. Let τ be the transposition exchanging 1 and 2, and let σ be the cycle $1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1$; the Cayley graph of S_3 with respect to the generating system $\{\tau, \sigma\}$ is depicted in Figure 3.5.

  Notice that the Cayley graph $\operatorname{Cay}(S_3, S_3)$ is a complete graph on six vertices; similarly, $\operatorname{Cay}(\mathbb{Z}/6\mathbb{Z}, \mathbb{Z}/6\mathbb{Z})$ is a complete graph on six vertices. In particular, we see that non-isomorphic groups may have isomorphic Cayley graphs with respect to certain generating sets.

  However, the isomorphism type of any Cayley graph of a finitely generated Abelian group is still rigid enough to remember the size of the torsion part [58].
- The Cayley graph of a free group with respect to a free generating set is a tree (see Theorem 3.3.1 below).

Remark 3.2.3 (Elementary properties of Cayley graphs)**.**

1. Cayley graphs are connected as every vertex g can be reached from the vertex of the neutral element by walking along the edges corresponding to a presentation of minimal length of g in terms of the given generators.
2. Cayley graphs are regular in the sense that every vertex has the same number $|(S \cup S^{-1}) \setminus \{e\}|$ of neighbours.
3. A Cayley graph is locally finite if and only if the generating set is finite; a graph is said to be *locally finite* if every vertex has only finitely many neighbours.

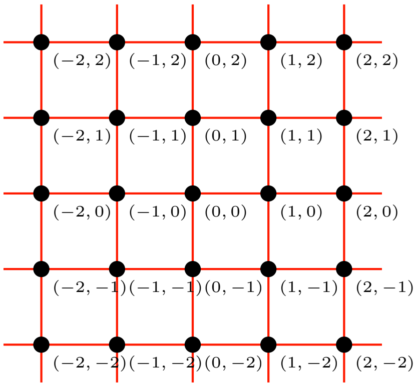

Figure 3.4.: The Cayley graph $\operatorname{Cay}\big(\mathbb{Z}^2, \{(1,0),(0,1)\}\big)$

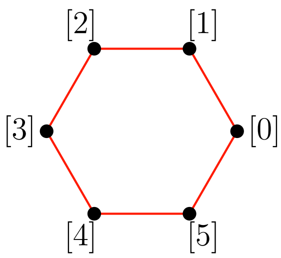

$\operatorname{Cay}\big(\mathbb{Z}/6\mathbb{Z}, \{[1]\}\big)$

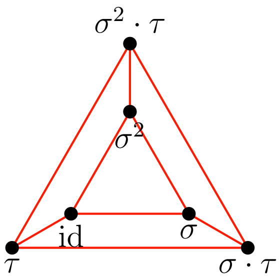

$\operatorname{Cay}\big(S_3, \{\tau, \sigma\}\big)$

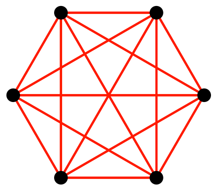

$\operatorname{Cay}(S_3, S_3)$
$\cong \operatorname{Cay}(\mathbb{Z}/6\mathbb{Z}, \mathbb{Z}/6\mathbb{Z})$

Figure 3.5.: Cayley graphs of $\mathbb{Z}/6\mathbb{Z}$ and S_3

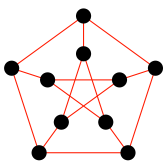

Figure 3.6.: The Petersen graph

Exercise 3.2.1 (Petersen graph)**.** Show that the Petersen graph (depicted in Figure 3.6), even though it is a highly regular graph, is *not* the Cayley graph of any group.

Remark 3.2.4 (Cayley complexes, classifying spaces)**.** There are higher dimensional analogues of Cayley graphs in topology: Associated with a presentation of a group, there is the so-called Cayley complex [17, Chapter 8A], which is a 2-dimensional object. More generally, every group admits a so-called classifying space (or Eilenberg-MacLane space of type $K(\,\cdot\,,1)$, a space whose fundamental group is the given group, and whose higher dimensional homotopy groups are trivial [44, Chapter I.B]. These spaces allow to model group theory (both groups and homomorphisms) in topology and play an important role in the study of group cohomology [18, 57].

Notice that Cayley graphs and Cayley complexes require additional data on the group (generating sets and presentations respectively) and, viewed as combinatorial objects, are only functorial with respect to maps/homomorphisms that respect this additional data.

So far, we considered only the combinatorial structure of Cayley graphs; later, we will also consider Cayley graphs from the point of view of group actions (most groups act freely on their Cayley graphs) (Chapter 4), and from the point of view of large-scale geometry, by introducing metric structures on Cayley graphs (Chapter 5).

3.3 Cayley graphs of free groups

A combinatorial characterisation of free groups can be given in terms of trees:

Theorem 3.3.1 (Cayley graphs of free groups)**.** *Let F be a free group, freely generated by $S \subset F$. Then the corresponding Cayley graph* $\operatorname{Cay}(F, S)$ *is a tree.*

The converse is *not* true in general:

Example 3.3.2 (Non-free groups with Cayley trees)**.**

- The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}/2\mathbb{Z}, [1]\bigr)$ consists of two vertices joined by an edge; clearly, this graph is tree, but the group $\mathbb{Z}/2\mathbb{Z}$ is not free.
- The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}, \{-1, 1\}\bigr)$ coincides with $\operatorname{Cay}\bigl(\mathbb{Z}, \{1\}\bigr)$, which is a tree (looking like a line). But $\{-1, 1\}$ is not a free generating set of $\mathbb{Z}$.

However, these are basically the only things that can go wrong:

Theorem 3.3.3 (Cayley trees and free groups)**.** *Let G be a group and let $S \subset G$ be a generating set satisfying $s \cdot t \neq e$ for all s, $t \in S$. If the Cayley graph* $\operatorname{Cay}(G, S)$ *is a tree, then S is a free generating set of G.*

While it might be intuitively clear that free generating sets do not lead to any cycles in the corresponding Cayley graphs and vice versa, a formal proof requires the description of free groups in terms of reduced words (Section 3.3.1). More generally, any explicit and complete description of the Cayley graph of a group G with respect to a generating set S basically requires to solve the word problem of G with respect to S.

3.3.1 Free groups and reduced words

The construction $F(S)$ of the free group generated by S consisted of taking the set of all words in elements of S and their formal inverses, and taking the quotient by a certain equivalence relation (proof of Theorem 2.2.7). While this construction is technically clean and simple, it has the disadvantage that getting hold of the precise nature of said equivalence relation is tedious.

In the following, we discuss an alternative construction of a group freely generated by S by means of reduced words; it is technically a little bit more cumbersome, but has the advantage that every group element is represented by a canonical word:

Definition 3.3.4 (Reduced word)**.** Let S be a set, and let $(S \cup \overline{S})^*$ be the set of words over S and formal inverses of elements of S.

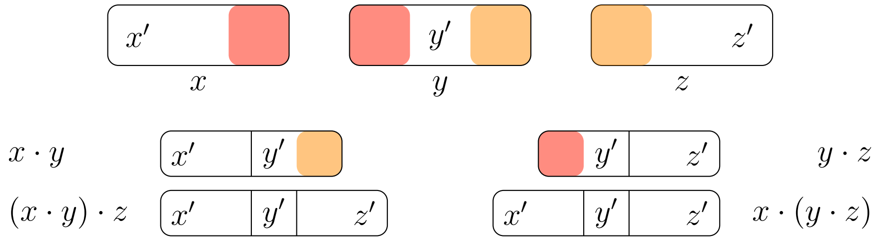

Figure 3.7.: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do not interfere

– Let $n \in \mathbb{N}$ and let $s_1, \dots, s_n \in S \cup \overline{S}$. The word $s_1 \dots s_n$ is *reduced* if

$$s_{j+1} \neq \overline{s_j} \qquad \text{and} \qquad \overline{s_{j+1}} \neq s_j$$

holds for all $j \in \{1, \dots, n-1\}$. (In particular, ε is reduced.)
– We write $F_{\text{red}}(S)$ for the set of all reduced words in $(S \cup \overline{S})^*$.

Proposition 3.3.5 (Free groups via reduced words)**.** *Let S be a set.*
1. *The set $F_{\text{red}}(S)$ of reduced words over $S \cup \overline{S}$ forms a group with respect to the composition $F_{\text{red}}(S) \times F_{\text{red}}(S) \longrightarrow F_{\text{red}}(S)$ given by*

$$(s_1 \dots s_n, s_{n+1} \dots s_m) \longmapsto (s_1 \dots s_{n-r} s_{n+1+r} \dots s_{n+m}),$$

where $s_1 \dots s_n$ and $s_{n+1} \dots s_m$ are elements of $F_{\text{red}}(S)$, and

$$r := \max\bigl\{k \in \{0, \dots, \min(n, m-1)\} \bigm| \forall_{j \in \{0,\dots,k-1\}} \quad s_{n-j} = \overline{s_{n+1+j}} \\ \lor \overline{s_{n-j}} = s_{n+1+j}\bigr\}$$

2. *The group $F_{\text{red}}(S)$ is freely generated by S.*

Sketch of proof. Ad. 1. The above composition is well-defined because if two reduced words are composed, then the composed word is reduced by construction. Moreover, the composition has the empty word ε (which is reduced!) as neutral element, and it is not difficult to show that every reduced word admits an inverse with respect to this composition (take the inverse sequence and flip the bar status of every element).

Thus it remains to prove that this composition is associative (which is the ugly part of this construction): Instead of giving a formal proof

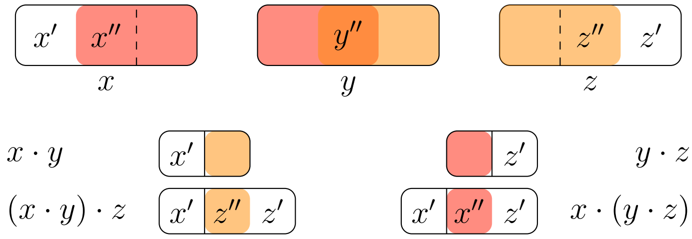

Figure 3.8.: Associativity of the composition in $F_{\mathrm{red}}(S)$; if the reduction areas of the outer elements do interfere

involving lots of indices, we sketch the argument graphically (Figures 3.7 and 3.8): Let x, y, $z \in F_{\mathrm{red}}(S)$; we want to show that $(x \cdot y) \cdot z = x \cdot (y \cdot z)$. By definition, when composing two reduced words, we have to remove the maximal reduction area where the two words meet.

- If the reduction areas of x, y and y, z have no intersection in y, then clearly $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ (Figure 3.7).
- If the reduction areas of x, y and y, z have a non-trivial intersection y'' in y, then the equality $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ follows by carefully inspecting the reduction areas in x and z and the neighbouring regions, as indicated in Figure 3.8; notice that because of the overlap in y'', we know that x'' and z'' coincide (they both are the inverse of y'').

Ad. 2. We show that S is a free generating set of $F_{\mathrm{red}}(S)$ by verifying that the universal property is satisfied: So let H be a group and let $\varphi \colon S \longrightarrow H$ be a map. Then a straightforward (but slightly technical) computation shows that

$$\overline{\varphi} := \varphi^*|_{F_{\mathrm{red}}(S)} \colon F_{\mathrm{red}}(S) \longrightarrow H$$

is a group homomorphism (recall that φ^* is the extension of φ to the set $(S \cup \overline{S})^*$ of all words). Clearly, $\overline{\varphi}|_S = \varphi$; because S generates $F_{\mathrm{red}}(S)$ it follows that $\overline{\varphi}$ is the only such homomorphism. Hence, $F_{\mathrm{red}}(S)$ is freely generated by S. □

As a corollary to the proof of the second part, we obtain:

Corollary 3.3.6. *Let S be a set. Any element of $F(S) = (S \cup \overline{S})^* / \sim$ can be represented by exactly one reduced word over $S \cup \overline{S}$.* □

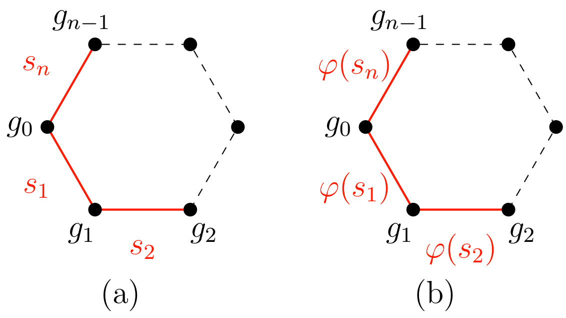

Figure 3.9.: Cycles lead to reduced words, and vice versa

Corollary 3.3.7. *The word problem in free groups with respect to free generating sets is solvable: we just need to consider and compare reduced words.* □

Remark 3.3.8 (Reduced words in free products etc.)**.** Using the same method of proof, one can describe free products $G_1 * G_2$ of groups G_1 and G_2 by reduced words; in this case, one calls a word

$$g_1 \dots g_n \in (G_1 \sqcup G_2)^*$$

with $n \in \mathbb{N}$ and $g_1, \dots, g_n \in G_1 \sqcup G_2$ *reduced*, if for all $j \in \{1, \dots, n-1\}$

- either $g_j \in G_1 \setminus \{e\}$ and $g_{j+1} \in G_2 \setminus \{e\}$,
- or $g_j \in G_2 \setminus \{e\}$ and $g_{j+1} \in G_1 \setminus \{e\}$.

Similarly, one can also describe free amalgamated products and HNN-extensions by suitable classes of reduced words [87, Chapter I].

3.3.2 Free groups → trees

Proof of Theorem 3.3.1. Suppose the group F is freely generated by S. By Proposition 3.3.5, the group F is isomorphic to $F_{\text{red}}(S)$ via an isomorphism that is the identity on S; without loss of generality we can therefore assume that F is $F_{\text{red}}(S)$.

We now show that the Cayley graph $\text{Cay}(F, S)$ is a tree: Because S generates F, the graph $\text{Cay}(F, S)$ is connected. *Assume* for a contradiction

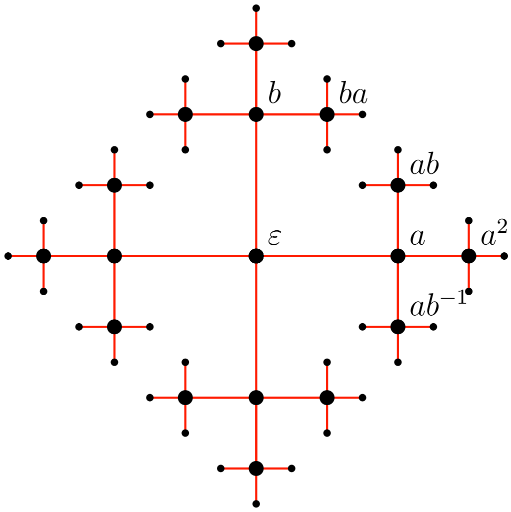

Figure 3.10.: Cayley graph of the free group of rank 2 with respect to a free generating set $\{a, b\}$

that $\operatorname{Cay}(F, S)$ contains a cycle $g_0, \dots, g_{n-1}$ of length n with $n \geq 3$; in particular, the elements $g_0, \dots, g_{n-1}$ are distinct, and

$$s_{j+1} := g_{j+1} \cdot {g_j}^{-1} \in S \cup S^{-1}$$

for all $j \in \{0, \dots, n-2\}$, as well as $s_n := g_0 \cdot {g_{n-1}}^{-1} \in S \cup S^{-1}$ (Figure 3.9 (a)). Because the vertices are distinct, the word $s_0 \dots s_{n-1}$ is reduced; on the other hand, we obtain

$$s_n \dots s_1 = g_0 \cdot {g_{n-1}}^{-1} \cdot \dots \cdot g_2 \cdot {g_1}^{-1} \cdot g_1 \cdot {g_0}^{-1} = e = \varepsilon$$

in $F = F_{\text{red}}(S)$, which is impossible. Therefore, $\operatorname{Cay}(F, S)$ cannot contain any cycles. So $\operatorname{Cay}(F, S)$ is a tree. □

Example 3.3.9 (Cayley graph of the free group of rank 2)**.** Let S be a set consisting of two different elements a and b. Then the corresponding Cayley graph $\operatorname{Cay}\big(F(S), \{a, b\}\big)$ is a regular tree whose vertices have exactly four neighbours (see Figure 3.10).

3.3.3 Trees → free groups

Proof of Theorem 3.3.3. Let G be a group and let $S \subset G$ be a generating set satisfying $s \cdot t \neq e$ for all $s, t \in S$ and such that the corresponding Cayley graph $\operatorname{Cay}(G, S)$ is a tree. In order to show that then S is a free generating set of G, in view of Proposition 3.3.5, it suffices to show that G is isomorphic to $F_{\mathrm{red}}(S)$ via an isomorphism that is the identity on S.

Because $F_{\mathrm{red}}(S)$ is freely generated by S, the universal property of free groups provides us with a group homomorphism $\varphi \colon F_{\mathrm{red}}(S) \longrightarrow G$ that is the identity on S. As S generates G, it follows that φ is surjective. *Assume* for a contradiction that φ is not injective. Let $s_1 \dots s_n \in F_{\mathrm{red}}(S) \setminus \{\varepsilon\}$ with $s_1, \dots, s_n \in S \cup \overline{S}$ be an element of minimal length that is mapped to e by φ. We consider the following cases:

- Because $\varphi|_S = \mathrm{id}_S$ is injective, it follows that $n > 1$.
- If $n = 2$, then it would follow that
$$e = \varphi(s_1 \cdot s_2) = \varphi(s_1) \cdot \varphi(s_2) = s_1 \cdot s_2$$
in G, contradicting that $s_1 \dots s_n$ is reduced and that $s \cdot t \neq e$ for all $s, t \in S$.
- If $n \geq 3$, we consider the sequence $g_0, \dots, g_{n-1}$ of elements of G given inductively by $g_0 := e$ and
$$g_{j+1} := g_j \cdot \varphi(s_{j+1})$$
for all $j \in \{0, \dots, n-2\}$ (Figure 3.9 (b)). The sequence $g_0, \dots, g_{n-1}$ is a cycle in $\operatorname{Cay}(G, S)$ because by minimality of the word $s_1 \dots s_n$, the elements $g_0, \dots, g_{n-1}$ are all distinct; moreover, $\operatorname{Cay}(G, S)$ contains the edges $\{g_0, g_1\}, \dots, \{g_{n-2}, g_{n-1}\}$, and the edge
$$\begin{aligned}\{g_{n-1}, g_0\} &= \{s_1 \cdot s_2 \cdot \dots \cdot s_{n-1}, e\} \\ &= \{s_1 \cdot s_2 \cdot \dots \cdot s_{n-1}, s_1 \cdot s_2 \cdot \dots \cdot s_n\}.\end{aligned}$$

  However, this contradicts that $\operatorname{Cay}(G, S)$ is a tree.

Hence, $\varphi \colon F_{\mathrm{red}}(S) \longrightarrow G$ is injective. $\square$

4 Groups $\to$ geometry, II: Group actions

In the previous chapter, we took the first step from groups to geometry by considering Cayley graphs. In the present chapter, we consider another geometric aspect of groups by looking at so-called group actions, which can be viewed as a generalisation of seeing groups as symmetry groups. We start by reviewing some basic concepts about group actions (Section 4.1). Further introductory material on group actions and symmetry can be found in Armstrong's book [2].

As we have seen, free groups can be characterised combinatorially as the groups admitting trees as Cayley graphs (Section 3.3). In Section 4.2, we will prove that this characterisation can be generalised to a first geometric characterisation of free groups: A group is free if and only if it admits a free action on a tree. An important consequence of this characterisation is that it leads to an elegant proof of the fact that subgroups of free groups are free – which is a purely algebraic statement! (Section 4.3).

Another tool helping us to recognise that certain groups are free is the so-called ping-pong lemma (Section 4.4); this is particularly useful to prove that certain matrix groups are free – which also is a purely algebraic statement (Section 4.5).

4.1 Review of group actions

Recall that for an object X in a category C the set $\mathrm{Aut}_C(X)$ of all C-automorphisms of X is a group with respect to composition in the category C (Exercise 2.1.1).

Definition 4.1.1 (Group action)**.** Let G be a group, let C be a category, and let X be an object in C. An *action of G on X in the category C* is a group homomorphism $G \longrightarrow \mathrm{Aut}_C(X)$. In other words, a group action of G on X consists of a family $(f_g)_{g\in G}$ of automorphisms of X such that

$$f_g \circ f_h = f_{g\cdot h}$$

holds for all $g, h \in G$.

Example 4.1.2 (Group actions, generic examples)**.**

– Every group G admits an action on any object X in any category C, namely the *trivial action*:

$$\begin{aligned} G &\longrightarrow \mathrm{Aut}_C(X) \\ g &\longmapsto \mathrm{id}_X . \end{aligned}$$

– If X is an object in a category C, the automorphism group $\mathrm{Aut}_C(X)$ canonically acts on X via the homomorphism

$$\mathrm{id}_{\mathrm{Aut}_C(X)} \colon \mathrm{Aut}_C(X) \longrightarrow \mathrm{Aut}_C(X).$$

In other words: group actions are a concept generalising automorphism/symmetry groups.

– Let G be a group and let X be a set. If $\varrho \colon G \longrightarrow \mathrm{Aut}_{\mathsf{Set}}(X)$ is an action of G on X by bijections, then we also use the notation

$$g \cdot x := \bigl(\varrho(g)\bigr)(x)$$

for $g \in G$ and $x \in X$, and we can view ϱ as a map $G \times X \longrightarrow X$. If $\varrho\colon G \longrightarrow \operatorname{Aut}_{\mathsf{Set}}(X)$ is a map, then ϱ is a G-action on X if and only if

$$\forall_{g,h\in G}\ \forall_{x\in X}\ \ (g \cdot h) \cdot x = g \cdot (h \cdot x).$$

More generally, a map $\cdot : G \times X \longrightarrow X$ defines a G-action on X if and only if

$$\begin{aligned} \forall_{g,h\in G}\ \forall_{x\in X}\ \ (g \cdot h) \cdot x &= g \cdot (h \cdot x) \\ \forall_{x\in X}\ \ e \cdot x &= x. \end{aligned}$$

More generally, we also use this notation whenever the group G acts on an object in a category, where objects are sets (with additional structure), where morphisms are (structure preserving) maps of sets and the composition of morphisms is nothing but composition of maps. This applies for example to

  - actions by isometries on a metric space (*isometric actions*),
  - actions by homeomorphisms on a topological space (*continuous actions*),
  - actions by linear isomorphisms on vector spaces (*representations*),
  - ...

- Further examples of group actions are actions of groups on a topological space by homotopy equivalences or actions on a metric space by quasi-isometries (see Chapter 5); notice however, that in these cases, automorphisms are equivalence classes of maps of sets and composition of morphisms is done by composing representatives of the corresponding equivalence classes.

On the one hand, group actions allow us to understand groups better by looking at suitable objects on which the groups act nicely; on the other hand, group actions also allow us to understand geometric objects better by looking at groups that can act nicely on these objects.

4.1.1 Free actions

The relation between groups and geometric objects acted upon is particularly strong if the group action is a so-called free action. Important examples of free actions are the natural actions of groups on their Cayley graphs (provided the group does not contain any elements of order 2), and the action of the fundamental group of a space on its universal covering.

Definition 4.1.3 (Free action on a set)**.** Let G be a group, let X be a set, and let $G \times X \longrightarrow X$ be an action of G on X. This action is *free* if

$$g \cdot x \neq x$$

holds for all $g \in G \setminus \{e\}$ and all $x \in X$. In other words, an action is free if and only if every non-trivial group element acts without fixed points.

Example 4.1.4 (Left translation action)**.** If G is a group, then the left translation action

$$\begin{aligned} G &\longrightarrow \operatorname{Aut}(G) \\ g &\longmapsto (h \mapsto g \cdot h) \end{aligned}$$

is a free action of G on itself by bijections.

Example 4.1.5 (Rotations on the circle)**.** Let $S^1 := \{z \in \mathbb{C} \mid |z| = 1\}$ be the unit circle in $\mathbb{C}$, and let $\alpha \in \mathbb{R}$. Then the rotation action

$$\begin{aligned} \mathbb{Z} \times S^1 &\longrightarrow S^1 \\ (n, z) &\longmapsto e^{2\pi i \cdot \alpha \cdot n} \cdot z \end{aligned}$$

of $\mathbb{Z}$ on S^1 is free if and only if α is irrational (because π is irrational).

Example 4.1.6 (Isometry groups)**.** In general, the action of an isometry group on its underlying geometric object is not necessarily free: For example, the isometry group of the unit square does not act freely on the unit square – e.g., the vertices of the unit square are fixed by reflection along the diagonal through the vertex in question.

Example 4.1.7 (Universal covering)**.** Let X be a "nice" path-connected topological space (e.g., a CW-complex). Associated with X there is a so-called universal covering space $\widetilde{X}$, a path-connected space covering X that has trivial fundamental group [64, Chapter V].

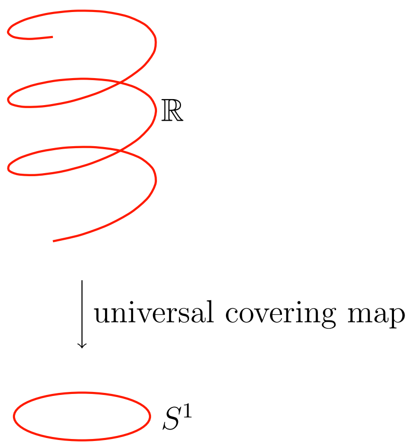

Figure 4.1.: Universal covering of S^1

The fundamental group $\pi_1(X)$ can be identified with the deck transformation group of the universal covering $\widetilde{X} \longrightarrow X$ and the action of $\pi_1(X)$ on $\widetilde{X}$ by deck transformations is free (and properly discontinuous) [64, Chapter V].

For instance:

– The fundamental group of S^1 is isomorphic to $\mathbb{Z}$, the universal covering of S^1 is the exponential map $\mathbb{R} \longrightarrow S^1$, and the action of the fundamental group of S^1 on $\mathbb{R}$ is given by translation (Figure 4.1).
– The fundamental group of the torus $S^1 \times S^1$ is isomorphic to $\mathbb{Z} \times \mathbb{Z}$, the universal covering of $S^1 \times S^1$ is the component-wise exponential map $\mathbb{R}^2 \longrightarrow S^1 \times S^1$, and the action of the fundamental group of $S^1 \times S^1$ on $\mathbb{R}^2$ is given by translation.
– The fundamental group of the figure eight $S^1 \vee S^1$ is isomorphic to F_2, the universal covering space of $S^1 \vee S^1$ is the CW-complex T whose underlying combinatorics is given by the regular tree of degree 4 (Figure 3.10), the universal covering map $T \longrightarrow S^1 \vee S^1$ collapses all 0-cells to the gluing point of $S^1 \vee S^1$ and wraps the "horizontal" and "vertical" edges around the two different circles.

There are two natural definitions of free actions on graphs – one that requires that no vertex and no edge is fixed by any non-trivial group element and one that only requires that no vertex is fixed. We will use the first, stronger, one:

Definition 4.1.8 (Free action on a graph)**.** Let G be a group acting on a graph (V, E) by graph isomorphisms via $\varrho\colon G \longrightarrow \operatorname{Aut}(V, E)$. The action ϱ is *free* if for all $g \in G \setminus \{e\}$ we have

$$\begin{aligned} \forall_{v \in V} \ \ \bigl(\varrho(g)\bigr)(v) &\neq v, \text{and} \\ \forall_{\{v,v'\} \in E} \ \ \bigl\{(\varrho(g))(v), (\varrho(g))(v')\bigr\} &\neq \{v, v'\}. \end{aligned}$$

Example 4.1.9 (Left translation action on Cayley graphs)**.** Let G be a group and let S be a generating set of G. Then the group G acts by graph isomorphisms on the Cayley graph $\operatorname{Cay}(G, S)$ via left translation:

$$\begin{aligned} G &\longrightarrow \operatorname{Aut}\bigl(\operatorname{Cay}(G, S)\bigr) \\ g &\longmapsto (h \mapsto g \cdot h); \end{aligned}$$

notice that this map is indeed well-defined and a group homomorphism.

Proposition 4.1.10 (Free actions on Cayley graphs)**.** *Let G be a group and let S be a generating set of G. Then the left translation action on the Cayley graph* $\operatorname{Cay}(G, S)$ *is free if and only if S does not contain any elements of order* 2*.*

Recall that the *order* of a group element g of a group G is the infimum of all $n \in \mathbb{N}_{>0}$ with $g^n = e$; here, we use the convention $\inf \emptyset := \infty$.

Proof. The action on the vertices is nothing but the left translation action by G on itself, which is free. It therefore suffices to study under which conditions the action of G on the edges is free:

If the action of G on the edges of the Cayley graph $\operatorname{Cay}(G, S)$ is not free, then S contains an element of order 2: Let $g \in G$, and let $\{v, v'\}$ be an edge of $\operatorname{Cay}(G, S)$ with $\{v, v'\} = g \cdot \{v, v'\} = \{g \cdot v, g \cdot v'\}$; by definition, we can write $v' = v \cdot s$ with $s \in S \cup S^{-1} \setminus \{e\}$. Then one of the following cases occurs:

1. We have $g \cdot v = v$ and $g \cdot v' = v'$. Because the action of G on the vertices is free, this is equivalent to $g = e$.
2. We have $g \cdot v = v'$ and $g \cdot v' = v$. Then in G we have

   $$v = g \cdot v' = g \cdot (v \cdot s) = (g \cdot v) \cdot s = v' \cdot s = (v \cdot s) \cdot s = v \cdot s^2$$

   and so $s^2 = e$. As $s \neq e$, it follows that S contains an element of order 2.

Conversely, if $s \in S$ has order 2, then s fixes the edge $\{e, s\} = \{s^2, s\}$ of $\mathrm{Cay}(G, S)$. $\Box$

4.1.2 Orbits and stabilisers

A group action can be disassembled into orbits, leading to the orbit space of the action. Conversely, one can try to understand the whole object by looking at the orbit space and the orbits/stabilisers.

Definition 4.1.11 (Orbit)**.** Let G be a group acting on a set X.

- The *orbit* of an element $x \in X$ with respect to this action is the set

$$G \cdot x := \{g \cdot x \mid g \in G\}.$$

- The *quotient* of X by the given G-action (or *orbit space*) is the set

$$G \setminus X := \{G \cdot x \mid x \in X\}$$

of orbits; we write "$G \setminus X$" because G acts "from the left."

In a sense, the orbit space describes the original object "up to symmetry" or "up to irrelevant transformations."

If a group does not only act by bijections on a set, but if the set is equipped with additional structure that is preserved by the action (e.g., an action by isometries on a metric space), then usually also the orbit space inherits additional structure similar to the one on the space acted upon. However, in general, the orbit space is not as well-behaved as the original space; e.g., the quotient space of an action on a metric space by isometries in general is only a pseudo-metric space – even if the action is free (e.g., this happens for irrational rotations on the circle).

Example 4.1.12 (Rotation on $\mathbb{C}$)**.** We consider the action of the unit circle S^1 (which is a group with respect to multiplication) on the complex numbers $\mathbb{C}$ given by multiplication of complex numbers. The orbit of the origin 0 is just $\{0\}$; the orbit of an element $z \in \mathbb{C} \setminus \{0\}$ is the circle around 0 passing through z (Figure 4.2). The quotient of $\mathbb{C}$ by this action can be identified with $\mathbb{R}_{\geq 0}$ (via the absolute value).

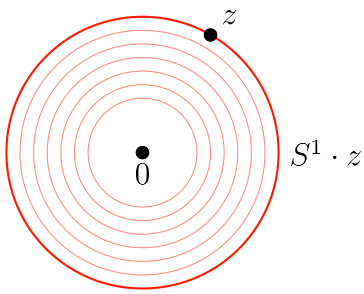

Figure 4.2.: Orbits of the rotation action of S^1 on $\mathbb{C}$

Example 4.1.13 (Universal covering)**.** Let X be a "nice" path-connected topological space (e.g., a CW-complex). The quotient of the universal covering $\widetilde{X}$ by the action of the fundamental group $\pi_1(X)$ by deck transformations is homeomorphic to X [64, Chapter V].

Definition 4.1.14 (Stabiliser, fixed set)**.** Let G be a group acting on a set X.

- The *stabiliser group* of an element $x \in X$ with respect to this action is given by
$$G_x := \{g \in G \mid g \cdot x = x\};$$
notice that G_x indeed is a group (a subgroup of G).
- The *fixed set* of an element $g \in G$ is given by
$$X^g := \{x \in X \mid g \cdot x = x\};$$
more generally, if $H \subset G$ is a subset, then we write $X^H := \bigcap_{h \in H} X^h$.
- We say that the action of G on X has *a global fixed point*, if $X^G \neq \emptyset$.

Example 4.1.15 (Isometries of the unit square)**.** Let Q be the (filled) unit square in $\mathbb{R}^2$, and let G be the isometry group of Q with respect to the Euclidean metric on $\mathbb{R}^2$. Then G naturally acts on Q by isometries and we know that $G \cong D_4$ (Example 2.2.20).

- Let $t \in G$ be the reflection along the diagonal passing through $(0,0)$ and $(1,1)$. Then
$$Q^t = \bigl\{(x,x) \bigm| x \in [0,1]\bigr\}.$$

– Let $s \in G$ be rotation around $2\pi/4$. Then
$$Q^s = \{(1/2, 1/2)\}.$$
– The orbit of $(0,0)$ are all four vertices of Q, and the stabiliser of $(0,0)$ is $G_{(0,0)} = \{\mathrm{id}_Q, t\}$.
– The stabiliser of $(1/3, 0)$ is the trivial group.
– The stabiliser of $(1/2, 1/2)$ is $G_{(1/2,1/2)} = G$, so $(1/2, 1/2)$ is a global fixed point of this action.

Proposition 4.1.16 (Actions of finite groups on trees)**.** *Any action of a finite group on a (non-empty) tree has a global fixed point (in the sense that there is a vertex fixed by all group elements or an edge fixed by all group elements).*

Proof. Exercise. [We look at a vertex such that the diameter of the corresponding orbit is minimal. If there is no globally fixed vertex or edge, then this minimal diameter is at least 2. We now consider the vertices of such a minimal orbit and all paths connecting them. Using the uniqueness of paths in trees (Proposition 3.1.9) one can then find a vertex on these paths whose orbit has smaller diameter, contradicting minimality.] □

Proposition 4.1.17 (Counting orbits)**.** *Let G be a group acting on a set X.*

1. *If $x \in X$, then the map*
$$\begin{aligned} A_x \colon G/G_x &\longrightarrow G \cdot x \\ g \cdot G_x &\longmapsto g \cdot x \end{aligned}$$
is well-defined and bijective. Here, G/G_x denotes the set of all G_x-cosets in G, i.e., $G/G_x = \{g \cdot G_x \mid g \in G\}$.
2. *Moreover, the number of distinct orbits equals the average number of points fixed by a group element: If G and X are finite, then*
$$|G \setminus X| = \frac{1}{|G|} \cdot \sum_{g \in G} |X^g|.$$

Proof. *Ad 1.* We start by showing that A_x is well-defined, i.e., that the definition does not depend on the chosen representatives in G/G_x: Let $g_1, g_2 \in G$ with $g_1 \cdot G_x = g_2 \cdot G_x$. Then there exists an $h \in G_x$ with $g_1 = g_2 \cdot h$. By definition of G_x, we then have $g_1 \cdot x = (g_2 \cdot h) \cdot x = g_2 \cdot (h \cdot x) = g_2 \cdot x$; thus, A_x is well-defined.

By construction, the map A_x is surjective. Why is A_x injective? Let $g_1, g_2 \in G$ with $g_1 \cdot x = g_2 \cdot x$. Then $(g_1^{-1} \cdot g_2) \cdot x = x$ and so $g_1^{-1} \cdot g_2 \in G_x$. Therefore, $g_1 \cdot G_1 = g_1 \cdot (g_1^{-1} \cdot g_2) \cdot G_x = g_2 \cdot G_x$. Hence, A_x is injective.

Ad 2. This equality is proved by double counting: More precisely, we consider the set

$$F := \{(g, x) \mid g \in G, x \in X, g \cdot x = x\} \subset G \times X.$$

By definition of stabiliser groups and fixed sets, we obtain

$$\sum_{x \in X} |G_x| = |F| = \sum_{g \in G} |X^g|.$$

We now transform the right hand side: Notice that $|G/G_x| \cdot |G_x| = |G|$ because every coset of G_x in G has the same size as G_x; therefore, using the first part, we obtain

$$\begin{aligned}
\sum_{x \in X} |G_x| &= \sum_{x \in X} \frac{|G|}{|G/G_x|} \\
&= \sum_{x \in X} \frac{|G|}{|G \cdot x|} \\
&= \sum_{G \cdot x \in G \setminus X} \sum_{y \in G \cdot x} \frac{|G|}{|G \cdot x|} \\
&= \sum_{G \cdot x \in G \setminus X} |G \cdot x| \cdot \frac{|G|}{|G \cdot x|} \\
&= |G \setminus X| \cdot |G|.
\end{aligned}$$

□

4.1.3 Application: Counting via group actions

Group actions can be used to solve certain counting problems. A standard example from algebra is the proof of the Sylow theorems in finite group

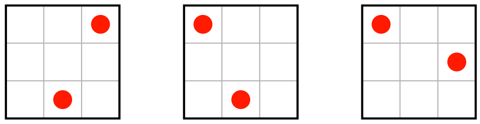

Figure 4.3.: The same punchcard seen from different sides/angles

theory: in this proof, the conjugation action of a group on the set of certain subgroups is considered [2, Chapter 20].

Another class of examples arises in combinatorics:

Example 4.1.18 (Counting punchcards [1])**.** How many 3×3-punchcards with exactly two holes are there, if front and back sides of the punchcards are not distinguishable, and also all vertices are indistinguishable? For example, the punchcards depicted in Figure 4.3 are all considered to be the same.

In terms of group actions, this question can be reformulated as follows: We consider the set of all configurations

$$\bigl\{(x, y) \bigm| x, y \in \{0, 1, 2\} \times \{0, 1, 2\},\ x \neq y\bigr\}$$

of two holes in a 3×3-square, on which the isometry group D_4 of the "square" $\{0, 1, 2\}^2$ acts, and we want to know how many different orbits this action has.

In view of Proposition 4.1.17, it suffices to determine for each element of $D_4 \cong \langle s, t \,|\, t^2, s^4, tst^{-1} = s^{-1}\rangle$ (Example 2.2.20) the number of configurations fixed by this element. Notice that conjugate elements have the same number of fixed points. We obtain the following table (where $\overline{s}$ and $\overline{t}$ denote the images of s and t respectively under the canonical map from $F(\{s, t\})$ to D_4):

conjugacy class in D_4	*number of fixed configurations*
e	36
$\overline{s}$, $\overline{s}^{-1}$	0
$\overline{s}^2$	4
$\overline{t}$, $\overline{s}^2 \cdot \overline{t}$	$3 + 3 = 6$
$\overline{s} \cdot \overline{t}$, $\overline{t} \cdot \overline{s}$	$3 + 3 = 6$

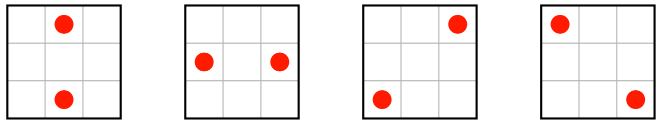

Figure 4.4.: The punchcards fixed by rotation around π

For example, the element $\overline{s}^2$, i.e., rotation around π, fixes exactly the four configurations shown in Figure 4.4. Using the formula of Propositon 4.1.17 we obtain that in total there are exactly

$$\frac{1}{8} \cdot (1 \cdot 36 + 2 \cdot 0 + 1 \cdot 4 + 2 \cdot 6 + 2 \cdot 6) = 8$$

different punchcards.

In this example, it is also possible to go through all 36 configurations and check by hand which of the configurations lead to the same punchcards; however, the argument given above, easily generalises to bigger punchcards – the formula of Proposition 4.1.17 provides a systematic way to count essentially different configurations.

4.2 Free groups and actions on trees

In this section, we show that free groups can be characterised geometrically via free actions on trees; recall that for a free action of a group on a graph no non-trivial group element is allowed to fix any vertices or edges (Definition 4.1.8).

Theorem 4.2.1 (Free groups and actions on trees)**.** *A group is free if and only if it admits a free action on a (non-empty) tree.*

Proof of Theorem 4.2.1, part I. Let F be a free group, freely generated by $S \subset F$; then the Cayley graph $\operatorname{Cay}(F, S)$ is a tree by Theorem 3.3.1. We now consider the left translation action of F on $\operatorname{Cay}(F, S)$.

Looking at the description of F in terms of reduced words (Proposition 3.3.5) or applying the universal property of F with respect to the free generating set S to maps $S \longrightarrow \mathbb{Z}$ it is easily seen that S cannot contain any element of order 2; therefore, the left translation action of F on $\operatorname{Cay}(F, S)$ is free by Proposition 4.1.10. $\square$

Conversely, suppose that a group G acts freely on a tree T. How can we prove that G has to be free? Roughly speaking, we will show that out of T and the G-action on T we can construct – by contracting certain subtrees – a tree that is a Cayley graph of G for a suitable generating set and such that the assumptions of Theorem 3.3.3 are satisfied. This allows us to deduce that the group G is free.

The subtrees that will be contracted are so-called (equivariant) spanning trees, which we will discuss in the following section.

4.2.1 Spanning trees

Trees are the simplest and best accessible type of graphs. In many applications, one therefore wants to approximate graphs by trees. One, crude, approximation is given by spanning trees:

Definition 4.2.2 (Spanning tree)**.** Let G be a group acting on a connected graph X by graph automorphisms. A *spanning tree* of this action is a subtree of X that contains exactly one vertex of every orbit of the induced G-action on the vertices of X.

A *subgraph* of a graph (V, E) is a graph (V', E') with $V' \subset V$ and $E' \subset E$; a subgraph is a *subtree* if it is in addition a tree.

Example 4.2.3 (Spanning trees)**.** We consider the action of $\mathbb{Z}$ by "horizontal" shifting on the (infinite) tree depicted in Figure 4.5. Then the red subgraph is a spanning tree for this action.

Theorem 4.2.4 (Existence of spanning trees)**.** *Every action of a group on a connected graph by graph automorphisms admits a spanning tree.*

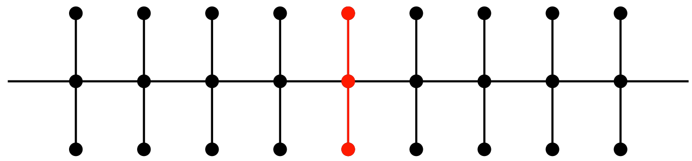

Figure 4.5.: A spanning tree (red) for a shift action of $\mathbb{Z}$

Proof. Let G be a group acting on a connected graph X. In the following, we may assume that X is non-empty (otherwise the empty tree is a spanning tree for the action). We consider the set T_G of all subtrees of X that contain at most one vertex from every G-orbit. We show that T_G contains an element T that is maximal with respect to the subtree relation. The set T_G is non-empty, e.g., the empty tree is an element of T_G. Clearly, the set T_G is partially ordered by the subgraph relation, and any totally ordered chain of T_G has an upper bound in T_G (namely the "union" over all trees in this chain). By the lemma of Zorn, there is a maximal element T in T_G; because X is non-empty, so is T.

We now show that T is a spanning tree for the G-action on X: *Assume* for a contradiction that T is not a spanning tree for the G-action on X. Then there is a vertex v such that none of the vertices of the orbit $G \cdot v$ is a vertex of T. We now show that there is such a vertex v such that one of the neighbours of v is a vertex of T:

As X is connected there is a path p connecting some vertex u of T with v. Let v' be the first vertex on p that is not in T. We distinguish the following two cases:

1. None of the vertices of the orbit $G \cdot v'$ is contained in T; then the vertex v' has the desired property.
2. There is a $g \in G$ such that $g \cdot v'$ is a vertex of T. If p' denotes the subpath of p starting in v' and ending in v, then $g \cdot p'$ is a path starting in the vertex $g \cdot v'$, which is a vertex of T, and ending in $g \cdot v$, a vertex such that none of the vertices in $G \cdot g \cdot v = G \cdot v$ is in T. Notice that the path p' is shorter than the path p, so that iterating this procedure produces eventually a vertex with the desired property.

Let v be a vertex such that none of the vertices of the orbit $G \cdot v$ is in T, and such that some neighbour u of v is in T. Then clearly, adding v and

Figure 4.6.: Contracting the spanning tree and all its copies to vertices (red), in the situation of Example 4.2.3

the edge $\{u, v\}$ to T produces a tree in T_G, which contains T as a proper subgraph. This contradicts the maximality of T. Hence, T is a spanning tree for the G-action on X. $\square$

4.2.2 Completing the proof

In the following, we use the letter "e" both for the neutral group element, and for edges in a graph; it will always be clear from the context which of the two is meant.

Proof of Theorem 4.2.1, part II. Let G be a group acting freely on a tree T by graph automorphisms. By Theorem 4.2.4 there exists a spanning tree T' for this action.

The idea is to think about the graph obtained from T by contracting T' and all its copies $g \cdot T'$ for $g \in G$ each to a single vertex (Figure 4.6 shows this in the situation of Example 4.2.3). Hence, the candidates for a generating set come from the edges joining these new vertices: An edge of T is called *essential* if it does not belong to T', but if one of the vertices of the edge in question belongs to T' (then the other vertex cannot belong to T' as well, by uniqueness of paths in trees (Proposition 3.1.9)).

As first step we construct a candidate $S \subset G$ for a free generating set of G: Let e be an essential edge of T, say $e = \{u, v\}$ with u a vertex of T' and v not a vertex of T'. Because T' is a spanning tree, there is an element $g_e \in G$ such that $g_e^{-1} \cdot v$ is a vertex of T'; equivalently, v is a vertex of $g_e \cdot T'$. Notice that the element g_e is uniquely determined by this property as the orbit $G \cdot v$ shares only a single vertex with T', and as G acts freely on T.

We define

$$\widetilde{S} := \{g_e \in G \mid e \text{ is an essential edge of } T\}.$$

This set $\widetilde{S}$ has the following properties:

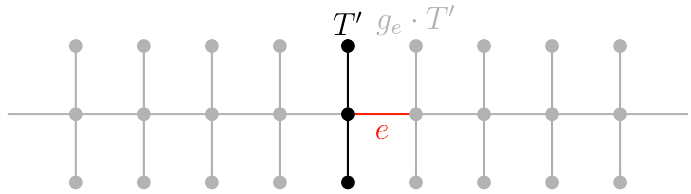

Figure 4.7.: An essential edge (red) for the shift action of $\mathbb{Z}$ (Example 4.2.3)

1. By definition, the neutral element is not contained in $\widetilde{S}$.
2. The set $\widetilde{S}$ does not contain an element of order 2 because G acts freely on a non-empty tree (and so cannot contain any non-trivial elements of finite order by Proposition 4.1.16).
3. If e and e' are essential edges with $g_e = g_{e'}$, then $e = e'$ (because T is a tree and therefore there cannot be two different edges connecting the connected subgraphs T' and $g_e \cdot T' = g_{e'} \cdot T'$).
4. If $g \in \widetilde{S}$, say $g = g_e$ for some essential edge e, then also $g^{-1} = g_{g^{-1} \cdot e}$ is in $\widetilde{S}$, because $g^{-1} \cdot e$ is easily seen to be an essential edge.

In particular, there is a subset $S \subset \widetilde{S}$ with

$$S \cap S^{-1} = \emptyset \qquad \text{and} \qquad |S| = \frac{|\widetilde{S}|}{2} = \frac{1}{2} \cdot \#\text{essential edges of } T.$$

The set $\widetilde{S}$ (and hence S) generates G: Let $g \in G$. We pick a vertex v of T'. Because T is connected, there is a path p in T connecting v and $g \cdot v$. The path p passes through several copies of T', say, $g_0 \cdot T', \dots, g_n \cdot T'$ of T' in this order, where $g_{j+1} \neq g_j$ for all $j \in \{0, \dots, n-1\}$, and $g_0 = e$, $g_n = g$ (Figure 4.8).

Let $j \in \{0, \dots, n-1\}$. Because T' is a spanning tree and $g_j \neq g_{j+1}$, the copies $g_j \cdot T'$ and $g_{j+1} \cdot T'$ are joined by an edge e_j. By definition, $g_j^{-1} \cdot e_j$ is an essential edge, and the corresponding group element

$$s_j := g_j^{-1} \cdot g_{j+1}$$

lies in $\widetilde{S}$. Therefore, we obtain that

$$\begin{aligned} g = g_n &= g_0^{-1} \cdot g_n \\ &= g_0^{-1} \cdot g_1 \cdot g_1^{-1} \cdot g_2 \cdot \dots \cdot g_{n-1}^{-1} \cdot g_n \\ &= s_0 \cdot \dots \cdot s_{n-1} \end{aligned}$$

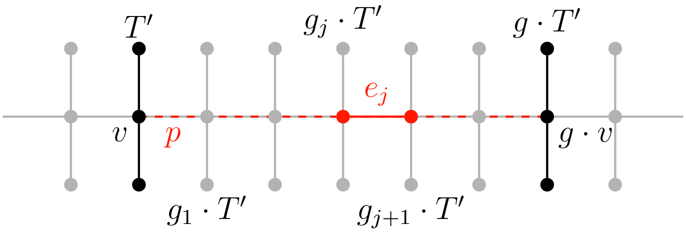

Figure 4.8.: The set $\widetilde{S}$ generates G

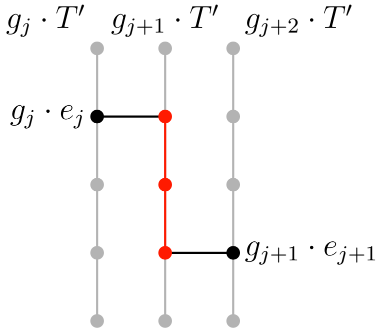

Figure 4.9.: Cycles in $\mathrm{Cay}(G, \widetilde{S})$ lead to cycles in T by connecting translates of essential edges in the corresponding translates of T' (red path)

is in the subgroup of G generated by $\widetilde{S}$. In other words, $\widetilde{S}$ is a generating set of G. (And we can view the graph obtained by collapsing each of the translates of T' in T to a vertex as the Cayley graph $\mathrm{Cay}(G, \widetilde{S})$).

The set S is a free generating set of G: In view of Theorem 3.3.3 it suffices to show that the Cayley graph $\mathrm{Cay}(G, S)$ does not contain any cycles. *Assume* for a contradiction that there is an $n \in \mathbb{N}_{\geq 3}$ and a cycle $g_0, \dots, g_{n-1}$ in $\mathrm{Cay}(G, S) = \mathrm{Cay}(G, \widetilde{S})$. By definition, the elements

$$\forall_{j \in \{0,\dots,n-2\}} \;\; s_{j+1} := g_j^{-1} \cdot g_{j+1}$$

and $s_n := g_{n-1}^{-1} \cdot g_0$ are in $\widetilde{S}$. For $j \in \{1, \dots, n\}$ let e_j be an essential edge joining T' and $s_j \cdot T'$.

Because each of the translates of T' is a connected subgraph, we can connect those vertices of the edges $g_j \cdot e_j$ and $g_j \cdot s_j \cdot e_{j+1} = g_{j+1} \cdot e_{j+1}$ that lie in $g_{j+1} \cdot T'$ by a path in $g_{j+1} \cdot T'$ (Figure 4.9). Using the fact that $g_0, \dots, g_{n-1}$ is a cycle in $\operatorname{Cay}(G, \widetilde{S})$, one sees that the resulting concatenation of paths is a cycle in T, which contradicts that T is a tree. □

Remark 4.2.5 (Topological proof)**.** Let G be a group acting freely on a tree; then G also acts freely, continuously, cellularly, and properly discontinuously on the CW-realisation X of this tree, which is contractible. Covering theory shows that the quotient space $G \setminus X$ is homeomorphic to a one-dimensional CW-complex and that $G \cong \pi_1(G \setminus X)$. The Seifert-van Kampen theorem then yields that G is a free group [64, Chapter VI].

Remark 4.2.6 (Characterising free products with amalgamation)**.** Similarly, also free products with amalgamations and HNN-extensions can be characterised geometrically by suitable actions on graphs [87]; generalisations of this type eventually lead to complexes of groups and Bass-Serre theory [87].

4.3 Application: Subgroups of free groups are free

The characterisation of free groups in terms of free actions on trees allows us to prove freeness of certain subgroups of suitable groups:

Corollary 4.3.1 (Nielsen-Schreier theorem)**.** *Subgroups of free groups are free.*

Proof. Let F be a free group, and let $G \subset F$ be a subgroup of F. Because F is free, the group F acts freely on a non-empty tree; hence, also the (sub)group G acts freely on this non-empty tree. Therefore, G is a free group by Theorem 4.2.1. □

Recall that the *index* of a subgroup $H \subset G$ of a group G is the number of cosets of H in G; we denote the index of H in G by $[G : H]$. For example, the subgroup $2 \cdot \mathbb{Z}$ of $\mathbb{Z}$ has index 2 in $\mathbb{Z}$.

Corollary 4.3.2 (Nielsen-Schreier theorem, quantitative version)**.** *Let F be a free group of rank $n \in \mathbb{N}$, and let $G \subset F$ be a subgroup of index $k \in \mathbb{N}$. Then G is a free group of rank $k \cdot (n-1) + 1$.*

In particular, finite index subgroups of free groups of finite rank are finitely generated.

Proof. Let S be a free generating set of F, and let $T := \mathrm{Cay}(F, S)$; so T is a tree and the left translation action of F on T is free. Therefore, also the left translation action of the subgroup G on T is free (and so G is free). Looking at the proof of Theorem 4.2.1 shows that the rank of G equals $E/2$, where E is the number of essential edges of the action of G on T.

We determine E by a counting argument: Let T' be a spanning tree of the action of G on T. From $[F : G] = k$ we deduce that T' has exactly k vertices. For a vertex v in T we denote by $d_T(v)$ the *degree of v in T*, i.e., the number of neighbours of v in T. Because T is a regular tree all of whose vertices have degree $2 \cdot |S| = 2 \cdot n$, we obtain (where $V(T')$ denotes the set of vertices of T')

$$2 \cdot n \cdot k = \sum_{v \in V(T')} d_T(v).$$

On the other hand, T' is a finite tree with k vertices and so has $k-1$ edges. Because the edges of T' are counted twice when summing up the degrees of the vertices of T', we obtain

$$2 \cdot n \cdot k = \sum_{v \in V(T')} d_T(v) = 2 \cdot (k-1) + E;$$

in other words, the G-action on T has $2 \cdot (k \cdot (n-1) + 1)$ essential edges, as desired. $\square$

Remark 4.3.3 (Topological proof of the Nielsen-Schreier theorem)**.** Another proof of the Nielsen-Schreier theorem can be given via covering theory [64, Chapter VI]: Let F be a free group of rank n; then F is the fundamental group of an n-fold bouquet X of circles. If G is a subgroup of F, we can look at the corresponding covering $\overline{X} \longrightarrow X$ of X. As X can be viewed as a one-dimensional CW-complex, also the covering space $\overline{X}$ is a one-dimensional CW-complex. On the other hand, any such space is homotopy equivalent to a bouquet of circles, and hence has free fundamental group. Because $\overline{X} \longrightarrow X$ is the covering corresponding to the subgroup G of F, it follows that $G \cong \pi_1(\overline{X})$ is a free group.

Taking into account that the Euler characteristic of finite CW-complexes is multiplicative with respect to finite coverings, one can also prove the quantitative version of the Nielsen-Schreier theorem via covering theory.

Corollary 4.3.4. *If F is a free group of rank at least 2, and $n \in \mathbb{N}$, then there is a subgroup of G that is free of finite rank at least n.*

Proof. Exercise. □

Remark 4.3.5 (Hanna Neumann conjecture)**.** Let F be a free group, and let G and H be subgroups of F. Hence, G and H are free. Suppose that the ranks m and n of G and H respectively are finite and non-zero. Then Hanna Neumann (Appendix A.2) conjectured that the rank r of $G \cap H$ (which as a subgroup of a free group again is free) satisfies

$$r - 1 \leq (m - 1) \cdot (n - 1).$$

This conjecture was proved to be correct by J. Friedman [35] and I. Mineyev [66, 67], in 2011.

Corollary 4.3.6. *Finite index subgroups of finitely generated groups are finitely generated.*

Proof. Let G be a finitely generated group, and let H be a finite index subgroup of G. If S is a finite generating set of G, then the universal property of the free group $F(S)$ freely generated by S provides us with a surjective homomorphism $\pi \colon F(S) \longrightarrow G$. Let H' be the preimage of H under π; so H' is a subgroup of $F(S)$, and a straightforward calculation shows that $[F(S) : H'] \leq [G : H]$.

By Corollary 4.3.2, the group H' is finitely generated; but then also the image $H = \pi(H')$ is finitely generated. □

We will later see an alternative proof of Corollary 4.3.6 via the Švarc-Milnor lemma (Corollary 5.3.10).

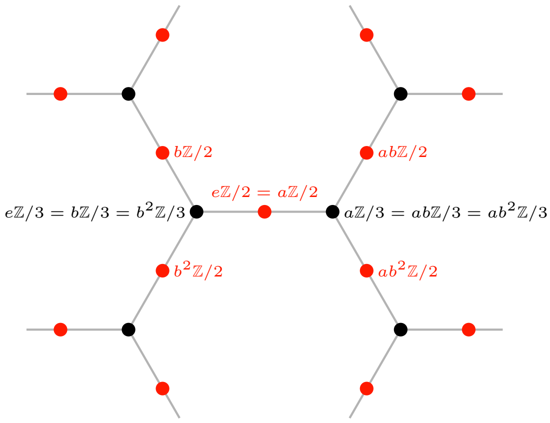

Figure 4.10.: The tree for the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z} \cong \langle a, b \,|\, a^2, b^3 \rangle$

Corollary 4.3.7 (Free subgroups of free products)**.** *Let G and H be finite groups. Then all torsion-free subgroups of the free product $G * H$ are free groups.*

Sketch of proof. Without loss of generality we may assume that G and H are non-trivial. We construct a tree on which the group $G * H$ acts with finite stabilisers:

Let X be the graph

- whose set of vertices is $V := \{x \cdot G \mid x \in G * H\} \cup \{x \cdot H \mid x \in G * H\}$ (where we view the vertices as subsets of $G * H$), and
- whose set of edges is

$$\bigl\{\{x \cdot G, x \cdot H\} \bigm| x \in G * H\bigr\}$$

(see Figure 4.10 for the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z}$). Using the description of the free product $G * H$ in terms of reduced words (Remark 3.3.8) one can show that the graph X is a tree.

The free product $G * H$ acts on the tree X by left translation, given on the vertices by

$$\begin{aligned} (G * H) \times V &\longrightarrow V \\ (y, x \cdot G) &\longmapsto (y \cdot x) \cdot G \\ (y, x \cdot H) &\longmapsto (y \cdot x) \cdot H. \end{aligned}$$

What are the stabilisers of this action? Let $x \in G * H$, and $y \in G * H$. Then y is in the stabiliser of $x \cdot G$ if and only if

$$x \cdot G = y \cdot (x \cdot G) = (y \cdot x) \cdot G,$$

which is equivalent to $y \in x \cdot G \cdot x^{-1}$. Analogously, y is in the stabiliser of the vertex $x \cdot H$ if and only if $y \in x \cdot H \cdot x^{-1}$. A similar computation shows that the stabiliser of an edge $\{x \cdot G, x \cdot H\}$ is $x \cdot G \cdot x^{-1} \cap x \cdot H \cdot x^{-1} = \{e\}$.

Because G and H are finite, all stabilisers of the above action of $G * H$ on the tree X are finite. Therefore, any torsion-free subgroup of $G * H$ acts freely on the tree X. Applying Theorem 4.2.1 finishes the proof. □

A similar technique as in the previous proof shows for all primes $p \in \mathbb{Z}$ that all torsion-free subgroups of $\mathrm{SL}(2, \mathbb{Q}_p)$ are free [87].

4.4 The ping-pong lemma

The following sufficient criterion for freeness via suitable actions is due to F. Klein:

Theorem 4.4.1 (Ping-pong lemma). *Let G be a group, generated by elements a and b. Suppose there is a G-action on a set X such that there are non-empty subsets A, $B \subset X$ with B not included in A and such that for all $n \in \mathbb{Z} \setminus \{0\}$ we have*

$$a^n \cdot B \subset A \qquad \text{and} \qquad b^n \cdot A \subset B.$$

Then G is free of rank 2, freely generated by $\{a, b\}$.

Proof. Let $\alpha \neq \beta$. It suffices to find an isomorphism $F_{\mathrm{red}}(\{\alpha, \beta\}) \cong G$ that maps $\{\alpha, \beta\}$ to $\{a, b\}$. By the universal property of the free group $F_{\mathrm{red}}(\{\alpha, \beta\})$ there is a group homomorphism $\varphi \colon F_{\mathrm{red}}(\{\alpha, \beta\}) \longrightarrow G$

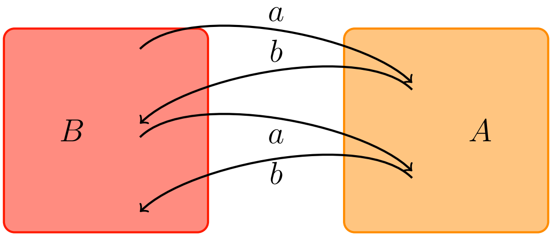

Figure 4.11.: The ping-pong lemma

mapping α to a and β to b. Because G is generated by $\{a, b\}$, the homomorphism φ is surjective.

Assume for a contradiction that φ is not injective; hence, there is a reduced word $w \in F_{\text{red}}(\{\alpha, \beta\}) \setminus \{\varepsilon\}$ with $\varphi(w) = e$. Depending on the first and last letter of w there are four cases:

1. The word w starts and ends with a (non-trivial) power of α, i.e., we can write $w = \alpha^{n_0}\beta^{m_1}\alpha^{n_1}\dots\beta^{m_k}\alpha^{n_k}$ for some $k \in \mathbb{N}$ and certain $n_0, \dots, n_k, m_1, \dots, m_k \in \mathbb{Z} \setminus \{0\}$. Then (ping-pong! – see Figure 4.11)

$$\begin{aligned} B = e \cdot B = \varphi(w) \cdot B & \\ = a^{n_0} \cdot b^{m_1} \cdot a^{n_1} \cdot \dots \cdot b^{m_k} \cdot a^{n_k} \cdot B & \\ \subset a^{n_0} \cdot b^{m_1} \cdot a^{n_1} \cdot \dots \cdot b^{m_k} \cdot A & \qquad \text{ping!} \\ \subset a^{n_0} \cdot b^{m_1} \cdot a^{n_1} \cdot \dots \cdot a^{n_{k-1}} \cdot B & \qquad \text{pong!} \\ \subset \cdots & \qquad \vdots \\ \subset a^{n_0} \cdot B & \\ \subset A, & \end{aligned}$$

which contradicts the assumption that B is not contained in A.

2. The word w starts and ends with non-trivial powers of β. Then $\alpha w \alpha^{-1}$ is a reduced word starting and ending in non-trivial powers of α. So

$$e = \varphi(\alpha) \cdot e \cdot \varphi(\alpha)^{-1} = \varphi(\alpha) \cdot \varphi(w) \cdot \varphi(\alpha^{-1}) = \varphi(\alpha w \alpha^{-1}),$$

contradicting what we already proved for the first case.

3. The word w starts with a non-trivial power of α and ends with a non-trivial power of β, say $w = \alpha^n w' \beta^m$ with $n, m \in \mathbb{Z} \setminus \{0\}$ and w' a reduced word not starting with a non-trivial power of α and not ending in a non-trivial power of β. Let $r \in \mathbb{Z} \setminus \{0, -n\}$. Then $\alpha^r w \alpha^{-r} = \alpha^{r+n} w' \beta^m \alpha^r$ starts and ends with a non-trivial power of α and
$$e = \varphi(\alpha^r w \alpha^{-r}),$$
contradicting what we already proved for the first case.
4. The word w starts with a non-trivial power of β and ends with a non-trivial power of α. Then the inverse of w falls into the third case and $\varphi(w^{-1}) = e$, which cannot be.

Therefore, φ is injective, and so $\varphi\colon F_{\mathrm{red}}(\{\alpha, \beta\}) \longrightarrow G$ is an isomorphism with $\varphi(\{\alpha, \beta\}) = \{a, b\}$, as was to be shown. □

Remark 4.4.2 (Ping-pong lemma for free products)**.** Similarly, using the description of free products in terms of reduced words (Remark 3.3.8), one can show the following [42, II.24]: Let G be a group, let G_1 and G_2 be two subgroups of G with $|G_1| \geq 3$ and $|G_2| \geq 2$, and suppose that G is generated by the union $G_1 \cup G_2$. If there is a G-action on a set X such that there are non-empty subsets $X_1, X_2 \subset X$ with X_2 not included in X_1 and such that
$$\forall_{g \in G_1 \setminus \{e\}} \;\; g \cdot X_2 \subset X_1 \qquad \text{and} \qquad \forall_{g \in G_2 \setminus \{e\}} \;\; g \cdot X_1 \subset X_2,$$
then $G \cong G_1 * G_2$.

4.5 Application: Free subgroups of matrix groups

Using the ping-pong lemma we can establish that certain matrix groups are free.

Example 4.5.1 (Free subgroups of $\mathrm{SL}(2,\mathbb{Z})$)**.** We consider the matrices a, $b \in \mathrm{SL}(2,\mathbb{Z})$ given by

$$a := \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad b := \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}.$$

We show that the subgroup of $\mathrm{SL}(2,\mathbb{Z})$ generated by $\{a,b\}$ is a free group of rank 2 (freely generated by $\{a,b\}$) via the ping-pong lemma:

The matrix group $\mathrm{SL}(2,\mathbb{Z})$ acts on $\mathbb{R}^2$ by matrix multiplication. We now consider the subsets

$$A := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2 \;\middle|\; |x| > |y| \right\} \quad \text{and} \quad B := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^2 \;\middle|\; |x| < |y| \right\}$$

of $\mathbb{R}^2$. Then A and B are non-empty and B is not contained in A. Moreover, for all $n \in \mathbb{Z} \setminus \{0\}$ and all $(x,y) \in B$ we have

$$a^n \cdot \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 1 & 2 \cdot n \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + 2 \cdot n \cdot y \\ y \end{pmatrix}$$

and

$$\begin{aligned} |x + 2 \cdot n \cdot y| &\geq |2 \cdot n \cdot y| - |x| \\ &\geq 2 \cdot |y| - |x| \\ &> 2 \cdot |y| - |y| \\ &= |y|; \end{aligned}$$

so $a^n \cdot B \subset A$. Similarly, we see that $b^n \cdot A \subset B$ for all $n \in \mathbb{Z} \setminus \{0\}$. Thus, we can apply the ping-pong lemma and deduce that the subgroup of $\mathrm{SL}(2,\mathbb{Z})$ generated by $\{a,b\}$ is freely generated by $\{a,b\}$. Notice that it would be rather awkward to prove this by hand, using only matrix calculations. A more careful analysis shows that this subgroup has index 12 in $\mathrm{SL}(2,\mathbb{Z})$ (exercise).

This example is related to the following fact:

Remark 4.5.2 (Tits alternative)**.** J. Tits [93] proved: For all fields K of characteristic 0 and all $n \in \mathbb{N}_{\geq 1}$ the following holds: If G is a subgroup of $\mathrm{GL}(n,K)$, then either G contains a free subgroup of rank 2 or G contains a finite index subgroup that is solvable.

Further examples are given in de la Harpe's book [42, Chapter II.B]; in particular, it can be shown that the group of homeomorphisms $\mathbb{R} \longrightarrow \mathbb{R}$ contains a free group of rank 2.

5

Groups $\longrightarrow$ geometry, III: Quasi-isometry

As explained in the introduction, one of the objectives of geometric group theory is to view groups as geo*metric* objects. We now add the metric layer to the combinatorics given by Cayley graphs:

If G is a group and S is a generating set of G, then the paths on the associated Cayley graph $\operatorname{Cay}(G,S)$ induce a metric on G, the so-called word metric with respect to the generating set S; unfortunately, in general, this metric depends on the chosen generating set.

In order to obtain a notion of distance on a group independent of a given generating set we pass to large scale geometry. Using the language of quasi-geometry, we arrive at such a notion for finitely generated groups – the so-called quasi-isometry type, which is one of the central objects of geometric group theory.

We start with some generalities on isometries, bilipschitz equivalences and quasi-isometries. As next step, we specialise to the case of finitely generated groups. The key to linking the geometry of groups to actual geometry is the Švarc-Milnor lemma (Section 5.3). Moreover, we discuss Gromov's dynamic criterion for quasi-isometry (Section 5.4) and give an outlook on geometric properties of groups and quasi-isometry invariants (Section 5.5).

5.1 Quasi-isometry types of metric spaces

In the following, we introduce different levels of similarity between metric spaces: isometries, bilipschitz equivalences and quasi-isometries. Intuitively, we want a large scale geometric notion of similarity – i.e., we want metric spaces to be equivalent if they seem to be the same when looked at from far away. A category theoretic framework will be explained in Remark 5.1.12. A guiding example to keep in mind is that we want the real line and the integers (with the induced metric from the real line) to be equivalent.

For the sake of completeness, we recall the definition of a metric space:

Definition 5.1.1 (Metric space)**.** A *metric space* is a set X together with a map $d\colon X \times X \longrightarrow \mathbb{R}_{\geq 0}$ satisfying the following conditions:

- For all $x, y \in X$ we have $d(x,y) = 0$ if and only if $x = y$.
- For all $x, y \in X$ we have $d(x,y) = d(y,x)$.
- For all $x, y, z \in X$ the *triangle inequality* holds:

$$d(x,z) \leq d(x,y) + d(y,z).$$

We start with the strongest type of similarity between metric spaces:

Definition 5.1.2 (Isometry)**.** Let $f\colon X \longrightarrow Y$ be a map between metric spaces (X, d_X) and (Y, d_Y).

- We say that f is an *isometric embedding* if

$$\forall_{x,x' \in X} \;\; d_Y\bigl(f(x), f(x')\bigr) = d_X(x,x').$$

- The map f is an *isometry* if it is an isometric embedding and if there is an isometric embedding $g\colon Y \longrightarrow X$ such that

$$f \circ g = \mathrm{id}_Y \quad \text{and} \quad g \circ f = \mathrm{id}_X\,.$$

– Two metric spaces are *isometric* if there exists an isometry between them.

Remark 5.1.3. Clearly, every isometric embedding is injective, and every isometry is a homeomorphism with respect to the topologies induced by the metrics. Moreover, an isometric embedding is an isometry if and only if it is bijective.

The notion of isometry is very rigid – too rigid for our purposes. We want a notion of "similarity" for metric spaces that only reflects the large scale shape of the space, but not the local details. A first step is to relax the isometry condition by allowing for a uniform multiplicative error:

Definition 5.1.4 (Bilipschitz equivalence)**.** Let $f\colon X \longrightarrow Y$ be a map between metric spaces (X, d_X) and (Y, d_Y).

– We say that f is a *bilipschitz embedding* if there is a constant $c \in \mathbb{R}_{>0}$ such that

$$\forall_{x,x' \in X} \quad \frac{1}{c} \cdot d_X(x,x') \leq d_Y\bigl(f(x), f(x')\bigr) \leq c \cdot d_X(x,x').$$

– The map f is a *bilipschitz equivalence* if it is a bilipschitz embedding and if there is a bilipschitz embedding $g\colon Y \longrightarrow X$ such that

$$f \circ g = \mathrm{id}_Y \quad \text{and} \quad g \circ f = \mathrm{id}_X \,.$$

– Two metric spaces are called *bilipschitz equivalent* if there exists a bilipschitz equivalence between them.

Remark 5.1.5. Clearly, every bilipschitz embedding is injective, and every bilipschitz equivalence is a homeomorphism with respect to the topologies induced by the metrics. Moreover, a bilipschitz embedding is a bilipschitz equivalence if and only if it is bijective.

Hence, also bilipschitz equivalences preserve the local information; so bilipschitz equivalences still remember too much detail for our purposes. As next – and final – step, we allow for a uniform additive error:

Definition 5.1.6 (Quasi-isometry)**.** Let $f\colon X \longrightarrow Y$ be a map between metric spaces (X, d_X) and (Y, d_Y).

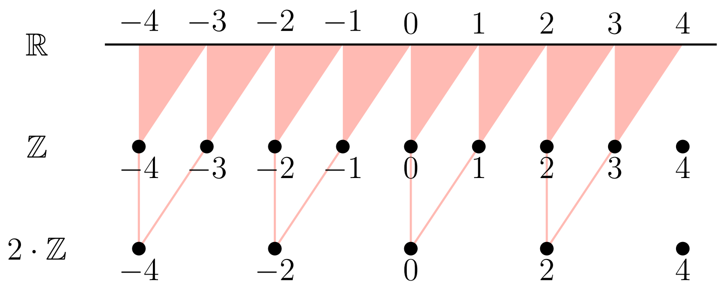

Figure 5.1.: The metric spaces $\mathbb{R}$, $\mathbb{Z}$, and $2 \cdot \mathbb{Z}$ are quasi-isometric

- The map f is a *quasi-isometric embedding* if there are constants $c \in \mathbb{R}_{>0}$ and $b \in \mathbb{R}_{>0}$ such that

$$\forall_{x,x' \in X} \ \frac{1}{c} \cdot d_X(x,x') - b \leq d_Y\bigl(f(x), f(x')\bigr) \leq c \cdot d_X(x,x') + b.$$

- A map $f' \colon X \longrightarrow Y$ has *finite distance from* f if there is a constant $c \in \mathbb{R}_{\geq 0}$ with

$$\forall_{x \in X} \ d_X\bigl(f(x), f'(x)\bigr) \leq c.$$

- The map f is a *quasi-isometry* if it is a quasi-isometric embedding for which there is a quasi-inverse quasi-isometric embedding, i.e., if there is a quasi-isometric embedding $g \colon Y \longrightarrow X$ such that $g \circ f$ has finite distance from id_Y and $f \circ g$ has finite distance from id_X.
- Two metric spaces X and Y are *quasi-isometric* if there exists a quasi-isometry $X \longrightarrow Y$; in this case, we write $X \sim_{\mathrm{QI}} Y$.

Example 5.1.7 (Isometries, bilipschitz equivalences and quasi-isometries)**.** Clearly, every isometry is a bilipschitz equivalence, and every bilipschitz equivalence is a quasi-isometry. In general, the converse does not hold:

We consider $\mathbb{R}$ as a metric space with respect to the distance function given by the absolute value of the difference of two numbers; moreover, we consider the subsets $\mathbb{Z} \subset \mathbb{R}$ and $2 \cdot \mathbb{Z} \subset \mathbb{R}$ with respect to the induced metrics (Figure 5.1).

Clearly, the inclusions $2 \cdot \mathbb{Z} \hookrightarrow \mathbb{Z}$ and $\mathbb{Z} \hookrightarrow \mathbb{R}$ are quasi-isometric embeddings but no bilipschitz equivalences (as they are not bijective). Moreover, the maps

$$\begin{aligned}
\mathbb{R} &\longrightarrow \mathbb{Z} \\
x &\longmapsto \lfloor x \rfloor, \\
\mathbb{Z} &\longrightarrow 2 \cdot \mathbb{Z} \\
n &\longmapsto \begin{cases} n & \text{if } n \in 2 \cdot \mathbb{Z}, \\ n+1 & \text{if } n \notin 2 \cdot \mathbb{Z} \end{cases}
\end{aligned}$$

are quasi-isometric embeddings that are quasi-inverse to the inclusions (here, $\lfloor x \rfloor$ denotes the integral part of x, i.e., the largest integer that is not larger than x).

The spaces $\mathbb{Z}$ and $2 \cdot \mathbb{Z}$ are bilipschitz equivalent (via the map given by multiplication by 2). However, $\mathbb{Z}$ and $2 \cdot \mathbb{Z}$ are not isometric – in $\mathbb{Z}$ there are points having distance 1, whereas in $2 \cdot \mathbb{Z}$ the minimal distance between two different points is 2.

Finally, because $\mathbb{R}$ is uncountable but $\mathbb{Z}$ and $2 \cdot \mathbb{Z}$ are countable, the metric space $\mathbb{R}$ cannot be isometric or bilipschitz equivalent to $\mathbb{Z}$ or $2 \cdot \mathbb{Z}$.

Caveat 5.1.8. In particular, we see that

- quasi-isometries in general are neither injective, nor surjective,
- quasi-isometries in general are not continuous at all,
- in general there is no isometry at finite distance of any given quasi-isometry,
- in general quasi-isometries do not preserve dimension locally.

Example 5.1.9 (More (non-)quasi-isometric spaces)**.**

- All non-empty metric spaces of finite diameter are quasi-isometric; the *diameter* of a metric space (X, d) is

$$\operatorname{diam} X := \sup_{x,y \in X} d(x, y).$$

- Conversely, if a space is quasi-isometric to a space of finite diameter, then it has finite diameter as well. So the metric space $\mathbb{Z}$ (with the metric induced from $\mathbb{R}$) is *not* quasi-isometric to a non-empty metric space of finite diameter.
- The metric spaces $\mathbb{R}$ and $\mathbb{R}^2$ (with respect to the Euclidean metric) are *not* quasi-isometric (exercise).

Exercise 5.1.1.

1. Are the metric spaces $\mathbb{Z}$ and $\mathbb{N}$ with the metrics induced from the standard metric on $\mathbb{R}$ quasi-isometric?
2. Are the metric spaces $\mathbb{Z}$ and $\{n^3 \mid n \in \mathbb{Z}\}$ with the metrics induced from the standard metric on $\mathbb{R}$ quasi-isometric?

Proposition 5.1.10 (Alternative characterisation of quasi-isometries)**.** *A map $f\colon X \longrightarrow Y$ between metric spaces (X, d_X) and (Y, d_Y) is a quasi-isometry if and only if it is a quasi-isometric embedding with quasi-dense image; a map $f\colon X \longrightarrow Y$ is said to have* quasi-dense image *if there is a constant $c \in \mathbb{R}_{>0}$ such that*

$$\forall_{y\in Y}\ \exists_{x\in X}\ d_Y\bigl(f(x), y\bigr) \leq c.$$

Proof. If $f\colon X \longrightarrow Y$ is a quasi-isometry, then, by definition, there exists a quasi-inverse quasi-isometric embedding $g\colon Y \longrightarrow X$. Hence, there is a $c \in \mathbb{R}_{>0}$ such that

$$\forall_{y\in Y}\ d_Y\bigl(f\circ g(y), y\bigr) \leq c;$$

in particular, f has quasi-dense image.

Conversely, suppose that $f\colon X \longrightarrow Y$ is a quasi-isometric embedding with quasi-dense image. Using the axiom of choice, we construct a quasi-inverse quasi-isometric embedding:

Because f is a quasi-isometric embedding with quasi-dense image, there is a constant $c \in \mathbb{R}_{>0}$ such that

$$\begin{aligned}\forall_{x,x'\in X}\ \frac{1}{c}\cdot d_X(x,x') - c \leq d_Y\bigl(f(x), f(x')\bigr) &\leq c\cdot d_X(x,x') + c,\\ \forall_{y\in Y}\ \exists_{x\in X}\ d_Y\bigl(f(x), y\bigr) &\leq c.\end{aligned}$$

By the axiom of choice, we can define a map

$$\begin{aligned} g\colon Y &\longrightarrow X\\ y &\longmapsto x_y,\end{aligned}$$

such that $d_Y(f(x_y), y) \leq c$ holds for all $y \in Y$.

The map g is quasi-inverse to f, because: By construction, for all $y \in Y$ we have

$$d_Y\bigl(f\circ g(y), y\bigr) = d_Y(f(x_y), y) \leq c;$$

conversely, for all $x \in X$ we obtain (using the fact that f is a quasi-isometric embedding)

$$d_X\bigl(g\circ f(x), x\bigr) = d_X(x_{f(x)}, x) \leq c\cdot d_Y\bigl(f(x_{f(x)}), f(x)\bigr)+c^2 \leq c\cdot c+c^2 = 2\cdot c^2.$$

So $f \circ g$ and $g \circ f$ have finite distance from the respective identity maps.

Moreover, g also is a quasi-isometric embedding, because: Let $y, y' \in Y$. Then

$$\begin{aligned}
d_X\bigl(g(y), g(y')\bigr) &= d_X(x_y, x_{y'}) \\
&\leq c \cdot d_Y\bigl(f(x_y), f(x_{y'})\bigr) + c^2 \\
&\leq c \cdot \Bigl(d_Y\bigl(f(x_y), y\bigr) + d_Y(y, y') + d_Y\bigl(f(x_{y'}, y')\bigr)\Bigr) + c^2 \\
&\leq c \cdot \bigl(d_Y(y, y') + 2 \cdot c\bigr) + c^2 \\
&= c \cdot d_Y(y, y') + 3 \cdot c^2,
\end{aligned}$$

and

$$\begin{aligned}
d_X\bigl(g(y), g(y')\bigr) &= d_X(x_y, x_{y'}) \\
&\geq \frac{1}{c} \cdot d_Y\bigl(f(x_y), f(x_{y'})\bigr) - 1 \\
&\geq \frac{1}{c} \cdot \Bigl(d_Y(y, y') - d_Y\bigl(f(x_y), y\bigr) - d_Y\bigl(f(x_{y'}), y'\bigr)\Bigr) - 1 \\
&\geq \frac{1}{c} \cdot d_Y(y, y') - \frac{2 \cdot c}{c} - 1.
\end{aligned}$$

(Clearly, the same argument shows that quasi-inverses of quasi-isometric embeddings are quasi-isometric embeddings). $\square$

When working with quasi-isometries, the following inheritance properties can be useful:

Proposition 5.1.11 (Inheritance properties of quasi-isometric embeddings)**.**

1. *Any map at finite distance of a quasi-isometric embedding is a quasi-isometric embedding.*
2. *Any map at finite distance of a quasi-isometry is a quasi-isometry.*
3. *Let X, Y, Z be metric spaces and let $f, f'\colon X \longrightarrow Y$ be maps that have finite distance from each other.*

   a) *If* $g\colon Z \longrightarrow X$ *is a map, then* $f \circ g$ *and* $f' \circ g$ *have finite distance from each other.*
   b) *If* $g\colon Y \longrightarrow Z$ *is a quasi-isometric embedding, then also* $g \circ f$ *and* $g \circ f'$ *have finite distance from each other.*
4. *Compositions of quasi-isometric [bilipschitz] embeddings are quasi-isometric [bilipschitz] embeddings.*
5. *Compositions of quasi-isometries [bilipschitz equivalences] are quasi-isometries [bilipschitz equivalences].*

Proof. All these properties follow via simple calculations from the respective definitions (exercise). □

In particular, we obtain the following more conceptual description of isometries, bilipschitz equivalences, and quasi-isometries:

Remark 5.1.12 (A category theoretic framework for quasi-isometry)**.** Let $\mathsf{Met}_{\mathsf{isom}}$ be the category whose objects are metric spaces, whose morphisms are isometric embeddings and where the composition is given by the ordinary composition of maps. Then isometries of metric spaces correspond to isomorphisms in the category $\mathsf{Met}_{\mathsf{isom}}$.

Let $\mathsf{Met}_{\mathsf{bilip}}$ be the category whose objects are metric spaces, whose morphisms are bilipschitz embeddings and where the composition is given by the ordinary composition of maps. Then bilipschitz equivalences of metric spaces correspond to isomorphisms in the category $\mathsf{Met}_{\mathsf{bilip}}$.

Let QMet' be the category whose objects are metric spaces, whose morphisms are quasi-isometric embeddings and where the composition is given by the ordinary composition of maps. For metric spaces X, Y the relation "having finite distance from" is an equivalence relation on $\mathrm{Mor}_{\mathsf{QMet'}}(X,Y)$ and this equivalence relation is compatible with composition (Proposition 5.1.11). Hence, we can define the corresponding homotopy category QMet as follows:
- Objects in QMet are metric spaces.
- For metric spaces X and Y, the set of morphisms from X to Y in QMet is given by

$$\mathrm{Mor}_{\mathsf{QMet}}(X,Y) := \mathrm{Mor}_{\mathsf{QMet'}}(X,Y) \bigm/ \text{finite distance}.$$

- For metric spaces X, Y, Z, the composition of morphisms in QMet is given by

$$\begin{aligned} \operatorname{Mor}_{\mathsf{QMet}}(Z,Y) \times \operatorname{Mor}_{\mathsf{QMet}}(X,Y) &\longrightarrow \operatorname{Mor}_{\mathsf{QMet}}(X,Z) \\ \bigl([g],[f]\bigr) &\longmapsto [g \circ f] \end{aligned}$$

Then quasi-isometries of metric spaces correspond to isomorphisms in the category QMet.

As quasi-isometries are not bijective in general, some care has to be taken when defining quasi-isometry groups of metric spaces; however, looking at the category QMet gives us a natural definition of quasi-isometry groups:

Definition 5.1.13 (Quasi-isometry group)**.** Let X be a metric space. Then the *quasi-isometry group of* X is defined by

$$\operatorname{QI}(X) := \operatorname{Aut}_{\mathsf{QMet}}(X),$$

i.e., the group of quasi-isometries $X \longrightarrow X$ modulo finite distance.

For example, the category theoretic framework immediately yields that quasi-isometric metric spaces have isomorphic quasi-isometry groups.

Having the notion of a quasi-isometry group of a metric space also allows to define what an *action of a group by quasi-isometries on a metric space* is – namely, a homomorphism from the group in question to the quasi-isometry group of the given space.

Example 5.1.14 (Quasi-isometry groups)**.**

- Clearly, the quasi-isometry group of a metric space of finite diameter is trivial.
- The quasi-isometry group of $\mathbb{Z}$ is huge; for example, it contains the multiplicative group $\mathbb{R} \setminus \{0\}$ as a subgroup via the injective homomorphism

$$\begin{aligned} \mathbb{R} \setminus \{0\} &\longrightarrow \operatorname{QI}(\mathbb{Z}) \\ \alpha &\longmapsto \bigl[n \mapsto \lfloor \alpha \cdot n \rfloor\bigr] \end{aligned}$$

together with many rather large and non-commutative groups [84].

5.2 Quasi-isometry types of groups

Any generating set of a group yields a metric on the group in question by looking at the lengths of paths in the corresponding Cayley graph. The large scale geometric notion of quasi-isometry then allows us to associate geometric types to finitely generated groups that do not depend on the choice of finite generating sets.

Definition 5.2.1 (Metric on a graph)**.** Let $G = (V, E)$ be a connected graph. Then the map

$$\begin{aligned} V \times V &\longrightarrow \mathbb{R}_{\geq 0} \\ (v, w) &\longmapsto \min\{n \in \mathbb{N} \mid \text{there is a path of length } n \\ &\qquad\qquad \text{connecting } v \text{ and } w \text{ in } G\} \end{aligned}$$

is a metric on V, the *metric on V associated with G*.

Remark 5.2.2. A map between the sets of vertices of graphs is an isometry with respect to the associated metrics if and only if the map is an isomorphism of graphs.

Definition 5.2.3 (Word metric)**.** Let G be a group and let $S \subset G$ be a generating set. The *word metric d_S on G with respect to S* is the metric on G associated with the Cayley graph $\operatorname{Cay}(G, S)$. In other words,

$$d_S(g, h) = \min\{n \in \mathbb{N} \mid \exists_{s_1, \dots, s_n \in S \cup S^{-1}} \;\; g^{-1} \cdot h = s_1 \cdot \dots \cdot s_n\}$$

for all g, $h \in G$.

Example 5.2.4 (Word metrics on $\mathbb{Z}$)**.** The word metric on $\mathbb{Z}$ corresponding to the generating set $\{1\}$ coincides with the metric on $\mathbb{Z}$ induced from the standard metric on $\mathbb{R}$. On the other hand, in the word metric on $\mathbb{Z}$ corresponding to the generating set $\mathbb{Z}$, all group elements have distance 1 from every other group element.

In general, word metrics on a given group do depend on the chosen set of generators. However, the difference is negligible when looking at the group from far away:

Proposition 5.2.5. *Let G be a finitely generated group, and let S and S' be finite generating sets of G.*

1. *Then the identity map id_G is a bilipschitz equivalence between (G, d_S) and $(G, d_{S'})$.*
2. *In particular, any metric space (X, d) that is bilipschitz equivalent (or quasi-isometric) to (G, d_S) is also bilipschitz equivalent (or quasi-isometric respectively) to $(G, d_{S'})$ (via the same maps).*

Proof. The second part directly follows from the first part because the composition of bilipschitz equivalences is a bilipschitz equivalence, and the composition of quasi-isometries is a quasi-isometry (Proposition 5.1.11).

Thus it remains to prove the first part: Because S is finite,

$$c := \max_{s \in S \cup S^{-1}} d_{S'}(e, s)$$

is finite. Let g, $h \in G$ and let $n := d_S(g, h)$. Then we can write $g^{-1} \cdot h = s_1 \cdot \dots \cdot s_n$ for certain $s_1, \dots, s_n \in S \cup S^{-1}$. Using the triangle inequality and the fact that the metric $d_{S'}$ is left-invariant by definition, we obtain

$$\begin{aligned}
d_{S'}(g, h) &= d_{S'}(g, g \cdot s_1 \cdot \dots \cdot s_n) \\
&\leq d_{S'}(g, g \cdot s_1) + d_{S'}(g \cdot s_1, g \cdot s_1 \cdot s_2) + \dots \\
&\quad + d_{S'}(g \cdot s_1 \cdot \dots \cdot s_{n-1}, g \cdot s_1 \cdot \dots \cdot s_n) \\
&= d_{S'}(e, s_1) + d_{S'}(e, s_2) + \dots + d_{S'}(e, s_n) \\
&\leq c \cdot n \\
&= c \cdot d_S(g, h).
\end{aligned}$$

Interchanging the roles of S and S' shows that also a similar estimate holds in the other direction and hence that $\mathrm{id}_G \colon (G, d_S) \longrightarrow (G, d_{S'})$ is a bilipschitz equivalence. $\square$

This proposition gives a precise meaning to the statement that Cayley graphs of a finitely generated group with respect to different finite generating sets seem to be the same when looked at from far away (see Figure 5.2 for an example for the additive group $\mathbb{Z}$).

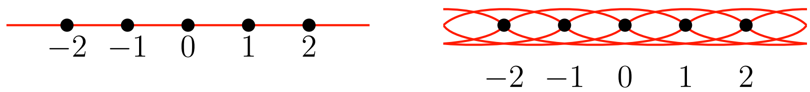

$\mathrm{Cay}(\mathbb{Z}, \{1\})$ $\mathrm{Cay}(\mathbb{Z}, \{2, 3\})$

Figure 5.2.: Cayley graphs of $\mathbb{Z}$ having the same large scale geometry

For infinite generating sets the first part of the above proposition does *not* hold in general; for example taking $\mathbb{Z}$ as a generating set for $\mathbb{Z}$ leads to the space $(\mathbb{Z}, d_{\mathbb{Z}})$ of finite diameter, while $(\mathbb{Z}, d_{\{1\}})$ does *not* have finite diameter (Example 5.2.4).

Definition 5.2.6 (Quasi-isometry type of finitely generated groups)**.** Let G be a finitely generated group.

- The group G is *bilipschitz equivalent* to a metric space X if for some (and hence all) finite generating set S of G the metric spaces (G, d_S) and X are bilipschitz equivalent.
- The group G is *quasi-isometric* to a metric space X if for some (and hence all) finite generating set S of G the metric spaces (G, d_S) and X are quasi-isometric. We write $G \sim_{\mathrm{QI}} X$ if G and X are quasi-isometric.

Analogously we define when two finitely generated groups are called bilipschitz equivalent or quasi-isometric.

Example 5.2.7. If $n \in \mathbb{N}$, then the group $\mathbb{Z}^n$ is quasi-isometric to Euclidean space $\mathbb{R}^n$ because the inclusion $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$ is a quasi-isometric embedding with quasi-dense image. In this sense, Cayley graphs of $\mathbb{Z}^n$ (with respect to finite generating sets) resemble the geometry of $\mathbb{R}^n$ (Figure 5.3).

At this point it might be more natural to consider bilipschitz equivalence of groups as a good geometric equivalence of finitely generated groups. However, later we will see why considering quasi-isometry types of groups is more appropriate: For instance, there is no suitable analogue of the Švarc-Milnor lemma for bilipschitz equivalence (Section 5.3).

The question of how quasi-isometry and bilipschitz equivalence are related for finitely generated groups leads to interesting problems and useful applications. A first step towards an answer is the following:

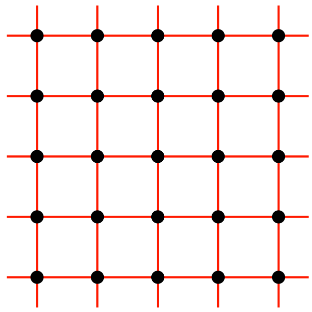

Figure 5.3.: The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ resembles the geometry of the Euclidean plane $\mathbb{R}^2$

Exercise 5.2.1 (Quasi-isometry vs. bilipschitz equivalence)**.**

1. Show that bijective quasi-isometries between finitely generated groups (with respect to the word metric of certain finite generating sets) are bilipschitz equivalences.
2. Does this also hold in general? I.e., are all bijective quasi-isometries between general metric spaces necessarily bilipschitz equivalences?

However, not all infinite fintely generated groups that are quasi-isometric are bilipschitz equivalent. We will study in Section **??** which quasi-isometric groups are bilipschitz equivalent; in particular, we will see then under which conditions free products of quasi-isometric groups lead to quasi-isometric groups.

5.2.1 First examples

As a simple example, we start with the quasi-isometry classification of finite groups:

Remark 5.2.8 (Properness of word metrics)**.** Let G be a group and let $S \subset G$ be a generating set. Then S is finite if and only if the word metric d_S on G is *proper* in the sense that all balls of finite radius in (G, d_S) are finite:

If S is infinite, then the ball of radius 1 around the neutral element of G contains $|S|$ elements, which is infinite. Conversely, if S is finite, then every

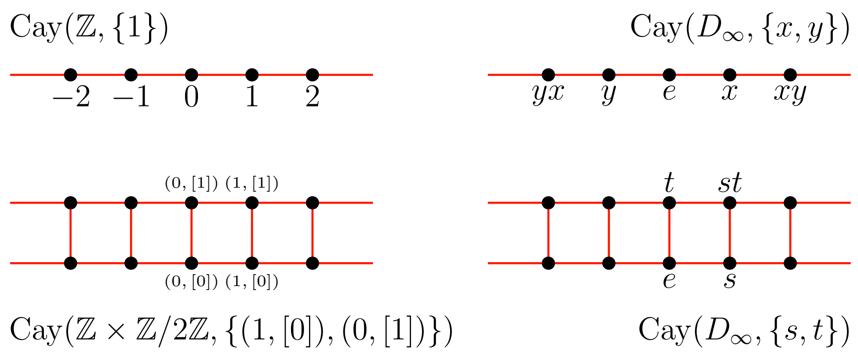

Figure 5.4.: The groups $\mathbb{Z}$, $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ and D_∞ are quasi-isometric

ball B of finite radius n around the neutral element contains only finitely many elements, because the set $(S \cup S^{-1})^n$ is finite and there is a surjective map $(S \cup S^{-1})^n \longrightarrow B$; because the metric d_S is invariant under the left translation action of G, it follows that all balls in (G, d_S) of finite radius are finite.

Example 5.2.9 (Quasi-isometry classification of finite groups)**.** A finitely generated group is quasi-isometric to a finite group if and only if it is finite: All finite groups lead to metric spaces of finite diameter and so all are quasi-isometric. Conversely, if a group is quasi-isometric to a finite group, then it has finite diameter with respect to some word metric; because balls of finite radius with respect to word metrics of finite generating sets are finite (Remark 5.2.8), it follows that the group in question has to be finite.

Notice that two finite groups are bilipschitz equivalent if and only if they have the same number of elements.

This explains why we drew the class of finite groups as a separate small spot of the universe of groups (Figure 1.2).

The next step is to look at groups (not) quasi-isometric to $\mathbb{Z}$:

Example 5.2.10 (Some groups quasi-isometric to $\mathbb{Z}$)**.** The groups $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ and D_∞ are bilipschitz equivalent and so in particular quasi-isometric (see Figure 5.4):

To this end we consider the following two presentations of the infinite dihedral group D_∞ by generators and relations:

$$\langle x, y \mid x^2, y^2 \rangle \cong D_\infty \cong \langle s, t \mid t^2, tst^{-1} = s^{-1} \rangle.$$

The Cayley graph $\operatorname{Cay}(D_\infty, \{x, y\})$ is isomorphic to $\operatorname{Cay}(\mathbb{Z}, \{1\})$; in particular, D_∞ and $\mathbb{Z}$ are bilipschitz equivalent. On the other hand, the Cayley graph $\operatorname{Cay}(D_\infty, s, t)$ is isomorphic to $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}, \{(1, [0]), (0, [1])\})$; in particular, D_∞ and $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ are bilipschitz equivalent. Because the word metrics on D_∞ corresponding to the generating sets $\{x, y\}$ and $\{s, t\}$ are bilipschitz equivalent, it follows that also $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ are bilipschitz equivalent.

Caveat 5.2.11 (isometry classification of finitely generated Abelian groups)**.** Even though $\mathbb{Z}$ and D_∞ as well as D_∞ and $\mathbb{Z} \times \mathbb{Z}/2$ admit finite generating sets with isomorphic Cayley graphs, the groups $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2$ do *not* admit finite generating sets with isomorphic Cayley graphs. More generally, finitely generated Abelian groups admit isomorphic Cayley graphs if and only if they have the same rank and if the torsion part has the same cardinality [58].

Exercise 5.2.2 (Groups not quasi-isometric to $\mathbb{Z}$)**.** Let $n \in \mathbb{N}_{\geq 2}$.

1. Show that any quasi-isometric embedding $\mathbb{Z} \longrightarrow \mathbb{Z}$ is a quasi-isometry.
2. Show that there is *no* quasi-isometric embedding $X \longrightarrow \mathbb{Z}$ where $X := (\mathbb{Z} \times \{0\}) \cup (\{0\} \times \mathbb{Z})$ is equipped with the ℓ^1-metric on $\mathbb{R}^2$.
3. Conclude that the groups $\mathbb{Z}$ and $\mathbb{Z}^n$ are *not* quasi-isometric. In particular, $\mathbb{R}$ is not quasi-isometric to $\mathbb{R}^n$ with the Euclidean metric (because these spaces are quasi-isometric to $\mathbb{Z}$ and $\mathbb{Z}^n$ respectively).
4. Show that the group $\mathbb{Z}$ is *not* quasi-isometric to a free group of rank n.

However, much more is true – the group $\mathbb{Z}$ is quasi-isometrically rigid in the following sense:

Theorem 5.2.12 (Quasi-isometry rigidity of $\mathbb{Z}$)**.** *A finitely generated group is quasi-isometric to $\mathbb{Z}$ if and only if it is virtually $\mathbb{Z}$. A group is called* virtually $\mathbb{Z}$ *if it contains a finite index subgroup isomorphic to $\mathbb{Z}$.*

In other words, the property of being virtually $\mathbb{Z}$ is a geometric property of groups. We will give a proof of this result later when we have more tools at hand (Section 6.3.6, Corollary 7.5.8).

More generally, it is one of the primary goals of geometric group theory to understand as much as possible of the quasi-isometry classification of finitely generated groups.

5.3 The Švarc-Milnor lemma

Why should we be interested in understanding how finitely generated groups look like up to quasi-isometry? A first answer to this question is given by the Švarc-Milnor lemma, which is one of the key ingredients linking the geometry of groups to the geometry of spaces arising naturally in geometry and topology.

The Švarc-Milnor lemma roughly says that given a "nice" action of a group on a "nice" metric space, we can already conclude that the group in question is finitely generated and that the group is quasi-isometric to the given metric space.

In practice, this result can be applied both ways: If we want to know more about the geometry of a group or if we want to know that a given group is finitely generated, it suffices to exhibit a nice action of this group on a suitable space. Conversely, if we want to know more about a metric space, it suffices to find a nice action of a suitable well-known group. Therefore, the Švarc-Milnor lemma is also called the "fundamental lemma of geometric group theory."

5.3.1 Quasi-geodesics and quasi-geodesic spaces

In order to state the Švarc-Milnor lemma, we have to specify what a "nice" metric space is; in particular, we need that the metric on our space can be realised (up to some uniform error) by paths:

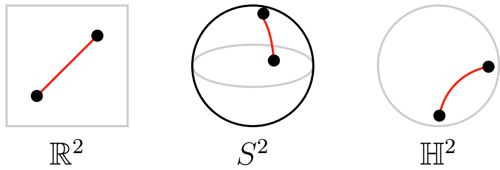

Figure 5.5.: Geodesic spaces and example geodesics (red)

Definition 5.3.1 (Geodesic space)**.** Let (X, d) be a metric space.

- Let $L \in \mathbb{R}_{\geq 0}$. A *geodesic of length* L in X is an isometric embedding $\gamma\colon [0, L] \longrightarrow X$, where the interval $[0, L]$ carries the metric induced from the standard metric on $\mathbb{R}$; the point $\gamma(0)$ is the *start point of* γ, and $\gamma(L)$ is the *end point of* γ.
- The metric space X is called *geodesic*, if for all $x, x' \in X$ there exists a geodesic in X with start point x and end point x'.

Example 5.3.2 (Geodesic spaces)**.** The following statements are illustrated in Figure 5.5.

- Let $n \in \mathbb{N}$. Geodesics in the Euclidean space $\mathbb{R}^n$ are precisely the Euclidean line segments (parametrised via a vector of unit length). As any two points in $\mathbb{R}^n$ can be joined by a line segment, Euclidean space $\mathbb{R}^n$ is geodesic.
- The space $\mathbb{R}^2 \setminus \{0\}$ endowed with the metric induced from the Euclidean metric on $\mathbb{R}^2$ is *not* geodesic (exercise).
- The sphere S^2 with the standard round Riemannian metric is a geodesic metric space.
- The hyperbolic plane $\mathbb{H}^2$ is a geodesic metric space.

Caveat 5.3.3. The notion of geodesic in differential geometry is related to the one above, but not quite the same; geodesics in differential geometry are only required to be locally isometric, not necessarily globally.

Finitely generated groups together with a word metric coming from a finite generating system are *not* geodesic (if the group in question is non-trivial), as the underlying metric space is discrete. However, they are geodesic in the sense of large scale geometry:

Definition 5.3.4 (Quasi-geodesic)**.** Let (X, d) be a metric space and let $c \in \mathbb{R}_{>0}$, $b \in \mathbb{R}_{\geq 0}$.

- Then a (c, b)-*quasi-geodesic* in X is a (c, b)-quasi-isometric embedding $\gamma\colon I \longrightarrow X$, where $I = [t, t'] \subset \mathbb{R}$ is some closed interval; the point $\gamma(t)$ is the *start point of* γ, and $\gamma(t')$ is the *end point of* γ.
- The space X is (c, b)-*quasi-geodesic*, if for all x, $x' \in X$ there exists a (c, b)-quasi-geodesic in X with start point x and end point x'.

Clearly, any geodesic space is also quasi-geodesic (namely, $(1, 0)$-quasi-geodesic); however, not every quasi-geodesic space is geodesic:

Example 5.3.5 (Quasi-geodesic spaces)**.**

- If $G = (V, E)$ is a connected graph, then the associated metric on V turns V into a $(1, 1)$-geodesic space: The distance between two vertices is realised as the length of some graph-theoretic path in the graph G, and any path in the graph G that realises the distance between two vertices yields a $(1, 1)$-quasi-geodesic (with respect to a suitable parametrisation).
- In particular: If G is a group and S is a generating set of G, then (G, d_S) is a $(1, 1)$-quasi-geodesic space.
- For every $\varepsilon \in \mathbb{R}_{>0}$ the space $\mathbb{R}^2 \setminus \{0\}$ is $(1, \varepsilon)$-quasi-geodesic with respect to the metric induced from the Euclidean metric on $\mathbb{R}^2$ (exercise).

Any quasi-geodesic space is quasi-isometric to a geodesic space (Proposition A.3.4). Notice however, that it is not always desirable to replace the original space by a geodesic space, because some control is lost during this replacement. Especially, in the context of graphs, one has to weigh up whether the rigidity of the combinatorial model or the flexibility of the geometric realisation is more useful for the situation at hand.

5.3.2 The Švarc-Milnor lemma

We now come to the Švarc-Milnor lemma. We start with a formulation using the language of quasi-geometry; in a second step, we deduce the topological version, the version usually used in applications.

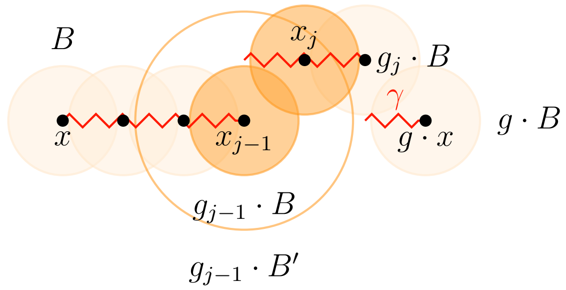

Figure 5.6.: Covering a quasi-geodesic by translates of B

Proposition 5.3.6 (Švarc-Milnor lemma)**.** *Let G be a group, and let G act on a (non-empty) metric space (X, d) by isometries. Suppose that there are constants $c,\ b \in \mathbb{R}_{>0}$ such that X is (c, b)-quasi-geodesic and suppose that there is a subset $B \subset X$ with the following properties:*

- *The diameter of B is finite.*
- *The G-translates of B cover all of X, i.e., $\bigcup_{g \in G} g \cdot B = X$.*
- *The set $S := \{g \in G \mid g \cdot B' \cap B' \neq \emptyset\}$ is finite, where*

$$B' := B_{2 \cdot b}(B) = \{x \in X \mid \exists_{y \in B}\ \ d(x, y) \leq 2 \cdot b\}.$$

Then the following holds:

1. *The group G is generated by S; in particular, G is finitely generated.*
2. *For all $x \in X$ the associated map*

$$\begin{aligned} G &\longrightarrow X \\ g &\longmapsto g \cdot x \end{aligned}$$

is a quasi-isometry (with respect to the word metric d_S on G).

Proof. *The set S generates G:* Let $g \in G$. We show that $g \in \langle S \rangle_G$ by using a suitable quasi-geodesic and following translates of B along this quasi-geodesic (Figure 5.6): Let $x \in B$. As X is (c, b)-quasi-geodesic, there is a (c, b)-quasi-geodesic $\gamma \colon [0, L] \longrightarrow X$ starting in x and ending in $g \cdot x$. We now look at close enough points on this quasi-geodesic:

Let $n := \lceil L \cdot c/b \rceil$. For $j \in \{0, \dots, n-1\}$ we define

$$t_j := j \cdot \frac{b}{c},$$

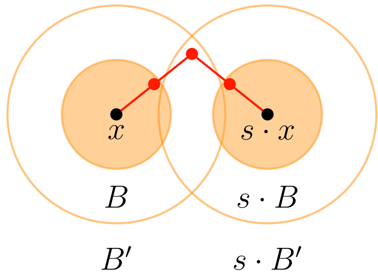

Figure 5.7.: If $s \in S$, then $d(x, s \cdot x) \leq 2 \cdot (\operatorname{diam} B + 2 \cdot b)$

and $t_n := L$, as well as

$$x_j := \gamma(t_j);$$

notice that $x_0 = \gamma(0) = x$ and $x_n = \gamma(L) = g \cdot x$. Because the translates of B cover all of X, there are group elements $g_j \in G$ with $x_j \in g_j \cdot B$; in particular, we can choose $g_0 := e$ and $g_n := g$.

For all $j \in \{1, \dots, n\}$, the group element $s_j := g_{j-1}^{-1} \cdot g_j$ lies in S because: As γ is a (c, b)-quasi-geodesic, we obtain

$$d(x_{j-1}, x_j) \leq c \cdot |t_{j-1} - t_j| + b \leq c \cdot \frac{b}{c} + b \leq 2 \cdot b.$$

Therefore, $x_j \in B_{2 \cdot b}(g_{j-1} \cdot B) = g_{j-1} \cdot B_{2 \cdot b}(B) = g_{j-1} \cdot B'$ (notice that in the second to last equality we used that G acts on X by isometries). On the other hand, $x_j \in g_j \cdot B \subset g_j \cdot B'$ and thus

$$g_{j-1} \cdot B' \cap g_j \cdot B' \neq \emptyset;$$

so, by definition of S, it follows that $s_j = g_{j-1}^{-1} \cdot g_j \in S$.

In particular,

$$g = g_n = g_{n-1} \cdot g_{n-1}^{-1} \cdot g_n = \dots = g_0 \cdot s_1 \cdot \dots \cdot s_n = s_1 \cdot \dots \cdot s_n$$

lies in the subgroup generated by S, as desired.

The group G is quasi-isometric to X: Let $x \in X$. We show that the map

$$\begin{aligned} \varphi \colon G &\longrightarrow X \\ g &\longmapsto g \cdot x \end{aligned}$$

is a quasi-isometry by showing that it is a quasi-isometric embedding with quasi-dense image. First notice that because G acts by isometries on X and because the G-translates of B cover all of X, we may assume that B contains x (so that we are in the same situation as in the first part of the proof).

The map φ has quasi-dense image, because: If $x' \in X$, then there is a $g \in G$ with $x' \in g \cdot B$. Then $g \cdot x \in g \cdot B$ yields

$$d\bigl(x', \varphi(g)\bigr) = d(x', g \cdot x) \leq \operatorname{diam} g \cdot B = \operatorname{diam} B,$$

which is assumed to be finite. Thus, φ has quasi-dense image.

The map φ is a quasi-isometric embedding, because: Let $g \in G$. We first give a uniform lower bound of $d(\varphi(e), \varphi(g))$ in terms of $d_S(e,g)$: Let $\gamma\colon [0, L] \longrightarrow X$ be as above a (c,b)-quasi-geodesic from x to $g \cdot x$. Then the argument from the first part of the proof (and the definition of n) shows that

$$\begin{aligned} d\bigl(\varphi(e), \varphi(g)\bigr) &= d(x, g \cdot x) = d\bigl(\gamma(0), \gamma(L)\bigr) \\ &\geq \frac{1}{c} \cdot L - b \\ &\geq \frac{1}{c} \cdot \frac{b \cdot (n-1)}{c} - b \\ &= \frac{b}{c^2} \cdot n - \frac{1}{c^2} - b \\ &\geq \frac{b}{c^2} \cdot d_S(e,g) - \frac{1}{c^2} - b. \end{aligned}$$

Conversely, we obtain a uniform upper bound of $d(\varphi(e), \varphi(g))$ in terms of $d_S(e,g)$ as follows: Suppose $d_S(e,g) = n$; so there are $s_1, \dots, s_n \in S \cup S^{-1} = S$ with $g = s_1 \cdot \dots \cdot s_n$. Hence, using the triangle inequality, the fact that G acts isometrically on X, and the fact that $s_j \cdot B' \cap B' \neq \emptyset$ for all $j \in \{1, \dots, n-1\}$ (see Figure 5.7) we obtain

$$\begin{aligned} d\bigl(\varphi(e), \varphi(g)\bigr) &= d(x, g \cdot x) \\ &\leq d(x, s_1 \cdot x) + d(s_1 \cdot x, s_1 \cdot s_2 \cdot x) + \cdots \\ &\quad + d(s_1 \cdot \dots \cdot s_{n-1} \cdot x, s_1 \cdot \dots \cdot s_n \cdot x) \\ &= d(x, s_1 \cdot x) + d(x, s_2 \cdot x) + \cdots + d(x, s_n \cdot x) \\ &\leq n \cdot 2 \cdot (\operatorname{diam} B + 2 \cdot b) \\ &= 2 \cdot (\operatorname{diam} B + 2 \cdot b) \cdot d_S(e,g). \end{aligned}$$

(Recall that diam B is assumed to be finite).

Because

$$d\bigl(\varphi(g),\varphi(h)\bigr) = d\bigl(\varphi(e),\varphi(g^{-1}\cdot h)\bigr) \quad \text{and} \quad d_S(g,h) = d_S(e,g^{-1}\cdot h)$$

holds for all g, $h \in G$, these bounds show that φ is a quasi-isometric embedding. $\square$

Notice that the proof of the Švarc-Milnor lemma does only give a quasi-isometry, not a bilipschitz equivalence. Indeed, the translation action of $\mathbb{Z}$ on $\mathbb{R}$ shows that there is no analogue of the Švarc-Milnor lemma for bilipschitz equivalence. Therefore, quasi-isometry of finitely generated groups is in geometric contexts considered to be the more appropriate notion than bilipschitz equivalence.

In many cases, the following, topological, formulation of the Švarc-Milnor lemma is used:

Corollary 5.3.7 (Švarc-Milnor lemma, topological formulation)**.** *Let G be a group acting by isometries on a (non-empty) proper geodesic metric space (X,d). Furthermore suppose that this action is proper and cocompact. Then G is finitely generated, and for all $x \in X$ the map*

$$\begin{aligned} G &\longrightarrow X \\ g &\longmapsto g\cdot x \end{aligned}$$

is a quasi-isometry.

Before deducing this version from the quasi-geometric version, we briefly recall the topological notions occuring in the statement:

- A metric space X is *proper* if for all $x \in X$ and all $r \in \mathbb{R}_{>0}$ the closed ball $\{y \in X \mid d(x,y) \leq r\}$ is compact with respect to the topology induced by the metric.
  Hence, proper metric spaces are locally compact.
- An action $G \times X \longrightarrow X$ of a group G on a topological space X (e.g., with the topology coming from a metric on X) is *proper* if for all compact sets $B \subset X$ the set $\{g \in G \mid g\cdot B \cap B \neq \emptyset\}$ is finite.
  Example 5.3.8 (Proper actions)**.**
  - The translation action of $\mathbb{Z}$ on $\mathbb{R}$ is proper (with respect to the standard topology on $\mathbb{R}$).

  - More generally, the action by deck transformations of the fundamental group of a locally compact path-connected topological space (that admits a universal covering) on its universal covering is proper [64, Chapter V].
  - Clearly, all stabiliser groups of a proper action are finite. The converse is *not* necessarily true: For example, the action of $\mathbb{Z}$ on the circle S^1 given by rotation around an irrational angle is free but not proper (because $\mathbb{Z}$ is infinite and S^1 is compact).
- An action $G \times X \longrightarrow X$ of a group G on a topological space X is *cocompact* if the quotient space $G \setminus X$ with respect to the quotient topology is compact.

  Example 5.3.9 (Cocompact actions)**.**
  - The translation action of $\mathbb{Z}$ on $\mathbb{R}$ is cocompact (with respect to the standard topology on $\mathbb{R}$), because the quotient is homeomorphic to the circle S^1, which is compact.
  - More generally, the action by deck transformations of the fundamental group of a compact path-connected topological space X (that admits a universal covering) on its universal covering is cocompact because the quotient is homeomorphic to X (Example 4.1.13).
  - The (horizontal) translation action of $\mathbb{Z}$ on $\mathbb{R}^2$ is *not* cocompact (with respect to the standard topology on $\mathbb{R}^2$), because the quotient is homeomorphic to the infinite cylinder $S^1 \times \mathbb{R}$, which is not compact.
  - The action of $\mathrm{SL}(2, \mathbb{Z})$ by Möbius transformations, i.e., via

$$\begin{aligned} \mathrm{SL}(2,\mathbb{Z}) \times H &\longrightarrow H \\ \left(\begin{pmatrix} a & b \\ c & d \end{pmatrix}, z \right) &\longmapsto \frac{a \cdot z + b}{c \cdot z + d}, \end{aligned}$$

    on the upper half plane $H := \{z \in \mathbb{C} \mid \operatorname{Re} z > 0\}$ is *not* cocompact (exercise).

Proof (of Corollary 5.3.7). Notice that under the given assumptions the space X is $(1, \varepsilon)$-quasi-geodesic for all $\varepsilon \in \mathbb{R}_{>0}$. In order to be able to apply the Švarc-Milnor lemma (Proposition 5.3.6), we need to find a suitable subset $B \subset X$.

Because the projection $\pi \colon X \longrightarrow G \setminus X$ associated with the action is an open map and because $G \setminus X$ is compact, one can easily find a closed subspace $B \subset X$ of finite diameter with $\pi(B) = G \setminus X$ (e.g., a suitable union of finitely many closed balls). In particular, $\bigcup_{g \in G} g \cdot B = X$ and

$$B' := B_{2 \cdot \varepsilon}(B)$$

has finite diameter. Because X is proper, the subset B' is compact; thus the action of G on X being proper implies that the set $\{g \in G \mid g \cdot B' \cap B' \neq \emptyset\}$ is finite.

Hence, we can apply the Švarc-Milnor lemma (Proposition 5.3.6). $\square$

5.3.3 Applications of the Švarc-Milnor lemma to group theory, geometry and topology

The Švarc-Milnor lemma has numerous applications to geometry, topology and group theory; we will give a few basic examples of this type, indicating the potential of the Švarc-Milnor lemma:

- Subgroups of finitely generated groups are finitely generated.
- (Weakly) commensurable groups are quasi-isometric.
- Certain groups are finitely generated (for instance, certain fundamental groups).
- Fundamental groups of nice compact metric spaces are quasi-isometric to the universal covering space.

As a first application of the Švarc-Milnor lemma, we give another proof of the fact that finite index subgroups of finitely generated groups are finitely generated:

Corollary 5.3.10. *Finite index subgroups of finitely generated groups are finitely generated and quasi-isometric to the ambient group.*

Proof. Let G be a finitely generated group, and let H be a subgroup of finite index. If S is a finite generating set of G, then the left translation action of H on (G, d_S) is an isometric action satisfying the conditions of the Švarc-Milnor lemma (Proposition 5.3.6): The space (G, d_S) is $(1, 1)$-quasi-geodesic. Moreover, we let $B \subset G$ be a finite set of representatives of $H \setminus G$ (in particular, the diameter of B is finite). Then $H \cdot B = G$, and the set $B' := B_2(B)$ is finite as well, and so the set

$$\{h \in H \mid h \cdot B' \cap B' \neq \emptyset\}$$

is finite.

Therefore, H is finitely generated and the inclusion $H \hookrightarrow G$ is a quasi-isometry (with respect to any word metrics on H and G coming from finite generating sets). $\square$

Pursuing this line of thought consequently leads to the notion of (weak) commensurability of groups:

Definition 5.3.11 ((Weak) commensurability)**.**
- Two groups G and H are *commensurable* if they contain finite index subgroups $G' \subset G$ and $H' \subset H$ with $G' \cong H'$.
- More generally, two groups G and H are *weakly commensurable* if they contain finite index subgroups $G' \subset G$ and $H' \subset H$ satisfying the following condition: There are finite normal subgroups N of G' and M of H' respectively such that the quotient groups G'/N and H'/M are isomorphic.

In fact, both commensurability and weak commensurability are equivalence relations on the class of groups (exercise).

Corollary 5.3.12 (Weak commensurability and quasi-isometry)**.** *Let G be a group.*
1. *Let G' be a finite index subgroup of G. Then G' is finitely generated if and only if G is finitely generated. If these groups are finitely generated, then $G \sim_{\mathrm{QI}} G'$.*
2. *Let N be a finite normal subgroup. Then G/N is finitely generated if and only if G is finitely generated. If these groups are finitely generated, then $G \sim_{\mathrm{QI}} G/N$.*

In particular, if G is finitely generated, then any group weakly commensurable to G is finitely generated and quasi-isometric to G.

Proof. Ad 1. In view of Corollary 5.3.10, it suffices to show that G is finitely generated if G' is; but clearly combining a finite generating set of G' with a finite set of representatives of the G'-cosets in G yields a finite generating set of G.

Ad 2. If G is finitely generated, then so is the quotient G/N; conversely, if G/N is finitely generated, then combining lifts with respect to the canonical projection $G \longrightarrow G/N$ of a finite generating set of G/N with the finite set N gives a finite generating set of G.

Let G and G/N be finitely generated, and let S be a finite generating set of G/N. Then the (pre-)composition of the left translation action of G/N on $(G/N, d_S)$ with the canonical projection $G \longrightarrow G/N$ gives an isometric action of G on G/N that satisfies the conditions of the Švarc-Milnor lemma (Proposition 5.3.6). Therefore, we obtain $G \sim_{\mathrm{QI}} G/N$. $\square$

Example 5.3.13.

- Let $n \in \mathbb{N}_{\geq 2}$. Then the free group of rank 2 contains a free group of rank n as finite index subgroup, and hence these groups are commensurable; in particular, all free groups of finite rank bigger than 1 are quasi-isometric.
- The subgroup of $\mathrm{SL}(2, \mathbb{Z})$ generated by the two matrices

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$$

  is free of rank 2 (Example 4.5.1); moreover, one can show that this subgroup index 12 in $\mathrm{SL}(2, \mathbb{Z})$. Thus, $\mathrm{SL}(2, \mathbb{Z})$ is finitely generated and commensurable to a free group of rank 2, and therefore quasi-isometric to a free group of rank 2 (and hence to all free groups of finite rank bigger than 1).
- Later we will find more examples of finitely generated groups that are not quasi-isometric. Hence, all these examples cannot be weakly commensurable (which might be rather difficult to check by hand).

Caveat 5.3.14. However, not all quasi-isometric groups are commensurable [42, p. 105f]: Let F_3 be a free group of rank 3, and let F_4 be a free group of rank 4. Then the finitely generated groups $(F_3 \times F_3) * F_3$ and $(F_3 \times F_3) * F_4$ are bilipschitz equivalent (and hence quasi-isometric) (Example 9.4.8).

On the other hand, the Euler characteristic of the corresponding classifying spaces is multiplicative under finite coverings. Hence, commensurable groups G and G' (that admit sufficiently finite models of classifying spaces) satisfy

$$\chi(G) = 0 \Longleftrightarrow \chi(G') = 0.$$

However,

$$\begin{aligned}\chi\bigl((F_3 \times F_3) * F_3\bigr) &= \chi(F_3) \cdot \chi(F_3) + \chi(F_3) - 1 \\ &= (1-3) \cdot (1-3) + (1-3) - 1 \\ &\neq 0 \\ &= (1-3) \cdot (1-3) + (1-4) - 1 \\ &= \chi(F_3) \cdot \chi(F_3) + \chi(F_4) - 1 \\ &= \chi\bigl((F_3 \times F_3) * F_4\bigr).\end{aligned}$$

So, $(F_3 \times F_3) * F_3$ and $(F_3 \times F_3) * F_4$ are *not* commensurable; moreover, because these groups are torsion-free, they also are not weakly commensurable.

Even more drastically, there also exist groups that are weakly commensurable but not commensurable [42, III.18(xi)].

As second topic, we look at applications of the Švarc-Milnor lemma in algebraic topology/Riemannian geometry via fundamental groups; a concise introduction to Riemannian geometry is given by Lee's book [54].

Corollary 5.3.15 (Fundamental groups and quasi-isometry)**.** *Let M be a closed (i.e., compact and without boundary) connected Riemannian manifold, and let $\widetilde{M}$ be its Riemannian universal covering manifold. Then the fundamental group $\pi_1(M)$ is finitely generated and for every $x \in \widetilde{M}$, the map*

$$\begin{aligned}\pi_1(M) &\longrightarrow \widetilde{M} \\ g &\longmapsto g \cdot x\end{aligned}$$

given by the action of the fundamental group $\pi_1(M)$ on $\widetilde{M}$ via deck transformations is a quasi-isometry. Here, M and $\widetilde{M}$ are equipped with the metrics induced from their Riemannian metrics.

Sketch of proof. Standard arguments from Riemannian geometry and topology show that in this case $\widetilde{M}$ is a proper geodesic metric space and that the action of $\pi_1(M)$ on $\widetilde{M}$ is isometric, proper, and cocompact (the quotient being the compact space M).

Applying the topological version of the Švarc-Milnor lemma (Corollary 5.3.7) finishes the proof. $\square$

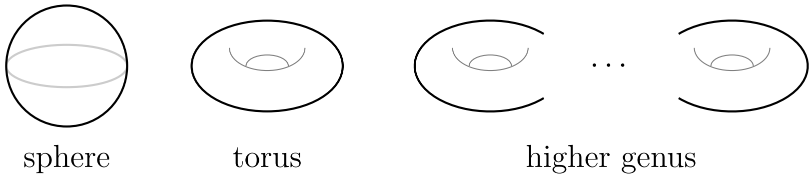

Figure 5.8.: Oriented closed connected surfaces

We give a sample application of this consequence of the Švarc-Milnor lemma to Riemannian geometry:

Definition 5.3.16 (Flat manifold, hyperbolic manifold)**.**
- A Riemannian manifold is called *flat* if its Riemannian universal covering is isometric to the Euclidean space of the same dimension.
- A Riemannaian manifold is called *hyperbolic* if its Riemannian universal covering is isometric to the hyperbolic space of the same dimension.

Example 5.3.17 (Surfaces)**.** Oriented closed connected surfaces are determined up to homeomorphism/diffeomorphism by their genus (i.e., the number of "handles", see Figure 5.8) [64, Chapter I].
- The oriented surface of genus 0 is the sphere of dimension 2; it is simply connected, and so coincides with its universal covering space. In particular, no Riemannian metric on S^2 is flat or hyperbolic.
- The oriented surface of genus 1 is the torus of dimension 2, which has fundamental group isomorphic to $\mathbb{Z}^2$. The torus admits a flat Riemannian metric: The translation action of $\mathbb{Z}^2$ on $\mathbb{R}^2$ is isometric with respect to the flat Riemannian metric on $\mathbb{R}^2$ and properly discontinuous; hence, the quotient space (i.e., the torus $S^1 \times S^1$) inherits a flat Riemannian metric.
- Oriented surfaces of genus $g \geq 2$ have fundamental group isomorphic to
$$\biggl\langle a_1, \dots, a_g, b_1, \dots, b_g \biggm| \prod_{j=1}^{g} [a_j, b_j] \biggr\rangle$$
and one can show that these surfaces admit hyperbolic Riemannian metrics [11, Chapter B.1, B.3].

Being flat is the same as having vanishing sectional curvature and being hyperbolic is the same as having constant sectional curvature -1 [54, Chapter 11].

As a direct consequence of Corollary 5.3.15 we obtain:

Corollary 5.3.18 ((Non-)Existence of flat/hyperbolic structures)**.**

1. *If M is a closed connected Riemannian n-manifold that is flat, then its fundamental group $\pi_1(M)$ is quasi-isometric to Euclidean space $\mathbb{R}^n$, and hence to $\mathbb{Z}^n$.*
2. *In other words: If the fundamental group of a closed connected smooth n-manifold is not quasi-isometric to $\mathbb{R}^n$ (or $\mathbb{Z}^n$), then this manifold does not admit a flat Riemannian metric.*
3. *If M is a closed connected Riemannian n-manifold that is hyperbolic, then its fundamental group $\pi_1(M)$ is quasi-isometric to the hyperbolic space $\mathbb{H}^n$.*
4. *In other words: If the fundamental group of a closed connected smooth n-manifold is not quasi-isometric to $\mathbb{H}^n$, then this manifold does not admit a hyperbolic Riemannian metric.*

So, classifying finitely generated groups up to quasi-isometry and studying the quasi-geometry of finitely generated groups gives insights into geometry and topology of smooth/Riemannian manifolds.

5.4 The dynamic criterion for quasi-isometry

The Švarc-Milnor lemma translates an action of a group into a quasi-isometry of the group in question to the metric space acted upon. Similarly, we can also use certain actions to compare two groups with each other:

Definition 5.4.1 (Set-theoretic coupling)**.** Let G and H be groups. A *set-theoretic coupling for G and H* is a non-empty set X together with a left action of G on X and a right action[1] of H on X that commute with each other (i.e., $(g \cdot x) \cdot h = g \cdot (x \cdot h)$ holds for all $x \in X$ and all $g \in G$, $h \in H$) such that X contains a subset K with the following properties:

1. The G- and H-translates of K cover X, i.e. $G \cdot K = X = K \cdot H$.
2. The sets
$$\begin{aligned} F_G &:= \{g \in G \mid g \cdot K \cap K \neq \emptyset\}, \\ F_H &:= \{h \in H \mid K \cdot h \cap K \neq \emptyset\} \end{aligned}$$
are finite.
3. For each $g \in G$ there is a finite subset $F_H(g) \subset H$ with $g \cdot K \subset K \cdot F_H(g)$, and for each $h \in H$ there is a finite subset $F_G(h) \subset G$ with $K \cdot h \subset F_G(h) \cdot K$.

Example 5.4.2. Let X be a group and let $G \subset X$ and $H \subset X$ be subgroups of finite index. Then the left action
$$\begin{aligned} G \times X &\longrightarrow X \\ (g, x) &\longmapsto g \cdot x \end{aligned}$$
of G on X and the right action
$$\begin{aligned} X \times H &\longrightarrow X \\ (x, h) &\longmapsto x \cdot h \end{aligned}$$
of H on X commute with each other (because multiplication in the group X is associative). The set X together with these actions is a set-theoretic coupling – a suitable subset $K \subset X$ can for example be obtained by taking the union of finite sets of representatives for G-cosets in X and H-cosets in X respectively.

Proposition 5.4.3 (Quasi-isometry and set-theoretic couplings)**.** *Let G and H be two finitely generated groups that admit a set-theoretic coupling. Then $G \sim_{\mathrm{QI}} H$.*

[1] A *right action* of a group H on a set X is a map $X \times H \longrightarrow H$ such that $x \cdot e = x$ and $(x \cdot h) \cdot h' = x \cdot (h \cdot h')$ holds for all $x \in X$ and all $h, h' \in H$. In other words, a right action is the same as an antihomomorphism $H \longrightarrow S_X$.

Proof. Let X be a set-theoretic coupling space for G and H with the corresponding commuting actions by G and H; in the following, we will use the notation from Definition 5.4.1. We prove $G \sim_{\mathrm{QI}} H$ by writing down a candidate for a quasi-isometry $G \longrightarrow H$ and by then verifying (similar to the proof of the Švarc-Milnor lemma) that this map indeed has quasi-dense image and is a quasi-isometric embedding: Let $x \in K \subset X$. Using the axiom of choice, we obtain a map $f \colon G \longrightarrow H$ satisfying

$$g^{-1} \cdot x \in K \cdot f(g)^{-1}$$

for all $g \in G$.

Moreover, we will use the following notation: Let $S \subset G$ be a finite generating set of G, and let $T \subset H$ be a finite generating set of H. For a subset $B \subset H$, we define

$$D_T(B) := \sup_{b \in B} d_T(e, b),$$

and similarly, we define $D_S(A)$ for subsets A of G.

The map f has quasi-dense image: Let $h \in H$. Using $G \cdot K = X$ we find a $g \in G$ with $x \cdot h \in g \cdot K$; because the actions of G and H commute with each other, it follows that $g^{-1} \cdot x \in K \cdot h^{-1}$. On the other hand, also $g^{-1} \cdot x \in K \cdot f(g)^{-1}$, by definition of f. In particular, $K \cdot h^{-1} \cap K \cdot f(g)^{-1} \neq \emptyset$, and so $h^{-1} \cdot f(g) \in F_H$. Therefore,

$$d_T\bigl(h, f(g)\bigr) \leq D_T F_H,$$

which is finite (the set F_H is finite by assumption) and independent of h; hence, f has quasi-dense image.

The map f is a quasi-isometric embedding: Notice that the sets

$$F_H(S) := \bigcup_{s \in S \cup S^{-1}} F_H(s) \qquad \text{and} \qquad F_G(T) := \bigcup_{t \in T \cup T^{-1}} F_G(t)$$

are finite by assumption. Let $g, g' \in G$.

- We first give an upper bound of $d_T(f(g), f(g'))$ in terms of $d_S(g, g')$: More precisely, we will show that

$$d_T\bigl(f(g), f(g')\bigr) \leq D_T F_H(S) \cdot d_S(g, g') + D_T F_H.$$

To this end let $n := d_S(g,g')$. As first step we show that the intersection $K \cdot f(g)^{-1} \cdot f(g') \cap K \cdot F_H(S)^n$ is non-empty: On the one hand,

$$g^{-1} \cdot x \cdot f(g') \in K \cdot f(g)^{-1} \cdot f(g')$$

by construction of f. On the other hand, because $d_S(g,g') = n$ we can write $g^{-1} \cdot g' = s_1 \cdot \dots \cdot s_n$ for certain $s_1, \dots, s_n \in S \cup S^{-1}$, and thus

$$\begin{aligned}
g^{-1} \cdot x \cdot f(g') &= g^{-1} \cdot g' \cdot g'^{-1} \cdot x \cdot f(g') \\
&\in g^{-1} \cdot g' \cdot K \cdot f(g')^{-1} \cdot f(g') \\
&= g^{-1} \cdot g' \cdot K \\
&= s_1 \cdot \dots \cdot s_{n-1} \cdot s_n \cdot K \\
&\subset s_1 \cdot \dots \cdot s_{n-1} \cdot K \cdot F_H(S) \\
&\vdots \\
&\subset K \cdot F_H(S)^n.
\end{aligned}$$

In particular, $K \cdot f(g)^{-1} \cdot f(g') \cap K \cdot F_H(S)^n \neq \emptyset$. Notice that in all these computations we used heavily that the actions of G and H on X commute with each other.

Using the definition of F_H, we see that

$$f(g)^{-1} \cdot f(g') \in F_H \cdot F_H(S)^n;$$

in particular, we obtain (via the triangle inequality)

$$\begin{aligned}
d_T\bigl(f(g), f(g')\bigr) &= d_T\bigl(e, f(g)^{-1} \cdot f(g')\bigr) \\
&\leq D_T(F_H \cdot F_H(S)^n) \\
&\leq n \cdot D_T F_H(S) + D_T F_H \\
&= d_S(g,g') \cdot D_T F_H(S) + D_T F_H,
\end{aligned}$$

as desired (the constants $D_T F_H(S)$ and $D_T F_H$ are finite because the sets $F_H(S)$ and F_H are finite by assumption).

– Moreover, there is a lower bound of $d_T(f(g), f(g'))$ in terms of $d_S(g,g')$: Let $m := d_T(f(g), f('g))$. Using similar arguments as above, one sees that

$$g^{-1} \cdot x \cdot f(g') \in F_G(T)^m \cdot K \cap g^{-1} \cdot g' \cdot K$$

and hence that this intersection is non-empty. Therefore, we can conclude that

$$d_S(g, g') \leq D_S F_G(T) \cdot d_T\bigl(f(g), f(g')\bigr) + D_S F_G,$$

which gives the desired lower bound. □

Remark 5.4.4 (Cocycles)**.** The construction of the map f in the proof above is an instance of a more general principle associating interesting maps with actions. Namely, suitable actions lead to so-called *cocycles* (which is an algebraic object); considering cocycles up to an appropriate equivalence relation ("being a coboundary") then gives rise to so-called *cohomology groups*. In this way, certain aspects of group actions on a space can be translated into an algebraic theory.

Moreover, we do not need to assume that both groups are finitely generated as being finitely generated is preserved by set-theoretic couplings:

Exercise 5.4.1 (Set-theoretic couplings and finite generation)**.** Let G and H be groups that admit a set-theoretic coupling. Show that if G is finitely generated, then so is H.

The converse of Proposition 5.4.3 also holds: whenever two finitely generated groups are quasi-isometric, then there exists a coupling (even a topological coupling) between them:

Definition 5.4.5 (Topological coupling)**.** Let G and H be groups. A *topological coupling for G and H* is a non-empty locally compact space X together with a proper cocompact left action of G on X by homeomorphisms and a proper cocompact right action of H on X by homeomorphisms that commute with each other.

A topological space X is called *locally compact*[2] if for every $x \in X$ and every open neighbourhood $U \subset X$ of x there exists a compact neighbourhood $K \subset X$ of x with $K \subset U$. For example, a metric space is locally compact if and only if it is proper.

We can now formulate Gromov's dynamic criterion for quasi-isometry:

[2] There are several *different* notions of local compactness in the literature!

Theorem 5.4.6 (Dynamic criterion for quasi-isometry)**.** *Let G and H be finitely generated groups. Then the following are equivalent:*

1. *There is a topological coupling for G and H.*
2. *There is a set-theoretic coupling for G and H.*
3. *The groups G and H are quasi-isometric.*

Sketch of proof. Ad "1 $\Longrightarrow$ 2". Let G and H be finitely generated groups that admit a topological coupling, i.e., there is a non-empty locally compact space X together with a proper cocompact action from G on the left and from H on the right such that these two actions commute with each other. We show that such a topological coupling gives rise to a set-theoretic coupling:

A standard argument from topology shows that in this situation there is a compact subset $K \subset X$ such that $G \cdot K = X = K \cdot H$. Because the actions of G and H on X are proper, the sets

$$\{g \in G \mid g \cdot K \cap K \neq \emptyset\}, \qquad \{h \in H \mid K \cdot h \cap K \neq \emptyset\}$$

are finite; moreover, compactness of the set K as well as the local compactness of X also give us that for every $g \in G$ there is a finite set $F_H(g) \subset H$ satisfying $g \cdot K \subset K \cdot F_H(g)$, and similarly for elements of H. Hence, we obtain a set-theoretic coupling for G and H.

Ad "2 $\Longrightarrow$ 3". This was proved in Proposition 5.4.3.

Ad "3 $\Longrightarrow$ 1". Suppose that the finitely generated groups G and H are quasi-isometric. We now sketch how this leads to a topological coupling of G and H:

Let $S \subset G$ and $T \subset H$ be finite generating sets of G and H respectively. As first step, we show that there is a finite group F and a constant $C \in \mathbb{R}_{>0}$ such that the set

$$\begin{aligned} X := \Bigl\{ f \colon G \longrightarrow H \times F \Bigm| \; & f \text{ has } C\text{-dense image in } H \times F, \text{ and} \\ & \forall_{g,g' \in G} \;\; \frac{1}{C} \cdot d_S(g,g') \leq d_{T \times F}\bigl(f(g), f(g')\bigr) \leq C \cdot d_S(g,g') \Bigr\} \end{aligned}$$

is non-empty: Let $f \colon G \longrightarrow H$ be a quasi-isometry. Because f is a quasi-isometry, there is a $c \in \mathbb{R}_{>0}$ such that f has c-dense image in H and

$$\forall_{g,g' \in G} \;\; \frac{1}{c} \cdot d_S(g,g') - c \leq d_T\bigl(f(g), f(g')\bigr) \leq c \cdot d_S(g,g') + c.$$

In particular, if g, $g' \in G$ with $f(g) = f(g')$, then $d_S(g,g') \leq c^2$. Let F be a finite group that has more elements than the d_S-ball of radius c^2 in G (around the neutral element). Then out of f we can construct an *injective* quasi-isometry $\overline{f}\colon G \longrightarrow H \times F$. Let $\overline{c} \in \mathbb{R}_{>0}$ be chosen in such a way that $\overline{f}$ is a $(\overline{c},\overline{c})$-quasi-isometric embedding with $\overline{c}$-dense image. Because $\overline{f}$ is injective, then $\overline{f}$ satisfies a $\max(2\cdot\overline{c},\overline{c}^2+\overline{c})$-bilipschitz estimate (this follows as in Exercise 5.2.1 from the fact that any two different elements of a finitely generated group have distance at least 1 in any word metric). Hence, F and $C := \max(2\cdot\overline{c},\overline{c}^2+\overline{c})$ have the desired property that the corresponding set X is non-empty.

We now consider the following left G-action and right H-action on X:

$$\begin{aligned} G \times X &\longrightarrow X \\ (g,f) &\longmapsto \bigl(x \mapsto f(g^{-1}\cdot x)\bigr) \end{aligned}$$

$$\begin{aligned} X \times H &\longrightarrow X \\ (f,h) &\longmapsto \bigl(x \mapsto f(x)\cdot(h,e)\bigr) \end{aligned}$$

Clearly, these two actions commute with each other.

Furthermore, we equip X with the topology of pointwise convergence (which coincides with the compact-open topology when viewing X as a subspace of all "continuous" functions $G \longrightarrow H \times F$). By the Arzelá-Ascoli theorem [47, Chapter 7], the space X is locally compact with respect to this topology; at this point it is crucial that the functions in X satisfy a uniform (bi)lipschitz condition (instead of a quasi-isometry condition) so that X is equicontinuous. A straightforward computation (also using the Arzelá-Ascoli theorem) shows that the actions of G and H on X are proper and cocompact [86]. $\square$

5.4.1 Applications of the dynamic criterion

A topological version of subgroups of finite index are uniform lattices; the dynamic criterion shows that finitely generated uniform lattices in the same ambient locally compact group are quasi-isometric (Corollary 5.4.8).

Definition 5.4.7 (Uniform lattices)**.** Let G be a locally compact topological group. A *uniform (or cocompact) lattice in G* is a discrete subgroup Γ

of G such that the left translation action (equivalently, the right translation action) of Γ on G is cocompact.

Recall that a *topological group* is a group G that in addition is a topological space such that the composition $G \times G \longrightarrow G$ in the group and the map $G \longrightarrow G$ given by taking inverses are continuous (on $G \times G$ we take the product topology). A subgroup Γ of a topological group G is *discrete* if there is an open neighbourhood U of the neutral element e in G such that $U \cap \Gamma = \{e\}$.

Corollary 5.4.8 (Uniform lattices and quasi-isometry)**.** *Let G be a locally compact topological group. Then all finitely generated uniform lattices in G are quasi-isometric.*

Proof. Let Γ and Λ be two finitely generated uniform lattices in G. Then the left action

$$\begin{aligned} \Gamma \times G &\longrightarrow G \\ (\gamma, g) &\longmapsto \gamma \cdot g \end{aligned}$$

of Γ on G and the right action

$$\begin{aligned} G \times \Lambda &\longrightarrow G \\ (g, \lambda) &\longmapsto g \cdot \lambda \end{aligned}$$

of Λ on G are continuous (because G is a topological group) and commute with each other, and these actions are cocompact and proper. Hence, Γ and Λ are quasi-isometric by the dynamic criterion (Theorem 5.4.6). □

Example 5.4.9 (Uniform lattices)**.**

- If G is a group, equipped with the discrete topology, then a subgroup of G is a uniform lattice if and only if it has finite index in G.
- Let $n \in \mathbb{N}$. Then $\mathbb{Z}^n$ is a discrete subgroup of the locally compact topological group $\mathbb{R}^n$, and $\mathbb{Z}^n \setminus \mathbb{R}^n$ is compact (namely, the n-torus); hence, $\mathbb{Z}^n$ is a uniform lattice in $\mathbb{R}^n$.
- The subgroup $\mathbb{Q} \subset \mathbb{R}$ is *not* discrete in $\mathbb{R}$.
- Because the quotient $\mathbb{Z} \times \{0\} \setminus \mathbb{R}^2$ is not compact, $\mathbb{Z} \times \{0\}$ is *not* a uniform lattice in $\mathbb{R}^2$. Notice that the above corollary would not hold in general without requiring that the lattices are uniform: the group $\mathbb{Z}$ is *not* quasi-isometric to $\mathbb{R}^2$ (Exercise 5.2.2).

– Let $H_{\mathbb{R}}$ be the *real Heisenberg group*, and let H be the *Heisenberg group*, i.e.,

$$H_{\mathbb{R}} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x,\ y,\ z \in \mathbb{R} \right\}, \qquad H := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x,\ y,\ z \in \mathbb{Z} \right\}.$$

Then $H_{\mathbb{R}}$ is a locally compact topological group (with respect to the topology given by convergence of all matrix coefficients), and H is a finitely generated uniform lattice in $H_{\mathbb{R}}$ (exercise).
So, any finitely generated group that is not quasi-isometric to H cannot be a uniform lattice in $H_{\mathbb{R}}$; for example, we will see in Chapter 6 that $\mathbb{Z}^3$ is not quasi-isometric to H, and that free groups of finite rank are not quasi-isometric to H.
– The subgroup $\mathrm{SL}(2,\mathbb{Z})$ of the matrix group $\mathrm{SL}(2,\mathbb{R})$ is discrete and the quotient $\mathrm{SL}(2,\mathbb{Z}) \setminus \mathrm{SL}(2,\mathbb{R})$ has finite invariant measure, but this quotient is not compact (this is similar to the fact that the action of $\mathrm{SL}(2,\mathbb{Z})$ on the upper half-plane is not cocompact); so $\mathrm{SL}(2,\mathbb{Z})$ is *not* a uniform lattice in $\mathrm{SL}(2,\mathbb{R})$.
– If M is a closed connected Riemannian manifold, then the isometry group $\mathrm{Isom}(\widetilde{M})$ of the Riemannian universal covering of M is a locally compact topological group (with respect to the compact-open topology). Because the fundamental group $\pi_1(M)$ acts by isometries (via deck transformations) on $\widetilde{M}$, we can view $\pi_1(M)$ as a subgroup of $\mathrm{Isom}(\widetilde{M})$. One can show that this subgroup is discrete and cocompact, so that $\pi_1(M)$ is a uniform lattice in $\mathrm{Isom}(\widetilde{M})$ [86, Theorem 2.35].

Remark 5.4.10 (Measure equivalence)**.** Another important aspect of the dynamic criterion for quasi-isometry is that it admits translations to other settings. For example, the corresponding measure-theoretic notion is *measure equivalence* of groups, which plays a central role in measurable group theory [36].

5.5 Preview: Quasi-isometry invariants and geometric properties

The central classification problem of geometric group theory is to classify all finitely generated groups up to quasi-isometry. As we have seen in the previous sections, knowing that certain groups are *not* quasi-isometric leads to interesting consequences in group theory, topology, and geometry.

5.5.1 Quasi-isometry invariants

While a complete classification of finitely generated groups up to quasi-isometry is far out of reach, partial results can be obtained. A general principle to obtain partial classification results in many mathematical fields is to construct suitable invariants. We start with the simplest case, namely set-valued quasi-isometry invariants:

Definition 5.5.1 (Quasi-isometry invariants)**.** Let V be a set. A *quasi-isometry invariant with values in* V is a map I from the class of all finitely generated groups to V such that

$$I(G) = I(H)$$

holds for all finitely generated groups G and H with $G \sim_{\mathrm{QI}} H$.

Proposition 5.5.2 (Using quasi-isometry invariants)**.** *Let V be a set, and let I be a quasi-isometry invariant with values in V, and let G and H be finitely generated groups with $I(G) \neq I(H)$. Then G and H are* not *quasi-isometric.*

Proof. Assume for a contradiction that G and H are quasi-isometric. Because I is a quasi-isometry invariant, this implies $I(G) = I(H)$, which contradicts the assumption $I(G) \neq I(H)$. Hence, G and H cannot be quasi-isometric. □

So, the more quasi-isometry invariants we can find, the more finitely generated groups we can distinguish up to quasi-isometry.

Caveat 5.5.3. If I is a quasi-isometry invariant of finitely generated groups, and G and H are finitely generated groups with $I(G) = I(H)$, then in general we *cannot* deduce that G and H are quasi-isometric, as the example of the trivial invariant shows (see below).

Some basic examples of quasi-isometry invariants are the following:

Example 5.5.4 (Quasi-isometry invariants)**.**

- *The trivial invariant.* Let V be a set containing exactly one element, and let I be the map associating with every finitely generated group this one element. Then clearly I is a quasi-isometry invariant – however, I does not contain any interesting information.
- *Finiteness.* Let $V := \{0, 1\}$, and let I be the map that sends all finite groups to 0 and all finitely generated infinite groups to 1. Then I is a quasi-isometry invariant, because a finitely generated group is quasi-isometric to a finite group if and only if it is finite (Example 5.2.9).
- *Rank of free groups.* Let $V := \mathbb{N}$, and let I be the map from the class of all finitely generated free groups to V that associates with a finitely generated free group its rank. Then I is *not* a quasi-isometry invariant on the class of all finitely generated free groups, because free groups of rank 2 and rank 3 are quasi-isometric (Example 5.3.13).

In order to obtain more interesting classification results we need further quasi-isometry invariants. In the following chapters, we will, for instance, study the growth of groups (Chapter 6), and ends of groups (i.e., geometry at infinity) (Chapter 8).

5.5.2 Functorial quasi-isometry invariants

More generally, quasi-isometry invariants can be conveniently formalised as functors between categories. Functors translate objects and morphisms between categories:

Definition 5.5.5 (Functor)**.** Let C and D be categories. A *(covariant) functor* $F\colon C \longrightarrow D$ consists of the following components:

– A map $F\colon \operatorname{Ob}(C) \longrightarrow \operatorname{Ob}(D)$ between the classes of objects.
– For all objects $X, Y \in \operatorname{Ob}(C)$ a map

$$F\colon \operatorname{Mor}_C(X,Y) \longrightarrow \operatorname{Mor}_D\bigl(F(X), F(Y)\bigr).$$

These maps are required to be compatible in the following sense:
– For all objects $X \in \operatorname{Ob}(C)$ we have $F(\mathrm{id}_X) = \mathrm{id}_{F(X)}$.
– For all $X, Y, Z \in \operatorname{Ob}(C)$, all $f \in \operatorname{Mor}_C(X,Y)$, and all $g \in \operatorname{Mor}_C(Y,Z)$ we have

$$F(g \circ f) = F(g) \circ F(f).$$

Functors, by definition, satisfy a fundamental invariance principle:

Proposition 5.5.6 (Functors preserve isomorphisms)**.** *Let C and D be categories, let $F\colon C \longrightarrow D$ be a functor, and let $X, Y \in \operatorname{Ob}(C)$.*
1. *If $f \in \operatorname{Mor}_C(X,Y)$ is an isomorphism in the category C, then $F(f) \in \operatorname{Mor}_D(F(X), F(Y))$ is an isomorphism in D.*
2. *If $X \cong_C Y$, then $F(X) \cong_D F(Y)$.*
3. *If $F(X) \not\cong_D F(Y)$, then $X \not\cong_C Y$.*

Proof. This is an immediate consequence of the definition of functors and of isomorphism in categories. □

Functors are ubiquitous in modern mathematics. For example, the fundamental group (Appendix A.1) is a functor from the (homotopy) category of pointed topological spaces to the category of groups; geometric realisation (Appendix A.3) can be viewed as a functor from the category of graphs to the category of metric spaces.

Definition 5.5.7 (functorial quasi-isometry invariant)**.** Let C be a category. A *functorial quasi-isometry invariant with values in C* is a functor from (a subcategory of) QMet to C.

Clearly, any set-valued quasi-isometry invariant can be viewed as a functorial quasi-isometry invariant (on the full subcategory of finitely generated groups with word metrics in QMet) by adding the trivial category structure on the set V of values (i.e., the elements of V are the objects of the category, and the only morphisms are the identity morphisms).

Basic examples of functorial quasi-isometry invariants are the ends functor from the subcategory of QMet generated by geodesic metric spaces to the category of topological spaces and the Gromov boundary functor from the subcategory of QMet generated by so-called quasi-hyperbolic spaces to the category of topological spaces (Chapter 8).

Moreover, there is a general principle turning functors from algebraic topology into quasi-isometry invariants, based on the coarsening construction by Higson and Roe [45, 81, 77]: For simplicity, we restrict ourselves to the domain category of uniformly discrete spaces of bounded geometry.

Definition 5.5.8 (UDBG space)**.** A metric space X is a *UDBG space* if it is *uniformly discrete* and of *bounded geometry*, i.e., if

$$\inf\bigl\{d(x,x') \bigm| x,x' \in X,\ x \neq x'\bigr\} > 0$$

and if there exists for all $r \in \mathbb{R}_{>0}$ a constant $K_r \in \mathbb{N}$ such that

$$\forall_{x \in X}\ \bigl|B_r(x)\bigr| \leq K_r.$$

The full subcategory of QMet generated by all UDBG spaces is denoted by UDBG.

For example, if G is a finitely generated group and $S \subset G$ is a finite generating set, then (G, d_S) is a UDBG space.

The coarsening of a functor is – in a certain sense – the maximal quasi-isometry invariant contained in the given functor:

Theorem 5.5.9 (Coarsening of functors)**.** *Let C be a category that is closed under direct limits, and let $F \colon \mathsf{Simp}^{\mathsf{lf}}_{\mathsf{h}} \longrightarrow C$ be a functor. Then there exists a functor $QF \colon \mathsf{UDBG} \longrightarrow C$ and a natural transformation $c_F \colon F \circ I \Longrightarrow QF \circ V$ with the following universal property: If $G \colon \mathsf{UDBG} \longrightarrow C$ is a functor and if $c_G \colon F \circ I \Longrightarrow G \circ V$ is a natural transformation, then there exists a unique natural transformation $c \colon QF \Longrightarrow G$ with*

$$c \circ c_F = c_G.$$

$$\begin{tikzcd} F \circ I \arrow[r, Rightarrow] \arrow[dr, Rightarrow] & G \circ V \\ & QF \circ V \arrow[u, Rightarrow, dashed] \end{tikzcd}$$

We call QF the coarsening of F *and c_F the* comparison map for F.

Before sketching the proof, we briefly explain the terms in the theorem: A natural transformation is a family of morphisms in the target category relating two given functors [62]. Let $\mathsf{Simp}^{\mathsf{lf}}$ denote the category of locally finite simplicial complexes and proper simplicial maps; let $\mathsf{Simp}^{\mathsf{lf}}_{\mathsf{h}}$ be the homotopy category of $\mathsf{Simp}^{\mathsf{lf}}$ with respect to proper simplicial homotopy of proper simplicial maps.

We write $\mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}}$ for the subcategory of $\mathsf{Simp}^{\mathsf{lf}}$ of uniformly locally finite simplicial complexes and simplicial maps that are not only proper but also quasi-isometric embeddings with respect to the metric on the vertices associated with the 1-skeleton of the simplicial complex in question; notice that the set of vertices is a UDBG space with respect to this metric. Moreover, we consider the corresponding canonical functors

$$\begin{aligned} I\colon \mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}} &\longrightarrow \mathsf{Simp}^{\mathsf{lf}}_{\mathsf{h}} \\ V\colon \mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}} &\longrightarrow \mathsf{UDBG}\,. \end{aligned}$$

Sketch of proof of Theorem 5.5.9. As first step we indicate how the functor $QF\colon \mathsf{UDBG} \longrightarrow C$ can be constructed: The basic idea is to "zoom out" of the given space using so-called Rips complexes: If X is a UDBG space and $r \in \mathbb{R}_{\geq 0}$, then we define the *Rips complex $R_r(X)$ of X with radius r* as the following simplicial complex: If $n \in \mathbb{N}$, then the n-simplices of $R_r(X)$ are the set

$$\bigl\{x \in X^{n+1} \bigm| \forall_{j,k \in \{0,\dots,n\}}\ d(x_j, x_k) \leq r\bigr\}.$$

Hence, the local structure of radius r is blurred in $R_r(X)$, and only the large-scale structure beyond radius r is preserved (Figure 5.9).

– *On objects:* Let X be a UDBG space. Then the family of Rips complexes $(R_r(X))_{r \in \mathbb{R}_{\geq 0}}$ with increasing radius forms a directed system of (uniformly) locally finite simplicial complexes with respect to the natural inclusions

$$R_r(X) \longrightarrow R_s(X)$$

for $r, s \in \mathbb{R}_{\geq 0}$ with $r \leq s$. We then define

$$QF(X) := \lim_{r \to \infty} F\bigl(R_r(X)\bigr),$$

where the direct limit is taken by applying F to the above inclusions.

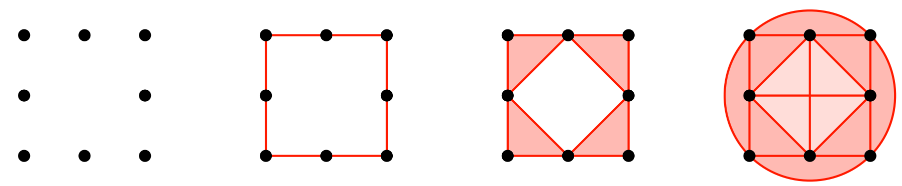

$X = R_0(X)$ $\quad$ $R_1(X)$ $\quad$ $R_{\sqrt{2}}(X)$ $\quad$ $R_2(X)$

Figure 5.9.: Toy example for Rips complexes of a discrete metric space X; here, X has the metric induced from $\mathbb{R}^2$ and the large square has sides of length 2.

– *On morphisms:* Let $f\colon X \longrightarrow Y$ be a quasi-isometric embedding of UDBG spaces. Then for every $r \in \mathbb{R}_{\geq 0}$ there exists an $s \in \mathbb{R}_{\geq 0}$ such that $f\colon X \longrightarrow Y$ induces a well-defined proper simplicial map $R_{r,s}(f)\colon R_r(X) \longrightarrow R_s(Y)$. Taking the direct limit of these maps yields a morphism

$$QF(f)\colon QF(X) \longrightarrow QF(Y).$$

Using the proper homotopy invariance of F, it is not hard to check that $QF(f) = QF(f')$ if f and f' are finite distance apart.

A straightforward calculation shows that QF indeed is a functor.

We can construct a natural transformation $c_F\colon F \circ I \Longrightarrow QF \circ V$ as follows: If $X \in \operatorname{Ob}(\mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}})$, then X is a subcomplex of $R_1(V(X))$, and so we obtain a natural morphism

$$F(X) \longrightarrow F\bigl(R_1(V(X))\bigr) \longrightarrow \lim_{r\to\infty} F\bigl(R_r(X)\bigr) = QF(X).$$

by composing F applied to the inclusion with the structure morphism in the direct limit.

Moreover, standard arguments show that QF and c_F have the claimed universality property: Let $G\colon \mathsf{QMet} \longrightarrow C$ be a functor. Then for all UDBG spaces X and all $r, s \in \mathbb{R}_{\geq 0}$ with $r \leq s$ the canonical inclusion $V(R_r(X)) \longrightarrow V(R_s(X))$ is a quasi-isometry, and so there is a canonical isomorphism

$$\lim_{r\to\infty} G\bigl(V(R_r(X))\bigr) \cong G\bigl(V(X)\bigr).$$

Therefore, a natural transformation $F \circ I \Longrightarrow G \circ V$ gives rise to the desired natural transformation $QF \Longrightarrow G$. It might seem at first that this construction only gives a natural transformation between $QF \circ V$ and $G \circ V$, but as any UDBG space is quasi-isometric to a space in the image of V, this natural transformation can be extended to UDBG. $\square$

For example, functors from algebraic topology such as the fundamental group, higher homotopy groups or singular/cellular/simplicial homology on the geometric realisation of simplicial complexes have trivial coarsenings because every element can be represented on a subspace of finite diameter (which will be killed by a later stage in the Rips construction).

In contrast, truly locally finite theories do have interesting coarsenings. For example, the coarsening of locally finite simplicial homology with uniformly bounded coefficients coincides with *uniformly finite homology*. Uniformly finite homology also has a suitable concrete description in terms of combinatorial chains that satisfy certain geometric finiteness conditions [12, 96, 77]. Applications of uniformly finite homology include the following:

- Uniformly finite homology allows to characterise so-called amenable groups [12] (Chapter 9.2.4).
- Uniformly finite homology measures the difference between bilipschitz equivalence and quasi-isometry of UDBG spaces [96] (Theorem **??**).
- Uniformly finite homology can be used to show the existence of certain aperiodic tilings [12].
- Uniformly finite homology is a tool that allows to study certain large-scale notions of dimension [26].
- Uniformly finite homology has applications in the context of existence of positive scalar curvature metrics [12].

5.5.3 Geometric properties of groups and rigidity

In geometric group theory, it is common to use the following term:

Definition 5.5.10 (Geometric property of groups)**.** Let P be a property of finitely generated groups (i.e., any finitely generated group either has P

or does not have P; more formally, P is a subclass of the class of finitely generated groups). We say that P is a *geometric property of groups*, in case the following holds for all finitely generated groups G and H: If G has P and H is quasi-isometric to G, then also H has P.

Example 5.5.11 (Geometric properties)**.**
- Being finite is a geometric property of groups (Example 5.2.9).
- Being Abelian is *not* a geometric property of groups: For example, the trivial group and the symmetric group S_3 are quasi-isometric (because they are both finite), but the trivial group is Abelian and S_3 is not Abelian.

Surprisingly, there are many interesting (many of them purely algebraic!) properties of groups that are geometric:
- Being virtually[3] infinite cyclic is a geometric property.
- More generally, for every $n \in \mathbb{N}$ the property of being virtually $\mathbb{Z}^n$ is geometric.
- Being finitely generated and virtually free is a geometric property.
- Being finitely generated and virtually nilpotent is a geometric property of groups.
- Being finitely presented is a geometric property of groups.

Proving that these properties are geometric is far from easy; some of the techniques and invariants needed to prove such statements are explained in later chapters.

That a certain algebraic property of groups turns out to be geometric is an instance of the *rigidity* phenomenon; so, for example, the fact that being virtually infinite cyclic is a geometric property can also be formulated as the group $\mathbb{Z}$ being *quasi-isometrically rigid*.

Conversely, in the following chapters, we will also study geometrically defined properties of finitely generated groups such as hyperbolicity (Chapter 7) and amenability (Chapter 9) and we will investigate how the geometry of these groups affects their algebraic structure.

[3]Let P be a property of groups. A group is *virtually* P if it contains a finite index subgroup that has property P.

6 Growth types of groups

The first quasi-isometry invariant we discuss is the growth type. Basically, we measure the "volume"/size of balls in a given finitely generated group and study the asymptotic behaviour when the radius tends to infinity.

We will start by introducing growth functions for finitely generated groups (with respect to finite generating sets); while these growth functions depend on the chosen finite generating set, a straightforward calculation shows that growth functions for different finite generating sets only differ by a small amount, and more generally that growth functions of quasi-isometric groups are asymptotically equivalent. This leads to the notion of growth type of a finitely generated group.

The quasi-isometry invariance of the growth type allows us to show for many groups that they are not quasi-isometric.

Surprisingly, having polynomial growth is a rather strong constraint for finitely generated groups: By a celebrated theorem of Gromov, all finitely generated groups of polynomial growth are virtually nilpotent(!). We will discuss this theorem in Section 6.3.

6.1 Growth functions of finitely generated groups

We start by introducing growth functions of groups with respect to finite generating sets:

Definition 6.1.1 (Growth function)**.** Let G be a finitely generated group and let $S \subset G$ be a finite generating set of G. Then

$$\begin{aligned} \beta_{G,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto \bigl|B_r^{G,S}(e)\bigr| \end{aligned}$$

is the *growth function of G with respect to S*; here, $B_r^{G,S}(e) := \{g \in G \mid d_S(g,e) \leq r\}$ denotes the (closed) ball of radius r around e with respect to the word metric d_S on G.

Notice that this definition makes sense because balls for word metrics with respect to finite generating sets are finite (Remark 5.2.8).

Example 6.1.2 (Growth functions of groups)**.**

– The growth function of the additive group $\mathbb{Z}$ with respect to the generating set $\{1\}$ clearly is given by

$$\begin{aligned} \beta_{\mathbb{Z},\{1\}} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto 2 \cdot r + 1. \end{aligned}$$

On the other hand, a straightforward induction shows that the growth function of $\mathbb{Z}$ with respect to the generating set $\{2,3\}$ is given by

$$\begin{aligned} \beta_{\mathbb{Z},\{2,3\}} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto \begin{cases} 1 & \text{if } r = 0 \\ 5 & \text{if } r = 1 \\ 6 \cdot r + 1 & \text{if } r > 1. \end{cases} \end{aligned}$$

So, in general, growth functions for different finite generating sets are different.

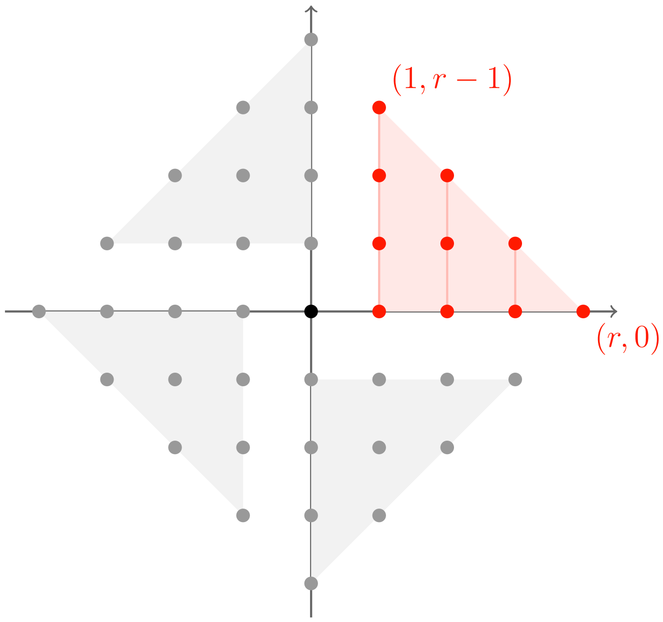

Figure 6.1.: An r-ball in $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$

- The growth function of $\mathbb{Z}^2$ with respect to the standard generating set $S := \{(1,0),(0,1)\}$ is quadratic (see Figure 6.1):

$$\begin{aligned} \beta_{\mathbb{Z}^2,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longrightarrow 1 + 4 \cdot \sum_{j=1}^{r} (r+1-j) = 2 \cdot r^2 + 2 \cdot r + 1. \end{aligned}$$

- More generally, if $n \in \mathbb{N}$, then the growth functions of $\mathbb{Z}^n$ grows like a polynomial of degree n.
- The growth function of the Heisenberg group

$$H = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x,\ y,\ z \in \mathbb{Z} \right\} \cong \bigl\langle x, y, z \bigm| [x,z], [y,z], [x,y] = z \bigr\rangle$$

with respect to the generating set $\{x, y, z\}$ grows like a polynomial of degree 4 (exercise).

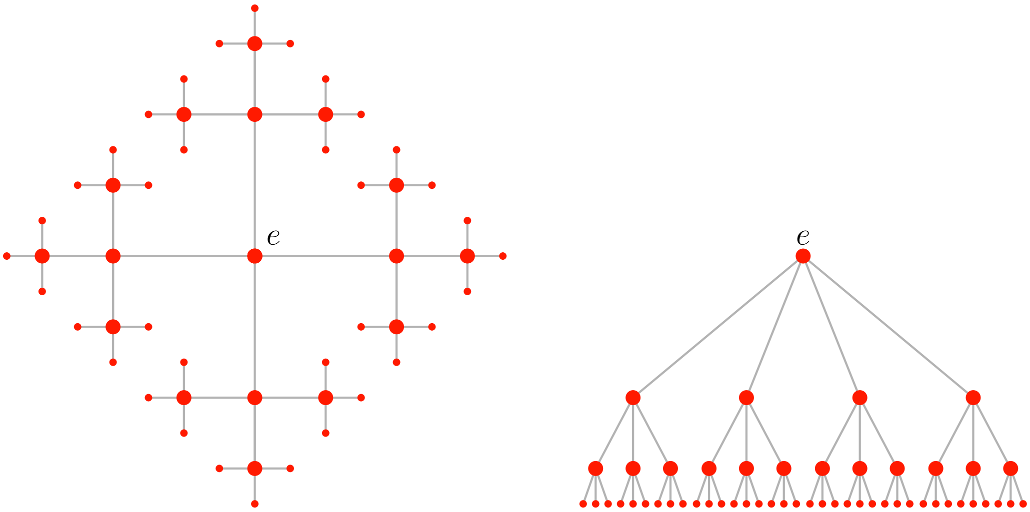

Figure 6.2.: A 3-ball in $\operatorname{Cay}\bigl(F(\{a,b\}),\{a,b\}\bigr)$, drawn in two ways

- The growth function of a free group F of finite rank $n \geq 2$ with respect to a free generating set S is exponential (see Figure 6.2):

$$\begin{aligned}\beta_{F,S} \colon \mathbb{N} &\longrightarrow \mathbb{N}\\ r &\longmapsto 1 + 2 \cdot n \cdot \sum_{j=0}^{r-1} (2 \cdot n - 1)^j = 1 + \frac{n}{n-1} \cdot \bigl((2 \cdot n - 1)^r - 1\bigr).\end{aligned}$$

Proposition 6.1.3 (Basic properties of growth functions)**.** *Let G be a finitely generated group, and let $S \subset G$ be a finite generating set.*

1. Sub-multiplicativity. *For all r, $r' \in \mathbb{N}$ we have*
$$\beta_{G,S}(r + r') \leq \beta_{G,S}(r) \cdot \beta_{G,S}(r').$$
2. A general lower bound. *Let G be infinite. Then $\beta_{G,S}$ is strictly increasing; in particular, $\beta_{G,S}(r) \geq r$ for all $r \in \mathbb{N}$.*
3. A general upper bound. *For all $r \in \mathbb{N}$ we have*
$$\beta_{G,S}(r) \leq \beta_{F(S),S}(r) = 1 + \frac{|S|}{|S| - 1} \cdot \bigl((2 \cdot |S| - 1)^r - 1\bigr).$$

Proof. Ad 1./2. This follows easily from the definition of the word metric d_S on G (exercise).

Ad 3. The homomorphism $\varphi\colon F(S) \longrightarrow G$ characterised by $\varphi|_S = \mathrm{id}_S$ is contracting with respect to the word metrics given by S on $F(S)$ and G respectively. Moreover, φ is surjective. Therefore, we obtain

$$\beta_{G,S}(r) = \left|B_r^{G,S}(e)\right| = \left|\varphi(B_r^{F(S),S}(e))\right| \leq \left|B_r^{F(S),S}(e)\right| = \beta_{F(S),S}(r)$$

for all $r \in \mathbb{N}$. The growth function $\beta_{F(S),S}$ is calculated in Example 6.1.2. □

6.2 Growth types of groups

As we have seen, different finite generating sets in general lead to different growth functions; however, one might suspect already that growth functions coming from different generating sets only differ by some "finite" terms. We therefore introduce the following notion of equivalence for growth functions:

6.2.1 Growth types

Definition 6.2.1 (Quasi-equivalence of (generalised) growth functions)**.**

- A *generalised growth function* is a function of type $\mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ that is increasing.
- Let $f, g\colon \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ be generalised growth functions. We say that g *quasi-dominates* f, if there exist $c, b \in \mathbb{R}_{>0}$ such that

  $$\forall_{r\in\mathbb{R}_{\geq 0}} \quad f(r) \leq c \cdot g(c \cdot r + b) + b.$$

  If g quasi-dominates f, then we write $f \prec g$.

– Two generalised growth functions f, $g\colon \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ are *quasi-equivalent* if both $f \prec g$ and $g \prec f$; if f and g are quasi-equivalent, then we write $f \sim g$.

A straightforward computation shows that quasi-equivalence is an equivalence relation on the set of all generalised growth functions. Quasi-domination then induces a partial order on the set of equivalence classes; however, this partial order clearly is *not* total.

Example 6.2.2 (Generalised growth functions)**.**

– *Monomials.* If $a \in \mathbb{R}_{\geq 0}$, then

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ x &\longmapsto x^a \end{aligned}$$

is a generalised growth function.
For all a, $a' \in \mathbb{R}_{\geq 0}$ we have

$$(x \mapsto x^a) \prec (x \mapsto x^{a'}) \Longleftrightarrow a \leq a',$$

because: If $a \leq a'$, then for all $r \in \mathbb{R}_{\geq 0}$

$$r^a \leq r^{a'} + 1,$$

and so $(x \mapsto x^a) \prec (x \mapsto x^{a'})$.
Conversely, if $a > a'$, then for all c, $b \in \mathbb{R}_{>0}$ we have

$$\lim_{r\to\infty} \frac{r^a}{c \cdot (c \cdot r + b)^{a'} + b} = \infty;$$

thus, for all c, $b \in \mathbb{R}_{>0}$ there is $r \in \mathbb{R}_{\geq 0}$ such that $r^a \geq c\cdot(c\cdot r+b)^{a'}+b$, and so $(x \mapsto x^a) \not\prec (x \mapsto x^{a'})$.
In particular, $(x \mapsto x^a) \sim (a \mapsto x^{a'})$ if and only if $a = a'$.

– *Exponential functions.* If $a \in \mathbb{R}_{>1}$, then

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ x &\longmapsto a^x \end{aligned}$$

is a generalised growth function. A straightforward calculation shows that

$$(x \mapsto a^x) \sim (x \mapsto a'^x)$$

holds for all a, $a' \in \mathbb{R}_{>1}$, as well as

$$(x \mapsto a^x) \succ (x \mapsto x^{a'}) \qquad \text{and} \qquad (x \mapsto a^x) \not\prec (x \mapsto x^{a'})$$

for all $a \in \mathbb{R}_{>1}$ and all $a' \in \mathbb{R}_{\geq 0}$ (exercise). Moreover, there exist generalised growth functions such that

$$\begin{aligned} f \prec (x \mapsto a^x) \quad &\text{and} \quad f \not\sim (x \mapsto a^x), \ \text{and} \\ f \succ (x \mapsto x^{a'}) \quad &\text{and} \quad f \not\sim (x \mapsto x^{a'}) \end{aligned}$$

holds for all $a \in \mathbb{R}_{>1}$, $a' \in \mathbb{R}_{\geq 0}$ (exercise).

Example 6.2.3 (Growth functions yield generalised growth functions)**.** Let G be a finitely generated group, and let $S \subset G$ be a finite generating set. Then the function

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ r &\longrightarrow \beta_{G,S}(\lceil r \rceil) \end{aligned}$$

associated with the growth function $\beta_{G,S} \colon \mathbb{N} \longrightarrow \mathbb{N}$ indeed is a generalised growth function (which is also sub-multiplicative).

If G and H are finitely generated groups with finite generating sets S and T respectively, then we say that the growth function $\beta_{G,S}$ is *quasi-dominated by*/*quasi-equivalent to* the growth function $\beta_{H,T}$ if the associated generalised growth functions are quasi-dominated by/quasi-equivalent to each other.

Notice that then $\beta_{G,S}$ is quasi-dominated by $\beta_{H,T}$ if and only if there exist c, $b \in \mathbb{N}$ such that

$$\forall_{r \in \mathbb{N}} \ \ \beta_{G,S}(r) \leq c \cdot \beta_{H,T}(c \cdot r + b) + b.$$

6.2.2 Growth types and quasi-isometry

We will now show that growth functions of different finite generating sets are quasi-equivalent; more generally, we show that the quasi-equivalence class of growth functions of finite generating sets is a quasi-isometry invariant:

Proposition 6.2.4 (Growth functions and quasi-isometry)**.** *Let G and H be finitely generated groups, and let $S \subset G$ and $T \subset H$ be finite generating sets of G and H respectively.*

1. *If there exists a quasi-isometric embedding $(G, d_S) \longrightarrow (H, d_T)$, then*
$$\beta_{G,S} \prec \beta_{H,T}.$$
2. *In particular, if G and H are quasi-isometric, then the growth functions $\beta_{G,S}$ and $\beta_{H,T}$ are quasi-equivalent.*

Proof. The second part follows directly from the first one (and the definition of quasi-isometry and quasi-equivalence of generalised growth functions).

For the first part, let $f\colon G \longrightarrow H$ be a quasi-isometric embedding; hence, there is a $c \in \mathbb{R}_{>0}$ such that
$$\forall_{g,g' \in G} \quad \frac{1}{c} \cdot d_S(g, g') - c \leq d_T\bigl(f(g), f(g')\bigr) \leq c \cdot d_S(g, g') + c.$$
We write $e' := f(e)$, and let $r \in \mathbb{N}$. Using the estimates above we obtain the following:

– If $g \in B_r^{G,S}(e)$, then $d_T(f(g), e') \leq c \cdot d_S(g, e) + c \leq c \cdot r + c$, and thus
$$f\bigl(B_r^{G,S}(e)\bigr) \subset B_{c\cdot r+c}^{H,T}(e').$$
– For all g, $g' \in G$ with $f(g) = f(g')$, we have
$$d_S(g, g') \leq c \cdot \bigl(d_T(f(g), f(g')) + c\bigr) = c^2.$$

Because the metrics d_S on G and d_T on H are invariant under left translation, it follows that
$$\begin{aligned} \beta_{G,S}(r) &\leq \bigl|B_{c^2}^{G,S}(e)\bigr| \cdot \bigl|B_{c\cdot r+c}^{H,T}(e')\bigr| \\ &= \bigl|B_{c^2}^{G,S}(e)\bigr| \cdot \bigl|B_{c\cdot r+c}^{H,T}(e)\bigr| \\ &= \beta_{G,S}(c^2) \cdot \beta_{H,T}(c \cdot r + c), \end{aligned}$$
which shows that $\beta_{G,S} \prec \beta_{H,T}$ (the term $\beta_{G,S}(c^2)$ does not depend on the radius r). $\square$

Proposition 6.2.4 shows in particular that quasi-equivalence classes of growth functions yield a quasi-isometry invariant with values in the set of quasi-equivalence classes of generalised growth functions. Moreover, this can also be viewed as a functorial quasi-isometry invariant with values in the category associated with the partially ordered set given by quasi-equivalence classes of generalised growth functions with respect to quasi-domination.

In particular, we can define the growth type of finitely generated groups:

Definition 6.2.5 (Growth types of finitely generated groups)**.** Let G be a finitely generated group.

- The *growth type of* G is the (common) quasi-equivalence class of all growth functions of G with respect to finite generating sets of G.
- The group G is of *exponential growth*, if it has the growth type of the exponential map $(x \mapsto e^x)$.
- The group G has *polynomial growth*, if for one (and hence every) finite generating set S of G there is an $a \in \mathbb{R}_{\geq 0}$ such that $\beta_{G,S} \prec (x \mapsto x^a)$.
- The group G is of *intermediate growth*, if it is neither of exponential nor of polynomial growth.

Recall that growth functions of finitely generated groups grow at most exponentially (Proposition 6.1.3), and that polynomials and exponential functions are not quasi-equivalent (Example 6.2.2); hence the term "intermediate growth" does make sense and a group cannot have exponential and polynomial growth at the same time.

We obtain from Proposition 6.2.4 and Example 6.2.2 that having exponential growth/polynomial growth/intermediate growth respectively is a geometric property of groups. More generally:

Corollary 6.2.6 (Quasi-isometry invariance of the growth type)**.** *By Proposition 6.2.4, the growth type of finitely generated groups is a quasi-isometry invariant, i.e., quasi-isometric finitely generated groups have the same growth type.*

In other words: Finitely generated groups having different growth types cannot be quasi-isometric. □

Example 6.2.7 (Growth types)**.**

- If $n \in \mathbb{N}$, then $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$.

– The Heisenberg group has the growth type of $(x \mapsto x^4)$ (exercise).
– Non-Abelian free groups of finite rank have the growth type of the exponential function $(x \mapsto e^x)$.

In other words: The groups $\mathbb{Z}^n$ and the Heisenberg group have polynomial growth, while non-Abelian free groups have exponential growth.

Example 6.2.8 (Quasi-isometry classification of finitely generated Abelian groups)**.** In analogy with topological invariance of dimension, the following holds: We can recover the rank of free Abelian groups from their quasi-isometry type, namely: For all $m, n \in \mathbb{N}$ we have

$$\mathbb{Z}^m \sim_{\mathrm{QI}} \mathbb{Z}^n \iff m = n;$$

this follows by combining Example 6.2.2, Example 6.2.7, and the above corollary.

More generally, combining this observation with Corollary 5.3.10, we obtain: For all finitely generated Abelian groups A and A' we have

$$A \sim_{\mathrm{QI}} A' \iff \mathrm{rk}_{\mathbb{Z}} A = \mathrm{rk}_{\mathbb{Z}} A'.$$

On the other hand, finitely generated Abelian groups admit equal growth functions if and only if they have the same rank and if their torsion subgroups have the same parity [59].

Example 6.2.9 (Distinguishing quasi-isometry types of basic groups)**.**
– Similarly, we obtain for the Heisenberg group H that $H \not\sim_{\mathrm{QI}} \mathbb{Z}^3$, which might be surprising because H fits into a short exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow H \longrightarrow \mathbb{Z}^2 \longrightarrow 1$$

of groups!
– Let F be a non-Abelian free group of finite rank, and let $n \in \mathbb{N}$. Because F grows exponentially, but $\mathbb{Z}^n$ has polynomial growth, we obtain

$$F \not\sim_{\mathrm{QI}} \mathbb{Z}^n \quad \text{and} \quad F \not\sim_{\mathrm{QI}} H.$$

Grigorchuk was the first to show that there indeed exist groups that have intermediate growth [37][42, Chapter VIII]:

Theorem 6.2.10 (Existence of groups of intermediate growth)**.** *There exists a finitely generated group of intermediate growth.*

This group constructed by Grigorchuk [37] can be described via automorphisms of trees or as a so-called automatic group. Furthermore, the group constructed by Grigorchuk also has several other interesting properties [42, Chapter VIII]; for example, it is a finitely generated infinite torsion group, and it is commensurable to the direct product with itself.

Corollary 6.2.11 (Growth of subgroups)**.** *Let G be a finitely generated group, and let H be a finitely generated subgroup of G. If T is a finite generating set of H, and S is a finite generating set of G, then*

$$\beta_{H,T} \prec \beta_{G,S}.$$

Proof. Let $S' := S \cup T$; then S' is a finite generating set of G. Let $r \in \mathbb{N}$; then for all $h \in B_r^{H,T}(e)$ we have

$$d_{S'}(h, e) \leq d_T(h, e) \leq r,$$

and so $B_r^{H,T}(e) \subset B_r^{G,S'}(e)$. In particular,

$$\beta_{H,T}(r) \leq \beta_{G,S'}(r),$$

and thus $\beta_{H,T} \prec \beta_{G,S'}$. Moreover, we know that (G, d_S) and $(G, d_{S'})$ are quasi-isometric, and hence the growth functions $\beta_{G,S'}$ and $\beta_{G,S}$ are quasi-equivalent by Proposition 6.2.4. Therefore, we obtain $\beta_{H,T} \prec \beta_{G,S}$, as desired. □

Example 6.2.12 (Subgroups of exponential growth)**.** Let G be a finitely generated group; if G contains a non-Abelian free subgroup, then G has exponential growth. For instance, it follows that the Heisenberg group does not contain a non-Abelian free subgroup. However, not every finitely generated group of exponential growth contains a non-Abelian free subgroup.

Caveat 6.2.13. The inclusion of a finitely generated subgroup of a finitely generated group into this ambient group in general is *not* a quasi-isometric embedding. For example the inclusion

$$\mathbb{Z} \longrightarrow \bigl\langle x, y, z \bigm| [x, z], [y, z], [x, y] = z \bigr\rangle$$

given by mapping 1 to the generator z of the Heisenberg group is *not* a quasi-isometric embedding: Let $S := \{x, y, z\}$. Then for all $n \in \mathbb{N}$ we have

$$d_S(e, z^{n^2}) = d_S\bigl(e, [x^n, y^n]\bigr) \leq 4 \cdot n;$$

hence, $(n \mapsto d_S(e, z^n))$ does not grow linearly, and so the above inclusion cannot be a quasi-isometric embedding.

6.2.3 Application: Volume growth of manifolds

Whenever we have a reasonable notion of volume on a metric space, we can define corresponding growth functions; in particular, any choice of base point in a Riemannian manifold leads to a growth function. Similarly to the Švarc-Milnor lemma, nice isometric actions of groups give a connection between the growth type of the group acting and the growth type of the metric space acted upon. One instance of this type of results is the following [68]:

Proposition 6.2.14 (Švarc-Milnor lemma for growth types)**.** *Let M be a closed connected Riemannian manifold, let $\widetilde{M}$ be its Riemannian universal covering, and let $x \in \widetilde{M}$. Then the Riemannian volume growth function*

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ r &\longmapsto \operatorname{vol}_{\widetilde{M}} B_r^{\widetilde{M}}(x) \end{aligned}$$

of $\widetilde{M}$ is quasi-equivalent to the growth functions (with respect to one (and hence any) finite generating set) of the fundamental group $\pi_1(M)$ (which is finitely generated by Corollary 5.3.15). Here $B_r^{\widetilde{M}}(x)$ denotes the closed ball in $\widetilde{M}$ of radius r around x with respect to the metric induced by the Riemannian metric on $\widetilde{M}$, and "$\operatorname{vol}_{\widetilde{M}}$" denotes the Riemannian volume with respect to the Riemannian metric on $\widetilde{M}$.

Sketch of proof. By the Švarc-Milnor lemma (Corollary 5.3.15), the map

$$\begin{aligned} \varphi\colon \pi_1(M) &\longrightarrow \widetilde{M} \\ g &\longmapsto g \cdot x \end{aligned}$$

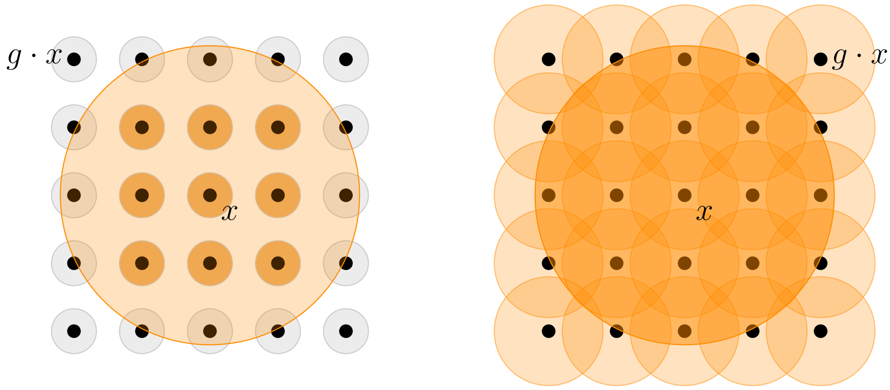

Figure 6.3.: Packing balls into balls, and covering balls by balls

given by the deck transformation action of $\pi_1(M)$ on $\widetilde{M}$ is a quasi-isometry. This allows to translate between radii for balls in $\pi_1(M)$ into radii for balls in $\widetilde{M}$, and vice versa.

- The Riemannian volume growth function of $\widetilde{M}$ at x quasi-dominates the growth functions of $\pi_1(M)$: Because $\pi_1(M)$ acts freely, isometrically, and properly discontinuously on $\widetilde{M}$ there is an $R \in \mathbb{R}_{>0}$ with $d_{\widetilde{M}}(h \cdot x, g \cdot x) \geq R$ for all $g, h \in \pi_1(M)$ with $g \neq h$. Hence, the balls $(B^{\widetilde{M}}_{R/3}(g \cdot x))_{g \in \pi_1(M)}$ are pairwise disjoint (Figure 6.3 (a)). Packing balls and a straightforward computation – using that φ is a quasi-isometric embedding – then proves this claim.
- The Riemannian volume growth function of $\widetilde{M}$ at x is quasi-dominated by the growth functions of $\pi_1(M)$: Because the image of φ is quasi-dense, there is an $R \in \mathbb{R}_{>0}$ such that the balls $(B^{\widetilde{M}}_{R}(g \cdot x))_{g \in \pi_1(M)}$ cover $\widetilde{M}$ (Figure 6.3 (b)). Covering balls and a straightforward compuation – using that φ is a quasi-isometry – then proves this claim.

More details are given in de la Harpes book [42, Proposition VI.36]. □

In particular, we obtain the following obstruction for existence of maps of non-zero degree:

Corollary 6.2.15 (Maps of non-zero degree and growth)**.** *Let M and N be oriented closed connected manifolds of the same dimension and suppose* *that there exists a continuous map $M \longrightarrow N$ of non-zero degree.*

1. *Then*
$$\beta_{\pi_1(M),S} \succ \beta_{\pi_1(N),T}$$
holds for all finite generating sets S and T of $\pi_1(M)$ and $\pi_1(N)$, respectively.
2. *In particular: If M and N are smooth and carry Riemannian metrics, then*
$$\bigl(r \mapsto \operatorname{vol}_{\widetilde{M}} B_r^{\widetilde{M}}(x)\bigr) \succ \bigl(r \mapsto \operatorname{vol}_{\widetilde{N}} B_r^{\widetilde{N}}(y)\bigr)$$
holdsfor all $x \in \widetilde{M}$ and all $y \in \widetilde{N}$.

Associated with a continuous map between oriented closed connected manifolds of the same dimension is an integer, the so-called *mapping degree*. Roughly speaking the mapping degree is the number of preimages (counted with multiplicity and sign) under the given map of a generic point in the target; more precisely, the mapping degree can be defined in terms of singular homology with integral coefficients and the fundamental classes of the manifolds in question [25, Chapter VIII.4].

Sketch of proof. In view of Proposition 6.2.14 it suffices to prove the first part.

Let us recall a standard (but essential) argument from algebraic topology: If $f \colon M \longrightarrow N$ has non-zero degree, then the image G of $\pi_1(M)$ in $\pi_1(N)$ under the induced group homomorphism $\pi_1(f) \colon \pi_1(M) \longrightarrow \pi_1(N)$ has finite index:

By covering theory, there is a connected covering $p \colon \overline{N} \longrightarrow N$ satisfying $\operatorname{im} \pi_1(p) = G$ [64, Theorem V.10.2 and V.4.2]; in particular, $\overline{N}$ also is a connected manifold of dimension $\dim N$ without boundary. By covering theory and construction of G, there exists a continuous map $\overline{f} \colon M \longrightarrow \overline{N}$ with $p \circ \overline{f} = f$ [64, Theorem V.5.1]:

$$\begin{tikzcd} & \overline{N} \arrow[d, "p"] \\ M \arrow[ur, dashed, "\overline{f}"] \arrow[r, "f"'] & N \end{tikzcd}$$

Looking at the induced diagram in singular homology with integral coefficients in top degree $n := \dim M = \dim N = \dim \overline{N}$ shows $H_n(\overline{N};\mathbb{Z}) \not\cong 0$:

$$\begin{tikzcd}
 & \overline{H}_n(N;\mathbb{Z}) \arrow[d, "H_n(p;\mathbb{Z})"] \\
H_n(M;\mathbb{Z}) \arrow[ur, dashed, "H_n(\overline{f};\mathbb{Z})"] \arrow[r, "H_n(f;\mathbb{Z})"'] & H_n(N;\mathbb{Z})
\end{tikzcd}$$

In particular, $\overline{N}$ is compact [25, Corollary VIII.3.4]. On the other hand, $|\deg p|$ coincides with the number of sheets of the covering p [25, Proposition VIII.4.7], which in turn equals the index of $\operatorname{im} \pi_1(p)$ in $\pi_1(N)$ [64, p. 133]. Hence,

$$\big[\pi_1(N) : G\big] = \big[\pi_1(N) : \operatorname{im} \pi_1(p)\big] = |\deg p| < \infty,$$

as claimed.

In particular, G is quasi-isometric to $\pi_1(N)$ and $\pi_1(f)$ provides a surjective homomorphism from $\pi_1(M)$ to G; so the growth functions of $\pi_1(N)$ are quasi-dominated by those of $\pi_1(M)$. ☐

Corollary 6.2.16 (Maps of non-zero degree to hyperbolic manifolds)**.** *In particular: If N is an oriented closed connected hyperbolic manifold and if M is an oriented closed connected Riemannian manifold of the same dimension whose Riemannian universal covering has polynomial or intermediate volume growth, then there is no continuous map $M \longrightarrow N$ of non-zero degree.*

Sketch of proof. This follows from the previous corollary by taking into account that the volume of balls in hyperbolic space $\mathbb{H}^{\dim N} = \widetilde{N}$ grows exponentially with the radius [80, Chapter 3.4]. ☐

In Chapter 7, we will discuss a concept of negative curvature for finitely generated groups, leading to generalisations of Corollary 6.2.16. Alternatively, Corollary 6.2.16 can also be obtained via simplicial volume [56, 57].

6.3 Groups of polynomial growth

One of the central theorems in geometric group theory is Gromov's discovery that groups of polynomial growth can be characterised algebraically as those groups that are virtually nilpotent. The original proof by Gromov [39] was subsequently simplified by van den Dries and Wilkie [27], Kleiner [48], and Shalom and Tao [88].

Theorem 6.3.1 (Gromov's polynomial growth theorem)**.** *Finitely generated groups have polynomial growth if and only if they are virtually nilpotent.*

In Sections 6.3.1 and 6.3.2 we briefly discuss nilpotent groups and their growth properties; in Section 6.3.3, we sketch Gromov's argument why groups of polynomial growth are virtually nilpotent. In the remaining sections, we give some applications of the polynomial growth theorem.

6.3.1 Nilpotent groups

There are two natural ways to inductively take commutator subgroups of a given group, leading to the notion of nilpotent and solvable groups respectively:

Definition 6.3.2 ((Virtually) nilpotent group)**.**

- Let G be a group. For $n \in \mathbb{N}$ we inductively define $C_n(G)$ by

$$C_0(G) := G \qquad \text{and} \qquad \forall_{n\in\mathbb{N}}\ C_{n+1}(G) := \bigl[G, C_n(G)\bigr].$$

  (The sequence $(C_n(G))_{n\in\mathbb{N}}$ is the *lower central series of* G). The group G is *nilpotent*, if there is an $n \in \mathbb{N}$ such that $C_n(G)$ is the trivial group.

– A group is *virtually nilpotent* if it contains a nilpotent subgroup of finite index.

If G is a group and $A, B \subset G$, then $[A, B]$ denotes the subgroup of G *generated* by the set $\{[a, b] \mid a \in A, b \in B\}$ of commutators.

Definition 6.3.3 ((Virtually) solvable group)**.**

– Let G be a group. For $n \in \mathbb{N}$ we inductively define $G^{(n)}$ by

$$G^{(0)} := G \qquad \text{and} \qquad \forall_{n\in\mathbb{N}}\ G^{(n+1)} := [G^{(n)}, G^{(n)}].$$

(The sequence $(G^{(n)})_{n\in\mathbb{N}}$ is the *derived series of* G). The group G is *solvable*, if there is an $n \in \mathbb{N}$ such that $G^{(n)}$ is the trivial group.

– A group is *virtually solvable* if it contains a solvable subgroup of finite index.

Solvable groups owe their name to the fact that a polynomial is solvable by radicals if and only if the corresponding Galois group is solvable [52, Chapter VI.7].

Clearly, the terms of the derived series of a group are subgroups of the corresponding stages of the lower central series; hence, every nilpotent group is solvable.

Example 6.3.4.

– All Abelian groups are nilpotent (and solvable) because their commutator subgroup is trivial.

– The Heisenberg group $H \cong \langle x, y, z \mid [x, z], [y, z], [x, y] = z \rangle$ is nilpotent: We have

$$C_1(H) = [H, H] \cong \langle z \rangle_H,$$

and hence $C_2(H) = [H, C_1(H)] \cong [H, \langle z \rangle_H] = \{e\}$.

– In general, virtually nilpotent groups need not be nilpotent or solvable: For example, every finite group is virtually nilpotent, but not every finite group is nilpotent (for instance the alternating groups A_n are simple and so not even solvable for $n \in \mathbb{N}_{\geq 5}$ [52, Theorem I.5.5]).

– There exist solvable groups that are *not* virtually nilpotent: For example, the semi-direct product

$$\mathbb{Z}^2 \rtimes_\alpha \mathbb{Z},$$

where $\alpha\colon \mathbb{Z} \longrightarrow \operatorname{Aut}(\mathbb{Z}^2)$ is given by the action of the matrix

$$\begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

on $\mathbb{Z}^2$ is such a group (exercise).

– Free groups of rank at least 2 are *not* virtually solvable (exercise).

Nilpotent groups and solvable groups are built up from Abelian groups in a nice way – this corresponds to walking the first steps along the top boundary of the universe of groups (Figure 1.2):

Proposition 6.3.5 (Disassembling nilpotent groups)**.** *Let G be a group, and let $j \in \mathbb{N}$.*

1. *Then $C_{j+1}(G) \subset C_j(G)$ and $C_{j+1}(G)$ is a normal subgroup of $C_j(G)$.*
2. *Moreover, the quotient group $C_j(G)/C_{j+1}(G)$ is Abelian; more precisely, $C_j(G)/C_{j+1}(G)$ is a central subgroup of $G/C_{j+1}(G)$.*

Proof. This follows via a straightforward induction from the definition of the lower central series. □

Except for the last statement on centrality, the analogous statements also holds for the derived series instead of the lower central series.

6.3.2 Growth of nilpotent groups

The growth type of finitely generated nilpotent (and hence of virtually nilpotent) groups can be determined precisely in terms of the lower central series:

Theorem 6.3.6 (Growth type of nilpotent groups)**.** *Let G be a finitely generated nilpotent group, and let $n \in \mathbb{N}$ be minimal with the property that $C_n(G)$ is the trivial group. Then G has polynomial growth of degree*

$$\sum_{j=0}^{n-1} (j+1) \cdot \operatorname{rk}_{\mathbb{Z}} C_j(G)/C_{j+1}(G).$$

Why does the term $\operatorname{rk}_{\mathbb{Z}} C_j(G)/C_{j+1}(G)$ make sense? By Proposition 6.3.5, the quotient group $C_j(G)/C_{j+1}(G)$ is Abelian. Moreover, it can be shown that it is finitely generated (because G is finitely generated [98, Lemma 3.7]). Hence, by the structure theorem of finitely generated Abelian groups, there is a unique $r \in \mathbb{N}$ and a finite Abelian group A (unique up to isomorphims) with

$$C_j(G)/C_{j+1}(G) \cong \mathbb{Z}^r \oplus A;$$

one then defines $\operatorname{rk}_{\mathbb{Z}} C_j(G)/C_{j+1}(G) := r$.

The proof of Theorem 6.3.6 proceeds by induction over the nilpotence degree n [8] and uses a suitable normal form of group elements in terms of the lower central series; the arguments from the computation of the growth rate of the Heisenberg group (Example 6.1.2) give a first impression of how this induction proof works.

Caveat 6.3.7 (Growth type of solvable groups)**.** Even though solvable groups are also built up inductively out of Abelian groups, in general they do *not* have polynomial growth. This follows, for example, from the polynomial growth theorem (Theorem 6.3.1) and the fact that there exist solvable groups that are not virtually nilpotent (Example 6.3.4); moreover, it can also be shown by elementary calculations that the group $\mathbb{Z}^2 \rtimes_A \mathbb{Z}$ given by the action of the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

on $\mathbb{Z}^2$ has exponential growth (exercise).

Wolf [98] and Milnor [68] used algebraic means (similar to the calculations in Caveat 6.3.7) to prove the following predecessor of the polynomial growth theorem:

Theorem 6.3.8 (Growth type of solvable groups)**.** *A finitely generated solvable group has polynomial growth if and only if it is virtually nilpotent.*

Notice that this theorem is needed in all proofs of the polynomial growth theorem known so far.

6.3.3 Groups of polynomial growth are virtually nilpotent

We will now sketch Gromov's argument why finitely generated groups of polynomial growth are virtually nilpotent, following the exposition by van den Dries and Wilkie [27]:

The basic idea behind the proof is to proceed by induction over the degree of polynomial growth. In the following, let G be a finitely generated group of polynomial growth, say of polynomial growth of degree at most d with $d \in \mathbb{N}$.

In the case $d = 0$ the growth functions of G are bounded functions, and so G must be finite. In particular, G is virtually trivial, and so virtually nilpotent.

For the induction step we assume $d > 0$ and that we know already that all finitely generated groups of polynomial growth of degree at most $d - 1$ are virtually nilpotent. Moreover, we may assume without loss of generality that G is infinite.

The key to the inductive argument is the following theorem by Gromov [39]:

Theorem 6.3.9. *If G is a finitely generated infinite group of polynomial growth, then there exists a subgroup G' of G of finite index that admits a surjective homomorphism $G' \longrightarrow \mathbb{Z}$.*

In fact, the proof of this theorem is the lion share of the proof of the polynomial growth theorem. The alternative proofs of van den Dries and Wilkie, Kleiner, and Tao and Shalom mainly give different proofs of Theorem 6.3.9. We will now briefly sketch Gromov's argument:

Sketch of proof of Theorem 6.3.9. Gromov's cunning proof roughly works as follows: Let $S \subset G$ be a finite generating set. We then consider the sequence

$$\Bigl(G, \frac{1}{n} \cdot d_S\Bigr)_{n\in\mathbb{N}}$$

of metric spaces; this sequence models what happens when we move far away from the group. If G has polynomial growth, then Gromov proves that this sequence has a subsequence converging in an appropriate sense to a "nice" metric space Y [39]. Using the solution of Hilbert's fifth problem [69, 92], one can show that the isometry group of Y is a Lie group, and so is closely related to $\mathrm{GL}(n,\mathbb{C})$. Moreover, it can be shown that some finite index subgroup G' of G acts on Y in such a way that results on Lie groups (e.g., the Tits alternative for $\mathrm{GL}(n,\mathbb{C})$ (Remark 4.5.2)) allow to construct a surjective homomorphism $G' \longrightarrow \mathbb{Z}$.

A detailed proof is given in the paper by van den Dries and Wilkie [27]. Gromov's considerations of the sequence $(G, 1/n \cdot d_S)_{n\in\mathbb{N}}$ are a precursor of so-called asymptotic cones [28]. □

In view of Theorem 6.3.9 we can assume without loss of generality that our group G admits a surjective homomorphism $\pi\colon G \longrightarrow \mathbb{Z}$. Using such a homomorphism, we find a subgroup of G of lower growth rate inside of G:

Proposition 6.3.10 (Subgroup of lower growth rate)**.** *Let $d \in \mathbb{N}$ and let G be a finitely generated group af polynomial growth at most d that admits a surjective homomorphism $\pi\colon G \longrightarrow \mathbb{Z}$. Let $K := \ker \pi$.*

1. *Then the subgroup K is finitely generated.*
2. *The subgroup K is of polynomial growth of degree at most $d-1$.*

Proof. Ad 1. Exercise.

Ad 2. In view of the first part we find a finite generating set $S \subset G$ that contains a finite generating set $T \subset K$ of K and an element $g \in S$ with $\pi(g) = 1 \in \mathbb{Z}$. Let $c \in \mathbb{R}_{>0}$ with

$$\forall_{r\in\mathbb{N}} \;\; \beta_{G,S}(r) \leq c \cdot r^d.$$

Now let $r \in \mathbb{N}$, let $N := \beta_{K,T}(\lfloor r/2 \rfloor)$, and let $k_1, \dots, k_N \in K$ be the N elements of the ball $B^{K,T}_{\lfloor r/2\rfloor}(e)$. Then the elements

$$k_j \cdot g^s \quad \text{with } j \in \{1, \dots, N\} \text{ and } s \in \{-\lfloor r/2 \rfloor, \dots, \lfloor r/2 \rfloor\}$$

of G are all distinct, and have S-length at most r. Hence,

$$\beta_{K,T}(\lfloor r/2 \rfloor) \leq \frac{\beta_{G,S}(r)}{r} \leq c \cdot r^{d-1},$$

and therefore, $\beta_{K,T} \prec (r \mapsto r^{d-1})$. □

By Proposition 6.3.10, the subgroup $K := \ker \pi$ of G has polynomial growth of degree at most $d - 1$. Hence, by induction, we can assume that K is virtually nilpotent. Now an algebraic argument shows that the extension G of $\mathbb{Z}$ by K is a virtually *solvable* group [27, Lemma 2.1]. Because G has polynomial growth, Theorem 6.3.8 lets us deduce that G indeed is virtually nilpotent, as desired.

6.3.4 Application: Being virtually nilpotent is a geometric property

As a first application we show that being virtually nilpotent is a geometric property of finitely generated groups in the sense of Definition 5.5.10. Notice that from the algebraic definition of being virtually nilpotent it is not clear at all that this property is preserved under quasi-isometries.

Corollary 6.3.11. *Being virtually nilpotent is a geometric property of finitely generated groups.*

Proof. In view of Gromov's polynomial growth theorem (Theorem 6.3.1) for finitely generated groups being virtually nilpotent and having polynomial growth are equivalent. On the other hand, having polynomial growth is a geometric property (Corollary 6.2.6). □

6.3.5 Application: More on polynomial growth

A priori it is not clear that a finitely generated group having polynomial growth has the growth type of $(r \mapsto r^d)$, where d is a *natural number*.

Corollary 6.3.12. *Let G be a finitely generated group of polynomial growth. Then there is a $d \in \mathbb{N}$ such that*

$$\beta_{G,S} \sim (r \mapsto r^d)$$

holds for all finite generating sets S of G.

Proof. By the polynomial growth theorem (Theorem 6.3.1), the group G is virtually nilpotent. Therefore, G has polynomial growth of degree

$$d := \sum_{j=0}^{n-1} (j+1) \cdot \operatorname{rk}_{\mathbb{Z}} C_j(G)/C_{j+1}(G)$$

by Theorem 6.3.6 (where n denotes the degree of nilpotency of G). In particular, $\beta_{G,S} \sim (r \mapsto r^d)$ for all finite generating sets S of G. $\square$

6.3.6 Application: Quasi-isometry rigidity of free Abelian groups

Gromov's polynomial growth theorem can be used to show that finitely generated free Abelian groups are quasi-isometrically rigid:

Corollary 6.3.13 (Quasi-isometry rigidity of $\mathbb{Z}$)**.** *Let G be a finitely generated group quasi-isometric to $\mathbb{Z}$. Then G is virtually infinite cyclic.*

Proof. Because G is quasi-isometric to $\mathbb{Z}$, the group G has linear growth. In particular, G is virtually nilpotent by the polynomial growth theorem (Theorem 6.3.1). Let $H \subset G$ be a nilpotent subgroup of finite index; so H has linear growth as well. By Bass's theorem on the growth rate of nilpotent groups (Theorem 6.3.6) it follows that

$$1 = \sum_{j=0}^{n-1} (j+1) \cdot \operatorname{rk}_{\mathbb{Z}} C_j(H)/C_{j+1}(H),$$

where n is the degree of nilpotency of H. Because $\operatorname{rk}_{\mathbb{Z}}$ takes values in $\mathbb{N}$, it follows that

$$1 = \operatorname{rk}_{\mathbb{Z}} C_0(H)/C_1(H) \qquad \text{and} \qquad \forall_{j \in \mathbb{N}_{\geq 1}} \ 0 = \operatorname{rk}_{\mathbb{Z}} C_j(H)/C_{j+1}(H).$$

The classification of finitely generated Abelian groups shows that finitely generated Abelian groups of rank 0 are finite and that finitely generated Abelian groups of rank 1 are virtually $\mathbb{Z}$. So $C_1(H)$ is finite, and the quotient $C_0(H)/C_1(H)$ is Abelian and virtually $\mathbb{Z}$. Then also $H = C_0(H)$ is virtually $\mathbb{Z}$. In particular, G is virtually $\mathbb{Z}$.

We will see a more elementary proof in Chapter 7 (Corollary 7.5.8). $\square$

Corollary 6.3.14 (Quasi-isometry rigidity of $\mathbb{Z}^n$)**.** *Let $n \in \mathbb{N}$. Then every finitely generated group quasi-isometric to $\mathbb{Z}^n$ is virtually $\mathbb{Z}^n$.*

Sketch of proof. The proof is similar to the proof of quasi-isometry rigidity of $\mathbb{Z}$ above, but it needs in addition a description of the growth rate of virtually nilpotent groups in terms of their Hirsch rank [17, p. 149f]. □

Notice that no elementary proof of Corollary 6.3.14 seems to be known so far and that a full quasi-isometry classification of virtually nilpotent groups is still out of reach.

6.3.7 Application: Expanding maps of manifolds

We conclude this chapter by discussing Gromov's geometric application [39, Geometric corollary on p. 55] of the polynomial growth theorem (Theorem 6.3.1) to infra-nil-endomorphisms:

Corollary 6.3.15. *Every expanding self-map of a compact Riemannian manifold is topologically conjugate to an infra-nil-endomorphism.*

Before sketching the proof of this strong geometric rigidity result, we briefly explain the geometric terms:

A map $f \colon X \longrightarrow Y$ between metric spaces (X, d_X) and (Y, d_Y) is *globally expanding* if

$$\forall_{x,x' \in X} \;\; x \neq x' \Longrightarrow d_Y\bigl(f(x), f(x')\bigr) > d_X(x, x').$$

A map $f \colon X \longrightarrow Y$ is *expanding* if every point of X has a neighbourhood U such that the restriction $f|_U \colon U \longrightarrow Y$ is expanding. As Riemannian manifolds can be viewed as metric spaces, we obtain a notion of expanding maps of Riemannian manifolds.

As a simple example, let us consider the n-dimensional torus $\mathbb{Z}^n \setminus \mathbb{R}^n$. A straightforward calculation shows that a linear map $f \colon \mathbb{R}^n \longrightarrow \mathbb{R}^n$ with $f(\mathbb{Z}^n) \subset \mathbb{Z}^n$ induces a self-map $\mathbb{Z}^n \setminus \mathbb{R}^n \longrightarrow \mathbb{Z}^n \setminus \mathbb{R}^n$ and that this self-map is expanding if and only if all complex eigenvalues of f have absolute value bigger than 1.

A *nil-manifold* is a compact Riemannian manifold that can be obtained as a quotient $\Gamma \setminus N$, where N is a simply connected nilpotent Lie group and $\Gamma \subset N$ is a cocompact lattice. More generally, an *infra-nil-manifold* is a compact Riemannian manifold that can be obtained as a quotient $\Gamma \setminus N$, where N is a simply connected nilpotent Lie group and Γ is a subgroup of the group of all isometries of N generated by left translations of N and all automorphisms of N. Clearly, all nil-manifolds are also infra-nil-manifolds, and it can be shown that every infra-nil-manifold is finitely covered by a nil-manifold.

Let $\Gamma \setminus N$ be such an infra-nil-manifold. An expanding *infra-nil-endomorphism* is an expanding map $\Gamma \setminus N \longrightarrow \Gamma \setminus N$ that is induced by an expanding automorphism $N \longrightarrow N$ of the Lie group N.

For example, all tori and the quotient $H \setminus H_{\mathbb{R}}$ of the Heisenberg group $H_{\mathbb{R}}$ with real coefficients by the Heisenberg group H are nil-manifolds (and so also infra-nil-manifolds). The expanding maps on tori mentioned above are examples of expanding infra-nil-endomorphisms.

Two self-maps $f\colon X \longrightarrow X$ and $g\colon Y \longrightarrow Y$ between topological spaces are *topologically conjugate* if there exists a homeomorphism $h\colon X \longrightarrow Y$ with $h \circ f = g \circ h$, i.e., which fits into a commutative diagram:

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, dashed, "h"'] & X \arrow[d, dashed, "h"] \\ Y \arrow[r, "g"'] & Y \end{tikzcd}$$

Sketch of proof of Corollary 6.3.15. By a theorem of Franks [34], if a compact Riemannian manifold M admits an expanding self-map, then the Riemannian universal covering $\widetilde{M}$ has polynomial volume growth; hence, by the Švarc-Milnor lemma (Proposition 6.2.14), the fundamental group $\pi_1(M)$ is finitely generated and has polynomial growth as well.

In view of the polynomial growth theorem (Theorem 6.3.1), we obtain that $\pi_1(M)$ is virtually nilpotent.

By a result of Shub [89, 34], this implies that any expanding self-map of M is topologically conjugate to an infra-nil-endomorphism. $\square$

7 Hyperbolic groups

On the side of the universe of groups (Figure 1.2) opposite to Abelian, nilpotent, solvable and amenable groups, we find free groups, and then further down negatively curved groups. In this chapter, we look at the world of negatively curved groups and their astonishing properties.

We start with a brief reminder of classical curvature of plane curves and of surfaces in $\mathbb{R}^3$ (Section 7.1).

In Section 7.2, we study Gromov's extension of the notion of negative curvature to large scale geometry via slim triangles; this requires a careful discussion of the relation between geodesics and quasi-geodesics in hyperbolic spaces (Section 7.2.3).

In particular, we obtain a notion of negatively curved finitely generated groups, so-called hyperbolic groups (Section 7.3).

We spend the rest of the chapter studying algebraic consequences of the geometric condition of groups of being hyperbolic: We show that the word problem is solvable for hyperbolic groups (Section 7.4). Moreover, we will investigate the role and properties of elements of infinite order in hyperbolic groups (Section 7.5).

7.1 Classical curvature, intuitively

A central type of invariants in Riemannian geometry are curvature invariants. Classically, curvatures in Riemannian geometry are defined in terms of *local* data; however, some types of curvature constraints also influence the *global* shape. An example of such a situation is having everywhere negative sectional curvature.

What is curvature? Roughly speaking, curvature measures how much a space bends at a given point, i.e., how far away it is from being a “flat” Euclidean space.

In the following, we give a brief introduction into curvature in Riemannian geometry; however, instead of going into the details of bundles, connections, curvature tensors and related machinery, we rely on a graphic and intuitive description based on curves. Readers interested in concise and mathematically precise definitions of the various types of curvature are referred to the literature on Riemannian geometry, for instance to the pleasant book *Riemannian manifolds. An introduction to curvature* by Lee [54].

7.1.1 Curvature of plane curves

As a first step, we briefly describe curvature of curves in the Euclidean plane $\mathbb{R}^2$: Let $\gamma\colon [0,L] \longrightarrow \mathbb{R}^2$ be a smooth curve, parametrised by arc-length, and let $t \in (0,L)$. Geometrically, the *curvature $\kappa_\gamma(t)$ of γ at t* can be described as follows (see also Figure 7.1): We consider the set of all (parametrised in mathematically positive orientation) circles in $\mathbb{R}^2$ that are tangent to $\gamma(t)$. It can be shown that this set contains exactly one circle that at the point $\gamma(t)$ has the same acceleration vector as γ; the so-called

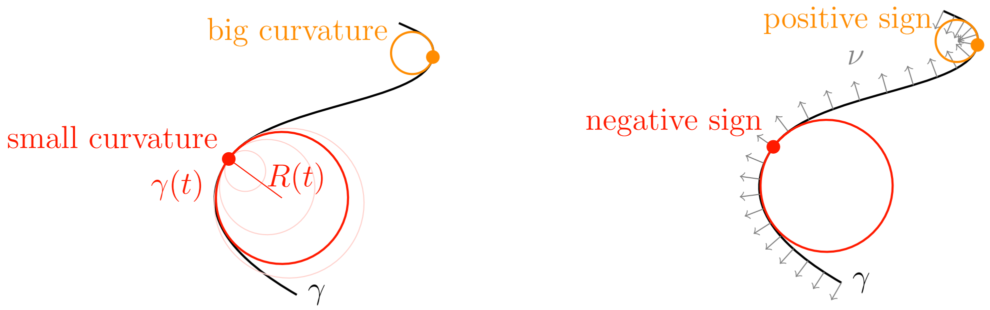

Figure 7.1.: Curvature of a plane curve

osculating circle of γ *at* t. Then the curvature of γ at t is defined as

$$\kappa_\gamma(t) := \frac{1}{R(t)},$$

where $R(t)$ is the radius of the osculating circle of γ at t. I.e., the smaller the curvature, the bigger is the osculating circle, and so the curve is rather close to being a straight line at this point; conversely, the bigger the curvature, the more the curve bends at this point.

More technically, if γ is parametrised by arc-length, then for all $t \in (0, L)$ the curvature of γ at t can be expressed as

$$\kappa_\gamma(t) = |\ddot{\gamma}(t)|.$$

By introducing a reference normal vector field along the curve γ, we can also add a sign to the curvature, describing in which direction the curve bends relative to the chosen normal vector field (see also Figure 7.1). Let $\nu\colon [0, L] \longrightarrow \mathbb{R}^2$ be a non-vanishing normal vector field along γ, i.e., ν is a smooth map satisfying

$$\nu(t) \perp \dot{\gamma}(t) \qquad \text{and} \qquad \nu(t) \neq 0$$

for all $t \in (0, L)$. The *signed curvature* $\widetilde{\kappa}_{\gamma,\nu}(t)$ *of* γ *at* t *with respect to* ν is

- defined to be $+\kappa_\gamma(t)$ if the normal vector $\nu(t)$ points from $\gamma(t)$ in the direction of the centre of the osculating circle of γ at t, and it is
- defined to be $-\kappa_\gamma(t)$ if the normal vector $\nu(t)$ points from $\gamma(t)$ away from the centre of the osculating circle of γ at t.

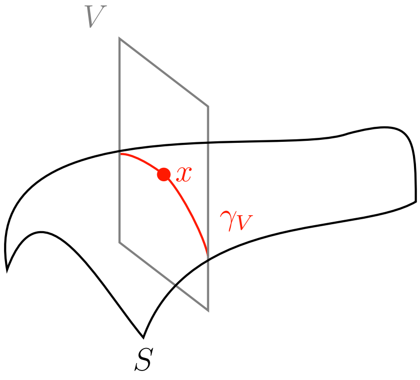

Figure 7.2.: Gaussian curvature of a surface via curves

7.1.2 Curvature of surfaces in $\mathbb{R}^3$

As second step, we have a look at the curvature of surfaces embedded in the Euclidean space $\mathbb{R}^3$: Let $S \subset \mathbb{R}^3$ be a smooth surface embedded into $\mathbb{R}^3$, and let $x \in S$ be a point on S. Then the *curvature of S at x* is defined as follows (see also Figure 7.2):

1. We choose a non-vanishing normal vector field ν on S in a neighbourhood U of x in S.
2. For every affine plane $V \subset \mathbb{R}^3$ containing x (spanned by $\nu(x)$ and a tangent vector at x), let γ_V be the component of "the" curve in S given by the intersection $U \cap V$ that passes through x, and let $\nu|_{\gamma_V}$ be the induced normal vector field on γ_V in V.
3. The *principal curvatures of S at x with respect to ν* are given by
$$\kappa^+_{S,\nu}(x) := \sup_V \widetilde{\kappa}_{\gamma_V,\nu|_{\gamma_V}}(x) \qquad \text{and} \qquad \kappa^-_{S,\nu}(x) := \inf_V \widetilde{\kappa}_{\gamma_V,\nu|_{\gamma_V}}(x).$$
4. The *(Gaussian) curvature of S at x* is defined as
$$\kappa_S(x) := \kappa^+_{S,\nu}(x) \cdot \kappa^-_{S,\nu}(x);$$
notice that $\kappa_S(x)$ does *not* depend on the choice of the normal vector field ν.

It can be shown that while the principal curvatures are not intrinsic invariants of a surface (i.e., they are in general *not* invariant under isometries) [54, p. 6], the Gaussian curvatures of a surface are intrinsic (Theorema Egregium [54, Chapter 8]).

Example 7.1.1. The following examples are illustrated in Figure 7.3.

- The sphere $S^2 \subset \mathbb{R}^3$ has everywhere positive Gaussian curvature because the principal curvatures at every point are non-zero and have the same sign.
- The plane $\mathbb{R}^2 \subset \mathbb{R}^3$ has everywhere vanishing Gaussian curvature (i.e., it is "flat") because all principal curvatures are 0.
  The cylinder $S^1 \times \mathbb{R} \subset \mathbb{R}^3$ has everywhere vanishing Gaussian curvature because at every point one of the principal curvatures is 0.
- Saddle-shapes in $\mathbb{R}^3$ have points with negative Gaussian curvature because at certain points the principal curvatures are non-zero and have opposite signs.
  Moreover, one can show that the hyperbolic plane has everywhere negative Gaussian curvature (Appendix **??**).
  Particularly nice illustrations of the geometry of the hyperbolic plane (which is everywhere negatively curved) are provided by Escher's *Cirkellimiet* works [32, 33].

Remark 7.1.2 (Curvature of Riemannian manifolds)**.** How can one define curvature of Riemannian manifolds? Let M be a Riemannian manifold, and let $x \in M$. For simplicity, let us assume that M is embedded into some Euclidean space $\mathbb{R}^n$. For every plane V tangent to M at x we can define a curvature: Taking all geodesics in M starting in x and tangent to V defines a surface S_V in M that inherits a Riemannian metric from the Riemannian metric on M; then the *sectional curvature of M at x with respect to V* is the Gaussian curvature of the surface S_V at x. Sectional curvature in fact is an intrinsic property of Riemannian manifold (i.e., independent of the choice of any embedding into Euclidean space) and can be described analytically in terms of certain tensors on M.

Taking suitable averages of sectional curvatures leads to the weaker curvature notions of *Ricci curvature* and *scalar curvature* respectively. For more details we refer to the book of Lee [54].

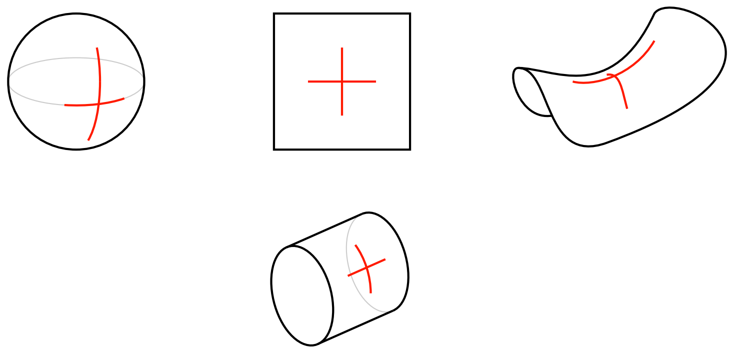

positive flat negative

Figure 7.3.: Examples of Gaussian curvatures of surfaces

However, by definition, the Gaussian curvatures are defined in terms of the local structure of a surface, and so are unsuited for a notion of curvature in large-scale geometry. However, surprisingly, e.g., negatively curved surfaces share certain global properties, and so it is conceivable that it is possible to define a notion of negative curvature that makes sense in large-scale geometry. In order to see how this can be done, we look at geodesic triangles in surfaces in $\mathbb{R}^3$ (Figure 7.4), i.e., at triangles in surfaces whose sides are goedesics in the surface in question.

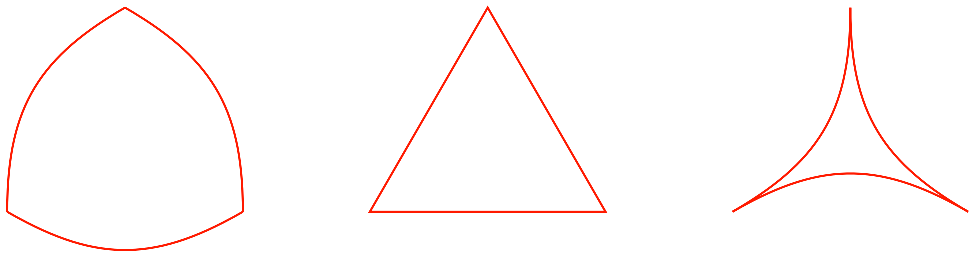

positive flat negative

Figure 7.4.: Geodesic triangles in surfaces

In positively curved spaces, goedesic triangles are "thicker" than in Euclidean space, while in negatively curved spaces, geodesic triangles are "thinner" than in Euclidean space.

7.2 (Quasi-)Hyperbolic spaces

Gromov and Rips realised that the global geometry of negatively curved spaces can be captured by the property that all geodesic triangles are slim [40]. Taking slim triangles as the defining property leads to the notion of (Gromov) hyperbolic spaces.

We start by explaining the notion of hyperbolicity for geodesic spaces (Section 7.2.1); as next step, we generalise this notion to quasi-hyperbolic spaces (Section 7.2.2). Finally, in Section 7.2.3, we discuss some technical aspects of geodesics in (quasi-)hyperbolic spaces.

7.2.1 Hyperbolic spaces

Taking slim geodesic triangles as defining property leads to the notion of (Gromov) hyperbolic spaces:

Definition 7.2.1 (δ-Slim geodesic triangle)**.** Let (X, d) be a metric space.

- A *geodesic triangle* in X is a triple $(\gamma_0, \gamma_1, \gamma_2)$ consisting of geodesics $\gamma_j \colon [0, L_j] \longrightarrow X$ in X such that

$$\gamma_0(L_0) = \gamma_1(0), \qquad \gamma_1(L_1) = \gamma_2(0), \qquad \gamma_2(L_2) = \gamma_0(0).$$

- A geodesic triangle $(\gamma_0, \gamma_1, \gamma_2)$ is δ-*slim* if (Figure 7.5)

$$\begin{aligned}
\operatorname{im}\gamma_0 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2\bigr),\\
\operatorname{im}\gamma_1 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2\bigr),\\
\operatorname{im}\gamma_2 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1\bigr).
\end{aligned}$$

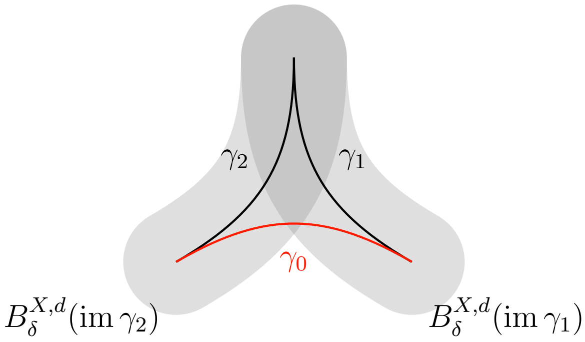

Figure 7.5.: A δ-slim triangle

Here, for $\gamma\colon [0,L] \longrightarrow X$ we use the abbreviation $\operatorname{im}\gamma := \gamma([0,L])$, and for $A \subset X$ we write

$$B_\delta^{X,d}(A) := \Bigl\{x \in X \Bigm| \inf_{a\in A} d(x,a) \leq \delta\Bigr\}.$$

Definition 7.2.2 (δ-Hyperbolic space)**.** Let X be a metric space.

- Let $\delta \in \mathbb{R}_{\geq 0}$. We say that X is δ-*hyperbolic* if X is geodesic and if all geodesic triangles in X are δ-slim.
- The space X is *hyperbolic* if there exists a $\delta \in \mathbb{R}_{\geq 0}$ such that X is δ-hyperbolic.

Example 7.2.3 (Hyperbolic spaces)**.**

- Any geodesic metric space X of finite diameter is $\operatorname{diam}(X)$-hyperbolic.
- The real line $\mathbb{R}$ is 0-hyperbolic because every geodesic triangle in $\mathbb{R}$ is degenerate.
- The Euclidean space $\mathbb{R}^2$ is *not* hyperbolic because for $\delta \in \mathbb{R}_{\geq 0}$, the Euclidean triangle with vertices $(0,0)$, $(0, 3\cdot\delta)$, and $(3\cdot\delta, 0)$ (where the sides are parametrised isometrically) is not δ-slim (Figure 7.6).
- The hyperbolic plane $\mathbb{H}^2$ is a hyperbolic metric space in the sense of Definition 7.2.2 [17]. More generally, if M is a closed connected compact Riemannian manifold of negative sectional curvature (e.g., hyperbolic manifolds), then the Riemannian universal covering of M is hyperbolic in the sense of Definition 7.2.2.
- If T is a tree, then the geometric realisation $|T|$ of T (Appendix A.3) is 0-hyperbolic because all geodesic triangles in $|T|$ are degenerate

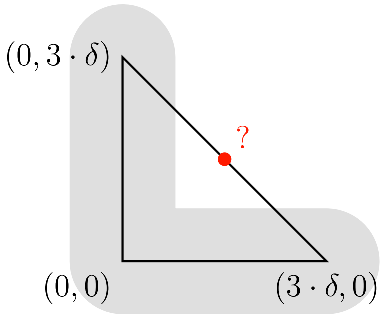

Figure 7.6.: The Euclidean plane $\mathbb{R}^2$ is *not* hyperbolic

tripods (exercise, see also Proposition 7.2.17 below); in a sense, hyperbolic spaces can be viewed as thickenings of metric trees.

Caveat 7.2.4 (Quasi-isometry invariance of hyperbolicity)**.** Notice that from the definition it is not clear that hyperbolicity is quasi-isometry invariant (among geodesic spaces), because the composition of a geodesic triangle with a quasi-isometry in general is only a quasi-geodesic triangle and not a geodesic triangle. However, we will see in Corollary 7.2.13 below that hyperbolicity indeed is quasi-isometry invariant among geodesic spaces.

7.2.2 Quasi-hyperbolic spaces

We will now translate the definition of hyperbolicity to the world of quasi-geometry; so instead of geodesics we consider quasi-geodesics. In particular, we thereby obtain a notion of quasi-hyperbolicity for quasi-geodesic spaces. On the one hand, it will be immediate from the definition that quasi-hyperbolicity indeed is a quasi-isometry invariant (Proposition 7.2.9); on the other hand, we will relate hyperbolicity of geodesic spaces to quasi-hyperbolicity (Theorem 7.2.10), which shows that also hyperbolicity is a quasi-isometry invariant in the class of geodesic spaces (Corollary 7.2.13).

Definition 7.2.5 (δ-Slim quasi-geodesic triangle)**.** Let (X, d) be a metric space, and let $c, b \in \mathbb{R}_{>0}$, $\delta \in \mathbb{R}_{\geq 0}$.

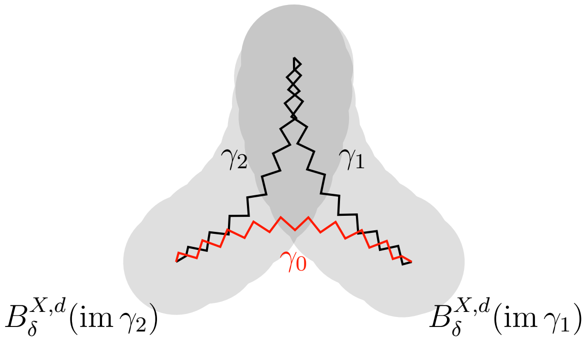

Figure 7.7.: A quasi-slim triangle

- A (c,b)-*quasi-geodesic triangle* in X is a triple $(\gamma_0, \gamma_1, \gamma_2)$ consisting of (c,b)-quasi-geodesics $\gamma_j\colon [0, L_j] \longrightarrow X$ in X such that

$$\gamma_0(L_0) = \gamma_1(0), \qquad \gamma_1(L_1) = \gamma_2(0), \qquad \gamma_2(L_2) = \gamma_0(0).$$

- A (c,b)-quasi-geodesic triangle $(\gamma_0, \gamma_1, \gamma_2)$ is δ-*slim* if (Figure 7.7)

$$\begin{aligned}\operatorname{im}\gamma_0 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_1 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_2 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1\bigr).\end{aligned}$$

Definition 7.2.6 (Quasi-hyperbolic space)**.** Let (X,d) be a metric space.
- Let c, $b \in \mathbb{R}_{>0}$, $\delta \in \mathbb{R}_{\geq 0}$. We say that X is (c,b,δ)-*quasi-hyperbolic* if X is (c,b)-quasi-hyperbolic and all (c,b)-quasi-geodesic triangles in X are δ-slim.
- Let c, $b \in \mathbb{R}_{>0}$. The space X is (c,b)-*quasi-hyperbolic* if for every c', $b' \in \mathbb{R}_{\geq 0}$ with $c' \geq c$ and $b' \geq b$ there exists a $\delta \in \mathbb{R}_{\geq 0}$ such that X is (c', b', δ)-quasi-hyperbolic.
- The space X is *quasi-hyperbolic* if there exist c, $b \in \mathbb{R}_{>0}$ such that X is (c,b)-quasi-hyperbolic.

Example 7.2.7. Clearly, all metric spaces of finite diameter are quasi-hyperbolic.

Caveat 7.2.8. In general, it is rather difficult to show that a space is quasi-hyperbolic by showing that all quasi-geodesic triangles in question are slim enough, because there are too many quasi-geodesics. Using Corollary 7.2.13 below simplifies this task considerably in case we know that the space in question is quasi-isometric to an accessible geodesic space. This will give rise to a large number of interesting quasi-hyperbolic spaces.

Proposition 7.2.9 (Quasi-isometry invariance of quasi-hyperbolicity)**.** *Let X and Y be metric spaces.*

1. *If Y is quasi-geodesic and if X and Y are quasi-isometric, then also X is quasi-geodesic.*
2. *If Y is quasi-hyperbolic and X is quasi-geodesic and if there exists a quasi-isometric embedding $X \longrightarrow Y$, then also X is quasi-hyperbolic.*
3. *In particular: If X and Y are quasi-isometric, then X is quasi-hyperbolic if and only if Y is quasi-hyperbolic.*

Proof. The proof consists of pulling back and pushing forward quasi-geodesics along quasi-isometric embeddings.

We start by proving the first part. Let Y be quasi-geodesic and suppose that $f\colon X \longrightarrow Y$ is a quasi-isometric embedding; let $c \in \mathbb{R}_{\geq 0}$ be so large that Y is (c, c)-quasi-geodesic and such that f is a (c, c)-quasi-isometric embedding with c-dense image. Furthermore, let x, $x' \in X$. Then there is a (c, c)-quasi-geodesic $\gamma\colon [0, L] \longrightarrow Y$ joining $f(x)$ and $f(x')$. Using the axiom of choice and the fact that f has c-dense image, we can find a map

$$\widetilde{\gamma}\colon [0, L] \longrightarrow X$$

such that $\widetilde{\gamma}(0) = x$, $\widetilde{\gamma}(L) = x'$, and

$$d_Y\bigl(f \circ \widetilde{\gamma}(t), \gamma(t)\bigr) \leq c$$

for all $t \in [0, L]$. The same arguments as in the proof of Proposition 5.1.10 show that $\widetilde{\gamma}$ is a $(c, \max(3 \cdot c^2, 3))$-quasi-geodesic joining x and x'. Hence, X is $(c, \max(3 \cdot c^2, 3))$-quasi-geodesic.

As for the second part, suppose that Y is quasi-hyperbolic, that X is quasi-geodesic, and that $f\colon X \longrightarrow Y$ is a quasi-isometric embedding. Hence, there are c, $b \in \mathbb{R}_{>0}$ such that Y is (c, b)-quasi-hyperbolic, such that X is (c, b)-quasi-geodesic, and such that f is a (c, b)-quasi-isometric

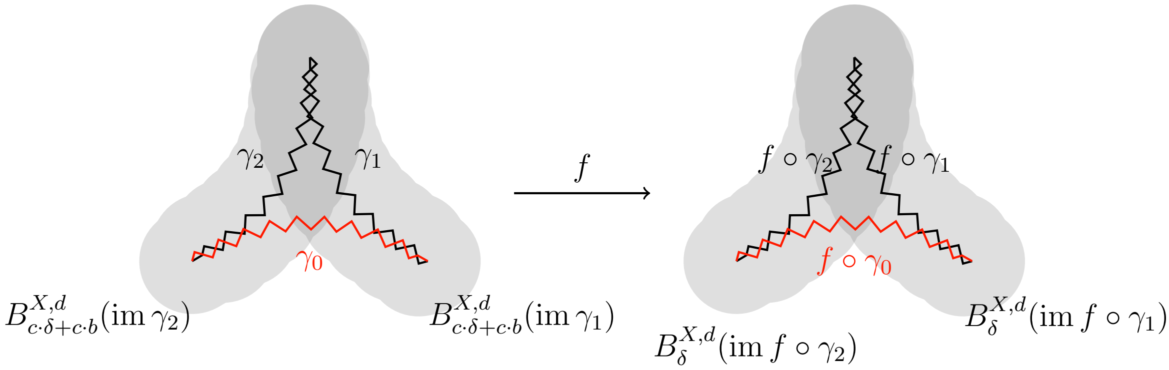

Figure 7.8.: Using a quasi-isometric embedding to translate quasi-geodesic triangles and neighbourhoods back and forth

embedding; we are allowed to choose common constants because of the built-in freedom of constants in the definition of quasi-hyperbolicity. We will show now that X is (c, b)-quasi-hyperbolic (see Figure 7.8 for an illustration of the notation).

Let c', $b' \in \mathbb{R}_{>0}$ with $c' \geq c$ and $b' \geq b$, and let $(\gamma_0, \gamma_1, \gamma_2)$ be a (c', b')-quasi-geodesic triangle in X. As $f\colon X \longrightarrow Y$ is a (c, b)-quasi-isometric embedding, $(f \circ \gamma_0, f \circ \gamma_1, f \circ \gamma_2)$ is a (c'', b'')-quasi-geodesic triangle in Y, where $c'' \in \mathbb{R}_{\geq c}$ and $b'' \in \mathbb{R}_{\geq b}$ are constants that depend only on c', b' and the quasi-isometry embedding constants c and b of f; without loss of generality, we may assume that $c'' \geq c$ and $b'' \geq b$.

Because Y is (c, b)-quasi-hyperbolic, there is a $\delta \in \mathbb{R}_{\geq 0}$ such that Y is (c'', b'', δ)-quasi-hyperbolic. In particular, the (c'', b'')-quasi-geodesic triangle $(f \circ \gamma_0, f \circ \gamma_1, f \circ \gamma_2)$ is δ-slim. Because f is a (c, b)-quasi-isometric embedding, a straightforward computation shows that

$$\begin{aligned}
\operatorname{im}\gamma_0 &\subset B^{X,d_X}_{c\cdot\delta+c\cdot b}(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2)\\
\operatorname{im}\gamma_1 &\subset B^{X,d_X}_{c\cdot\delta+c\cdot b}(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2)\\
\operatorname{im}\gamma_2 &\subset B^{X,d_X}_{c\cdot\delta+c\cdot b}(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1).
\end{aligned}$$

Therefore, X is $(c', b', c \cdot \delta + c \cdot b)$-quasi-hyperbolic, as was to be shown.

Clearly, the third part is a direct consequence of the first two parts. $\square$

7.2.3 Quasi-geodesics in hyperbolic spaces

Our next goal is to show that hyperbolicity is a quasi-isometry invariant in the class of geodesic spaces (Corollary 7.2.13). To this end we first compare hyperbolicity and quasi-hyperbolicity on geodesic spaces (Theorem 7.2.10), and then apply quasi-isometry invariance of quasi-hyperbolicity.

Theorem 7.2.10 (Hyperbolicity vs. quasi-hyperbolicity)**.** *Let X be a geodesic metric space. Then X is hyperbolic if and only if X is quasi-hyperbolic.*

In order to show that hyperbolic spaces indeed are quasi-hyperbolic, we need to understand how quasi-geodesics (and hence quasi-geodesic triangles) in hyperbolic spaces can be approximated by geodesics (and hence geodesic triangles).

Theorem 7.2.11 (Stability of quasi-geodesics in hyperbolic spaces)**.** *Let $\delta,\ c,\ b \in \mathbb{R}_{\geq 0}$. Then there exists a $\Delta \in \mathbb{R}_{\geq 0}$ with the following property: If X is a δ-hyperbolic metric space, if*[1] *$\gamma\colon [0, L] \longrightarrow X$ is a (c, b)-quasi-geodesic and $\gamma'\colon [0, L'] \longrightarrow X$ is a geodesic with $\gamma'(0) = \gamma(0)$ and $\gamma'(L') = \gamma(L)$, then*

$$\operatorname{im}\gamma' \subset B^{X,d}_{\Delta}(\operatorname{im}\gamma) \qquad \textit{and} \qquad \operatorname{im}\gamma \subset B^{X,d}_{\Delta}(\operatorname{im}\gamma').$$

Caveat 7.2.12. In general, the stability theorem for quasi-geodesics does *not* hold in non-hyperbolic spaces: For example, the logarithmic spiral (Figure 7.9)

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}^2 \\ t &\longmapsto t \cdot \bigl(\sin(\ln(1+t)), \cos(\ln(1+t))\bigr) \end{aligned}$$

is a quasi-isometric embedding with respect to the standard metrics on $\mathbb{R}$ and $\mathbb{R}^2$ (exercise), but this quasi-geodesic ray does not have bounded distance from any geodesic ray, and so quasi-geodesics in $\mathbb{R}^2$ cannot be uniformly approximated by geodesics.

---

[1] More precisely: if $L \in \mathbb{R}_{\geq 0}$ and if $\gamma\colon [0, L] \longrightarrow X$ is a (c, b)-quasi-geodesic, etc.

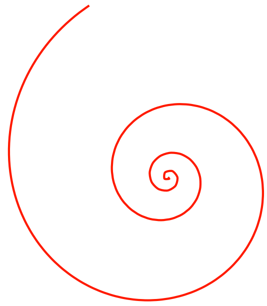

Figure 7.9.: The logarithmic spiral

We defer the proof of the stability theorem and first show how we can apply it to prove quasi-isometry invariance of hyperbolicity in the class of geodesic spaces:

Proof of Theoren 7.2.10. Clearly, if X is quasi-hyperbolic, then X is also hyperbolic (because every geodesic triangle is a quasi-geodesic triangle and so is slim enough by quasi-hyperbolicity).

Conversely, suppose that X is hyperbolic, say δ-hyperbolic for a suitable $\delta \in \mathbb{R}_{\geq 0}$. Moreover, let $c, b \in \mathbb{R}_{\geq 0}$, and let $\Delta \in \mathbb{R}_{\geq 0}$ be as provided by the stability theorem (Theorem 7.2.11) for the constants c, b, δ. We show now that X is $(c, b, 2 \cdot \Delta + \delta)$-quasi-hyperbolic.

To this end let $(\gamma_0, \gamma_1, \gamma_2)$ be a (c, b)-quasi-geodesic triangle in X. Because X is geodesic, we find geodesics γ_0', γ_1', γ_2' in X that have the same start and end points as the corresponding quasi-geodesics γ_0, γ_1, and γ_2, respectively (Figure 7.10). In particular, $(\gamma_0', \gamma_1', \gamma_2')$ is a geodesic triangle in X, and

$$\operatorname{im} \gamma_j' \subset B_\Delta^{X,d}(\operatorname{im} \gamma_j) \qquad \text{and} \qquad \operatorname{im} \gamma_j \subset B_\Delta^{X,d}(\operatorname{im} \gamma_j')$$

for all $j \in \{0, 1, 2\}$. Because X is δ-hyperbolic, it follows that

$$\begin{aligned}
\operatorname{im} \gamma_0' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_1' \cup \operatorname{im} \gamma_2') \\
\operatorname{im} \gamma_1' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_0' \cup \operatorname{im} \gamma_2') \\
\operatorname{im} \gamma_2' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_0' \cup \operatorname{im} \gamma_1'),
\end{aligned}$$

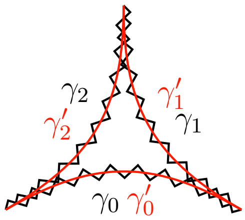

Figure 7.10.: Approximating quasi-geodesic triangles by geodesic triangles

and so

$$\begin{aligned}
\operatorname{im}\gamma_0 &\subset B^{X,d}_{\Delta+\delta+\Delta}(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2)\\
\operatorname{im}\gamma_1 &\subset B^{X,d}_{\Delta+\delta+\Delta}(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2)\\
\operatorname{im}\gamma_2 &\subset B^{X,d}_{\Delta+\delta+\Delta}(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1).
\end{aligned}$$

Therefore, X is $(c, b, 2 \cdot \Delta + \delta)$-quasi-hyperbolic. □

Corollary 7.2.13 (Quasi-isometry invariance of hyperbolicity)**.** *Let X and Y be metric spaces.*

1. *If Y is hyperbolic and X is quasi-geodesic and there is a quasi-isometric embedding $X \longrightarrow Y$, then X is quasi-hyperbolic.*
2. *If Y is geodesic and X is quasi-isometric to Y, then X is quasi-hyperbolic if and only if Y is hyperbolic.*
3. *If X and Y are geodesic and quasi-isometric, then X is hyperbolic if and only if Y is hyperbolic.*

Proof. Ad 1. In this case, by Theorem 7.2.10, the space Y is also quasi-hyperbolic. Therefore, X is quasi-hyperbolic as well (Proposition 7.2.9).

Ad 2. If Y is hyperbolic, then X is quasi-hyperbolic by the first part. Conversely, if X is quasi-hyperbolic, then Y is quasi-hyperbolic by Proposition 7.2.9. In particular, Y is hyperbolic.

Ad 3. This follows easily from the previous parts and Theorem 7.2.10. □

It remains to prove the stability theorem (Theorem 7.2.11). In order to do so, we need two facts about approximating curves by geodesics:

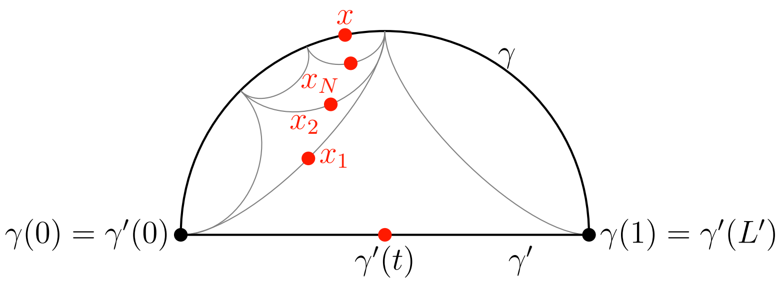

Figure 7.11.: Distance from geodesics to curves in hyperbolic spaces

Lemma 7.2.14 (Christmas tree lemma: Distance from geodesics to curves in hyperbolic spaces)**.** *Let $\delta \in \mathbb{R}_{\geq 0}$ and let X be a δ-hyperbolic space. If $\gamma\colon [0, L] \longrightarrow X$ is a continuous curve and if $\gamma'\colon [0, L'] \longrightarrow X$ is a geodesic with $\gamma'(0) = \gamma(0)$ and $\gamma'(L') = \gamma(L)$, then*

$$d\bigl(\gamma'(t), \operatorname{im}\gamma\bigr) \leq \delta \cdot \bigl|\ln_2(\ell(\gamma))\bigr| + 1$$

for all $t \in [0, L']$. Here, $\ell(\gamma)$ denotes the length of γ*, which is defined by*

$$\ell(\gamma) := \sup\biggl\{\sum_{j=0}^{k-1} d\bigl(\gamma(t_j), \gamma(t_{j+1})\bigr) \biggm| k \in \mathbb{N},\ t_0, \dots, t_k \in [0, L],\ t_0 \leq t_1 \leq \dots \leq t_k\biggr\}$$
$$\in \mathbb{R}_{\geq 0} \cup \{\infty\}.$$

Sketch of proof. Without loss of generality, we can assume that $\ell(\gamma) > 1$, that $\ell(\gamma) < \infty$, and that $\gamma\colon [0, 1] \longrightarrow X$ is parametrised by arc length; the latter is possible, because γ is continuous. These assumptions are notationally convenient, when filling in the details for the following arguments. Let $N \in \mathbb{N}$ with

$$\frac{\ell(\gamma)}{2^{N+1}} < 1 \leq \frac{\ell(\gamma)}{2^N}.$$

Let $t \in [0, L']$. Using the fact that X is δ-hyperbolic, we inductively construct a sequence $x_1, \dots, x_N \in X$ of points and geodesic triangles such that

$$d\bigl(\gamma'(t), x_1\bigr) \leq \delta, \qquad d(x_1, x_2) \leq \delta, \qquad d(x_2, x_3) \leq \delta, \qquad \dots$$

and such that x_N lies on a geodesic of length $\ell(\gamma)/2^N$ whose endpoints lie on $\operatorname{im}\gamma$ (see Figure 7.11). In particular, there is an $x \in \operatorname{im}\gamma$ with

$$\begin{aligned} d\bigl(\gamma'(t), x\bigr) &\leq d\bigl(\gamma'(t), x_N\bigr) + d(x_N, x) \\ &\leq \delta \cdot N + \frac{\ell(\gamma)}{2^{N+1}} \\ &\leq \delta \cdot |\ln_2(\gamma)| + 1. \end{aligned}$$

□

Lemma 7.2.15 (Taming quasi-geodesics in geodesic spaces)**.** *Let $c, b \in \mathbb{R}_{>0}$. Then there exist c', $b' \in \mathbb{R}_{\geq 0}$ with the following property: If X is a geodesic metric space and $\gamma\colon [0, L] \longrightarrow X$ is a (c, b)-quasi-geodesic, then there exists a* continuous *(c', b')-quasi-geodesic $\gamma'\colon [0, L] \longrightarrow X$ with $\gamma'(0) = \gamma(0)$ and $\gamma'(L) = \gamma(L)$ that satisfies the following properties:*

1. *For all s, $t \in [0, L]$ with $s \leq t$ we have*

$$\ell(\gamma'|_{[s,t]}) \leq c' \cdot d\bigl(\gamma'(s), \gamma'(t)\bigr) + b'.$$

2. *Moreover,*

$$\operatorname{im}\gamma' \subset B^{X,d}_{c+b}(\operatorname{im}\gamma) \qquad \textit{and} \qquad \operatorname{im}\gamma \subset B^{X,d}_{c+b}(\operatorname{im}\gamma').$$

Sketch of proof. Let $I := [0, L) \cap \mathbb{Z} \cup \{L\}$. As first step, we define

$$\gamma'|_I := \gamma|_I.$$

We then extend the definition of γ' to all of $[0, L]$ by inserting geodesic segments between the images of successive points in I. Then γ' is continuous, and a straightforward calculation shows that the conditions in the lemma are satisfied. □

Proof of Theorem 7.2.11. Let γ be a quasi-geodesic and let γ' be a geodesic as in the statement of the stability theorem (Theorem 7.2.11). In view of Lemma 7.2.15, by recplacing c, b if necessary with larger constants (depending only on c and b) we can assume without loss of generality that γ is a continuous (c, b)-quasi-geodesic satisfying the length condition of said lemma, i.e.,

$$\ell(\gamma|_{[r,s]}) \leq c \cdot d\bigl(\gamma(r), \gamma(s)\bigr) + b$$

for all r, $s \in [0, L]$ with $r \leq s$.

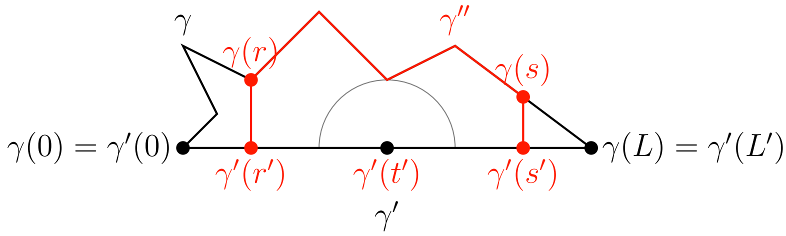

Figure 7.12.: In hyperbolic spaces, geodesics are close to quasi-geodesics

As *first step*, we give an upper estimate for $\sup_{t\in[0,L]}(\gamma'(t), \operatorname{im}\gamma)$ in terms of c, b, δ, i.e., we show that the geodesic γ' is close to the quasi-geodesic γ: Let

$$\Delta := \sup\bigl\{d(\gamma'(t'), \operatorname{im}\gamma) \bigm| t' \in [0, L']\bigr\};$$

as γ is continuous, a straightforward topological argument shows that there is a $t' \in [0, L]$ at which this supremum is attained. We now deduce an upper bound for Δ (see Figure 7.12 for an illustration of the notation):

Let

$$r' := \max(0, t' - 2\cdot\Delta) \leq \Delta \qquad \text{and} \qquad s' := \min(L', t' + 2\cdot\Delta) \geq L' - \Delta;$$

by construction of Δ, there exist r, $s \in [0, L]$ with

$$d\bigl(\gamma(r), \gamma'(r')\bigr) \leq \Delta \qquad \text{and} \qquad d\bigl(\gamma(s), \gamma'(s')\bigr) \leq \Delta.$$

We consider the curve γ'' in X given by starting in $\gamma'(r')$, then following a geodesic to $\gamma(r)$, then following γ until $\gamma(s)$, and finally following a geodesic to $\gamma'(s')$. From Lemma 7.2.14 and the construction of γ'' we obtain that

$$\Delta \leq d\bigl(\gamma'(t'), \operatorname{im}\gamma''\bigr) \leq \delta \cdot \bigl|\ln_2 \ell(\gamma'')\bigr| + 1;$$

moreover, because γ satisfies the length estimate from Lemma 7.2.15 as described above, we have

$$\begin{aligned}\ell(\gamma'') &\leq \ell(\gamma|_{[r,s]}) + 2\cdot\Delta\\ &\leq c\cdot d\bigl(\gamma(r), \gamma(s)\bigr) + b + 2\cdot\Delta\\ &\leq c\cdot(\Delta + 2\cdot\Delta + 2\cdot\Delta + \Delta)\cdot\Delta + b + 2\cdot\Delta.\end{aligned}$$

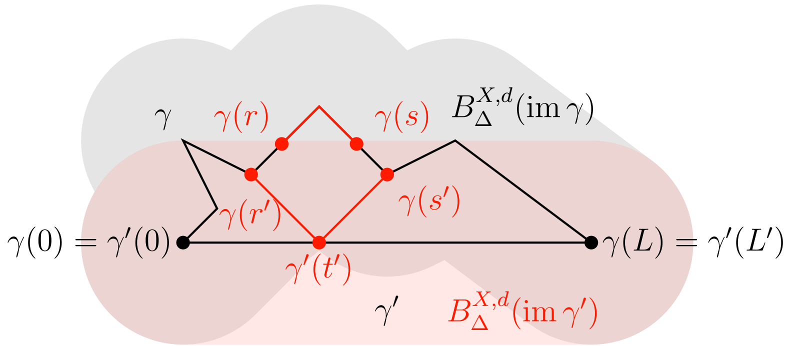

Figure 7.13.: In hyperbolic spaces, quasi-geodesics are close to geodesics

Hence,

$$\Delta \leq \delta \cdot \bigl|\ln_2((6 \cdot c + 2) \cdot \Delta + b)\bigr| + 1.$$

Because the logarithm function $\ln_2$ grows slower than linearly, this gives an upper bound for Δ in terms of c, b, and δ.

As *second step*, we give an upper estimate for $\sup_{t \in [0,L]}(\gamma(t), \operatorname{im}\gamma')$ in terms of c, b, δ, i.e., we show that the quasi-geodesic γ is close to the geodesic γ':

Let $\Delta := \sup\{d(\gamma'(t'), \operatorname{im}\gamma) \mid t' \in [0, L']\}$, as above. The idea is to show that "not much" of the quasi-geodesic γ lies outside $B^{X,d}_\Delta(\operatorname{im}\gamma')$. To this end, let r, $s \in [0, L]$ such that $[r, s]$ is a (with respect to inclusion) maximal interval with

$$\gamma((s,t)) \subset X \setminus B^{X,d}_\Delta(\operatorname{im}\gamma').$$

If there is no such non-empty interval, then there remains nothing to prove; hence, we assume that $r \neq s$. By choice of Δ, we know that $\operatorname{im}\gamma' \subset B^{X,d}_\Delta(\operatorname{im}\gamma)$, and so $\operatorname{im}\gamma' \subset B^{X,d}_\Delta(\operatorname{im}\gamma|_{[0,r]} \cup \operatorname{im}\gamma|_{[s,L]})$. Because γ is continuous and the interval $[0, L']$ is connected there exists a $t' \in [0, L']$ such that there are $r' \in [0, r]$, $s' \in [s, L]$ with

$$d\bigl(\gamma'(t'), \gamma(r')\bigr) \leq \Delta \qquad \text{and} \qquad d\bigl(\gamma'(t'), \gamma(s')\bigr) \leq \Delta$$

(see also Figure 7.13). Hence, we obtain

$$\begin{aligned} \ell(\gamma|_{[r,s]}) &\leq \ell(\gamma|_{[r',s']}) \leq c \cdot d\bigl(\gamma(r'), \gamma(s')\bigr) + b \\ &\leq 2 \cdot c \cdot \Delta + b, \end{aligned}$$

and thus

$$\gamma([r,s]) \subset B^{X,d}_{c\cdot\Delta+b+\Delta}(\operatorname{im}\gamma').$$

Applying the same reasoning to all components of γ lying outside of the neighbourhood $B^{X,d}_{\Delta}(\operatorname{im}\gamma')$, we can conclude that

$$\operatorname{im}\gamma \subset B^{X,d}_{c\cdot\Delta+b+\Delta}(\operatorname{im}\gamma').$$

Using the first part of the proof, we can bound Δ from above in terms of c, b, δ. Hence, this gives the desired estimate. □

7.2.4 Hyperbolic graphs

In particular, Corollary 7.2.13 applies to graphs: Graphs, viewed as metric spaces, are not geodesic (unless they have at most one vertex). However, in practice, it is often more convenient to be able to argue via geodesics than via quasi-geodesics. Therefore, in this context, it is common to work with geometric realisations of graphs.

Roughly speaking the *geometric realisation of a graph* is obtained by gluing a unit interval between every two vertices that are connected by an edge in the given graph. This construction can be turned into a geodesic metric space by combining the standard metric on the unit interval with the combinatorially defined metric on the vertices given by the graph structure. The details of this construction are explained in Appendix A.3.

Corollary 7.2.16 (Hyperbolicity of graphs)**.** *Let X be a connected graph. Then X is quasi-hyerbolic if and only if the geometric realisation $|X|$ is hyperbolic.*

Proof. This is an immediate consequence of Corollary 7.2.13, the fact that connected graphs are $(1,1)$-quasi-geodesic, and that the canonical inclusion of the vertices induces a quasi-isometry $X \sim_{\mathrm{QI}} |X|$ (Proposition A.3.3). □

Proposition 7.2.17 (Metric trees are hyperbolic)**.** *If T is a tree, then the geometric realisation $|T|$ of T is 0-hyperbolic. Hence, T is quasi-hyperbolic.*

Sketch of proof. Because the graph T does not contain any graph-theoretic cycles, one can show that any geodesic triangle in $|T|$ looks like a tripod (Figure 7.14). (Exercise). □

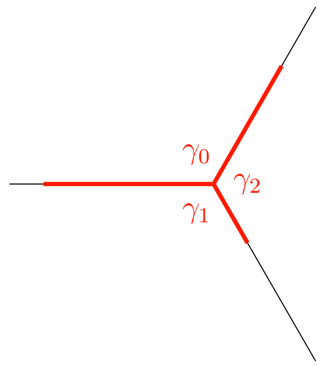

Figure 7.14.: Geodesic triangles in trees are tripods

7.3 Hyperbolic groups

Because (quasi-)hyperbolicity is a quasi-isometry invariant notion and because different finite generating sets of finitely generated groups give rise to canonically quasi-isometric word metrics/Cayley graphs, we obtain a sensible notion of hyperbolic groups [40]:

Definition 7.3.1 (Hyperbolic group)**.** A finitely generated group G is *hyperbolic* if for some (and hence any) finite generating set S of G the Cayley graph $\operatorname{Cay}(G, S)$ is quasi-hyperbolic.

In view of Corollary 7.2.16, we can check hyperbolicity of a finitely generated group also by checking that the geometric realisations of Cayley graphs are hyperbolic (which might be a more accessible problem).

Clearly, hyperbolicity of finitely generated groups is a geometric property:

Proposition 7.3.2 (Hyperbolicity is quasi-isometry invariant)**.** *Let G and H be finitely generated groups.*

1. *If H is hyperbolic and if there exist finite generating sets S and T of G and H respectively such that there is a quasi-isometric embedding $(G, d_S) \longrightarrow (H, d_T)$, then G is hyperbolic as well.*
2. *In particular: If G and H are quasi-isometric, then G is hyperbolic if and only if H is hyperbolic.*

Proof. This follows directly from the corresponding properties of quasi-hyperbolic spaces (Proposition 7.2.9) and the fact that Cayley graphs of groups are quasi-geodesic. □

Example 7.3.3 (Hyperbolic groups)**.**

- All finite groups are hyperbolic because the associated metric spaces have finite diameter.
- The group $\mathbb{Z}$ is hyperbolic, because it is quasi-isometric to the hyperbolic metric space $\mathbb{R}$.
- Finitely generated free groups are hyperbolic, because the Cayley graphs of free groups with respect to free generating sets are trees and hence hyperbolic by Proposition 7.2.17.
- In particular, $\mathrm{SL}(2, \mathbb{Z})$ is hyperbolic, because $\mathrm{SL}(2, \mathbb{Z})$ is quasi-isometric to a free group of rank 2 (Example 5.3.13).
- Let M be a compact Riemannian manifold of negative sectional curvature (e.g., a hyperbolic manifold in the sense of Definition 5.3.16). Then the fundamental group $\pi_1(M)$ is hyperbolic, because by the Švarc-Milnor lemma (Corollary 5.3.15) $\pi_1(M)$ is quasi-isometric to the Riemannian universal covering of M, which is hyperbolic (Example 7.2.3).

  In particular, the fundamental groups of oriented closed connected surfaces of genus at least 2 are hyperbolic (Example 5.3.17).
- The group $\mathbb{Z}^2$ is *not* hyperbolic, because it is quasi-isometric to the Euclidean plane $\mathbb{R}^2$, which is a goedesic metric space that is not hyperbolic (Example 7.2.3).
- We will see that the Heisenberg group is *not* hyperbolic (Example 7.5.21).

Even though $\mathrm{SL}(2, \mathbb{Z})$ is hyperbolic in the sense of geometric group theory and has a very close relation to the isometry group of the hyperbolic plane, there is no direct connection between these two properties:

Caveat 7.3.4. Let $z \in \mathbb{H}^2$. The canonical isometric action of $\mathrm{SL}(2,\mathbb{Z})$ on the hyperbolic plane $\mathbb{H}^2$ (in the half-plane model given by Möbius transformations) induces a map

$$\begin{aligned} \mathrm{SL}(2,\mathbb{Z}) &\longrightarrow \mathbb{H}^2 \\ A &\longmapsto A \cdot z \end{aligned}$$

with finite kernel (the kernel consists of E_2 and $-E_2$). This map is contracting with respect to the word metrics on $\mathrm{SL}(2,\mathbb{Z})$ and the hyperbolic metric on $\mathbb{H}^2$; however, this map is *not* a quasi-isometric embedding: We consider the matrix

$$A := \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}(2,\mathbb{Z}).$$

Then the word length of A^n (with respect to some finite generating set of $\mathrm{SL}(2,\mathbb{Z})$) grows linearly in $n \in \mathbb{N}$ (as can be seen by looking at the free subgroup of $\mathrm{SL}(2,\mathbb{Z})$ freely generated by A^2 and $(A^2)^T$), while the hyperbolic distance from the point $A^n \cdot z$ to z grows like $O(\ln n)$, as can be easily verified in the half-plane model (Appendix **??**): For all $n \in \mathbb{Z}$ we have

$$A^n \cdot z = \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} \cdot z = \frac{1 \cdot z + n}{0 \cdot z + 1} = z + n$$

and hence

$$\begin{aligned} d_{\mathbb{H}^2}(z, A^n \cdot z) &= d_{\mathbb{H}^2}(z, z+n) \\ &= \operatorname{arcosh}\Bigl(1 + \frac{n^2}{2 \cdot (\operatorname{Im} z)^2}\Bigr) \\ &\in O\bigl(\operatorname{arcosh}(n^2)\bigr) \\ &= O\bigl(\ln(n^2 + \sqrt{n^2 - 1})\bigr) \\ &= O(\ln n). \end{aligned}$$

Geometrically, the points $\{A^n \cdot z \mid n \in \mathbb{Z}\}$ lie all on a common horocycle, and the distance of consecutive points of this sequence is constant.

Similar arguments show that the standard embedding of a regular tree of degree 4 into $\mathbb{H}^2$ (given by viewing the free group of rank 2 as a subgroup of finite index in $\mathrm{SL}(2,\mathbb{Z})$) is *not* a quasi-isometric embedding. Nevertheless, in many situations, one can think of the geometry of the free group of rank 2 as an analogue of the hyperbolic plane in group theory.

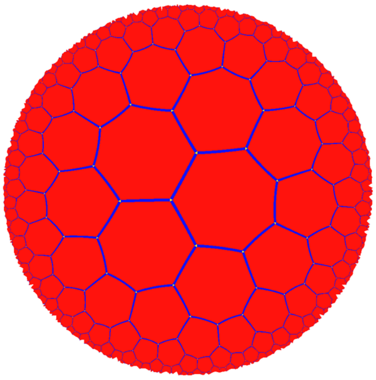

Figure 7.15.: A regular tiling of the hyperbolic plane (in the Poincaré disk model) by regular heptagons, where in each vertex exactly three heptagons meet [97]

Example 7.3.5 (Baking cookies via reflection groups)**.** When baking *Ausstecherle*, it is desirable to have cookie cutters that tile the dough in such a way that no rest of the dough remains (which would require recursive rolling out etc.). However, when baking in a Euclidean kitchen, there is only a very limited set of regular polygons that tile the plane (squares, regular triangles, regular hexagons).

Hyperbolic kitchens offer more flexibility: For example, it is possible to tile the hyperbolic plane $\mathbb{H}^2$ with isometric regular heptagons, where in each vertex exactly three heptagons meet (Figure 7.15). More generally, elementary hyperbolic geometry shows that the hyperbolic plane $\mathbb{H}^2$ admits regular tilings by isometric regular k-gons, where in each vertex exactly d of these k-gons meet, if and only if $k, d \in \mathbb{N}_{\geq 3}$ satisfy

$$\frac{1}{k} + \frac{1}{d} < \frac{1}{2}.$$

The *symmetry group* of the tiling above is the set of all isometries of $\mathbb{H}^2$ that preserve the tiling. It can be shown that the symmetry group of this tiling has the presentation

$$G := \bigl\langle r_1, r_2, r_3 \bigm| r_1^2, r_2^2, r_3^2, (r_1 r_2)^7, (r_1 r_3)^2, (r_2 r_3)^3 \bigr\rangle;$$

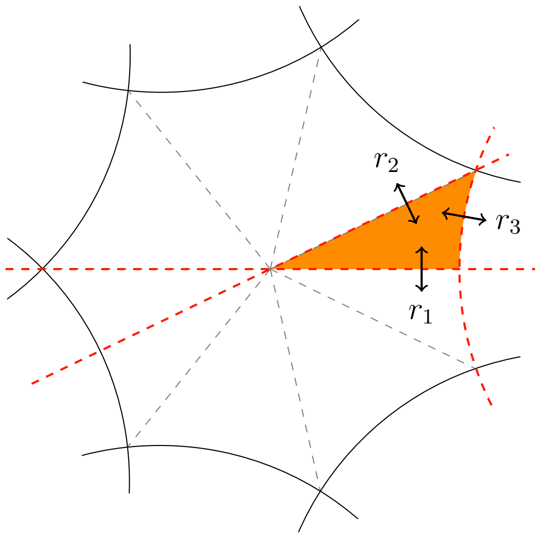

Figure 7.16.: Fundamental triangle of a regular hyperbolic heptagon and its reflections

geometrically, the generators r_1, r_2, r_3 correspond to the reflections at the geodesic lines bounding a fundamental triangle of this tiling (Figure 7.16).

Using the Švarc-Milnor lemma (Corollary 5.3.7) one can easily show that G is quasi-isometric to $\mathbb{H}^2$; hence, G is a hyperbolic group.

The group G is an example of a reflection group. More abstractly, reflection groups are special cases of Coxeter groups [19, 23]. Coxeter groups play an important role in various areas of mathematics. For example, the *reflection group trick* by Davis and Januszkiewicz [23] in algebraic and geometric topology allows to construct aspherical manifolds with exotic properties.

Beautiful tilings of $\mathbb{H}^2$ based on quasi-regular tilings of $\mathbb{H}^2$ can be found in the work of Escher [32, 33].

Surprisingly, the geometric property of being hyperbolic does not only have interesting geometric consequences mimicking the behaviour of manifolds of negative sectional curvature in Riemannian geometry for hyperbolic groups, but also non-trivial algebraic consequences (Section 7.4 and 7.5).

7.4 Application: "Solving" the word problem

In the following, we show that the geometric condition of being hyperbolic implies solvability of the word problem.

Definition 7.4.1. Let $\langle S \mid R \rangle$ be a finite presentation of a group. The *word problem is solvable for the presentation* $\langle S \mid R \rangle$, if there is an algorithm terminating on every input from $(S \cup S^{-1})^*$ that decides for every word w in $(S \cup S^{-1})^*$ whether w represents the trivial element of the group $\langle S \mid R \rangle$.

More precisely: The word problem is solvable for the presentation $\langle S \mid R \rangle$, if the sets

$$\{w \in (S \cup S^{-1})^* \mid w \text{ represents the neutral element of } \langle S \mid R \rangle\},$$
$$\{w \in (S \cup S^{-1})^* \mid w \text{ does not represent the neutral element of } \langle S \mid R \rangle\}$$

are recursively enumerable subsets of $(S \cup S^{-1})^*$.

The notion of being recursively enumerable or being algorithmically solvable can be formalised in several, equivalent, ways, e.g., using Turing machines, using μ-recursive functions, or using lambda calculus [14, 13, 6].

For example, it is not difficult to see that $\langle x, y \mid \rangle$ and $\langle x, y \mid [x, y] \rangle$ have solvable word problem. However, not all finite presentations have solvable word problem [82, Chapter 12]:

Theorem 7.4.2. *There exist finitely presented groups such that* no *finite presentation has solvable word problem.*

How can one prove such a theorem? The basic underlying arguments are self-referentiality and diagonalisation: One of the most prominent problems that cannot be solved algorithmically is the so-called *halting problem* for Turing machines: Roughly speaking, every Turing machine can be encoded by an integer (self-referentiality). Using a diagonalisation argument, one can show that there cannot exist a Turing machine that given two integers decides whether the Turing machine given by the first integer stops when applied to the second integer as input.

It is possible to encode the halting problem into group theory, thereby obtaining a finite presentation with unsolvable word problem.

Notice that the existence of finite presentations with unsolvable word problem has consequences in many other fields in mathematics; for example, reducing classification problems for manifolds to group theoretic questions shows that many classification problems in topology are unsolvable (Caveat 2.2.24).

Gromov noticed that hyperbolic groups have solvable word problem [40], thereby generalising and unifying previous work in combinatorial group theory and on fundamental groups of negatively curved manifolds:

Theorem 7.4.3 (Hyperbolic groups have solvable word problem)**.** *Let G be a hyperbolic group, and let S be a finite generating set of G. Then there exists a finite set $R \subset (S \cup S^{-1})^*$ such that $G \cong \langle S \mid R \rangle$ (in particular, G is finitely presented) and such that $\langle S \mid R \rangle$ has solvable word problem.*

Before proving this theorem, let us put this result in perspective: Gromov [41, 73] and Ol'shanskii [74] proved the following:

Theorem 7.4.4 (Generic groups are hyperbolic)**.** *In a certain well-defined statistical sense, almost all finite presentations of groups represent hyperbolic groups.*

So, statistically, the word problem for almost all finitely presented groups is solvable; however, one should keep in mind that there are interesting classes of groups that are *not* hyperbolic and so do not necessarily have solvable word problem.

The proof of Theorem 7.4.3 relies on a basic idea due to Dehn:

Definition 7.4.5 (Dehn presentation)**.** A finite presentation $\langle S \mid R \rangle$ is a *Dehn presentation* if there is an $n \in \mathbb{N}_{>0}$ and words $u_1, \dots, u_n$, $v_1, \dots, v_n$ such that

- we have $R = \{u_1 v_1^{-1}, \dots, u_n v_n^{-1}\}$,
- for all $j \in \{1, \dots, n\}$ the word v_j is shorter than u_j,
- and for all $w \in (S \cup S^{-1})^* \setminus \{\varepsilon\}$ that represent the neutral element of the group $\langle S \mid R \rangle$ there exists a $j \in \{1, \dots, n\}$ such that u_j is a subword of w.

Example 7.4.6. Looking at the characterisation of free groups in terms of reduced words shows that $\langle x, y \mid xx^{-1}\varepsilon, yy^{-1}\varepsilon, x^{-1}x\varepsilon, yy^{-1}\varepsilon\rangle$ is a Dehn presentation of the free group of rank 2. On the other hand, $\langle x, y \mid [x,y]\rangle$ is *not* a Dehn presentation for $\mathbb{Z}^2$.

The key property of Dehn presentations is the third one, as it allows to replace words by shorter words that represent the same group element:

Proposition 7.4.7 (Dehn's algorithm)**.** *If $\langle S \mid R\rangle$ is a Dehn presentation, then the word problem for $\langle S \mid R\rangle$ is solvable.*

Proof. We write $R = \{u_1 v_1^{-1}, \dots, u_n v_n^{-1}\}$, as given by the definition of Dehn presentations.

Given a word $w \in (S \cup S^{-1})^*$ we proceed as follows:

- If $w = \varepsilon$, then w represents the trivial element of the group $\langle S \mid R\rangle$.
- If $w \neq \varepsilon$, then:
  - If none of the words $u_1, \dots, u_n$ is a subword of w, then w does not represent the trivial element of the group $\langle S \mid R\rangle$ (by the third property of Dehn presentations).
  - If there is a $j \in \{1, \dots, n\}$ such that u_j is a subword of w, then we can write $w = w' u_j w''$ for certain words w', $w'' \in (S \cup S^{-1})^*$. Because $u_j v_j^{-1} \in R$, the words w and $w' v_j w''$ represent the same group element in $\langle S \mid R\rangle$; hence, w represents the trivial element of the group $\langle S \mid R\rangle$ if and only if the shorter word $w' v_j w''$ represents the trivial element of the group $\langle S \mid R\rangle$ (which we can check by applying this algorithm recursively to $w' v_j w''$).

Clearly, this algorithm terminates on all inputs from $(S \cup S^{-1})^*$ and decides whether the given word represents the trivial element in $\langle S \mid R\rangle$ or not. Hence, the word problem for $\langle S \mid R\rangle$ is solvable. □

In the setting of hyperbolic groups, short-cuts as required in the definition of Dehn presentations are enforced by the negative curvature (see Lemma 7.4.9 below).

Theorem 7.4.8 (Dehn presentations and hyperbolic groups)**.** *Let G be a hyperbolic group and let S be a finite generating set of G. Then there exists a finite set $R \subset (S \cup S^{-1})^*$ such that $\langle S \mid R\rangle$ is a Dehn presentation and $G \cong \langle S \mid R\rangle$.*

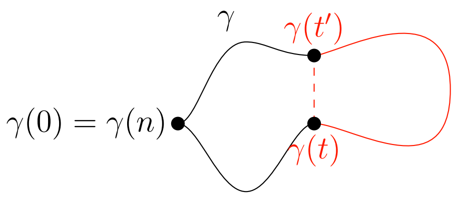

Figure 7.17.: Short-cuts in hyperbolic groups

The proof of this theorem relies on the existence of short-cuts in cycles in hyperbolic groups (Figure 7.17), which then give rise to a nice set of relations:

Lemma 7.4.9 (Short-cuts in cycles in hyperbolic groups)**.** *Let G be a hyperbolic group, let S be a finite generating set of G, and assume that $|\mathrm{Cay}(G,S)|$ is δ-hyperbolic with $\delta > 0$. If $\gamma\colon [0,n] \longrightarrow |\mathrm{Cay}(G,S)|$ is the geometric (piecewise linear) realisation of a graph-theoretic cycle in the Cayley graph $\mathrm{Cay}(G,S)$ of length $n > 0$, then there exist t, $t' \in [0,n]$ such that*

$$\ell(\gamma|_{[t,t']}) \leq 8 \cdot \delta$$

and such that the restriction $\gamma|_{[t,t']}$ is not *geodesic.*

A proof of this lemma will be given below.

Proof of Theorem 7.4.8. Because G is hyperbolic, there is a $\delta \in \mathbb{R}_{>0}$ such that $|\mathrm{Cay}(G,S)|$ is δ-hyperbolic. Let

$$D := \lceil 8 \cdot \delta \rceil + 1,$$

and let $\pi\colon F(S) \longrightarrow G$ be the canonical homomorphism. Modelling the short-cut lemma above in terms of group theory, leads to the (finite) set

$$\begin{aligned} R := &\bigl\{uv^{-1} \bigm| u,v \in (S \cup S^{-1})^*, |u| \leq D, |u| > d_S(e,\pi(u)),\\ &\qquad \pi(v) = \pi(u), |v| = d_S(e,\pi(u))\bigr\}\\ &\cup \{ss^{-1}\varepsilon \mid s \in S \cup S^{-1}\}. \end{aligned}$$

Clearly, the canonical homomorphism $\langle S \mid R\rangle \longrightarrow G$ is surjective (because S generates G).

This homomorphism is also injective, and $\langle S \mid R\rangle$ is a Dehn presentation for G, because: Let $w \in (S \cup S^{-1})^*$ be a word such that $\pi(w) = e$. We prove $w \in \langle R\rangle^{\triangleleft}_{F(S)}$ and the existence of a subword as required by the definition of Dehn presentations by induction over the length of the word w.

If w has length zero, then $w = \varepsilon$.

We now suppose that w has non-zero length and that the claim holds for all words in $(S \cup S^{-1})^*$ that are shorter than w.

If w is not reduced, then we find the desired subword in w by looking at the second part of the construction of R.

If w is reduced, then the word w (or a non-empty subword of w) translates into a graph-theoretic cycle in $\operatorname{Cay}(G, S)$ (see Definition 3.1.5). Applying the short-cut lemma (Lemma 7.4.9) to the geometric realisation of this cycle in $|\operatorname{Cay}(G, S)|$ shows that in $(S \cup S^{-1})^*$ we can decompose

$$w = w'uw''$$

into subwords w', u, w'' such that

$$d_S\bigl(e, \pi(u)\bigr) < |u| \leq D.$$

Let $v \in (S \cup S^{-1})^*$ be chosen in such a way that $\pi(v) = \pi(u)$ and $|v| = d_S(e, \pi(u)) < |u|$ (Figure 7.18). By construction of R, we find that

$$e = \pi(w) = \pi(w') \cdot \pi(u) \cdot \pi(w'') = \pi(w') \cdot \pi(v) \cdot \pi(w'') = \pi(w'vw'').$$

Because the word $w'vw''$ is shorter than $w'uw'' = w$, by induction, we obtain $w'vw'' \in \langle R\rangle^{\triangleleft}_{F(S)}$ and that $w'vw''$ (if non-trivial) contains a subword as in the definition of Dehn presentations. Therefore, also $w \in \langle R\rangle^{\triangleleft}_{F(S)}$ and w contains such a subword, as desired. $\square$

Corollary 7.4.10. *Every hyperbolic group admits a finite presentation.*

Proof. By the theorem, every hyperbolic group possesses a finite Dehn presentation, and finite Dehn presentations are examples of finite presentations. $\square$

In particular, we can now prove Theorem 7.4.3:

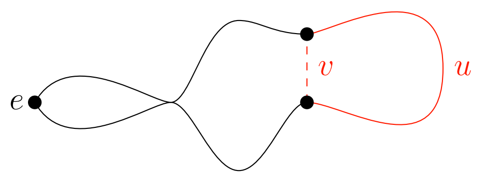

Figure 7.18.: Short-cuts in hyperbolic groups

Proof (of Theorem 7.4.3). In view of Theorem 7.4.8, every finite generating set S of a hyperbolic group can be extended to a Dehn presentation $\langle S \mid R \rangle$ of the group in question. By Proposition 7.4.7, the word problem for $\langle S \mid R \rangle$ is solvable. $\square$

It remains to prove the short-cut lemma: A key step in the proof uses that local geodesics in hyperbolic spaces stay close to actual geodesics:

Lemma 7.4.11 (Local geodesics in hyperbolic spaces)**.** *Let $\delta \in \mathbb{R}_{\geq 0}$, let (X, d) be a δ-hyperbolic space, and let $c \in \mathbb{R}_{>8\cdot\delta}$. Let $\gamma\colon [0, L] \longrightarrow X$ be a* c-local geodesic, *i.e., for all $t, t' \in [0, L]$ with $|t - t'| \leq c$ we have*

$$d\bigl(\gamma(t), \gamma(t')\bigr) = |t - t'|.$$

If $\gamma'\colon [0, L'] \longrightarrow X$ is a geodesic with $\gamma'(0) = \gamma(0)$ and $\gamma'(L') = \gamma(L)$, then

$$\operatorname{im}\gamma \subset B^{X,d}_{2\cdot\delta}(\operatorname{im}\gamma').$$

Proof. Exercise. $\square$

Proof of Lemma 7.4.9. We first show that γ cannot be a c-local geodesic for any $c \in \mathbb{R}_{>8\cdot\delta}$: *Assume* for a contradiction that there is a $c \in \mathbb{R}_{>8\cdot\delta}$ such that γ is a c-local geodesic in $|\mathrm{Cay}(G, S)|$. Because of $\gamma(0) = \gamma(n)$ and $c > 8 \cdot \delta$, it is clear that $n > 8 \cdot \delta$.

In view of Lemma 7.4.11, γ is $2 \cdot \delta$-close to any geodesic starting in $\gamma(0)$ and ending in $\gamma(n) = \gamma(0)$; because the constant map at $\gamma(0)$ is such a geodesic, it follows that

$$\operatorname{im}\gamma \subset B^{|\mathrm{Cay}(G,S)|, d_S}_{2\cdot\delta}\bigl(\gamma(0)\bigr).$$

Therefore, we obtain

$$4 \cdot \delta \geq \operatorname{diam} B_{2\cdot\delta}^{|\mathrm{Cay}(G,S)|,d_S}\bigl(\gamma(0)\bigr) \geq d_S\bigl(\gamma(0), \gamma(5\cdot\delta)\bigr) = 5 \cdot \delta,$$

which is a contradiction. So γ is not a c-local geodesic for any $c \in \mathbb{R}_{>8\cdot\delta}$.

Hence, there exist $t, t' \in [0, n]$ such that $|t - t'| \leq 8 \cdot \delta$ and $d\bigl(\gamma(t), \gamma(t')\bigr) \neq |t - t'|$; in particular $\gamma|_{[t,t']}$ is not geodesic. Moreover, because γ is the geometric realisation of a cycle in $\mathrm{Cay}(G, S)$, it follows that

$$\ell(\gamma|_{[t,t']}) = |t - t'| \leq 8 \cdot \delta. \qquad \square$$

7.5 Elements of infinite order in hyperbolic groups

In the following, we study elements of infinite order in hyperbolic groups; in particular, we show that any infinite hyperbolic group contains an element of infinite order (Section 7.5.1), and we show that centralisers of elements of infinite order in hyperbolic groups are "small" (Section 7.5.2). Consequently, hyperbolic groups cannot contain $\mathbb{Z}^2$ as a subgroup. Moreover, some generalisations and applications to topology and Riemannian geometry are discussed.

7.5.1 Existence of elements of infinite order in hyperbolic groups

In general, an infinite finitely generated group does not necessarily contain an element of infinite order; for example, the Grigorchuk group (p. 135) is of this type. However, in the case of hyperbolic groups, the geometry forces the existence of elements of infinite order:

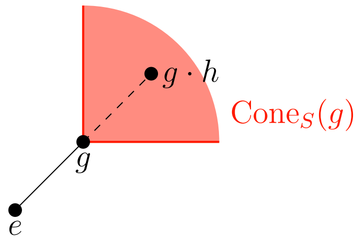

Figure 7.19.: Cone type, schematically; the drawing actually shows the more intuitive visualisation $g \cdot \mathrm{Cone}_S(g)$ of the cone type instead of $\mathrm{Cone}_S(g)$ itself.

Theorem 7.5.1. *Every infinite hyperbolic group contains an element of infinite order.*

In order to prove this theorem, we introduce the notion of cone types of group elements:

Definition 7.5.2 (Cone type)**.** Let G be a finitely generated group, let $S \subset G$ be a finite generating set, and let $g \in G$. The *cone type of g with respect to S* is the set (Figure 7.19)

$$\mathrm{Cone}_S(g) := \bigl\{h \in G \bigm| d_S(e, g \cdot h) \geq d_S(e, g) + d_S(e, h)\bigr\}.$$

Example 7.5.3.

- Let F be a finitely generated free group of rank n, and let S be a free generating set of F. Then F has exactly $2 \cdot n + 1$ cone types with respect to S, namely $\mathrm{Cone}_S(e) = F$, and

$$\begin{aligned} \mathrm{Cone}_S(s) = \{w \mid w \text{ is a reduced word over } S \cup S^{-1} \\ \text{that does not start with } s^{-1}\} \end{aligned}$$

  for all $s \in S \cup S^{-1}$.
- The group $\mathbb{Z}^2$ has only finitely many different cone types with respect to the genereating set $\{(1,0), (0,1)\}$ (Figure 7.20).

Theorem 7.5.1 is proved by verifying that hyperbolic groups have only finitely many cone types (Proposition 7.5.4) and that infinite groups that have only finitely many cone types must contain an element of infinite order (Proposition 7.5.6).

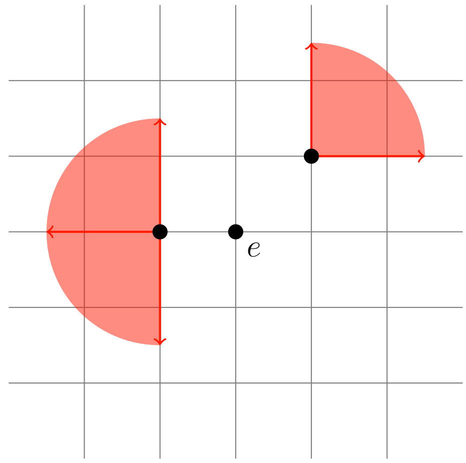

Figure 7.20.: Cone types of $\mathbb{Z}^2$

Proposition 7.5.4 (Cone types of hyperbolic groups)**.** *Let G be a hyperbolic group, and let S be a finite generating set of G. Then G has only finitely many cone types with respect to S.*

Proof. The idea of the proof is to show that the cone type of a given element depends only on the set of group elements close to g. More precisely, for $g \in G$ and $r \in \mathbb{R}_{\geq 0}$ we call

$$P_r^S(g) := \bigl\{ h \in G \bigm| d_S(e, g \cdot h) \leq d_S(e, g), h \in B_r^{G,S}(g) \bigr\}$$

the *r-past of g with respect to S* (Figure 7.21). Because G is hyperbolic there is a $\delta \in \mathbb{R}_{\geq 0}$ such that $|\mathrm{Cay}(G, S)|$ is δ-hyeprbolic. Let

$$r := 2 \cdot \delta + 2.$$

We will prove now that the r-past of an element determines its cone type, i.e.: for all g, $g' \in G$ with $P_r^S(g) = P_r^S(g')$ we have $\mathrm{Cone}_S(g) = \mathrm{Cone}_S(g')$.

Let g, $g' \in G$ with $P_r^S(g) = P_r^S(g')$, and let $h \in \mathrm{Cone}_S(g)$. We prove the assertion $h \in \mathrm{Cone}_S(g')$ by induction over $d_S(e, h)$:

If $d_S(e, h) = 0$, then $h = e$, and so the claim trivially holds.

If $d_S(e, h) = 1$, then $h \in \mathrm{Cone}_S(g)$ implies that $h \notin P_r^S(g) = P_r^S(g')$ (by definition of the cone type and the past); hence, in this case $h \in \mathrm{Cone}_S(g')$.

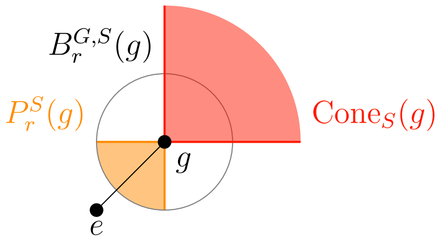

Figure 7.21.: The past of an element, schematically

Now suppose that $h = h' \cdot s$ with $s \in S \cup S^{-1}$ and $d_S(e, h') = d_S(e, h) - 1 > 0$, and that the claim holds for all group elements in $B^{G,S}_{d_S(e,h)-1}(e)$. Notice that $h \in \mathrm{Cone}_S(g)$ implies that $h' \in \mathrm{Cone}_S(g)$ as well; therefore, by induction, $h' \in \mathrm{Cone}_S(g')$. *Assume* for a contradiction that $h \not\in \mathrm{Cone}_S(g')$. Then

$$d_S(e, g' \cdot h) < d_S(e, g') + d_S(e, h);$$

without loss of generality we may assume that $d_S(e, g' \cdot h) \geq d_S(e, g')$. Thus, we can write

$$g' \cdot h = k_1 \cdot k_2$$

for certain group elements $k_1,\ k_2 \in G$ that in addition satisfy

$$\begin{aligned} d_S(e, g' \cdot h) &= d_S(e, k_1) + d_S(e, k_2) \\ d_S(e, k_1) &= d_S(e, g') \\ d_S(e, k_2) &\leq d_S(e, h) - 1 \end{aligned}$$

(such a decomposition of $g' \cdot h$ can, for instance, be obtained by looking at a shortest path in $\mathrm{Cay}(G, S)$ from e to $g' \cdot h$; see also Figure 7.22). We now consider the element

$$h'' := g'^{-1} \cdot k_1.$$

The element h'' lies in the past $P_r^S(g')$ of g', because: On the one hand, we have (by choice of k_1)

$$d_S(e, g' \cdot h'') = d_S(e, k_1) \leq d_S(e, g').$$

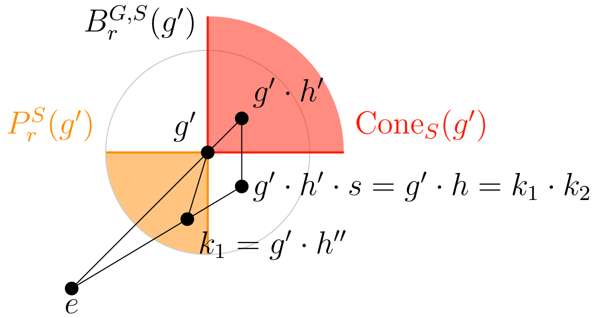

Figure 7.22.: The past of an element determines its cone type

On the other hand,

$$\begin{aligned} d_S(e, h'') &= d_S(e, g'^{-1} \cdot k_1) \\ &= d_S(g', k_1) \\ &\leq 2 \cdot \delta + 2 \\ &\leq r. \end{aligned}$$

For the penultimate inequality, we used the fact that g' and k_1 lie at the same time parameter (namely, $d_S(e, g') = d_S(e, k_1)$) on two geodesics in the δ-hyperbolic space $|\mathrm{Cay}(G, S)|$ that both start in e, and that end distance 1 apart (namely in $g' \cdot h'$ and $g' \cdot h$ respectively) (see Lemma 7.5.5 below); such geodesics indeed exist because $h' \in \mathrm{Cone}_S(g')$ shows that $d_S(e, g' \cdot h') \geq d_S(e, g') + d_S(e, h')$, and $d_S(e, g' \cdot h) = d_S(e, k_1) + d_S(e, k_2)$ by choice of k_1 and k_2. So $h'' \in P_r^S(g') = P_r^S(g)$.

Using the fact that $h \in \mathrm{Cone}_S(g)$, the choice of k_1 and k_2, as well as the fact that $h'' \in P_r^S(g)$, we obtain

$$\begin{aligned} d_S(e, g) + d_S(e, h) &\leq d_S(e, g \cdot h) \\ &= d_S(e, g \cdot g'^{-1} \cdot g' \cdot h) \\ &= d_S(e, g \cdot g'^{-1} \cdot k_1 \cdot k_2) \\ &\leq d_S(e, g \cdot h'') + d_S(e, k_2) \\ &\leq d_S(e, g) + d_S(e, h) - 1, \end{aligned}$$

which is a contradiction. Therefore, $h \in \mathrm{Cone}_S(g')$, completing the induction.

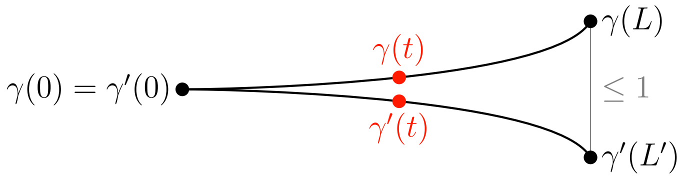

Figure 7.23.: Geodesics in hyperbolic spaces starting at the same point are uniformly close

Hence, the cone type of a group element g of G is determined by the r-past $P_r^S(g)$. By definition, $P_r^S(g) \subset B_r^{G,S}(e)$, which is a finite set. In particular, there are only finitely many possible different r-pasts with respect to S. Therefore, there are only finitely many cone types with respect to S in G, as desired. $\square$

Lemma 7.5.5 (Geodesics in hyperbolic spaces starting at the same point)**.** *Let $\delta \in \mathbb{R}_{\geq 0}$ and let (X, d) be a δ-hyperbolic space. Let $\gamma\colon [0, L] \longrightarrow X$ and $\gamma'\colon [0, L'] \longrightarrow X$ be geodesics in X with*

$$\gamma(0) = \gamma'(0) \qquad \textit{and} \qquad d\bigl(\gamma(L), \gamma'(L')\bigr) \leq 1.$$

Then γ and γ' are uniformly $(2 \cdot \delta + 2)$-close, i.e.,

$$\forall_{t \in [0, \min(L, L')]} \ d\bigl(\gamma(t), \gamma'(t)\bigr) \leq 2 \cdot \delta + 2.$$

Proof. Exercise (see also Figure 7.23). $\square$

Proposition 7.5.6 (Cone types and elements of infinite order)**.** *Let G be a finitely generated infinite group that has only finitely many cone types with respect to some finite generating set. Then G contains an element of infinite order.*

In particular, finitely generated infinite groups all of whose elements have finite order have infinitely many cone types.

Proof. Let $S \subset G$ be a finite generating set of G and suppose that G has only finitely many cone types with respect to S, say

$$k := \bigl|\bigl\{\mathrm{Cone}_S(g) \bigm| g \in G\bigr\}\bigr|.$$

Because G is infinite and $\mathrm{Cay}(G, S)$ is a a proper metric space with respect to the word metric d_S there exists a $g \in G$ with $d_S(e, g) > k$. In particular, choosing a shortest path from e to g in $\mathrm{Cay}(G, S)$ and applying the pigeonhole principle to the group elements on this path shows that we can write

$$g = g' \cdot h \cdot g''$$

such that $h \neq e$,

$$d_S(e, g) = d_S(e, g') + d_S(e, h) + d_S(e, g''),$$

and such that $\mathrm{Cone}_S(g') = \mathrm{Cone}_S(g' \cdot h)$ holds; in particular, we have $d_S(e, g' \cdot h) = d_S(e, g') + d_S(e, h)$ and so $h \in \mathrm{Cone}_S(g')$.

Then the element h has infinite order, because: We will show by induction over $n \in \mathbb{N}_{>0}$ that

$$d_S(e, g' \cdot h^n) \geq d_S(e, g') + n \cdot d_S(e, h).$$

In the case $n = 1$, this claim follows from the choice of g', h, and g'' above. Let now $n \in \mathbb{N}_{>0}$, and suppose that

$$d_S(e, g' \cdot h^n) \geq d_S(e, g') + n \cdot d_S(e, h);$$

in particular, $d_S(e, h^n) = n \cdot d_S(e, h)$. By definition of the cone type of g', it follows that $h^n \in \mathrm{Cone}_S(g') = \mathrm{Cone}_S(g' \cdot h)$. Hence,

$$\begin{aligned} d_S(e, g' \cdot h^{n+1}) &= d_S(e, g' \cdot h \cdot h^n) \\ &\geq d_S(e, g' \cdot h) + d_S(e, h^n) \\ &= d_S(e, g') + d_S(e, h) + n \cdot d_S(e, h) \\ &= d_S(e, g') + (n+1) \cdot d_S(e, h), \end{aligned}$$

which completes the induction step.

In particular, we obtain that for all $n \in \mathbb{N}_{>0}$ the elements $g' \cdot h^n$ and g' must be different. Thus, h has infinite order. $\square$

Proof of Theorem 7.5.1. Let G be an infinite hyperbolic group, and let $S \subset G$ be a finite generating set of G. Then G has only finitely many cone types (Proposition 7.5.4), and so contains an element of infinite order (Proposition 7.5.6). $\square$

As a first application of the existence of elements of infinite order, we give another proof of the fact that $\mathbb{Z}$ is quasi-isometrically rigid; we split the argument into two parts:

Corollary 7.5.7. *Any finitely generated group quasi-isometric to $\mathbb{Z}$ contains an element of infinite order.*

Proof. Let G be a finitely generated group quasi-isometric to $\mathbb{Z}$. Then G is infinite and hyperbolic, because $\mathbb{Z}$ is infinite and hyperbolic. In particular, G contains an element $g \in G$ of infinite order (by Theorem 7.5.1). □

Corollary 7.5.8 (Quasi-isometry rigidity of $\mathbb{Z}$)**.** *Any finitely generated group quasi-isometric to $\mathbb{Z}$ contains a finite index subgroup isomorphic to $\mathbb{Z}$.*

Proof. Let G be a finitely generated group quasi-isometric to $\mathbb{Z}$. By Corollary 7.5.7, G contains an element g of infinite order. One then shows:

- The subgroup $\langle g \rangle_G$ generated by g is quasi-dense in G because G and $\mathbb{Z}$ are quasi-isometric (exercise).
- Hence, $\langle g \rangle_G$ has finite index in G (exercise). □

7.5.2 Centralisers of elements of infinite order in hyperbolic groups

Because (quasi-)isometrically embedded (geodesic) subspaces of hyperbolic spaces are hyperbolic (Proposition 7.2.9), no hyperbolic space can contain the flat Euclidean plane $\mathbb{R}^2$ as a (quasi-)isometrically embedded subspace. The geometric group theoretic analogue is that $\mathbb{Z}^2$ cannot be quasi-isometrically embedded into a hyperbolic group.

In the following, we will show that also the, stronger, algebraic analogue holds: A hyperbolic group cannot contain $\mathbb{Z}^2$ as a subgroup (Corollary 7.5.19).

How can we prove such a statement? If a group contains $\mathbb{Z}^2$ as a subgroup, then it contains an element of infinite order whose centraliser contains a subgroup isomorphic to $\mathbb{Z}^2$; in particular, there are elements of infinite order with "large" centralisers.

In the case of hyperbolic groups, the geometry forces centralisers of elements of infinite order to be small:

Theorem 7.5.9 (Centraliseres in hyperbolic groups)**.** *Let G be a hyperbolic group and let $g \in G$ be an element of infinite order.*

1. *Then the map*

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longmapsto g^n \end{aligned}$$

   is a quasi-isometric embedding.
2. *The subgroup $\langle g \rangle_G$ has finite index in the centraliser $C_G(g)$ of g in G; in particular, $C_G(g)$ is virtually $\mathbb{Z}$.*

For sake of completeness, we briefly recall the notion of centraliser:

Definition 7.5.10 (Centraliser)**.** Let G be a group, and let $g \in G$. The *centraliser of g in G* is the set of all group elements commuting with g, i.e.,

$$C_G(g) := \{h \in G \mid h \cdot g = g \cdot h\}.$$

Notice that the centraliser of a group element always is a subgroup of the group in question.

Example 7.5.11.

- If G is Abelian, then $C_G(g) = G$ for all $g \in G$.
- If F is a free group and $g \in F \setminus \{e\}$, then $C_F(g) = \langle g \rangle_F \cong \mathbb{Z}$.
- If G and H are groups, then

$$G \times \langle h \rangle_H \subset C_{G \times H}(e, h)$$

  for all $h \in H$.
- If S is a generating set of a group G, then

$$\bigcap_{s \in S} C_G(s) = \bigcap_{g \in G} C_G(g) = \{h \in G \mid \forall_{g \in G} \;\; h \cdot g = g \cdot h\}$$

  is the so-called *centre of* G.

Caveat 7.5.12. If G is a finitely generated group with finite generating set S, and if $g \in G$ is an element of infinite order, then the map

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longrightarrow g^n \end{aligned}$$

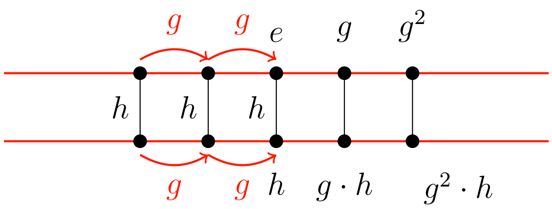

Figure 7.24.: Elements in the centraliser $C_G(g)$ lead to flat strips

is not necessarily a quasi-isometric embedding with respect to the standard metric on $\mathbb{Z}$ and the word metric d_S on G. For example, this occurs in the Heisenberg group (Caveat 6.2.13).

We now prove Theorem 7.5.9:

Proof of Theorem 7.5.9, part 2. We start by showing that the first part of the theorem implies the second part: This is a geometric argument involving invariant geodesics, inspired by Preissmann's theorem (and its proof) in Riemannian geometry [31, Theorem 10.2.2]:

In view of the first part of the theorem, there is a "geodesic line" in G that is left invariant under translation by g, namely $n \mapsto g^n$. If $h \in C_G(g)$, then translation by g also leaves the "geodesic line" in G given by $n \mapsto h \cdot g^n = g^n \cdot h$ invariant. Using the fact that $g^n \cdot h = h \cdot g^n$ for all $n \in \mathbb{Z}$, one can show that these two "geodesic lines" span a flat strip in G (Figure 7.24). However, as G is hyperbolic, this flat strip cannot be too wide; in particular, h has to be close to $\langle g \rangle_G$. So $C_G(g)$ is virtually cyclic.

We now explain this argument in detail: Let $S \subset G$ be a finite generating set and let $\delta \in \mathbb{R}_{\geq 0}$ be chosen in such a way that $|\mathrm{Cay}(G, S)|$ is δ-hyperbolic.

As first step, we show that without loss of generality we may assume that g is *not* conjugate to an element of $B^{G,S}_{4\cdot\delta+2}(e)$; to this end, we show that the elements $\{g^n \mid n \in \mathbb{Z}\}$ fall into infinitely many conjugacy classes: Suppose there are $n, m \in \mathbb{N}$ such that $n < m$ and g^n is conjugate to g^m. Then there exists an $h \in G$ with

$$h \cdot g^n \cdot h^{-1} = g^m.$$

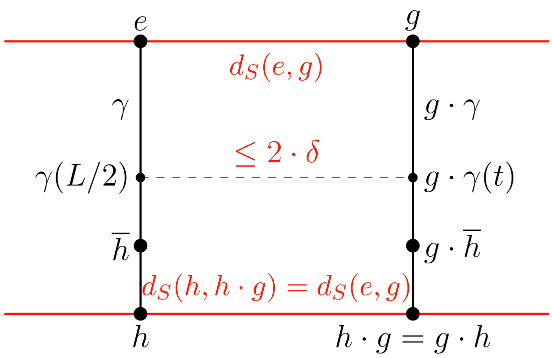

Figure 7.25.: Elements in the centraliser $C_G(g)$ are close to $\langle g \rangle_G$

A straightforward induction shows that then

$$h^k \cdot g^{n^k} \cdot h^{-k} = g^{m^k}$$

holds for all $k \in \mathbb{N}$. In particular, for all $k \in \mathbb{N}$ we obtain

$$d_S(e, g^{m^k}) \leq 2 \cdot k \cdot d_S(e,h) + n^k \cdot d_S(e,g).$$

However, if k is large enough, this contradicts the first part of the theorem, namely that the map $n \mapsto g^n$ is a quasi-isometric embedding. Hence, the elements $\{g^n \mid n \in \mathbb{Z}\}$ fall into infinitely many conjugacy classes, and so some power of g is not conjugate to an element in the finite set $B^{G,S}_{4\cdot\delta+2}(e)$.

Because $C_G(g) \subset C_G(g^n)$ and because $\langle g^n \rangle_G$ has finite index in $\langle g \rangle_G$ for all $n \in \mathbb{Z}$ we can replace g by such a power, and so may assume that g is not conjugate to an element of $B^{G,S}_{4\cdot\delta+2}(e)$.

We now consider

$$r := 2 \cdot d_S(e,g) + 4 \cdot \delta + 2,$$

and we will show that $C_G(g) \subset B^{G,S}_r(\langle g \rangle_G)$: Let $h \in C_G(g)$. *Assume* for a contradiction that $d_S(h, \langle g \rangle_G) > r$. Replacing h by $g^n \cdot h$ for a suitable $n \in \mathbb{Z}$ we may assume that

$$d_S(h,e) = d_S(h, \langle g \rangle_G) > r.$$

Let $\gamma \colon [0,L] \longrightarrow |\mathrm{Cay}(G,S)|$ be a geodesic starting in e and ending in h; then $g \cdot \gamma \colon [0,L] \longrightarrow |\mathrm{Cay}(G,S)|$ is a geodesic starting in g and ending in $g \cdot h$

(Figure 7.25). We now apply δ-hyperbolicity of $|\mathrm{Cay}(G,S)|$ to a geodesic quadrilateral whose sides are γ, $g\cdot\gamma$, some geodesic from e to g, and some geodesic from h to $g\cdot h = h\cdot g$. Because $d_S(e,h) > r = 2\cdot d_S(e,g)+4\cdot\delta+2$, it follows that $\gamma(L/2)$ is $2\cdot\delta$-close to $\operatorname{im} g\cdot\gamma$, say $d_S(\gamma(L/2), g\cdot\gamma(t)) \leq 2\cdot\delta$; here, the condition that $h\cdot g = g\cdot h$ enters in the form that it allows us to conclude that

$$d_S(h, g\cdot h) = d_S(h, h\cdot g) = d_S(e,g)$$

and thus that $\gamma(L/2)$ is "far away" from any geodesic from h to $g\cdot h$. Because $\mathrm{Cay}(G,S)$ is 1-dense in $|\mathrm{Cay}(G,S)|$, there is a group element $\overline{h}\in G$ with $d_S(\overline{h},\gamma(L/2)) \leq 1$. Hence, we obtain

$$\begin{aligned}
d_S(e,\overline{h}^{-1}\cdot g\cdot\overline{h}) &= d_S(\overline{h}, g\cdot\overline{h})\\
&\leq d_S\bigl(\overline{h},\gamma(L/2)\bigr) + d_S\bigl(\gamma(L/2), g\cdot\gamma(t)\bigr) + d_S\bigl(g\cdot\gamma(t), g\cdot\overline{h}\bigr)\\
&\leq 1 + 2\cdot\delta + d_S\bigl(\gamma(t),\overline{h}\bigr)\\
&\leq 1 + 2\cdot\delta + d_S\bigl(\gamma(t),\gamma(L/2)\bigr) + d_S\bigl(\gamma(L/2),\overline{h}\bigr)\\
&\leq 2 + 2\cdot\delta + |t - L/2|;
\end{aligned}$$

because $d_S(h,e) = d_S(h,\langle g\rangle_G)$ and because γ is a geodesic starting in e and ending in h, we have

$$\begin{aligned}
|t - L/2| &= \bigl|d_S\bigl(\gamma(t),\langle g\rangle_G\bigr) - d_S\bigl(\gamma(L/2),\langle g\rangle_G\bigr)\bigr|\\
&= \bigl|d_S\bigl(g\cdot\gamma(t),\langle g\rangle_G\bigr) - d_S\bigl(\gamma(L/2),\langle g\rangle_G\bigr)\bigr|\\
&\leq d_S\bigl(g\cdot\gamma(t),\gamma(L/2)\bigr)\\
&\leq 2\cdot\delta.
\end{aligned}$$

Therefore,

$$d_S(e,\overline{h}^{-1}\cdot g\cdot\overline{h}) \leq 4\cdot\delta + 2,$$

contradicting our assumption that g is not conjugate to an element of the ball $B^{G,S}_{4\cdot\delta+2}(e)$. Thus, $d_S(e,h) \leq r$, and so $C_G(g)\subset B^{G,S}_r(\langle g\rangle_G)$.

Now a simple calculation shows that $\langle g\rangle_G$ has finite index in $C_G(g)$. $\square$

The previous arguments were based on the fact that the group element in question acted by translation on a quasi-geodesic line. This property is already central in classical classification of isometries of the hyperbolic plane:

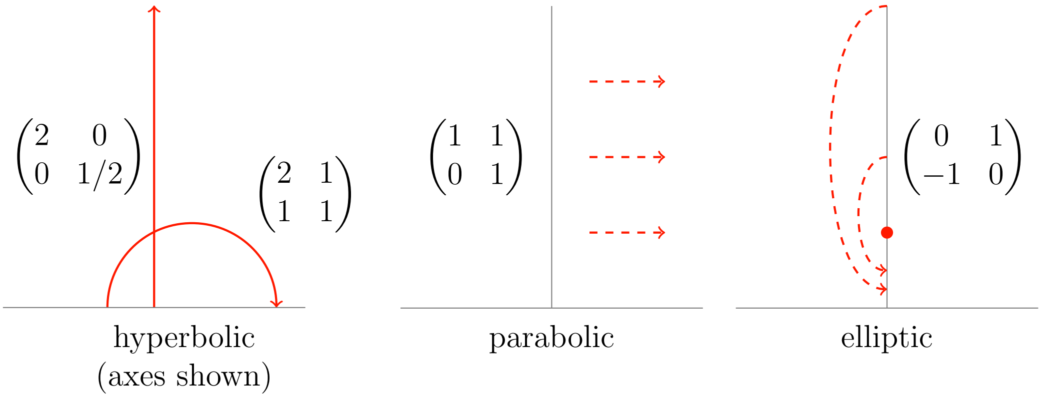

Figure 7.26.: The three types of orientation preserving isometries of $\mathbb{H}^2$, in the half-plane model

Remark 7.5.13 (The conic trichotomy)**.** Orientation preserving isometries (except the identity) of the hyperbolic plane $\mathbb{H}^2$ can be classified into the following three types [11, Proposition A.5.14ff]:

- *Hyperbolic.* A non-trivial orientation preserving isometry of $\mathbb{H}^2$ is *hyperbolic* if it has no fixed point in $\mathbb{H}^2$ and if it admits an *axis*, i.e., a geodesic line on which this geodesic acts by translation. Arguments similar to the ones in the previous proof show that such an axis is then unique.
- *Parabolic.* A non-trivial orientation preserving isometry of $\mathbb{H}^2$ is *parabolic* if it has no fixed point in $\mathbb{H}^2$ and if it admits no axis.
- *Elliptic.* A non-trivial orientation preserving isometry of $\mathbb{H}^2$ is *elliptic* if it has a fixed point in $\mathbb{H}^2$

The names derive from the fact that the group of orientation preserving isometries of $\mathbb{H}^2$ is isomorphic to the matrix group

$$\mathrm{PSL}(2,\mathbb{R}) = \mathrm{SL}(2,\mathbb{R})/\{E_2, -E_2\}$$

and that the above classification translates into quadratic equations for the matrix coefficients, which in turn are related to the classification of (non-degenerate) conic sections.

In the upper half-plane model for $\mathbb{H}^2$, prototypical examples of these three types of isometries are (Figure 7.26):

– *Hyperbolic.* For example, the matrices

$$\begin{pmatrix} 2 & 0 \\ 0 & 1/2 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$

induce the Möbius transformations

$$z \longmapsto 4 \cdot z \quad \text{and} \quad z \longmapsto \frac{2 \cdot z + 1}{z + 1},$$

respectively. Both of these isometries are hyperbolic. The first one has the imaginary axis as axis; the latter one has the geodesic line in $\mathbb{H}^2$ whose image is the half-circle that meets the real line in the points $1/2 \cdot (1 + \sqrt{5})$ and $1/2 \cdot (1 - \sqrt{5})$ as axis.

– *Parabolic.* The matrix

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

yields the parabolic Möbius transformation

$$z \longmapsto z + 1.$$

– *Elliptic.* The matrix

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

induces the elliptic Möbius transformation

$$z \longmapsto -\frac{1}{z}.$$

The fixed point of this Möbius transformation in $\mathbb{H}^2$ is i and this Möbius transformation can be viewed as rotation around π.

Using the language of boundaries of hyperbolic groups (Chapter 8.3), one can express this trichotomy also purely in terms of fixed points.

The proof of the first part in turn also heavily uses centralisers. To this end, we need some preparations: In order to gain control over the centralisers, we need a quasi-version of convexity. Recall that a subset of a geodesic space is *convex* if any geodesic joining whose endpoints lie in this subset must already be contained completely in this subset (Figure 7.27).

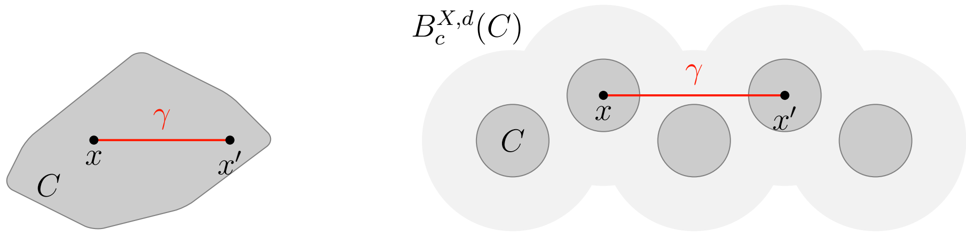

Figure 7.27.: Convexity and quasi-convexity, schematically

Definition 7.5.14 (Quasi-convex subspace)**.** Let X be a geodesic metric space. A subspace $C \subset X$ is *quasi-convex* if there is a $c \in \mathbb{R}_{\geq 0}$ with the following property: For all x, $x' \in C$ and all geodesics γ in X joining x with x' we have

$$\operatorname{im} \gamma \subset B_c^{X,d}(C).$$

Definition 7.5.15 (Quasi-convex subgroup)**.** Let G be a finitely generated group, let S be a finite generating set of G, and let $H \subset G$ be a subgroup of G. Then H is *quasi-convex with respect to* S, if H, viewed as a subset of $|\operatorname{Cay}(G, S)|$, is a quasi-convex subspace of $|\operatorname{Cay}(G, S)|$.

Proposition 7.5.16 (Properties of quasi-convex subgroups)**.** *Let G be a finitely generated group, and let S be a finite generating set of G.*

1. *If $H \subset G$ is a quasi-convex subgroup of G with respect to S, then H is finitely generated, and the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding.*
2. *If H and H' are quasi-convex subgroups of G with respect to S, then also the intersection $H \cap H'$ is a quasi-convex subgroup of G with respect to S.*

Proof. Exercise. $\square$

Caveat 7.5.17. The intersection of two convex subspaces of a geodesic space is always convex. However, the intersection of two quasi-convex subspaces of a geodesic space does not need to be quasi-convex again – this can even happen in hyperbolic metric spaces!

For example, the subsets

$$C := 2 \cdot \mathbb{Z},$$
$$C' := (2 \cdot \mathbb{Z} + 1) \cup \{n^2 \mid n \in \mathbb{Z}\}$$

are quasi-convex subsets of $\mathbb{R}$. However, their intersection

$$C \cap C' = \{n^2 \mid n \in 2 \cdot \mathbb{Z}\}$$

clearly is *not* quasi-convex in $\mathbb{R}$.

Proposition 7.5.18. *Let G be a hyperbolic group and let $g \in G$. Then the centraliser $C_G(g)$ is quasi-convex in G (with respect to any finite generating set of G).*

Proof. Let $S \subset G$ be a finite generating set of G, and let $\delta \in \mathbb{R}_{\geq 0}$ be chosen in such a way that $|\mathrm{Cay}(G, S)|$ is δ-hyperbolic. Let $\gamma\colon [0, L] \longrightarrow |\mathrm{Cay}(G, S)|$ be a geodesic with $\gamma(0) \in C_G(g)$ and $\gamma(L) \in C_G(g)$. We have to show that for all $t \in [0, L]$ the point $\gamma(t)$ is "uniformly close" to $C_G(g)$.

Without loss of generality, we may assume that $\gamma(0) = e$. We write $h := \gamma(L) \in C_G(g)$. Moreover, let $t \in [0, L]$; without loss of generality, we may assume that $\overline{h} := \gamma(t) \in G$ (otherwise, we pick a group element that is at most distance 1 away).

Clearly, $g \cdot \gamma\colon [0, L] \longrightarrow |\mathrm{Cay}(G, S)|$ is a geodesic starting in g, and ending in $g \cdot h = h \cdot g$. Because $|\mathrm{Cay}(G, S)|$ is δ-hyperbolic and because

$$d_S(h, g \cdot h) = d_S(h, h \cdot g) = d_S(e, g),$$

there is a constant $c \in \mathbb{R}_{\geq 0}$ depending only on $d_S(e, g)$ and δ such that γ and $g \cdot \gamma$ are uniformly c-close (apply a scaled version of Lemma 7.5.5 twice). In particular, we obtain

$$d_S(e, \overline{h}^{-1} \cdot g \cdot \overline{h}) = d_S(\overline{h}, g \cdot \overline{h}) = d_S\bigl(\gamma(t), g \cdot \gamma(t)\bigr) \leq c.$$

As next step we show that there is a "small" element $\overline{\overline{h}}$ satisfying

$$\overline{\overline{h}}^{-1} \cdot g \cdot \overline{\overline{h}} = \overline{h}^{-1} \cdot g \cdot \overline{h}.$$

To this end, we consider the following geodesics in $|\mathrm{Cay}(G, S)|$: Let $\overline{\gamma} := \gamma|_{[0,t]}\colon [0, t] \longrightarrow |\mathrm{Cay}(G, S)|$; i.e., $\overline{\gamma}$ is a geodesic starting in e and ending

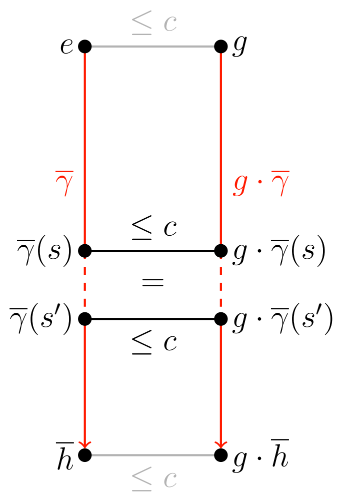

Figure 7.28.: Finding a shorter conjugating element

in $\overline{h}$. Furthermore, $g \cdot \overline{\gamma}$ is a geodesic starting in g and ending in $g \cdot \overline{h}$. By construction of c, the geodesics $\overline{\gamma}$ and $g \cdot \overline{\gamma}$ are uniformly c-close.

If $d_S(e, \overline{h}) > |B_c^{G,S}(e)|$, then by the pigeon-hole principle and the fact that goedesics in $|\operatorname{Cay}(G, S)|$ basically follow a shortest path in $\operatorname{Cay}(G, S)$ there exist parameters s, $s' \in [0, t]$ such that $s < s'$ and $\overline{\gamma}(s)$, $\overline{\gamma}(s') \in G$, as well as

$$\overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) = \overline{\gamma}(s')^{-1} \cdot g \cdot \overline{\gamma}(s')$$

(see Figure 7.28). Then the element

$$\overline{\overline{h}} := \overline{\gamma}(s) \cdot \overline{\gamma}(s')^{-1} \cdot \overline{h}$$

satisfies $d_S(e, \overline{\overline{h}}) < d_S(e, \overline{h})$ (because these elements are lined up on the same geodesic) and

$$\begin{aligned} \overline{\overline{h}}^{-1} \cdot g \cdot \overline{\overline{h}} &= \overline{h}^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) \cdot \overline{\gamma}(s')^{-1} \cdot \overline{h} \\ &= \overline{h}^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot g \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot \overline{h} \\ &= \overline{h}^{-1} \cdot g \cdot \overline{h}. \end{aligned}$$

Hence, inductively, we can find an element $\overline{\overline{h}} \in G$ with $d_S(e, \overline{\overline{h}}) \leq |B_c^{G,S}(e)|$ and
$$\overline{\overline{h}}^{-1} \cdot g \cdot \overline{\overline{h}} = \overline{h}^{-1} \cdot g \cdot \overline{h}.$$

The element $k := \overline{h} \cdot \overline{\overline{h}}^{-1}$ now witnesses that $\gamma(t)$ is $|B_c^{G,S}(e)|$-close to $C_G(g)$: On the one hand we have

$$\begin{aligned} d_S\bigl(k, \gamma(t)\bigr) &= d_S(\overline{h} \cdot \overline{\overline{h}}^{-1}, \overline{h}) = d_S(e, \overline{\overline{h}}^{-1}) = d_S(e, \overline{\overline{h}}) \\ &\leq \bigl|B_c^{G,S}(e)\bigr|. \end{aligned}$$

On the other hand, $k \in C_G(g)$ because

$$\begin{aligned} k \cdot g &= \overline{h} \cdot \overline{\overline{h}}^{-1} \cdot g \\ &= \overline{h} \cdot \overline{\overline{h}}^{-1} \cdot g \cdot \overline{\overline{h}} \cdot \overline{\overline{h}}^{-1} \\ &= \overline{h} \cdot \overline{h}^{-1} \cdot g \cdot \overline{h} \cdot \overline{\overline{h}}^{-1} \\ &= g \cdot k. \end{aligned}$$

Hence, γ is $|B_c^{G,S}(e)|$-close to $C_G(g)$, which shows that the centraliser $C_G(g)$ is quasi-convex in G with respect to S. $\square$

As promised, these quasi-convexity considerations allow us to complete the proof of the first part of Theorem 7.5.9:

Proof of Theorem 7.5.9, part 1. In view of Proposition 7.5.18, the centraliser $C_G(g)$ is a quasi-convex subgroup of G. In particular, $C_G(g)$ is finitely generated by Proposition 7.5.16, say by a finite generating set T. Then also the intersection
$$\bigcap_{t \in T} C_G(t) = C\bigl(C_G(g)\bigr),$$
which is the centre of $C_G(g)$, is a quasi-convex subgroup of G; so $C(C_G(g))$ is finitely generated and the inclusion $C(C_G(g)) \hookrightarrow G$ is a quasi-isometric embedding (Proposition 7.5.16).

On the other hand, $C(C_G(g))$ is Abelian and contains $\langle g \rangle_G \cong \mathbb{Z}$; because G is hyperbolic, it follows that $C(C_G(g))$ must be virtually $\mathbb{Z}$. Hence the infinite cyclic subgroup $\langle g \rangle_G$ has finite index in $C(C_G(g))$; in particular, the inclusion $\langle g \rangle_G \hookrightarrow C(C_G(g))$ is a quasi-isometric embedding.

Putting it all together, we obtain that the inclusion

$$\langle g \rangle_G \hookrightarrow C\bigl(C_G(g)\bigr) \hookrightarrow G$$

is a quasi-isometric embedding, as was to be shown. □

Corollary 7.5.19. *Let G be a hyperbolic group. Then G does not contain a subgroup isomorphic to $\mathbb{Z}^2$.*

Proof. Assume for a contradiction that G contains a subgroup H isomorphic to $\mathbb{Z}^2$. Let $h \in H \setminus \{e\}$; then h has infinite order and

$$\mathbb{Z}^2 \cong H = C_H(h) \subset C_G(h),$$

contradicting that the centraliser $C_G(h)$ is virtually $\langle h \rangle_G$ (Theorem 7.5.9). □

Corollary 7.5.20. *Let M be a closed connected smooth manifold. If the fundamental group $\pi_1(M)$ contains a subgroup isomorphic to $\mathbb{Z}^2$, then M does not admit a Riemannian metric of negative sectional curvature (in particular, M does not admit a hyperbolic structure).*

Proof. If M admits a Riemannian metric of negative sectional curvature, then its fundamental group $\pi_1(M)$ is hyperbolic (Example 7.3.3); hence, we can apply Corollary 7.5.19. □

Example 7.5.21 (Heisenberg group, $\mathrm{SL}(n,\mathbb{Z})$ and hyperbolicity)**.** By Corollary 7.5.19, the Heisenberg group is *not* hyperbolic: the subgroup of the Heisenberg group generated by the matrices

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

is isomorphic to $\mathbb{Z}^2$. Thus, for all $n \in \mathbb{N}_{\geq 3}$, the matrix groups $\mathrm{SL}(n,\mathbb{Z})$ also cannot be hyperbolic.

In particular, the closed connected smooth manifold given as the quotient of the Heisenberg group with $\mathbb{R}$-coefficients by the Heisenberg group does not admit a Riemannian metric of negative sectional curvature.

7.5.3 Application: Products and negative curvature

In view of Corollary 7.5.19, most non-trivial products of finitely generated groups are *not* hyperbolic.

A geometric counterpart of this consequence is the following classical theorem by Gromoll and Wolf [38] (for simplicity, we only state the version where the fundamental group has trivial centre):

Theorem 7.5.22 (Splitting theorem in non-positive curvature)**.** *Let M be a closed connected Riemannian manifold of non-positive sectional curvature whose fundamental group $\pi_1(M)$ has trivial centre.*

1. *If $\pi_1(M)$ is isomorphic to a product $G_1 \times G_2$ of non-trivial groups, then M is isometric to a product $N_1 \times N_2$ of closed connected Riemannian manifolds satisfying $\pi_1(N_1) \cong G_1$ and $\pi_1(N_2) \cong G_2$.*
2. *In particular: If M has negative sectional curvature, then M does not split as a non-trivial Riemannian product.*

In a more topological direction, the knowledge about centralisers in hyperbolic groups (together with standard arguments from algebraic topology) shows that manifolds of negative sectional curvature cannot even be dominated by non-trivial products [50, 51]:

Theorem 7.5.23 (Negatively curved manifolds are not presentable by products)**.** *Let M be an oriented closed connected Riemannian manifold of negative sectional curvature. Then there are no oriented closed connected manifolds N_1 and N_2 of non-zero dimension admitting a continuous map $N_1 \times N_2 \longrightarrow M$ of non-zero degree.*

7.5.4 Free subgroups of hyperbolic groups

Using similar arguments as in the discussion about centralisers of elements of infinite order in hyperbolic groups together with the ping-pong lemma (Theorem 4.4.1) shows that "generic" elements in hyperbolic groups fail to commute in the strongest possible way [17, Proposition III.Γ.3.20]:

Theorem 7.5.24 (Free subgroups in hyperbolic groups)**.** *Let G be a hyperbolic group, and suppose there are elements g, $h \in G$ of infinite order satisfying $\langle g\rangle_G \cap \langle h\rangle_G = \{e\}$.*

1. *Then the subgroup $\langle g, h\rangle_G$ of G generated by g and h is a free group of rank* 2.
2. *In particular, G has exponential growth (by Example 6.2.12).*

Proof. One considers the quasi-geodesic lines in G given by g and h. As a first step, one shows that one may assume that this lines are sufficiently far away from each other. In a second step, one then applies the ping-pong lemma to suitable subsets (defined in terms of these quasi-geodesic lines) to prove freeness of the subgroup generated by g and h. (Exercise.) □

The proofs of Theorem 7.5.9 and Theorem 7.5.24 are based on the geometry of (quasi-)geodesic lines in hyperbolic metric spaces. A systematic study of the geometry of (quasi-)geodesic rays in hyperbolic spaces leads to the Gromov boundary of hyperbolic groups (Chapter 8.3).

8 Ends and boundaries

In this chapter, we give a brief introduction into the geometry "at infinity" and how it can be applied to study finitely generated groups. Roughly speaking, a suitable notion of geometry at infinity (or a "boundary") should assign "nice" topological spaces to given metric spaces that reflect the behaviour of the given metric spaces "far out," and it should turn quasi-isometries of metric spaces into homeomorphisms of the corresponding topological spaces; so, a boundary mechanism can be thought of as a functor promoting maps and properties from the very wild world of quasi-isometries to the tamer world of topology.

In particular, boundaries yield quasi-isometry invariants; surprisingly, in many cases, the boundaries know enough about the underlying metric spaces to allow for interesting rigidity results.

In Section 8.1, we will describe what a boundary notion should be capable of and we will outline a basic construction principle. In Sections 8.2 and 8.3, we will discuss two such constructions in a bit more detail, namely ends of groups and the Gromov boundary of hyperbolic groups. As sample applications, we will look at Stallings's decomposition theorem for groups with infinitely many ends, and at Mostow rigidity (Section 8.4).

8.1 Geometry at infinity

What should a boundary object be capable of? A notion of boundary should be a functor $\cdot_\infty$ from the category QMet of metric spaces whose morphisms are quasi-isometric embeddings up to finite distance to the category of (preferably compact) topological spaces. I.e., a notion of boundary should be a way to associate

- to any metric space X a (preferably compact) topological space X_∞,
- and to any quasi-isometric embedding $f\colon X \longrightarrow Y$ between metric spaces a continuous map $f_\infty\colon X_\infty \longrightarrow Y_\infty$ between the corresponding topological spaces,

such that the following conditions are satisfied:

- If $f, g\colon X \longrightarrow Y$ are quasi-isometric embeddings that are finite distance apart, then $f_\infty = g_\infty$.
- If $f\colon X \longrightarrow Y$ and $g\colon Y \longrightarrow Z$ are quasi-isometric embeddings, then
$$(g \circ f)_\infty = g_\infty \circ f_\infty.$$
- Moreover, $(\mathrm{id}_X)_\infty = \mathrm{id}_{X_\infty}$ for all metric spaces X.

In particular: If $f\colon X \longrightarrow Y$ is a quasi-isometry, then $f_\infty\colon X_\infty \longrightarrow Y_\infty$ is a homeomorphism.

A first example is the trivial example: we can assign to every metric space the topological space consisting of exactly one element, and to every quasi-isometric embedding the unique map of this topological space to itself. Clearly, this construction does not contain any interesting geometric information.

A general construction principle that leads to interesting boundaries is given by thinking about points "at infinity" as the "endpoints" of rays: If X is a metric space, then one defines

$$X_\infty := \frac{\text{a suitable set of rays } [0,\infty) \longrightarrow X}{\text{a suitable equivalence relation on these rays}},$$

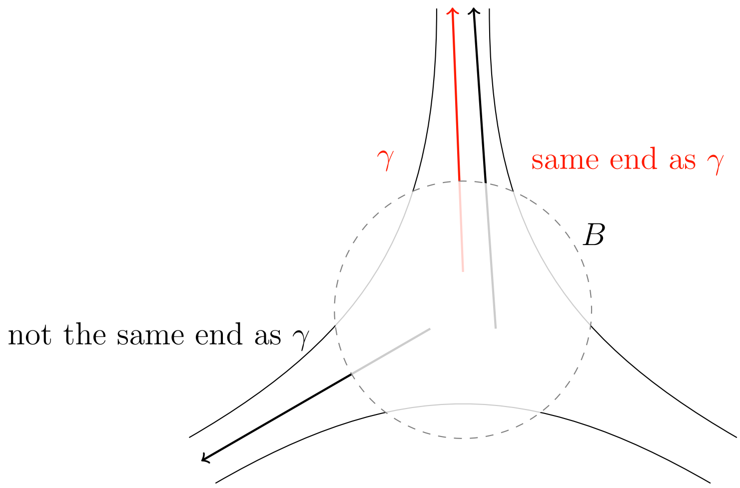

Figure 8.1.: The ends of a space, schematically

and if $f\colon X \longrightarrow Y$ is a quasi-isometric embedding, one sets

$$\begin{aligned} f_\infty\colon X_\infty &\longrightarrow Y_\infty \\ \bigl[\gamma\colon [0,\infty) \to X\bigr] &\longmapsto [f \circ \gamma]. \end{aligned}$$

In the following, we will briefly discuss two instances of this principle, namely ends of groups (Section 8.2) and the Gromov boundary of hyperbolic groups (Section 8.3).

8.2 Ends of groups

The ends of a space can be viewed as the set of principal "directions" leading "to infinity." Formally, the concept is described via rays and connected components when bounded pieces are removed (Figure 8.1):

Definition 8.2.1 (Ends of a metric space)**.** Let X be a geodesic metric space.

- A *proper ray* in X is a continuous map $\gamma\colon [0,\infty) \longrightarrow X$ such that for all bounded sets $B \subset X$ the pre-image $\gamma^{-1}(B) \subset [0,\infty)$ is bounded.
- Two proper rays γ, $\gamma'\colon [0,\infty) \longrightarrow X$ *represent the same end of* X if for all bounded subsets $B \subset X$ there exists a $t \in [0,\infty)$ such that $\gamma([t,\infty))$ and $\gamma'([t,\infty))$ lie in the same path-connected component of $X \setminus B$.
- If $\gamma\colon [0,\infty) \longrightarrow X$ is a proper ray, then we write $\operatorname{end}(\gamma)$ for the set of all proper rays that represent the same end as γ.
- We call

  $$\operatorname{Ends}(X) := \bigl\{\operatorname{end}(\gamma) \bigm| \gamma\colon [0,\infty) \longrightarrow X \text{ is a proper ray in } X\bigr\}$$

  the *space of ends of* X.
- We define a topology on $\operatorname{Ends}(X)$ by defining when a sequence of ends in X converges to a point in $\operatorname{Ends}(X)$: Let $(x_n)_{n\in\mathbb{N}} \subset \operatorname{Ends}(X)$, and let $x \in \operatorname{Ends}(X)$. We say that $(x_n)_{n\in\mathbb{N}}$ *converges to* x if there exist proper rays $(\gamma_n)_{n\in\mathbb{N}}$ and γ in X representing $(x_n)_{n\in\mathbb{N}}$ and x respectively such that the following condition is satisfied: For every bounded set $B \subset X$ there exists a sequence $(t_n)_{n\in\mathbb{N}} \subset [0,\infty)$ such that for all large enough $n \in \mathbb{N}$ the images $\gamma_n([t_n,\infty))$ and $\gamma([t_n,\infty))$ lie in the same path-connected component of $X \setminus B$.

Caveat 8.2.2. There are several common definitions of ends; while not all of them are equivalent on all classes of spaces, all of them give the same result when applied to finitely generated groups.

Example 8.2.3.

- If X is a metric space of finite diameter, then $\operatorname{Ends}(X) = \emptyset$.
- The proper rays

  $$\begin{aligned} [0,\infty) &\longrightarrow \mathbb{R} \\ t &\longmapsto t \\ t &\longmapsto -t \end{aligned}$$

  do not represent the same end. Moreover, a straightforward topological argument shows that every end of $\mathbb{R}$ can be represented by one of these two rays; hence,

  $$\operatorname{Ends}(\mathbb{R}) = \bigl\{\operatorname{end}(t \mapsto t), \operatorname{end}(t \mapsto -t)\bigr\}.$$

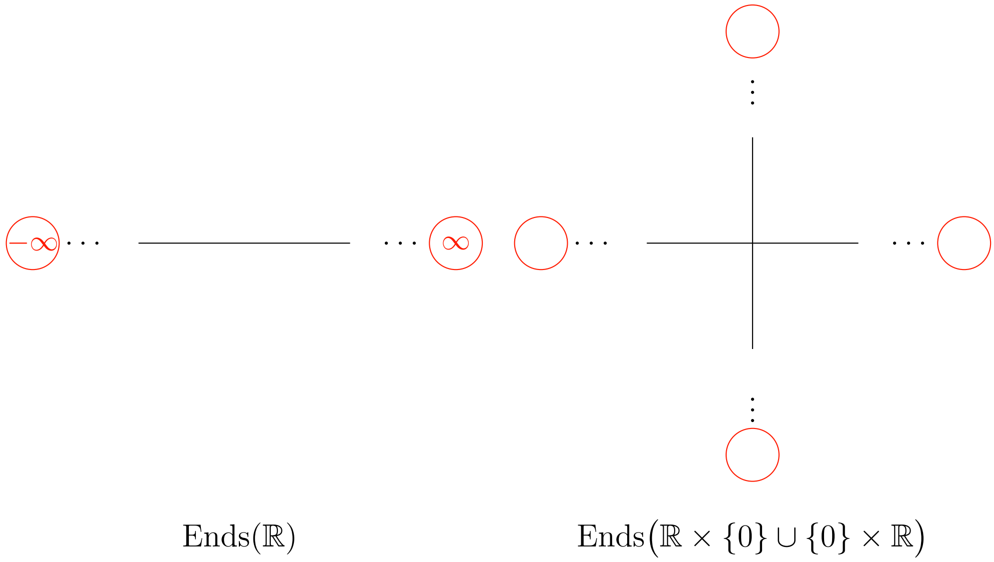

Figure 8.2.: Examples of ends (indicated by circles)

- If $n \in \mathbb{N}_{\geq 2}$, then Euclidean space $\mathbb{R}^n$ has only one end: If $B \subset \mathbb{R}^n$ is a bounded set, then there is an $r \in \mathbb{R}_{>0}$ such that $B \subset B_r^{\mathbb{R}^n,d}(0)$, and $\mathbb{R}^n \setminus B_r^{\mathbb{R}^n,d}(0)$ has exactly one path-component.
- Similarly, $|\operatorname{Ends}(\mathbb{H}^n)| = 1$ for all $n \in \mathbb{N}_{\geq 2}$ (Figure 8.2).
- The subspace $\mathbb{R} \times \{0\} \cup \{0\} \times \mathbb{R}$ of the Euclidean plane $\mathbb{R}^2$ has exactly four ends (Figure 8.2).
- Let $d \in \mathbb{N}_{\geq 3}$. If T is a (non-empty) tree in which every vertex has degree d, then T has infinitely many ends; more precisely, as a topological space, $\operatorname{Ends}(T)$ is a Cantor set [**?**, Chapter 13.4, 13.5][17, Exercise 8.31].

Remark 8.2.4 (Ends of quasi-geodesic spaces)**.** We can extend the definition of ends to quasi-geodesic spaces by allowing not only proper rays but also quasi-geodesic rays; quasi-geodesic rays then are defined to represent the same end, if, far out, they lie in the same quasi-path-components of the space in question minus bounded sets. Similarly, also the definition of the topology can be translated into the quasi-setting. Then a straightforward argument shows that these notions coincide for geodesic spaces [17, Lemma I.8.28, Proposition I.8.29].

Proposition 8.2.5 (Quasi-isometry invariance of ends)**.** *Let X and Y be quasi-geodesic metric spaces.*

1. *If $f\colon X \longrightarrow Y$ is a quasi-isometric embedding, then the map*

$$\begin{aligned} \operatorname{Ends}(f)\colon \operatorname{Ends}(X) &\longrightarrow \operatorname{Ends}(Y) \\ \operatorname{end}(\gamma) &\longmapsto \operatorname{end}(f \circ \gamma) \end{aligned}$$

   is well-defined and continuous.
2. *If $f, g\colon X \longrightarrow Y$ are quasi-isometric embeddings that are finite distance apart, then* $\operatorname{Ends}(f) = \operatorname{Ends}(g)$*.*

Hence, Ends *defines a functor from the full subcategory of* QMet *given by quasi-geodesic spaces to the category of topological spaces.*

In particular: If $f\colon X \longrightarrow Y$ is a quasi-isometry, then the induced map $\operatorname{Ends}(f)\colon \operatorname{Ends}(X) \longrightarrow \operatorname{Ends}(Y)$ *is a homeomorphism.*

Proof. This is a straightforward computation. Notice that in order to prove well-definedness in the first part, we need the freedom to represent ends by quasi-geodesic rays – the composition of a quasi-isometric embedding with a continuous ray in general is not continuous; so even if we were only interested in ends of geodesic spaces, we would still need to know that we can represent ends by quasi-geodesic rays. □

In particular, we obtain a notion of ends for finitely generated groups:

Definition 8.2.6 (Ends of a group)**.** Let G be a finitely generated group. The *space* $\operatorname{Ends}(G)$ *of ends of* G is defined as $\operatorname{Ends}(\operatorname{Cay}(G, S))$, where $S \subset G$ is some finite generating set of G. (Notice that, up to canonical homeomorphism, this does *not* depend on the choice of the finite generating set.)

In view of the proposition above, the space of ends of a finitely generated group is a quasi-isometry invariant. From Example 8.2.3 we obtain:

Example 8.2.7.

- If G is a finite group, then $\operatorname{Ends}(G) = \emptyset$.
- The group $\mathbb{Z}$ has exactly two ends (because $\mathbb{Z}$ is quasi-isometric to $\mathbb{R}$).
- Free groups of rank at least 2 have infinitely many ends; as a topological space the space of ends of a free group of rank at least 2 is a Cantor set.

– The group $\mathbb{Z}^2$ has only one end (because $\mathbb{Z}^2$ is quasi-isometric to $\mathbb{R}^2$).
– If M is a closed connected hyperbolic manifold of dimension at least 2, then the fundamental group $\pi_1(M)$ is quasi-isometric to $\mathbb{H}^n$, and so $\pi_1(M)$ has only one end.

Theorem 8.2.8 (Possible numbers of ends of groups). *Let G be a finitely generated group. Then G has 0, 1, 2 or infinitely many ends.*

Proof. By the functoriality of the ends construction, $\operatorname{Ends}(G)$ inherits a G-action from the left translation action of G on itself. *Assume* for a contradiction that G has a finite number $n \geq 3$ of ends. Then there is a finite index subgroup $H \subset G$ of G that acts trivially on $\operatorname{Ends}(G)$. Moreover, one can arrange for the vertices of one the rays to lie in H. Then one lets act these vertices on the other rays. (Exercise). $\square$

Moreover, using the Arzelá-Ascoli theorem, one can show that the space of ends of a finitely generated group is compact [17, Theorem I.8.32]. If a group has infinitely many ends, then the space of ends is uncountable, and every end is an accumulation point of ends [17, Theorem I.8.32].

As we have seen above, any finitely generated group has 0, 1, 2 or infinitely many ends. Conversely, we can use the number of ends of a group to learn something about the algebraic structure of the group in question [17, Theorem I.8.32][**?**, Chapter 13.6]:

Theorem 8.2.9 (Recognising groups via ends).

1. *A finitely generated group has no ends if and only if it is finite.*
2. *A finitely generated group has exactly two ends if and only if it is virtually $\mathbb{Z}$.*
3. Stallings's decomposition theorem. *A finitely generated group has infinitely many ends if and only if it is a non-trivial free amalgamated product over a finite group or if it is a non-trivial HNN-extension over a finite group.*

Notice that the second statement in particular gives yet another argument proving that $\mathbb{Z}$ is quasi-isometrically rigid.

The most interesting part is the third statement: If for some reason we know that a group has infinitely many ends, then we know that we can decompose the group into "smaller" pieces.

8.3 Boundary of hyperbolic groups

The space of ends is a rather crude invariant – many interesting groups have only one end. Therefore, we are interested in constructing finer boundary invariants. In the case of hyperbolic groups, such a refinement is possible:

Definition 8.3.1 (Gromov boundary)**.** Let X be a quasi-hyperbolic metric space.

- The *(Gromov) boundary* of X is defined as

  $$\partial X := \bigl\{\gamma\colon [0,\infty) \longrightarrow X \bigm| \gamma \text{ is a quasi-geodesic ray}\bigr\} /{\sim},$$

  where two quasi-geodesic rays γ, $\gamma'\colon [0,\infty) \longrightarrow X$ are equivalent if there exists a $c \in \mathbb{R}_{\geq 0}$ such that

  $$\operatorname{im}\gamma \subset B_c^{X,d}(\operatorname{im}\gamma') \qquad \text{and} \qquad \operatorname{im}\gamma' \subset B_c^{X,d}(\operatorname{im}\gamma)$$

  (i.e., $\operatorname{im}\gamma$ and $\operatorname{im}\gamma'$ have finite Hausdorff distance).
- We define a topology on ∂X by defining when a sequence in ∂X converges to a point in ∂X: Let $(x_n)_{n\in\mathbb{N}} \subset \partial X$, and let $x \in \partial X$. We say that $(x_n)_{n\in\mathbb{N}}$ *converges to* x if there exist quasi-geodesic rays $(\gamma_n)_{n\in\mathbb{N}}$ and γ representing $(x_n)_{n\in\mathbb{N}}$ and x respectively such that every subsequence of $(\gamma_n)_{n\in\mathbb{N}}$ contains a subsequence that converges (uniformly on compact subsets of $[0,\infty)$) to γ.

While the definition above technically would make sense for any (quasi-geoedesic) metric space, in general, such a boundary would be far too big (as there are too many quasi-geodesics) and rather difficult to keep under control.

Remark 8.3.2 (finite Hausdorff distance via QMet)**.** Let X be a metric space. The set of quasi-geodesic rays in X is nothing but $\operatorname{Mor}_{\mathsf{QMet}}([0,\infty), X)$ and the quasi-isometry group $\operatorname{QI}([0,\infty)) = \operatorname{Aut}_{\mathsf{QMet}}([0,\infty))$ acts from the right on $\mathrm{Mor}_{\mathsf{QMet}}([0,\infty), X)$ via pre-composition. Moreover, the canonical map

$$\begin{aligned} \mathrm{Mor}_{\mathsf{QMet}}([0,\infty), X) \big/ \mathrm{QI}([0,\infty)) &\longrightarrow \partial X \\ [\gamma] &\longmapsto [\gamma] \end{aligned}$$

is a bijection: Clearly, this map is well-defined and surjective. Moreover, it is injective: Let $\gamma, \gamma' \colon [0,\infty) \longrightarrow X$ be quasi-geodesic rays and suppose that there is a $c \in \mathbb{R}_{\geq} 0$ satisfying

$$\operatorname{im}\gamma \subset B_c^{X,d}(\operatorname{im}\gamma') \qquad \text{and} \qquad \operatorname{im}\gamma' \subset B_c^{X,d}(\operatorname{im}\gamma).$$

We then define a map $f \colon [0,\infty) \longrightarrow [0,\infty)$ by choosing for every $t \in [0,\infty)$ an $f(t) \in [0,\infty)$ with

$$d\bigl(\gamma(t), \gamma'(f(t))\bigr) \leq c.$$

A straightforward calculation shows that $f \colon [0,\infty) \longrightarrow [0,\infty)$ is a quasi-isometry with $\sup_{t\in[0,\infty)} d(\gamma(t), \gamma' \circ f(t)) \leq c$.

For proper hyperbolic metric spaces, the Gromov boundary can be expressed in terms of geodesic rays instead of quasi-geodesic rays:

Theorem 8.3.3 (A geodesic description of the boundary)**.** *Let X be a proper hyperbolic metric space.*

1. *Let $\gamma \colon [0,\infty) \longrightarrow X$ be a quasi-geodesic ray in X. Then there is a geodesic ray $\gamma' \colon [0,\infty) \longrightarrow X$ and a $c \in \mathbb{R}_{\geq 0}$ satisfying*

$$\operatorname{im}\gamma \subset B_c^{X,d}(\operatorname{im}\gamma') \qquad \textit{and} \qquad \operatorname{im}\gamma' \subset B_c^{X,d}(\operatorname{im}\gamma).$$

2. *Let $\gamma, \gamma' \colon [0,\infty) \longrightarrow X$ be geodesic rays in X with finite Hausdorff distance. Then $\sup_{t\in[0,\infty)} d(\gamma(t), \gamma(t')) < \infty$.*
3. *Let $x \in X$ and let $\gamma \colon [0,\infty) \longrightarrow X$ be a geodesic ray. Then there is a geodesic ray $\gamma' \colon [0,\infty) \longrightarrow X$ satisfying*

$$\gamma'(0) = x \quad \textit{and} \quad \sup_{t\in[0,\infty)} d\bigl(\gamma(t), \gamma'(t)\bigr) < \infty.$$

4. *In particular, for all $x \in X$ the canonical maps*

$$\begin{aligned} \textit{geodesic rays in } X \big/ \textit{finite distance} &\longrightarrow \partial X \\ \textit{geodesic rays in } X \textit{ starting in } x \big/ \textit{finite distance} &\longrightarrow \partial X \end{aligned}$$

are bijective.

Proof. The second part follows by applying scaled versions of Lemma 7.5.5 several times. The first and third part can be shown as follows: In hyperbolic spaces, quasi-geodesics stay close to geodesics (Theorem 7.2.11). Applying the Arzelá-Ascoli theorem to finite pieces of the quasi-geodesic rays in question proves the theorem. (Exercise). The last part just subsumes the other parts. □

Example 8.3.4.

- If X is a metric space of finite diameter, then $\partial X = \emptyset$.
- The Gromov boundary of the real line $\mathbb{R}$ consists of exactly two points corresponding to going to $+\infty$ and going to $-\infty$; hence, the Gromov boundary of $\mathbb{R}$ coincides with the space of ends of $\mathbb{R}$.
- The Gromov boundary of a regular tree of degree at least 3 is a Cantor set [**?**]
- One can show that for all $n \in \mathbb{N}_{\geq 2}$, the Gromov boundary of $\mathbb{H}^n$ is homeomorphic to the $(n-1)$-dimensional sphere S^{n-1} [**?**].

Proposition 8.3.5 (Quasi-isometry invariance of Gromov boundary)**.** *Let X and Y be quasi-hyperbolic spaces.*

1. *If $f \colon X \longrightarrow Y$ is a quasi-isometric embedding, then*
$$\begin{aligned} \partial f \colon \partial X &\longrightarrow \partial Y \\ [\gamma] &\longmapsto [f \circ \gamma] \end{aligned}$$
is well-defined and continuous.
2. *If $f,\, g \colon X \longrightarrow Y$ are quasi-isometric embeddings that are finite distance apart, then $\partial f = \partial g$.*

Hence, ∂ defines a functor from the full subcategory of QMet *given by quasi-hyperbolic spaces to the category of topological spaces. In particular: If $f \colon X \longrightarrow Y$ is a quasi-isometry, then the induced map $\partial f \colon \partial X \longrightarrow \partial Y$ is a homeomorphism.*

Proof. The proof is a straightforward computation. □

Corollary 8.3.6 (Quasi-isometry invariance of hyperbolic dimension)**.** *Let $n,\, m \in \mathbb{N}_{\geq 1}$. Then $\mathbb{H}^n \sim_{\mathrm{QI}} \mathbb{H}^m$ if and only if $n = m$.*

Proof. If $\mathbb{H}^n \sim_{\mathrm{QI}} \mathbb{H}^m$, then in view of the example above and the quasi-isometry invariance of the Gromov boundary we obtain homeomorphisms
$$S^{n-1} \cong \partial \mathbb{H}^n \cong \partial \mathbb{H}^m \cong S^{m-1}.$$

By a classical result in algebraic topology two spheres are homeomorphic if and only if they have the same dimension [25, Corollary IV.2.3]; hence, $n-1 = m-1$, and so $n = m$. □

Remark 8.3.7 (The conic trichotomy via fixed points)**.** The topology on the Gromov boundary of proper hyperbolic metric spaces X admits a canonical extension to

$$\overline{X} := X \cup \partial X$$

that is compatible with the metric topology on X and the topology on ∂X from above.

For example, there is a homeomorphism $\overline{\mathbb{H}^2} \longrightarrow D^2$ mapping $\partial\mathbb{H}^2$ to the boundary S^1 of the unit disk D^2 [**?**], and any isometry f of $\mathbb{H}^2$ yields a homeomorphism $\overline{f}\colon D^2 \longrightarrow D^2$. By the Brouwer fixed point theorem [25, Corollary IV.2.6], the latter map always has a fixed point. It is then possible to reformulate the conic trichotomy (Remark 7.5.13) for non-trivial orientation preserving isometries of $\mathbb{H}^2$ as follows [**?**]:

- Such an isometry f is *hyperbolic* if and only if $\overline{f}$ has exactly two fixed points and these fixed points lie on the boundary (namely the "endpoints" of the axis).
- Such an isometry f is *parabolic* if and only if $\overline{f}$ has exactly one fixed point and this fixed point lies on the boundary.
- Such an isometry f is *elliptic* if and only if $\overline{f}$ has exactly one fixed point and this fixed point does not lie on the boundary.

Moreover, the quasi-isometry invariance of Gromov boundary allows us to define the Gromov boundary for hyperbolic groups:

Definition 8.3.8 (Gromov boundary of a hyperbolic group)**.** Let G be a hyperbolic group. The *Gromov boundary* of G is defined as $\partial G := \partial\operatorname{Cay}(G, S)$, where $S \subset G$ is some finite generating set of G; up to canonical homeomorphism, this definition is independent of the choice of the finite generating set S.

Clearly, Proposition 8.3.5 also shows that the Gromov boundary of hyperbolic groups is a quasi-isometry invariant in the class of hyperbolic groups and that the Gromov boundary of group coincides with the Gromov boundary of the geometric realisations of its Cayley graphs (which in turn can be described via geodesic rays by Theorem 8.3.3). Moreover, Example 8.3.4 yields:

Example 8.3.9.

– If G is a finite group, then $\partial G = \emptyset$.
– The Gromov boundary $\partial \mathbb{Z}$ of $\mathbb{Z}$ consists of exactly two points because $\mathbb{Z}$ is quasi-isometric to $\mathbb{R}$.
– If M is a closed connected hyperbolic manifold of dimension n, then

$$\partial \pi_1(M) \cong \partial \mathbb{H}^n \cong S^{n-1}.$$

Conversely, it can be shown that if G is a torsion-free hyperbolic group whose boundary is a sphere of dimension $n - 1 \geq 5$, then G is the fundamental group of a closed connected aspherical manifold of dimension n [7].
– Let F be a finitely generated free group of rank at least 2. Then ∂F is a Cantor set; in particular, F and $\mathrm{SL}(2, \mathbb{Z})$ are *not* quasi-isometric to the hyperbolic plane $\mathbb{H}^2$.
– More generally, the Gromov boundary of a free product $G * H$ of two hyperbolic groups has the structure of a Cantor-like set, built from the Gromov boundaries of G and H respectively [91].

8.4 Application: Mostow rigidity

We briefly illustrate the power of boundary methods at the example of Mostow rigidity. Roughly speaking, Mostow rigidity says that certain manifolds that are equivalent in a rather weak, topological, sense (homotopy equivalent) must be equivalent in a rather strong, geometric, sense (isometric).

For the sake of simplicity, we discuss only the simplest version of Mostow rigidity, namely Mostow rigidity for closed hyperbolic manifolds:

Theorem 8.4.1 (Mostow rigidity – geometric version)**.** *Let $n \in \mathbb{N}_{\geq 3}$, and let M and N be closed connected hyperbolic manifolds of dimension n. If M and N are homotopy equivalent, then M and N are isometric.*

Theorem 8.4.2 (Mostow rigidity – algebraic version)**.** *Let $n \in \mathbb{N}_{\geq 3}$, and let Γ and Λ be cocompact lattices in* $\operatorname{Isom}(\mathbb{H}^n)$*. If Γ and Λ are isomorphic, then they are conjugate in* $\operatorname{Isom}(\mathbb{H}^n)$*, i.e., there exists a $g \in \operatorname{Isom}(\mathbb{H}^n)$ with $g \cdot \Gamma \cdot g^{-1} = \Lambda$.*

Notice that the corresponding statement for flat manifolds does not hold: Scaling the flat metric on a torus of any dimension still gives a flat metric on the same torus, but even though the underlying manifolds are homotopy equivalent (even homeomorphic), they are *not* isometric (e.g., scaling changes the volume).

Caveat 8.4.3. Notice that Mostow rigidity does *not* hold in dimension 2; in fact, in the case of surfaces of higher genus, the moduli space of hyperbolic structures is a rich and interesting mathematical object [11, Chapter B.4]. Together with Teichmüller space, this moduli space is the central object of study in Teichmüller theory.

Sketch proof of Mostow rigidity. We sketch the proof of the geometric version; that the geometric version and the algebraic version are equivalent follows from standard arguments in algebraic topology concerning the homotopy theory of classifying spaces/Eilenberg-Mac Lane spaces.

Let $f\colon M \longrightarrow N$ be a homotopy equivalence. Using covering theory and the fact that the Riemannian universal coverings of M and N are isometric to $\mathbb{H}^n$, we obtain a lift

$$\widetilde{f}\colon \mathbb{H}^n \longrightarrow \mathbb{H}^n$$

of f; in particular, $\widetilde{f}$ is compatible with the actions of $\pi_1(M)$ and $\pi_1(N)$ on $\mathbb{H}^n$ by deck transformations. Similar arguments as in the proof of the Švarc-Milnor lemma show that $\widetilde{f}$ is a quasi-isometry.

Hence, we obtain a homeomorphism

$$\partial\widetilde{f}\colon \partial\mathbb{H}^n \longrightarrow \partial\mathbb{H}^n$$

on the boundary of $\mathbb{H}^n$ that is compatible with the actions of $\pi_1(M)$ and $\pi_1(N)$ on $\partial\mathbb{H}^n$ induced by the deck transformation actions.

The main step of the proof is to show that this map $\partial\widetilde{f}$ on $\mathbb{H}^n$ is conformal (i.e., locally angle-preserving) with respect to the canonical homeoomorphism $\partial\mathbb{H}^n \cong S^{n-1}$; here, the condition that $n \geq 3$ enters. One way to show that $\partial\widetilde{f}$ is conformal is Gromov's proof via the simplicial volume and regular ideal simplices [**?**, **?**, **?**, **?**].

By a classical result from hyperbolic geometric, any conformal map on $\partial\mathbb{H}^n$ can be obtained as the boundary map of an isometry of $\mathbb{H}^n$ [**?**]; moreover, in our situation, it is also true that such an isometry $\overline{f} \in \operatorname{Isom}(\mathbb{H}^n)$ can be chosen in such a way that it is compatible with the deck transformation actions of $\pi_1(M)$ and $\pi_1(N)$ on $\mathbb{H}^n$.

Now a simple argument from covering theory shows that $\overline{f}$ induces an isometry $M \longrightarrow N$. (In view of homotopy theory of classifying spaces, it also follows that this isometry is homotopic to the original map f). □

Using finer boundary constructions (e.g., asymptotic cones), rigidity results can also be obtained in other situations [**?**].

Remark 8.4.4 (Borel conjecture)**.** In a more topological direction, a topological version of Mostow rigidity is formulated in the *Borel conjecture*: Closed connected manifolds with contractible universal covering space are homotopy equivalent if and only if they are homeomorphic. This conjecture is wide open in general, but many special cases are known to be true, including cases whose fundamental groups have a geometric meaning [60].

9 Amenable groups

The notion of amenability revolves around the leitmotiv of (almost) invariance. Different interpretations of this leitmotiv lead to different characterisations of amenable groups, e.g., via invariant means, Følner sets (i.e., almost invariant finite subsets), decomposition properties, or fixed point properties.

We will introduce amenable groups via invariant means (Chapter 9.1) and discuss first properties and examples. We will then study equivalent characterisations of amenability and their use cases (Chapter 9.2).

The different descriptions of amenability lead to various applications of amenability. For example, we will discuss the Banach-Tarski paradox (Chapter 9.2.3), Bilipschitz equivalence rigidity (Chapter 9.4), and the von Neumann problem (Chapter **??**). Moreover, we will briefly sketch (co)homological characterisations and impact of amenability.

9.1 Amenability via means

We will introduce amenable groups via invariant means. A mean can be viewed of as a generalised averaging operation for bounded functions. If X is a set, then $\ell^\infty(X, \mathbb{R})$ denotes the set of all bounded functions of type $X \longrightarrow \mathbb{R}$. Pointwise addition and scalar multiplication turns $\ell^\infty(X, \mathbb{R})$ into a real vector space. If G is a group, then any left G-action on X induces a left G-action on $\ell^\infty(G, \mathbb{R})$ via

$$\begin{aligned} G \times \ell^\infty(X, \mathbb{R}) &\longrightarrow \ell^\infty(X, \mathbb{R}) \\ (g, f) &\longmapsto \bigl(x \mapsto f(g^{-1} \cdot x)\bigr). \end{aligned}$$

Definition 9.1.1 (Amenable group)**.** A group G is *amenable*[1] if it admits a G-invariant mean on $\ell^\infty(G, \mathbb{R})$, i.e., an $\mathbb{R}$-linear map $m \colon \ell^\infty(G, \mathbb{R}) \longrightarrow \mathbb{R}$ with the following properties:

- *Normalisation.* We have $m(1) = 1$.
- *Positivity.* We have $m(f) \geq 0$ for all $f \in \ell^\infty(G, \mathbb{R})$ that satisfy $f \geq 0$ pointwise.
- *Left-invariance.* For all $g \in G$ and all $f \in \ell^\infty(G, \mathbb{R})$ we have

  $$m(g \cdot f) = m(f)$$

  with respect to the left G-action on $\ell^\infty(G, \mathbb{R})$ induced from the left translation action of G on G.

[1] in German: mittelbar/amenabel

9.1.1 First examples of amenable groups

Example 9.1.2. Clearly, finite groups are amenable: If G is a finite group, then

$$\begin{aligned} \ell^\infty(G,\mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto \frac{1}{|G|} \cdot \sum_{g \in G} f(g) \end{aligned}$$

is a G-invariant mean on $\ell^\infty(G,\mathbb{R})$.

Proposition 9.1.3. *Every Abelian group is amenable.*

The proof relies on the Markov-Kakutani fixed point theorem from functional analysis [79, Proposition 0.14]:

Theorem 9.1.4 (Markov-Kakutani fixed point theorem)**.** *Let V be a locally convex $\mathbb{R}$-vector space, let $A \times V \longrightarrow V$ be an action of an Abelian group A by continuous linear maps on V, and let $K \subset V$ be a non-empty compact convex subset with $A \cdot K \subset K$. Then there is an $x \in K$ such that*

$$\forall_{g \in A} \;\; g \cdot x = x.$$

Proof of Proposition 9.1.3. Let G be an Abelian group. The function space $\ell^\infty(G,\mathbb{R})$ is a normed real vector space with respect to the supremum norm, and the left G-action on $\ell^\infty(G,\mathbb{R})$ is isometric. We now consider the topological dual L of $\ell^\infty(G,\mathbb{R})$ with respect to the weak*-topology. Then L inherits a left G-action by continuous linear maps from the isometric left G-action on $\ell^\infty(G,\mathbb{R})$.

We now consider

$$M := \bigl\{ m \in L \bigm| m(1) = 1 \text{ and } m \text{ is positive} \bigr\}.$$

Clearly, M is a closed, convex subspace of L, and $G \cdot M \subset M$. Moreover, M is non-empty by the Hahn-Banach theorem [53, Chapter IV.1] and compact by the Banach-Alaoglu theorem [53, Chapter IV.1]. Hence, M contains a G-fixed point m, by the Markov-Kakutani theorem (Theorem 9.1.4). By construction, m is then a G-invariant mean on G, and so G is amenable. $\square$

Proposition 9.1.5. *The free group F_2 of rank 2 ist* not *amenable.*

Proof. *Assume* for a contradiction that F_2 is amenable, say via an invariant mean m on F_2. Let $\{a, b\} \subset F_2$ be a free generating set, and consider the set $A \subset F_2$ of reduced words that start with a non-trivial power of a. Then

$$A \cup a \cdot A = F_2.$$

Denoting characteristic functions of subsets with $\chi_{\dots}$, we obtain

$$\begin{aligned} 1 = m(1) = m(\chi_{F_2}) &= m(\chi_A + \chi_{a \cdot A}) \\ &\leq m(\chi_A) + m(a \cdot \chi_A) \\ &= 2 \cdot m(\chi_A), \end{aligned}$$

and hence $m(\chi_A) \geq 1/2$.

On the other hand, the sets $A, b \cdot A, b^2 \cdot A$ are disjoint, and so

$$\begin{aligned} 1 \geq m(\chi_{A \cup b \cdot A, b^2 \cdot A}) &= m(\chi_A) + m(\chi_{b \cdot A}) + m(\chi_{b^2} \cdot A) \\ &= m(\chi_A) + m(b \cdot \chi_A) + m(b^2 \cdot \chi_A) \\ &= 3 \cdot m(\chi_A) \\ &\geq \frac{3}{2}, \end{aligned}$$

which cannot be. Hence, F_2 is *not* amenable. □

9.1.2 Inheritance properties

We will now familiarise ourselves with the basic inheritance properties of amenable groups:

Proposition 9.1.6 (Inheritance properties of amenable groups)**.**

1. *Subgroups of amenable groups are amenable.*
2. *Homomorphic images of amenable groups are amenable.*
3. *Let*
$$1 \longrightarrow N \xrightarrow{\ i\ } G \xrightarrow{\ \pi\ } Q \longrightarrow 1$$
be an extension of groups. Then G is amenable if and only if N and Q are amenable.

4. *Let G be a group, and let $(G_i)_{i\in I}$ be a directed set of ascending amenable subgroups of G with $G = \bigcup_{n\in\mathbb{N}} G_n$. Then G is amenable.*

Proof. Ad 1. Let G be an amenable group with left-invariant mean m and let $H \subset G$ be a subgroup. Moreover, let $R \subset G$ be a set of representatives for $H \setminus G$, and let $s\colon G \longrightarrow H$ be a map with

$$g \in s(g) \cdot R$$

for all $g \in G$. Then a small calculation shows that

$$\begin{aligned} \ell^\infty(H, \mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto m(f \circ s) \end{aligned}$$

is a left-invariant mean on H.

Ad 2. Let G be an amenable group with left-invariant mean m and let $\pi\colon G \longrightarrow Q$ be a surjective group homomorphism. Then

$$\begin{aligned} \ell^\infty(Q, \mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto m(f \circ \pi) \end{aligned}$$

is a left-invariant mean on Q.

Ad 3. If G is amenable, then N and Q are amenable by the previous parts.

Conversely, suppose that N and Q are amenable with left-invariant means m_N and m_Q, respectively. Then

$$\begin{aligned} \ell^\infty(G, \mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto m_Q\bigl(g \cdot N \mapsto m_N(n \mapsto f(g \cdot n))\bigr) \end{aligned}$$

is a left-invariant mean on G.

Ad 4. For each $i \in I$ let m_i be a left-invariant mean on G_i. We then set

$$\begin{aligned} \widetilde{m}_i\colon \ell^\infty(G, \mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto m_i(f|_{G_i}). \end{aligned}$$

In view of the Banach-Alaoglu theorem, there is a subnet of $(\widetilde{m}_i)_{i\in I}$ that converges to a functional m on $\ell^\infty(G, \mathbb{R})$. One then easily checks that this limit m is a left-invariant mean on G. $\square$

Corollary 9.1.7 (Locally amenable groups are amenable)**.** *Let G be a group. Then G is amenable if and only if all finitely generated subgroups of G are amenable.*

Proof. If G is amenable, then all its subgroups are amenable.

Conversely, suppose that all finitely generated subgroups are amenable. Because the finitely generated subgroups of G form an ascending directed system of subgroups of G that cover all of G, the last part of Proposition 9.1.6 shows that G is amenable. □

Moreover, the inheritance properties give us some indication for the location of the class of amenable groups in the universe of groups (Figure 1.2):

Corollary 9.1.8 (Solvable groups are amenable)**.** *Every solvable group is amenable.*

Proof. By Proposition 9.1.3, every Abelian group is amenable. By induction along the derived series, we obtain with help of Proposition 9.1.6 that every solvable group is amenable. □

Conversely, obviously not every amenable group is solvable; for example, the finite group S_5 is amenable (Example 9.1.2), but it is well-known that S_5 is *not* solvable.

Remark 9.1.9 (Elementary amenable groups)**.** The class of *elementary amenable* groups is the smallest class of groups that contains all Abelian and all finite groups and that is closed under taking subgroups, quotients, extensions and directed ascending unions. By Example 9.1.2, Proposition 9.1.3, and Proposition 9.1.6, every elementary amenable group is amenable. Notice however, that not every amenable group is elementary amenable; this can for example be seen via the Grigorchuk groups [37].

Corollary 9.1.10. *Groups that contain a free subgroup of rank 2 are* not *amenable.*

Proof. This follows from Proposition 9.1.5 and Proposition 9.1.6. □

Remark 9.1.11 (The von Neumann problem)**.** The notion of amenability was originally introduced by John von Neumann [71]. In view of Corollary 9.1.10, he asked whether every non-amenable group contained a free subgroup of rank 2.

This question was answered negatively by Ol'shanskii [75] who constructed a non-amenable torsion group; in particular, these groups cannot contain a free subgroup of rank 2.

In Chapter **??** we discuss versions of the von Neumann problem for group actions and in measurable group theory.

9.2 Further characterisations of amenability

The notion of amenability revolves around the leitmotiv of (almost) invariance. We have seen the definition via invariant means in Chapter 9.1. In the following, we will study equivalent characterisations of amenability and their use cases, focusing on geometric properties. A more thorough treatment of amenable groups can be found in the books by Paterson [79] and Runde [83].

9.2.1 Følner sequences

We begin with a geometric characterisation of amenability via so-called Følner sets. We first formulate the notion of Følner sets for UDBG spaces (Definition 5.5.8).

Definition 9.2.1 (Følner sequence)**.** Let X be a UDBG space.

- If $F \subset X$ and $r \in \mathbb{N}$, then we denote the *r-boundary of F in X* by
$$\partial_r^X F := \{x \in X \setminus F \mid \exists_{f \in F} \; d(x,f) \leq r\}$$
(Figure 9.1).
- A *Følner sequence for X* is a sequence $(F_n)_{n \in \mathbb{N}}$ of non-empty finite subsets of X with the following property: For all $r \in \mathbb{N}$ we have
$$\lim_{n \to \infty} \frac{|\partial_r^X F_n|}{|F_n|} = 0.$$

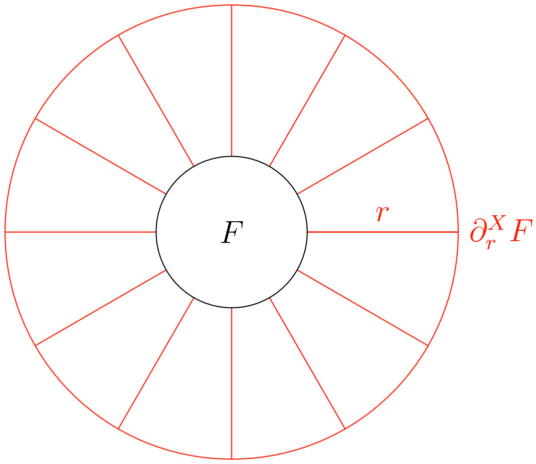

Figure 9.1.: The r-boundary of a set, schematically

Følner sets have two interpretations:

- Geometrically, Følner sets are efficient in the sense that they have "small" boundary, but "large" volume.
- More algebraically, for finitely generated groups, Følner sets can be thought of as (almost) invariant finite subsets: Let G be a finitely generated group and let $S \subset G$ be a finite generating set. Then for all $F \subset G$ we have

$$\partial_1^{G,d_S}(F) = \bigl\{ g \in G \setminus F \bigm| \exists_{s \in S \cup S^{-1}} \;\; g \cdot s \in F \bigr\}.$$

  Hence, $\partial_1^{G,d_S}(F)$ being "small" in comparison with F means that $(S \cup S^{-1}) \cdot f \in f$ for "many" $f \in F$.

Example 9.2.2. From the calculations in Example 6.1.2 we obtain:

- Let $n \in \mathbb{N}$ and let $S := \{e_1, \dots, e_n\}$ be the standard generating set of $\mathbb{Z}^n$. Then $(\{-k, \dots, k\}^n)_{k \in \mathbb{N}}$ is a Følner sequence for $(\mathbb{Z}^n, d_S)$.
- Let S be a free generating set of F_2. Then the balls $(B_n^{F,S})_{n \in \mathbb{N}}$ do *not* form a Følner sequence of (F_2, d_S). In fact, we will see that F_2 does *not* admit any Følner sequence.

Proposition 9.2.3 (Subexponential growth yields Følner sequences)**.** *Let G be a finitely generated group of subexponential growth. Let $S \subset G$ be a finite generating set of G. Then (G, d_S) admits a Følner sequence.*

Proof. For $n \in \mathbb{N}$ let
$$F_n := B_n^{G,S}(e) \subset G.$$
Then $(F_n)_{n\in\mathbb{N}}$ is a Følner sequence for G with respect to S: *Assume* for a contradiction that there is an $r \in \mathbb{N}_{>0}$ and an $\varepsilon \in \mathbb{R}_{>0}$ as well as an $N \in \mathbb{N}$ with
$$\forall_{n\in\mathbb{N}_{\geq N}} \quad \frac{|\partial_r^{G,d_S} F_n}{|F_n|} \geq \varepsilon.$$
By definition of the word metric, we have $|F_n| = \beta_{G,S}(n)$ and $|\partial_r^{G,d_S}(F_n)| = |B_{n+r}^{G,S} \setminus B_n^{G,S}| = \beta_{G,S}(n+r) - \beta_{G,S}(n)$. Hence, for all $n \in \mathbb{N}_{\geq N}$ we obtain
$$\frac{\beta_{G,S}(n+r)}{\beta_{G,S}(n)} \geq 1 + \varepsilon.$$
Inductively, we derive $\beta_{G,S} \succ (x \mapsto (1+\varepsilon)^x$, which contradicts that G has subexponential growth. Therefore, $(F_n)_{n\in\mathbb{N}}$ is a Følner sequence for G, which means that G is amenable. □

There are several variations of the notion of Følner sets or Følner sequences; in the end, they all lead to the same class of finitely generated groups. For instance, we have the following:

Proposition 9.2.4. *Let X be a UDBG space. Then X admits a Følner sequence if and only if for every $r \in \mathbb{N}$ and every $\varepsilon \in \mathbb{R}_{>0}$ there is a non-empty finite subset $F \subset X$ with*
$$\frac{|\partial_r^X F|}{|F|} \leq \varepsilon.$$

Proof. Any Følner sequence clearly leads to subsets with the desired properties. Conversely, suppose that for every $n \in \mathbb{N}$ there is a non-empty finite subset $F_n \subset X$ with
$$\frac{|\partial_n^X F_n|}{|F_n|} \leq \frac{1}{n}.$$
Then $(F_n)_{n\in\mathbb{N}}$ is easily seen to be a Følner sequence for X. □

Theorem 9.2.5 (Amenability via Følner sequences)**.** *Let G be a finitely generated group and let $S \subset G$ be a finite generating set. Then G is amenable if and only if (G, d_S) admits a Følner sequence.*

Sketch of proof. Let $(F_n)_{n\in\mathbb{N}}$ be a Følner sequence for G (with respect to the finite generating set S). Moreover, let ω be a non-principal ultrafilter on $\mathbb{N}$. Then

$$\begin{aligned} \ell^\infty(G,\mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto \lim_{n\in\omega} \frac{1}{|F_n|}\cdot \sum_{g\in F_n} f(g) \end{aligned}$$

is a left-invariant mean on G; here, $\lim_{n\in\omega}$ denotes the limit along ω, which is defined for all bounded sequences in $\mathbb{R}$ and picks one of the accumulation points of the argument sequence [21, Chapter J]. The almost invariance of the Følner sets in the limit translates into invarince. Hence, G is amenable. Alternatively, such a mean can also be obtained using weak*-limits in $\ell^\infty(G,\mathbb{R})'$ [**?**].

Conversely, let G be amenable. Recall that $\ell^1(G,\mathbb{R})$ is weak*-dense in the double dual $\ell^1(G,\mathbb{R})'' = \ell^\infty(G,\mathbb{R})'$ [**?**]. Then any invariant mean on $\ell^\infty(G,\mathbb{R})$ is an element of $\ell^1(G,\mathbb{R})'$ and thus can be approximated by ℓ^1-functions. These ℓ^1-functions in turn can be approximated by ℓ^1-functions with finite support. The invariance of the mean then translates into almost invariance of these finite supports, which yields a Følner sequence [**?**]. □

Because we know an explicit Følner sequence for $\mathbb{Z}$ we could use the "construction" in the proof of Theorem 9.2.5 to produce an invariant mean on $\ell^\infty(\mathbb{Z},\mathbb{R})$. However, non-principal ultrafilters on $\mathbb{N}$ cannot be made explicit, and also the resulting invariant means on $\ell^\infty(\mathbb{Z},\mathbb{R})$ cannot be made explicit [**?**].

Corollary 9.2.6. *Every finitely generated group of subexponential growth is amenable.*

Proof. This follows by combining Theorem 9.2.5 with Proposition 9.2.3. □

The converse of this corollary does not hold in general:

Caveat 9.2.7 (An amenable group of exponential growth)**.** The semi-direct product $\mathbb{Z}^2 \rtimes_A \mathbb{Z}$ with

$$A := \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

has exponential growth and is solvable (Caveat 6.3.7), whence amenable (Corollar 9.1.8).

9.2.2 Paradoxical decompositions

Another, geometric, characterisation of amenability is based on decomposition paradoxes. These decomposition properties are the main ingredient in the Banach-Tarski paradox (Chapter 9.2.3).

Definition 9.2.8 (Paradoxical group)**.** A group G is *paradoxical* if it admits a paradoxical decomposition. A *paradoxical decomposition of* G is a triple $\bigl(K, (A_g)_{g\in K}, (B_g)_{g\in K}\bigr)$ where $K \subset G$ is finite and $(A_g)_{g\in K}$, $(B_g)_{g\in K}$ are families of subsets of G with the property that

$$G = \Bigl(\bigcup_{g\in K} A_g\Bigr) \cup \Bigl(\bigcup_{g\in K} B_g\Bigr), \qquad G = \bigcup_{g\in K} g \cdot A_g, \qquad G = \bigcup_{g\in K} g \cdot B_g$$

are *disjoint* unions.

In other words, a paradoxical decomposition of a group is a decomposition into disjoint subsets such that these subsets can be rearranged into two copies of the group.

Proposition 9.2.9 (Non-Abelian free groups are paradoxical)**.** *The free group of rank* 2 *is paradoxical.*

Proof. Exercise. □

The proof of the previous proposition similar to the proof of Proposition 9.1.5; more generally, paradoxical groups are not amenable (and vice versa):

Theorem 9.2.10 (Tarski's theorem)**.** *Let G be a group. Then G is paradoxical if and only if G is* not *amenable.*

Proof. Let G be paradoxical and let $(K, (A_g)_{g\in K}, (B_g)_{g\in K})$ be a paradoxical decomposition of G. *Assume* for a contradiction that G is amenable, and let m be an invariant mean for G. Because the corresponding unions all are disjoint and m is left-invariant, we obtain

$$1 = m(\chi_G) = \sum_{g \in K} m(\chi_{g \cdot A_g}) = \sum_{g \in K} m(g \cdot \chi_{A_g}) = \sum_{g \in K} m(\chi_{A_g}),$$
$$1 = \sum_{g \in K} m(\chi_{B_g}),$$

and hence

$$1 = m(\chi_G) = \sum_{g \in K} m(\chi_{A_g}) + \sum_{g \in K} m(\chi_{B_g}) = 1 + 1 = 2,$$

which cannot be.

Conversely, let G be non-amenable. Then there is "enough space" in G to perform a combinatorial construction – based a version of the marriage theorem (Theorem 9.4.3) – leading to a paradoxical decomposition [21, Theorem 4.9.1]. □

9.2.3 Application: The Banach-Tarski paradox

We will now discuss how actions of paradoxical groups lead to paradoxical decomposition of sets, the most famous example being the Banach-Tarski paradox (Theorem 9.2.15).

Definition 9.2.11 (Paradoxical decomposition)**.** Let G be a group and consider an action of G on a set X. Then a subset $Y \subset X$ is *G-paradoxical* if Y admits a G-paradoxical decomposition. A *G-paradoxical decomposition* is a triple $\bigl(K, (A_g)_{g \in K}, (B_g)_{g \in K}\bigr)$ where $K \subset G$ is finite and $(A_g)_{g \in K}$, $(B_g)_{g \in K}$ are families of subsets of Y with the property that

$$Y = \Bigl(\bigcup_{g \in K} A_g\Bigr) \cup \Bigl(\bigcup_{g \in K} B_g\Bigr), \qquad Y = \bigcup_{g \in K} g \cdot A_g, \qquad Y = \bigcup_{g \in K} g \cdot B_g$$

are *disjoint* unions.

As in the group case, a paradoxical decomposition of a set is a decomposition into disjoint subsets such that these subsets can be rearranged using

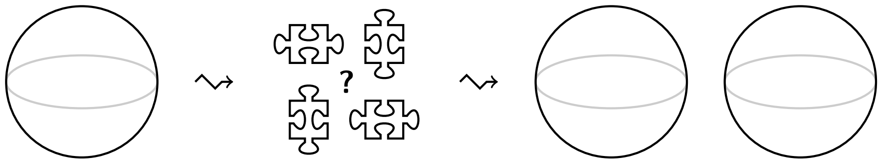

Figure 9.2.: The Hausdorff paradox, schematically

the group action into two copies of the group. Notice that in the literature several different notions of paradoxical decompositions are used.

In the presence of the axiom of choice, we can translate paradoxical decomposition of groups into paradoxical decomposition of sets:

Proposition 9.2.12 (Paradoxical groups induce paradoxical decomposition)**.** *Let G be a paradoxical group acting freely on a non-empty set X. Then X is G-paradoxical.*

Proof. Let $(K, (A_g)_{g\in G}, (B_g)_{g\in G})$ be a paradoxical decomposition of the group G. Using the axiom of choice, we find a subset $R \subset X$ that contains exactly one point from every G-orbit. Then a straightforward calculation shows that

$$\bigl((A_g \cdot R)_{g\in K}, (B_g \cdot R)_{g\in K}\bigr)$$

is a G-paradoxical decomposition of X. $\square$

The special orthogonal group $\mathrm{SO}(3)$ acts on $\mathbb{R}^3$ by matrix multiplication. This action is isometric and hence induces a well-defined action on the unit sphere $S^2 \subset \mathbb{R}^3$. Moreover, $\mathrm{SO}(3)$ contains a free subgroup of rank 2 (Proposition 9.2.13), which leads to the Hausdorff paradox (Figure 9.2).

Proposition 9.2.13. *The subgroup of the special orthogonal group* $\mathrm{SO}(3)$ *generated by*

$$\begin{pmatrix} \frac{3}{5} & \frac{4}{5} & 0 \\ -\frac{4}{5} & \frac{3}{5} & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad \textit{and} \qquad \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{3}{5} & -\frac{4}{5} \\ 0 & \frac{4}{5} & \frac{3}{5} \end{pmatrix}$$

is free of rank 2.

Proof. Exercise. $\square$

Corollary 9.2.14 (Hausdorff paradox)**.** *There is a countable set $D \subset S^2$ such that $S^2 \setminus D$ is* SO(3)*-paradoxical with respect to the canonical* SO(3)*-action on S^2.*

Proof. By Proposition 9.2.13, the group SO(3) contains a free subgroup G of rank 2. All matrices of SO(3) act on S^2 by a rotation around a line and so every non-trivial element g of SO(3) has exactly two fixed points $x_{g,1}, x_{g,2}$ on S^2. Therefore,

$$D := \bigl\{ g \cdot x_{g,j} \bigm| g \in G \setminus \{e\}, j \in \{1,2\} \bigr\}$$

is a countable set and the SO(3)-action on S^2 restricts to a free G-action on $S^2 \setminus D$. By Proposition 9.2.12 and Proposition 9.2.9, the complement $S^2 \setminus D$ is G-paradoxical. Because G is a subgroup of SO(3), this implies that $S^2 \setminus D \subset S^2$ is also SO(3)-paradoxical. □

The Hausdorff paradox can be improved as follows [83, Chapter 0.1]:

Theorem 9.2.15 (Banach-Tarski paradox for the sphere)**.** *The sphere S^2 is paradoxical with respect to the canonical* SO(3)*-action on S^2 in the following sense: There exist $n, m \in \mathbb{N}$ and pairwise disjoint subsets $A_1, \dots, A_n$, $B_1, \dots, B_m \subset S^2$ as well as group elements $g_1, \dots, g_n, h_1, \dots, h_m \in \mathrm{SO}(3)$ satisfying*

$$\bigcup_{j=1}^{n} g_j \cdot A_j = S^2 = \bigcup_{j=1}^{m} h_j \cdot B_j.$$

Clearly, the pieces of any such paradoxical decomposition of S^2 are *not* Lebesgue measurable, and hence are rather strange sets. Notice that the Hausdorff paradox and the Banach-Tarski paradox rely on the axiom of choice (for the set S^2) [**?**].

Further generalisations inlcude the Banach-Tarski paradoxes for balls in $\mathbb{R}^3$ and for bounded subsets in $\mathbb{R}^3$ [83, Chapter 0.1].

9.2.4 (Co)Homological characterisations of amenability

Amenability admits several characterisations in terms of suitable (co)-homology theories. We will briefly sketch characterisations in terms of uniformly finite homology and bounded cohomology, respectively.

We begin with the characterisation via uniformly finite homology [12, ?].

Definition 9.2.16 (Fundamental class)**.** Let R be a normed ring with unit and let X be a UDBG space. Then $\sum_{x \in X} 1 \cdot x$ is a cycle in $C_0^{\mathrm{uf}}(X; R)$. The corresponding class

$$[X]_R := \Bigl[\sum_{x \in X} 1 \cdot x\Bigr] \in H_0^{\mathrm{uf}}(X; R)$$

is the *fundamental class of X in $H_0^{\mathrm{uf}}(\,\cdot\,, R)$*.

Theorem 9.2.17 (Amenability and the fundamental class)**.** *Let X be a UDBG space. Then the following are equivalent:*

1. *The space X is* not *amenable.*
2. *We have $[X]_{\mathbb{Z}} = 0$ in $H_0^{\mathrm{uf}}(X; \mathbb{Z})$.*
3. *We have $[X]_{\mathbb{R}} = 0$ in $H_0^{\mathrm{uf}}(X; \mathbb{R})$.*

Moreover, there is also a characterisation of amenability in terms of bounded cohomology [46, 72] and ℓ^1-homology [55]:

Theorem 9.2.18 (Amenability and bounded cohomology/ℓ^1-homology)**.** *Let G be a group. Then the following are equivalent:*

1. *The group G is amenable.*
2. *For all Banach G-modules V and all $k \in \mathbb{N}$ we have $H_b^k(G; V') \cong 0$.*
3. *For all Banach G-modules V and all $k \in \mathbb{N}$ we have $H_k^{\ell^1}(G; V) \cong 0$.*

9.3 Quasi-isometry invariance of amenability

Amenability is a geometric property of finitely generated groups; more generally, amenability is a quasi-isometry invariant property of UDBG spaces:

Theorem 9.3.1 (Quasi-isometry invariance of amenability)**.** *Let X and Y be UDBG spaces. Then X is amenable if and only if Y is amenable.*

Proof. Let $f\colon X \longrightarrow Y$ be a quasi-isometry and let Y be amenable. Because f is a quasi-isometry and X and Y are UDBG spaces, there are constants $c, C \in \mathbb{R}_{>0}$ with the following properties:

- The map $f\colon X \longrightarrow Y$ has c-dense image.
- For all finite sets $F \subset Y$ we have

$$\bigl|f^{-1}(B^Y_c(F))\bigr| \geq \frac{1}{C} \cdot |F| \quad \text{and} \quad \bigl|f^{-1}(F)\bigr| \leq C \cdot |F|.$$

Moreover, let $(F_n)_{n\in\mathbb{N}}$ be a Følner sequence for Y. Then $(\widetilde{F}_n)_{n\in\mathbb{N}}$ defined by

$$\widetilde{F}_n := f^{-1}\bigl(B^Y_c(F_n)\bigr)$$

for all $n \in \mathbb{N}$ is a Følner sequence for X: Let $r \in \mathbb{N}$. Then a straightforward calculation shows that

$$f\bigl(\partial^X_r(\widetilde{F}_n)\bigr) \subset \partial^Y_{c\cdot(r+2)}(F_n)$$

holds for all $n \in \mathbb{N}$. Hence, we obtain

$$\begin{aligned} \bigl|\partial^X_r(\widetilde{F}_n)\bigr| \leq \bigl|f^{-1}(f(\partial^X_r(\widetilde{F}_n)))\bigr| &\leq \bigl|f^{-1}(\partial^Y_{c\cdot(r+2)}(F_n))\bigr| \\ &\leq C \cdot \bigl|\partial^Y_{c\cdot(r+2)}(F_n)\bigr| \end{aligned}$$

and

$$|\widetilde{F}_n| = \bigl|f^{-1}(B^Y_c(F_n))\bigr| \geq \frac{1}{C} \cdot |F_n|.$$

Combining these estimates, yields

$$\frac{\bigl|\partial^X_r(\widetilde{F}_n)\bigr|}{|\widetilde{F}_n|} \leq C^2 \cdot \frac{\bigl|\partial^Y_{c\cdot(r+2)}(F_n)\bigr|}{|F_n|}.$$

Because $(F_n)_{n\in\mathbb{N}}$ is a Følner sequence for Y, the right hand side converges to 0 for $n \to \infty$. Thus, $(\widetilde{F}_n)_{n\in\mathbb{N}}$ is a Følner sequence for X. $\Box$

In particular, amenability is a geometric property of finitely generated groups:

Corollary 9.3.2. *Let G and H be finitely generated groups with $G \sim_{\mathrm{QI}} H$. Then G is amenable if and only if H is amenable.*

Proof. This follows directly from the characterisation of amenable groups via Følner sequences (Theorem 9.2.5) and Theorem 9.3.1. □

Alternatively, quasi-isometry invariance of amenability of finitely generated groups or UDBG spaces can also be derived from the quasi-isometry invariance of uniformly finite homology and the characterisation of amenability via uniformly finite homology (Theorem 9.2.17).

9.4 Quasi-isometry and bilipschitz equivalence

We will now investigate the difference between quasi-isometry and bilipschitz equivalence for finitely generated groups. It can be shown that there exist finitely generated groups that are quasi-isometry invariant but not bilipschitz equivalent [29, 30]. However, for non-amenable groups we have a rigidity phenomenon: Every quasi-isometry between finitely generated non-amenable groups is at finite distance of a bilipschitz equivalence.

In Section 9.4.1, we give an elementary proof of this rigidity result; in Section **??**, we sketch a proof using uniformly finite homology.

9.4.1 Bilipschitz equivalence rigidity

We will now prove bilipschitz equivalence rigidity for non-amenable UDBG spaces and finitely generated groups by elementary means:

Theorem 9.4.1 (Bilipschitz equivalence rigidity for UDBG spaces)**.** *Let X and Y be UDBG spaces, let Y be non-amenable, and let $f\colon X \longrightarrow Y$ be a quasi-isometry. Then f is at finite distance of a bilipschitz equivalence $X \longrightarrow Y$.*

The idea of the proof is as follows:

- As bijective quasi-isometries between UDBG spaces are bilipschitz equivalences, it suffices to deform f into a bijective quasi-isometry.
- Because X (and hence also Y) are non-amenable, Hall's marriage theorem guarantees that there is enough space to deform f and a quasi-isometry inverse $Y \longrightarrow X$ into injective quasi-isometries.
- These injective maps can be glued via the Schröder-Bernstein theorem to form a bijective quasi-isometry.

We now explain the ingredients and the steps of this proof in more detail:

Proposition 9.4.2 (Bijective quasi-isometries vs. bilipschitz equivalences)**.** *Let X and Y be UDBG spaces, and let $f\colon X \longrightarrow Y$ be a bijective quasi-isometry. Then f is a bilipschitz equivalence.*

Proof. This is a straightforward calculation, analogously to Exercise 5.2.1. □

Theorem 9.4.3 (Hall's marriage theorem)**.** *Let W and M be non-empty sets and let $F\colon W \longrightarrow P^{\text{fin}}(M)$ be a map satisfying the* marriage condition

$$\forall_{V \in P^{\text{fin}}(W)} \quad \Bigl| \bigcup_{w \in V} F(w) \Bigr| \geq |V|;$$

here, $P^{\text{fin}}(W)$ denotes the set of all finite subsets of W. Thene there exists a (W, M, F)-marriage, *i.e., an injective map $\mu\colon W \longrightarrow M$ with*

$$\forall_{v \in W} \quad \mu(w) \in F(w).$$

The name *marriage theorem* is derived from the interpretation where W represents a set of women, M represents a set of men, and F models which men appear attractive to which women. By the theorem, there always exists a marriage $W \longrightarrow M$ that makes all women happy, provided that the obvious necessary condition is satisfied.

The marriage theorem for finite sets admits a beautiful proof by induction [**?**]; the general marriage theorem can then be obtained by applying Zorn's lemma to suitably extendable partial marriages [**?**].

The last ingredient is the Schröder-Bernstein theorem [**?**] from set theory:

Theorem 9.4.4 (Schröder-Bernstein)**.** *Let X and Y be sets admitting injections $f\colon X \longrightarrow Y$ and $g\colon Y \longrightarrow X$. Then there exists a bijection $X \longrightarrow$*

Y. More precisely, there is a disjoint decomposition $X = X' \sqcup X''$ of X such that

$$f|_{X'} \sqcup g^{-1}|_{X''} \colon X \longrightarrow Y$$

is a well-defined bijection between X and Y.

Using these tools, we can now prove the rigidity theorem Theorem 9.4.1:

Proof of Theorem 9.4.1. Let $f \colon X \longrightarrow Y$ be a quasi-isometry and let $g \colon Y \longrightarrow X$ be a quasi-isometry quasi-inverse to f.

We construct an injective map $X \longrightarrow Y$ at finite distance of f: Because X is a UDBG space and $f \colon X \longrightarrow Y$ is a quasi-isometry, there is an $E \in \mathbb{R}_{\geq 1}$ with

$$\forall_{V \in P^{\text{fin}}(X)} \ \ |V| \leq E \cdot \bigl|f(V)\bigr|.$$

Moreover, because Y is non-amenable, there is an $r \in \mathbb{N}$ such that

$$\forall_{V \in P^{\text{fin}}(X)} \ \ |V| \leq E \cdot \bigl|f(V)\bigr| \leq \bigl|B_r^X(f(V))\bigr|$$

(Exercise). Therefore, the map

$$\begin{aligned} X &\longrightarrow P^{\text{fin}}(Y) \\ x &\longmapsto B_r^Y\bigl(f(x)\bigr) \end{aligned}$$

satisfies the marriage condition. So, by the marriage theorem (Theorem 9.4.3) there exists an injection $f' \colon X \longrightarrow Y$ with

$$\forall_{x \in X} \ \ f'(x) \in B_r^Y\bigl(f(x)\bigr).$$

In particular, f' has finite distance from f.

Because Y is non-amenable also X is non-amenable (Theorem 9.3.1. Hence, we can construct with the same arguments as above an injective map $g' \colon Y \longrightarrow X$ at finite distance of g.

Because f' and g' have finite distance from quasi-isometric embeddings, they are also quasi-isometric embeddings.

Applying the Schröder-Bernstein theorem (Theorem 9.4.4) to the injections $f' \colon X \longrightarrow Y$ and $g' \colon Y \longrightarrow X$ we obtain a disjoint decomposition $X = X' \sqcup X''$ of X such that

$$\overline{f} := f'|_{X'} \sqcup g'^{-1}|_{X''} \colon X \longrightarrow Y$$

is a well-defined bijection. By construction, $\overline{f}$ is at finite distance from f. Hence, $\overline{f}$ is a quasi-isometric embedding. Proposition 9.4.2 then implies that $\overline{f}\colon X \longrightarrow Y$ is a bilipschitz equivalence. $\square$

In particular, we obtain:

Corollary 9.4.5 (Bilipschitz equivalence rigidity for groups)**.** *Let G and H be finitely generated groups, let H be non-amenable, and let $f\colon G \longrightarrow H$ be a quasi-isometry. Then f is at finite distance of a bilipschitz equivalence $G \longrightarrow H$.*

Remark 9.4.6. The groups F_2 and F_3 are quasi-isometric; because these groups are non-amenable, the rigidity theorem tells us that F_2 and F_3 are bilipschitz equivalent. Hence, the rank of free groups is *not* invariant under bilipschitz equivalence.

As an application, we discuss quasi-isometry stability of taking free products:

Corollary 9.4.7 (Quasi-isometry stability of taking free products)**.** *Let G, G', H be finitely generated groups with $G \sim_{\mathrm{QI}} G'$ and let G be non-amenable. Then*
$$G * H \sim_{\mathrm{QI}} G' * H.$$

Proof. Because G is non-amenable and $G \sim_{\mathrm{QI}} G'$, bilipschitz equivalence rigidity (Theorem 9.4.1) implies that there is even a bilipschitz equivalence $f\colon G \longrightarrow G'$. A straightforward calculation shows that the induced map
$$f * \mathrm{id}_H\colon G * H \longrightarrow G' * H$$
is also a bilipschitz equivalence (notice that the corresponding statement for quasi-isometries does *not* hold in general, because one loses control over the additive error term when taking free products). In particular, $G * H \sim_{\mathrm{QI}} G' * H$. $\square$

Example 9.4.8. The groups $(F_3 \times F_3) * F_3$ and $(F_3 \times F_3) * F_4$ are bilipschitz equivalent (and hence quasi-isometric) because F_3 is non-amenable and $F_3 \sim_{\mathrm{QI}} F_4$. This example can be used to separate commensurability and quasi-isometry (Caveat 5.3.14).

Caveat 9.4.9. If G, G', H are finitely generated groups with $G \sim_{\mathrm{QI}} G'$, then in general $G * H \sim_{\mathrm{QI}} G' * H$ does *not* hold. For example, $1 \sim_{\mathrm{QI}} \mathbb{Z}/2$, but $1 * \mathbb{Z}/2 \cong \mathbb{Z}/2$ is finite and hence *not* quasi-isometric to the infinite group $\mathbb{Z}/2 * \mathbb{Z}/2$. Another example of this type is that $\mathbb{Z}/3 * \mathbb{Z}/2$ is *not* quasi-isometric to $\mathbb{Z}/2 * \mathbb{Z}/2$, but $\mathbb{Z}/3 \sim_{\mathrm{QI}} \mathbb{Z}/2$. However, these two examples are the only cases where quasi-isometry is not inherited through free products [78].

A Appendix

A.1 The fundamental group, a primer

The fundamental group is a concept from algebraic topology: The fundamental group provides a translation from topological spaces into groups that ignores homotopies. In more technical terms, the fundamental group is a functor

$$\begin{aligned} \pi_1 \colon \text{path-connected pointed topological spaces} &\longrightarrow \text{groups} \\ \text{basepoint preserving continuous maps} &\longrightarrow \text{group homomorphisms} \end{aligned}$$

that is homotopy invariant. I.e., if $f, g \colon (X, x_0) \longrightarrow (Y, y_0)$ are basepoint preserving continuous maps that are homotopic (via a basepoint preserving homotopy), then the induced group homomorphisms

$$\pi_1(f),\ \pi_1(g) \colon \pi_1(X, x_0) \longrightarrow \pi_1(Y, y_0)$$

coincide.

Homotopy invariance and functoriality of π_1 show that if (X, x_0) and (Y, y_0) are homotopy equivalent pointed topological spaces, then $\pi_1(X, x_0)$ and $\pi_1(Y, y_0)$ are isomorphic groups. In particular, π_1 can detect that certain spaces are *not* homotopy equivalent: If $\pi_1(X, x_0) \not\cong \pi_1(Y, y_0)$, then (X, x_0) and (Y, y_0) cannot be homotopy equivalent (and so also cannot be homeomorphic).

Caveat A.1.1. In general, if spaces have isomorphic fundamental groups, then they need not be homotopy equivalent.

Geometrically, the fundamental group measures how many "holes" a space has that can be detected by loops in the space in question: The *fundamental group* $\pi_1(X, x_0)$ of a pointed space (X, x_0) is constructed as the set of all loops in X based at x_0 modulo basepoint preserving homotopies; the composition of two classes of loops is given by taking the class represented by concatenation of these loops. If $f \colon (X, x_0) \longrightarrow (Y, y_0)$ is a

Topological space/map	*Group/homomorphism*
$\bullet$	trivial group
$S^1 =$ 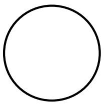	$\mathbb{Z}$
$S^1 \vee S^1 =$ 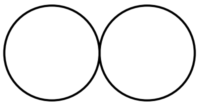	$\mathbb{Z} * \mathbb{Z}$
$S^1 \longrightarrow S^1,\ z \mapsto z^2$	$\mathbb{Z} \longrightarrow \mathbb{Z}, z \mapsto 2 \cdot z$
torus $= S^1 \times S^1 =$ 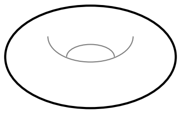	$\mathbb{Z} \times \mathbb{Z}$
sphere $= S^2 =$ 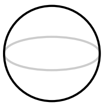	trivial group
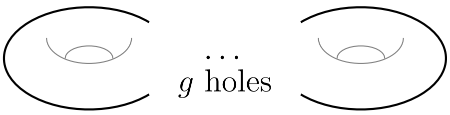	$\Bigl\langle a_1, \dots, a_g, b_1, \dots, b_g \Bigm\vert \prod_{j=1}^{g} [a_j, b_j] \Bigr\rangle$
$\mathbb{R}\mathrm{P}^2$	$\mathbb{Z}/2\mathbb{Z}$
products $(X, x_0) \times (Y, y_0)$	$\pi_1(X, x_0) \times \pi_1(Y, y_0)$
glueings $(X, a_0) \cup_{(A,a_0)} (Y, a_0)$ (where $X \cap Y = A$ is path-connected)	$\pi_1(X, a_0) *_{\pi_1(A,a_0)} \pi_1(Y, y_0)$
mapping tori	HNN-extensions
fibrations	extensions of groups

Figure A.1.: Basic examples of fundamental groups continuous map, then the induced homomorphism $\pi_1(f)$ is defined by composing f with loops in X. Some basic examples are given in Figure A.1.

In geometric group theory, the most important aspect of the fundamental group is that it is a source of convenient group actions and its connection with covering theory: For any "nice"[1] path-connected pointed topological space (X, x_0) there is a path-connected pointed space $(\widetilde{X}, \widetilde{x}_0)$, the so-called *universal covering space of* (X, x_0) with the following properties:

- The space $(\widetilde{X}, \widetilde{x}_0)$ has trivial fundamental group.
- The space $\widetilde{X}$ admits a free, properly discontinuous group action of $\pi_1(X, x_0)$ such that there is a homeomorphism from the quotient $\pi_1(X, x_0) \setminus \widetilde{X}$ to X that sends the class of $\widetilde{x}_0$ to x_0.

If X is a metric space, then $\widetilde{X}$ carries an induced metric (namely, the induced path-metric) and the action of the fundamental group is isometric with respect to this metric.

For example, the universal covering of the circle S^1 can be identified with the real line $\mathbb{R}$ together with the translation action of $\pi_1(S^1, 1) \cong \mathbb{Z}$ (Example 4.1.7 and Figure 4.1).

Moreover, taking quotients of the universal covering space by subgroups of the fundamental group yields intermediate coverings associated with these subgroups. These intermediate coverings together with the lifting properties of covering maps turn out to be a useful tool to study groups and spaces; one simple example of this type is given in the sketch proof of Corollary 6.2.15.

Conversely, for every group there is a path-connected topological space that has the given group as fundamental group; moreover, this can be arranged in such a way that the universal covering of this space is contractible. Spaces with this property are called *classifying spaces*. Classifying spaces provide a means to model group theory in terms of topological spaces, and are a central concept in the study of group (co)homology.

A precise definition of the fundamental group is given in most textbooks on algebraic topology; for example, an excellent treatment of the fundamental group and covering theory is given in Massey's book [64].

More generally, the paradigm of algebraic topology is to find good homotopy invariants of topological spaces; the fundamental group is just one example of this type. Further examples include, for instance, singular homology and cohomology, ...

[1] i.e., path-connected, locally path-connected, and semi-locally simply connected

A.2 The (von) Neumann forest

The name “Neumann” is ubiquitous in geometric group theory. On the one hand, there is the Neumann family (Hanna Neumann, Bernhard Neumann, Peter Neumann, Walter Neumann were/are all involved in geometric group theory and related fields); on the other hand, there is also John von Neumann, who – among many other disciplines – shaped geometric group theory:

(Jo)Hanna Neumann (née von Caemmerer) 1914–1971	∞	Bernhard Neumann 1909–2002	John von Neumann 1903–1957

Irene	Peter	Barbara	Walter	Daniel
1939–	1940–	1943–	1946–	1951–

The contributions to geometric group theory of the (von) Neumanns are too numerous to be listed here [5, 4, 76]; for the topics in these lecture notes, the most important ones are:

- Bernhard Neumann and Hanna Neumann developed and applied together with Higman the theory of a class of groups that is now accordingly named HNN extensions (Definition 2.3.10).
- There is/was the Hanna Neumann conjecture on ranks of certain subgroups of free groups (Remark 4.3.5).
- There is a joint article by Bernhard, Hanna, and Peter Neumann [70].
- There is/was the von Neumann conjecture on the relation between non-amenability and free subgroups (Remark 9.1.11).

A.3 Geometric realisation of graphs

Sometimes it is more convenient to be able to argue via geodesics than via quasi-geodesics. Therefore, we discuss in the following how we can associate a geodesic space with a connected graph and how quasi-geodesic spaces can be replaced by geodesic spaces:

Roughly speaking the geometric realisation of a graph is obtained by gluing a unit interval between every two vertices that are connected by an edge in the given graph. This construction can be turned into a metric space by combining the standard metric on the unit interval with the combinatorially defined metric on the vertices given by the graph structure.

A small technical point is that the unit interval is directed while our graphs are not. One alternative would be to choose an orientation of the given graph (and then prove that the realisation does not depend on the chosen orientation); we resolve this issue by replacing every undirected edge by the corresponding two directed edges and then identifying the corresponding intervals accordingly (Figure A.2):

Definition A.3.1 (Geometric realisation of graphs)**.** Let $X = (V, E)$ be a connected graph. The *geometric realisation of* X is the metric space

$$\bigl(|X|, d_{|X|}\bigr)$$

defined as follows:

If $E = \emptyset$, then X being connected implies that $|V| \leq 1$; in this case, we define $|X| := V$, and set $d_{|X|} := 0$.

If $E \neq \emptyset$, every vertex of X lies on at least one edge, and we define

$$|X| := \widetilde{E} \times [0,1] \bigm/ \!\sim .$$

Here,

$$\widetilde{E} := \bigl\{ (u,v) \bigm| u,v \in V,\ \{u,v\} \in E \bigr\}$$

$$
\begin{array}{ccccc}
\big(\{x,y,z\},\big\{\{x,y\},\{y,z\}\big\}\big) & \rightsquigarrow & \begin{array}{c} (u,v)\times[0,1] \qquad (v,w)\times[0,1] \\ u \rightleftarrows v \rightleftarrows w \\ (v,u)\times[0,1] \qquad (w,v)\times[0,1] \end{array} & \rightsquigarrow & u \quad v \quad w \\
(V,E) & \rightsquigarrow & \widetilde{E}\times[0,1] & \rightsquigarrow & \big|(V,E)\big|
\end{array}
$$

Figure A.2.: Geometric realisation of graphs

is the set of all directed edges (notice that for every unoriented edge $\{u,v\} = \{v,u\}$ we obtain two directed edges (u,v) and (v,u)), and the equivalence relation "$\sim$" is given as follows: For all $((u,v),t)$, $((u',v'),t') \in \widetilde{E}$ we have $((u,v),t) \sim ((u',v'),t')$ if and only if (see Figure A.3)

- the elements coincide, i.e., $((u,v),t) = ((u',v'),t')$, or
- the elements describe the same vertex lying on both edges, i.e.,
  - $u = u'$ and $t = 0 = t'$, or
  - $u = v'$ and $t = 0$ and $t' = 1$, or
  - $v = v'$ and $t = 1 = t'$, or
  - $v = u'$ and $t = 1$ and $t' = 0$,

  or
- the elements describe the same point on an edge but using different orientations, i.e., $(u,v) = (v',u')$ and $t = 1 - t'$.

The metric $d_{|X|}$ on $|X|$ is given by

$$
\begin{aligned}
&d_{|X|}\big(\big[((u,v),t)\big],\big[((u'v'),t')\big]\big)\\
&:= \begin{cases}
|t-t'| & \text{if } (u,v) = (u',v')\\
|t-(1-t')| & \text{if } (u,v) = (v',u')\\
\min\big(t + d_X(u,u') + t' & \\
\quad\ , t + d_X(u,v') + 1 - t' & \\
\quad\ , 1 - t + d_X(v,u') + t' & \\
\quad\ , 1 - t + d_X(v,v') + 1 - t'\big) & \text{if } \{u,v\} \neq \{u',v'\}
\end{cases}
\end{aligned}
$$

for all $[((u,v),t)]$, $[((u',v'),t')] \in |X|$, where d_X denotes the metric on V induced from the graph structure (see Definition 5.2.1).

Figure A.3.: Parametrisations describing the same points in the geometric realisation

Clearly, this construction can be extended to a functor from the category of graphs to the category of spaces. Hence, any action of a group on a graph induces a corresponding piecewise linear continuous action of the group on the geometric realisation. It is not difficult to see that the induced action is free [has a global fixed point] if and only if the original action on the graph is free [has a global fixed point] in the sense of Definition 4.1.8 and Proposition 4.1.16.

Example A.3.2.

- The geometric realisation of the graph $(\{0,1\},\{\{0,1\}\})$ consisting of two vertices and an edge joining them is isometric to the unit interval (Figure A.4).
- The geometric realisation of $\mathrm{Cay}(\mathbb{Z},\{1\})$ is isometric to the real line $\mathbb{R}$ with the standard metric.
- The geometric realisation of $\mathrm{Cay}(\mathbb{Z}^2,\{(1,0),(0,1)\})$ is isometric to the square lattice $\mathbb{R}\times\mathbb{Z}\cup\mathbb{Z}\times\mathbb{R}\subset\mathbb{R}^2$ with the metric induced from the ℓ^1-metric on $\mathbb{R}^2$.

Proposition A.3.3 (Geometric realisation of graphs)**.** *Let $X=(V,E)$ be a connected graph.*

1. *Then the geometric realisation $(|X|,d_{|X|})$ is a geodesic metric space.*
2. *There exists a canonical inclusion $V\hookrightarrow|X|$ and this map is an isometric embedding and a quasi-isometry.*

Proof. Exercise. $\square$

More generally, any quasi-geodesic space can be approximated by a geodesic space:

Proposition A.3.4 (Approximation of quasi-geodesic spaces by geoedesic spaces)**.** *Let X be a quasi-geodesic metric space. Then there exists a geodesic metric space that is quasi-isometric to X.*

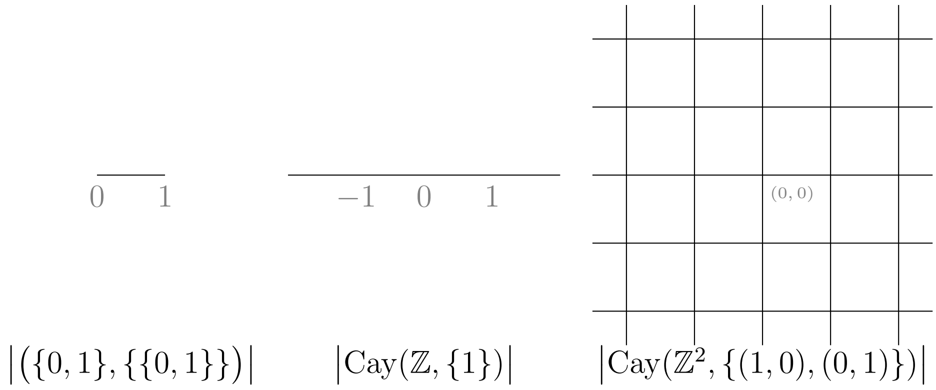

Figure A.4.: Geometric realisations (Example A.3.2)

Sketch of proof. Out of X we can define a graph Y as follows:

- The vertices of Y are the points of X,
- and two points of X are joined by an edge in Y if they are "close enough" together.

Then mapping points in X to the corresponding vertices of Y is a quasi-isometry; on the other hand, $Y \sim_{\mathrm{QI}} |Y|$ (Proposition A.3.3). Hence, X and $|Y|$ are quasi-isometric. Moreover, $|Y|$ is a geodesic metric space by Proposition A.3.3. We leave the details as an exercise. $\square$

Notice however, that it is not always desirable to replace the original space by a geodesic space, because some control is lost during this replacement. Especially, in the context of graphs, one has to weigh up whether the rigidity of the combinatorial model or the flexibility of the geometric realisation is more useful for the situation at hand.

Bibliography

Please note that the numbering of the references will change while the lecture notes grow and develop.

[1] Martin Aigner. *Diskrete Mathematatik*, fifth edition, Vieweg Verlag, 2004. Cited on page: 63

[2] Margaret A. Armstrong. *Groups and Symmetry*, Undergraduate Texts in Mathematics, Springer, 1988. Cited on page: 53, 63

[3] Tim Austin. Rational group ring elements with kernels having irrational dimension, preprint, available online at arXiv:0909.2360 [math.GR], 2009. Cited on page: 30

[4] Australian Academy of Science. *Bernhard Hermann Neumann 1909–2002*, http://sciencearchive.org.au/fellows/memoirs/neumann-b.html Cited on page: A.5

[5] Australian Academy of Science. *Professor Bernhard Neumann (1909–2002), Mathematician*, http://www.sciencearchive.org.au/scientists/interviews/n/bn.html Cited on page: A.5

[6] Henk P. Barendregt. *The lambda calculus, its syntax and semantics*, with addenda for the 6th imprinting, volume 40 of *Studies in*

Logic, Mathematical Logic and Foundations, College Publications, 2012. Cited on page: 176

[7] Arthur Bartels, Wolfgang Lück, Shmuel Weinberger. On hyperbolic groups with spheres as boundary, to appear in *Journal of Differential Geometry*, available online at arXiv:0911.3725, 2009. Cited on page: 214

[8] Hyman Bass. The degree of polynomial growth of finitely generated nilpotent groups, *Proc. London Math. Soc.*, 25, pp. 603–614, 1972. Cited on page: 142

[9] Gilbert Baumslag, Donald Solitar. Some two-generator one-relator non-Hopfian groups, *Bull. Amer. Math. Soc.*, 68, pp. 199-201, 1962. Cited on page: 26

[10] Gilbert Baumslag. Wreath products and finitely presented groups, *Math. Z.*, 75, pp. 22–28, 1961. Cited on page: 28

[11] Riccardo Benedetti, Carlo Petronio. *Lectures on Hyperbolic Geometry*, *Universitext*, Springer, 1992. Cited on page: 106, 193, 215

[12] Jonathan Block, Shmuel Weinberger. Aperiodic tilings, positive scalar curvature and amenability of spaces, *J. Amer. Math. Soc.*, 5(4), pp. 907–918, 1992. Cited on page: 122, 231

[13] George S. Boolos, John P. Burgess, Richard C. Jeffrey. *Computability and logic*, fifth edition, Cambridge University Press, 2007. Cited on page: 176

[14] Douglas S. Bridges. *Computability. A mathematical sketchbook*, volume 146 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 176

[15] Martin R. Bridson. Non-positive curvature and complexity for finitely presented groups, volume 2 of *International Congress of Mathematicians*, pp. 961–987, EMS, 2006. Cited on page: 2, 3, 5

[16] Martin R. Bridson. Geometric and combinatorial group theory, in *The Princeton Companion to Mathematics*, pp. 431–448, Princeton University Press, 2008. Cited on page: 5

[17] Martin R. Bridson, André Haefliger. *Metric Spaces of Non-positive Curvature*, Band 319 der *Grundlehren der Mathematischen Wissenschaften*, Springer, 1999. Cited on page: 5, 46, 147, 158, 201, 207, 209

[18] Kenneth S. Brown. *Cohomology of Groups*, volume 87 of *Graduate Texts in Mathematics*, Springer, 1982. Cited on page: 32, 46

[19] Kenneth S. Brown. *Buildings*, reprint of the 1989 original. *Springer Monographs in Mathematics*, Springer, 1998. Cited on page: 175

[20] James W. Cannon, William J. Floyd, Walter R. Parry. Introductory notes on Richard Thompson's groups, *Enseign. Math. (2)*, 42, pp. 215–256, 1996. Cited on page: 26

[21] Tullio Ceccherini-Silberstein, Michel Coornaert. *Cellular Automata and Groups*, *Springer Monographs in Mathematics*, Springer, 2010. Cited on page: 226, 228

[22] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein. *Introduction to algorithms*, third edition. MIT Press, 2009. Cited on page: 39

[23] Michael W. Davis. *The geometry and topology of Coxeter groups*, volume 32 of *London Mathematical Society Monographs Series*, Princeton University Press, 2008. Cited on page: 175

[24] Reinhard Diestel. *Graph theory*, third edition, volume 173 of *Graduate Texts in Mathematics*, Springer, 2005. Cited on page: 39

[25] Albrecht Dold. *Lectures on Algebraic Topology*, reprint of the 1972 edition, *Classics in Mathematics*, Springer, 1995. Cited on page: 138, 139, 213

[26] Alexander Dranishnikov. On macroscopic dimension or rationally essential manifolds, *Geom. Topol.*, 15(2), pp. 1107–1124, 2011. Cited on page: 122

[27] Lou van den Dries, Alex J. Wilkie. Gromov's theorem on groups of polynomial growth and elementary logic, *J. Algebra*, 89(2), pp. 349–374, 1984. Cited on page: 139, 143, 144, 145

[28] Cornelia Druţu, Mark Sapir. Tree-graded spaces and asymptotic cones of groups. With an appendix by Denis Osin and Mark Sapir. *Topology*, 44(5), pp.959–1058, 2005. Cited on page: 144

[29] Tullia Dymarz. Bijective quasi-isometries of amenable groups, *Geometric methods in group theory*, 181–188, *Contemp. Math.*, 372, Amer. Math. Soc., 2005. Cited on page: 233

[30] Tullia Dymarz. Bilipschitz equivalence is not equivalent to quasi-isometric equivalence for finitely generated groups, *Duke Math. J.*, 154(3), pp. 509–526, 2010. Cited on page: 233

[31] Patrick B. Eberlein. *Geometry of Nonpositively Curved Manifolds*, University of Chicago Press, 1996. Cited on page: 191

[32] Maurits C. Escher. *Cirkellimiet III*, 1960. Available online at http://www.mcescher.nl/galerij/erkenning-succes/cirkellimiet-iii/ Cited on page: 155, 175

[33] Maurits C. Escher. *Cirkellimiet IV*, 1960. Available online at http://www.mcescher.nl/galerij/erkenning-succes/cirkellimiet-iv/ Cited on page: 155, 175

[34] John Franks. Anosov diffeomorphisms. *1970 Global Analysis*, Volume XIV of *Proc. Sympos. Pure Math.*, Berkeley, 1968, pp. 61–93, American Mathematical Society. Cited on page: 149

[35] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture, preprint, available online at arXiv:1105.0129 [math.CO] Cited on page: 72

[36] Alex Furman. A survey of measure group theory, in *Geometry, Rigidity, and Group Actions* (edited by B. Farb and D. Fisher), Chicago University Press, 2011. Cited on page: 115

[37] Rostislav I. Grigorchuk. Degrees of growth of finitely generated groups and the theory of invariant means. *Izv. Akad. Nauk SSSR Ser. Mat.*, 48(5), pp. 939–985, 1984. Cited on page: 134, 135, 222

[38] Detlef Gromoll and Joseph A. Wolf. Some relations between the metric structure and the algebraic structure of the fundamental group in manifolds of nonpositive curvature, *Bull. Amer. Math. Soc.*, 77, pp. 545–552, 1971. Cited on page: 200

[39] Mikhael Gromov. Groups of polynomial growth and expanding maps, with an Appendix by Jacques Tits, *Publ. Math. IHES*, 53, pp. 53–78, 1981. Cited on page: 139, 144, 147

[40] Misha Gromov. Hyperbolic groups. in *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pp. 75–263, Springer, 1987. Cited on page: 5, 157, 171, 177

[41] Misha Gromov. Random walk in random groups, *Geom. Funct. Anal.*, 13(1), pp. 73–146, 2003. Cited on page: 177

[42] Pierre de la Harpe. *Topics in Geometric Group Theory*, Chicago University Press, 2000. Cited on page: 5, 28, 76, 78, 104, 105, 134, 135, 137

[43] John M. Harris, Jeffrey L. Hirst, Michael J. Mossinghoff. *Combinatorics and Graph Theory*, second edition, *Undergraduate Texts in Mathematics, Springer*, 2008. Cited on page: 39

[44] Allen Hatcher. *Algebraic Topology*, Cambridge University Press, 2002. http://www.math.cornell.edu/∼hatcher/AT/ATpage.html Cited on page: 46

[45] Nigel Higson, John Roe. The Baum-Connes conjecture in coarse geometry. In S. Ferry, A. Ranicki, J. Rosenberg, editors, *Proceedings of the 1993 Oberwolfach conference on the Novikov Conjecture*, volume 227 of *LMS Lecture Notes*, 1995. Cited on page: 119

[46] Barry E. Johnson. *Cohomology in Banach Algebras*, volume 127 of *Memoirs of the AMS*, AMS, 1972. Cited on page: 231

[47] John L. Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*, Springer, 1975. Cited on page: 113

[48] Bruce Kleiner. A new proof of Gromov's theorem on groups of polynomial growth, preprint, available online at arXiv:0710.4593, 2007. Cited on page: 139

[49] Johannes Köbler, Uwe Schöning, Jacobo Torán. *Graph Isomorphism Problem: The Structural Complexity*, Birkhäuser, 1993. Cited on page: 41

[50] Dieter Kotschick, Clara Löh. Fundamental classes not representable by products, *J. Lond. Math. Soc.*, 79, no. 3, 545–561, 2009. Cited on page: 201

[51] Dieter Kotschick, Clara Löh. Groups not presentable by products, preprint, available online at arXiv:1003.3853, 2010. Cited on page: 201

[52] Serge Lang. *Algebra*, Revised third edition. Volume 211 of *Graduate Texts in Mathematics*, Springer, 2002. Cited on page: 8, 141

[53] Serge Lang. *Real and Functional Analysis*, Volume 142 of *Graduate Texts in Mathematics*, Springer, 1993. Cited on page: 219

[54] John M. Lee. *Riemannian Manifolds. An Introduction to Curvature.* Volume 176 of *Graduate Texts in Mathematics*, Springer, 1997. Cited on page: 5, 105, 107, 152, 155

[55] Clara Löh. Isomorphisms in ℓ^1-homology, *Münster J. Math.*, 1, pp. 237–266, 2008. Cited on page: 231

[56] Clara Löh. Simplicial volume, *Bull. Man. Atl.*, pp. 7–18, 2011. Cited on page: 139

[57] Clara Löh. *Group Cohomology and Bounded Cohomology. An Introduction for Topologists.* Lecture notes, 2010. Available online at http://www.mathematik.uni-regensburg.de/loeh/teaching/topologie3_ws0910 Cited on page: 32, 46, 139

[58] Clara Löh. Which finitely generated Abelian groups admit isomorphic Cayley graphs? *Geom. Dedicata*, 164(1), pp. 97–111, 2013. Cited on page: 44, 93

[59] Clara Löh, Matthias Mann. Which finitely generated Abelian groups admit equal growth functions? Preprint, available online at arXiv: 1309.3381 [math.GR], 2013. Cited on page: 134

[60] Wolfgang Lück. Aspherical Manifolds, *Bull. Man. Atl.*, pp. 1-17, 2012. Cited on page: 216

[61] Roger C. Lyndon, Paul E. Schupp. *Combinatorial Group Theory*, Classics in Mathematics, reprint of the 1977 edition, Springer, 2001. Cited on page: 37

[62] Saunders MacLane. *Categories for the Working Mathematician*, second edition. Volume 5 of *Graduate Texts in Mathematics*, Springer, 1998. Cited on page: 120

[63] Andrei A. Markov. Insolubility of the problem of homeomorphy, (Russian) *Proc. Internat. Congress Math. 1958*, S. 300–306, Cambridge University Press, 1960.
(Inofficial) English translation (of an untraceable German translation): http://www.cs.dartmouth.edu/ afra/goodies/markov.pdf Cited on page: 27

[64] William S. Massey. *Algebraic Topology: An Introduction*, volume 56 of *Graduate Texts in Mathematics*, Springer, 1977. Cited on page: 23, 35, 56, 57, 60, 70, 71, 101, 106, 138, 139, A.4

[65] Donald W. Miller. On a theorem of Hölder, *The American Mathematical Monthly*, 65(4), pp. 252–254, 1958. Cited on page: 15

[66] Igor Mineyev. Submultiplicativity and the Hanna Neumann conjecture, *Ann. of Math.*, 175(1), pp. 393–414, 2012. Cited on page: 72

[67] Igor Mineyev. Groups, graphs, and the Hanna Neumann conjecture, *J. Topol. Anal.*, 4(1), pp. 1–12, 2012. Cited on page: 72

[68] John Milnor. A note on curvature and the fundamental group, *Journal of Differential Geometry*, 2, pp. 1–7, 1968. Cited on page: 136, 143

[69] Deane Montgomery, Leo Zippin. Small subgroups of finite-dimensional groups, *Ann. of Math.*, 56, pp. 213–241, 1952. Cited on page: 144

[70] Bernhard H. Neumann, Hanna Neumann, Peter Neumann. Wreath products and varieties of groups, *Math. Z.*, 80(1), pp. 44–62, 1962. Cited on page: A.5

[71] John von Neumann. Zur allgemeinen Theorie des Maes, *Fund. Math.*, 13(1), pp. 73–111, 1929. Cited on page: 222

[72] Gennady A. Noskov. Bounded cohomology of discrete groups with coefficients, *Algebra i Analiz*, 2(5), pp. 146–164, 1990. Translation in: *Leningrad Math. J.*, 2(5), pp. 1067–1084, 1991. Cited on page: 231

[73] Yann Ollivier. Sharp phase transition theorems for hyperbolicity of random groups, *Geom. Funct. Anal.*, 14(3), pp. 595–679, 2004. Cited on page: 177

[74] Alexander Yu. Olshanskiĭ. Almost every group is hyperbolic, *Internat. J. Algebra Comput.*, 2(1), pp. 1–17, 1992. Cited on page: 177

[75] Alexander Ju. Olshanskii. On the question of the existence of an invariant mean on a group, *Uspekhi Mat. Nauk*, 35 (4), pp. 199–200, 1980. Cited on page: 223

[76] J.C. Oxtoby, B.J. Pettis, G.B. Price (eds.). John von Neumann 1903–1957, *Bull. Amer. Math. Soc.*, 64(2), 1958. Cited on page: A.5

[77] Piotr W. Nowak, Guoliang Yu. *Large Scale Geometry*, *EMS Textbooks in Mathematics*, EMS, 2012. Cited on page: 119, 122

[78] Panos Papasoglu, Kevin Whyte. Quasi-isometries between groups with infinitely many ends, *Comment. Math. Helv.*, 77(1), pp. 133–144, 2002. Cited on page: 237

[79] Alan L.T. Paterson. *Amenability*. Volume 29 of *Mathematical Surveys and Monographs*, AMS, 1988. Cited on page: 219, 223

[80] John G. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Volume 149 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 139

[81] John Roe. *Index Theory, Coarse Geometry, and Topology of Manifolds*. Volume 90 of *CBMS Regional Conf. Ser. in Math.*, AMS, 1996. Cited on page: 119

[82] Joseph J. Rotman. *An Introduction to the Theory of Groups*, 4th edition. Volume 148 of *Graduate Texts in Mathematics*, Springer, 1999. Cited on page: 8, 27, 37, 176

[83] Volker Runde. *Amenability*, volume 1774 of *Springer Lecture Notes in Mathematics*, Springer, 2002. Cited on page: 223, 230

[84] Parameswaran Sankaran. On homeomorphisms and quasi-isometries of the real line, *Proc. Amer. Math. Soc.*, 134, no. 7, pp. 1875–1880, 2006. Cited on page: 87

[85] Mark Sapir. Another false proof of nonamenability of the R. Thompson group F, http://marksapir.wordpress.com/2014/08/16/another-false-proof-of-nonamenability-of-the-r-thompson-group-f/ Cited on page: 26

[86] Roman Sauer. *L^2-Invariants of groups and discrete measured groupoids*, PhD thesis, WWU Münster, 2003. http://nbn-resolving.de/urn:nbn:de:hbz:6-85659549583 Cited on page: 113, 115

[87] Jean-Pierre Serre. *Trees*, corrected 2nd printing of the 1980 English translation, *Springer Monographs in Mathematics*, Springer, 2003. Cited on page: 5, 34, 38, 50, 70, 74

[88] Yehuda Shalom, Terence Tao. A finitary version of Gromov's polynomial growth theorem, *Geom. Funct. Anal.*, 20(6), pp. 1502–1547, 2010. Cited on page: 139

[89] Michael Shub. Expanding maps. *1970 Global Analysis*, Volume XIV of *Proc. Sympos. Pure Math.*, Berkeley, 1968, pp. 273276, American Mathematical Society. Cited on page: 149

[90] John R. Stallings. On torsion-free groups with infinitely many ends, *Annals of Mathematics*, 88, pp. 312–334, 1968. Cited on page: 37

[91] Jacek Swiatkowski. The dense amalgam of metric compacta and topological characterization of boundaries of free products of groups, preprint, available online at arXiv:1410.4989 [math.GT], 2014. Cited on page: 214

[92] Terence Tao. *Hilbert's fifth problem and related topics*. Volume 153 of *Graduate Studies in Mathematics*, American Mathematical Society, 2014. Cited on page: 144

[93] Jacques Tits. Free subgroups in linear groups, *J. Algebra*, 20, pp. 250–270, 1972. Cited on page: 77

[94] Karen Vogtmann. Automorphisms of free groups and outer space. *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, *Geom. Dedicata*, 94, pp. 1–31, 2002. Cited on page: 15

[95] Shmuel Weinberger. *Computers, Rigidity, and Moduli. The large-scale fractal geometry of Riemannian moduli space.* M. B. Porter Lectures. Princeton University Press, 2005. Cited on page: 27

[96] Kevin Whyte. Amenability, bi-Lipschitz equivalence, and the von Neumann conjecture, *Duke Math. J.*, 99, no. 1, pp. 93–112, 1999. Cited on page: 122

[97] Wikipedia (Authour: Tomruen; Access: December 23rd, 2014). *The order-3 heptagonal tiling of the hyperbolic plane, drawn in the Poincaré disk model*, http://en.wikipedia.org/wiki/File:Uniform_tiling_73-t0.png Cited on page: 174

[98] Joseph A. Wolf. Growth of finitely generated solvable groups and curvature of Riemanniann manifolds, *J. Differential Geometry*, 2, pp. 421–446, 1968. Cited on page: 142, 143