Clara Löh

Geometric Group Theory

An Introduction

this is a draft version!

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

Fakultät für Mathematik
Universität Regensburg
93040 Regensburg
Germany

All drawings by Clara Löh

für A ∗ A

About this book

This book is an introduction into *geometric group theory*. It is certainly not an encyclopedic treatment of geometric group theory, but hopefully it will prepare and encourage the reader to take the next step and learn more advanced aspects of geometric group theory.

The core material of the book should be accessible to third year students, requiring only a basic acquaintance with group theory, metric spaces, and point-set topology. I tried to keep the level of the exposition as elementary as possible, preferring elementary proofs over arguments that require more machinery in topology or geometry. I refrained from adding complete proofs for some of the deeper theorems and instead included sketch proofs, highlighting the key ideas and the view towards applications. However, many of the applications will need a more extensive background in algebraic topology, Riemannian geometry, and algebra.

The exercises are rated in difficulty, from easy* over medium** to hard***. And very hard$^{\infty}$* (usually, open problems of some sort). The core exercises should be accessible to third year students, but some of the exercises aim at applications in other fields and hence require a background in these fields. Moreover, there are exercise sections that develop additional theory in a series of exercises; these exercise sections are marked with $^{+}$.

This book covers slightly more than a one-semester course. Most of the material originates from various courses and seminars I taught at the Universität Regensburg: the geometric group theory courses (2010 and 2014), the seminar on amenable groups (2011), the course on linear groups and heights (2015, together with Walter Gubler), and an elementary course on geometry (2016). Most of the students had a background in real and complex analysis, in linear algebra, algebra, and some basic geometry of manifolds; some of the students also had experience in algebraic topology and Riemannian geometry. I would like to thank the participants of these courses and seminars for their interest in the subject and their patience.

I am particularly grateful to Toni Annala, Matthias Blank, Luigi Caputi, Francesca Diana, Alexander Engel, Daniel Fauser, Stefan Friedl, Walter Gubler, Michał Marcinkowski, Andreas Thom, Johannes Witzig, and the anonymous referees for many valuable suggestions and corrections. This work was supported by the GRK 1692 *Curvature, Cycles, and Cohomology* (Universität Regensburg, funded by the DFG).

Regensburg, September 2017 Clara Löh

Contents

1 Introduction

Groups are an abstract concept from algebra, formalising the study of symmetries of various mathematical objects.

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 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.

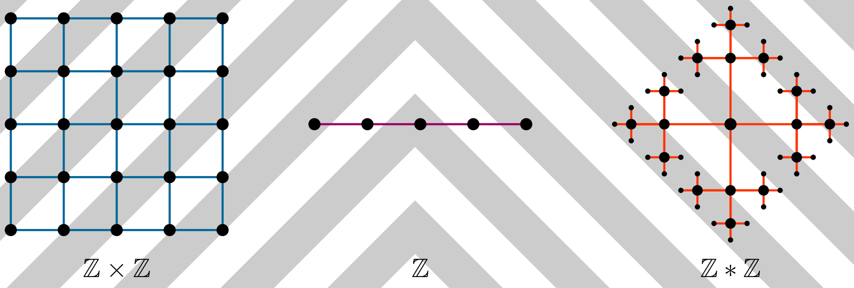

Figure 1.1.: Basic examples of Cayley graphs

A key example of geometric objects associated with a group are Cayley graphs (with respect to a chosen generating set) together with the corresponding word metrics. For instance, from the point of view of large scale geometry, Cayley graphs of $\mathbb{Z}$ resemble the geometry of the real line, Cayley graphs of $\mathbb{Z} \times \mathbb{Z}$ resemble the geometry of the Euclidean plane, while Cayley graphs of the free group $\mathbb{Z} * \mathbb{Z}$ on two generators have 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.” [29]
“A statement that holds for all finitely generated groups
has to be either trivial or wrong.” [attributed to M. Gromov]

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, Riemannian geometry, topology, and number theory.

For example, 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:

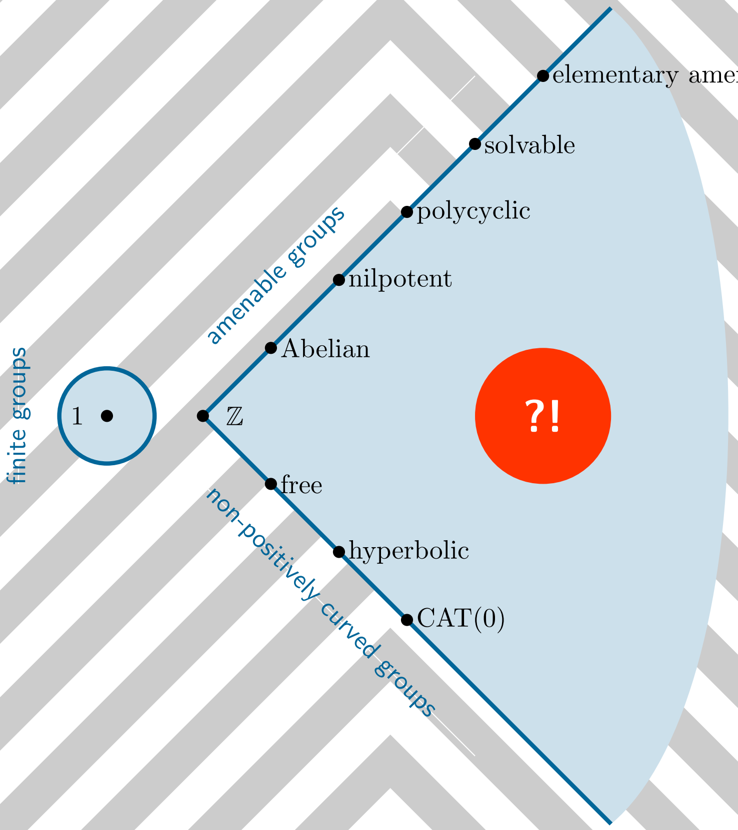

Figure 1.2.: The universe of groups (simplified version of Bridson's universe of groups [29])

- *Recognising that certain matrix groups are free groups;* there is a geometric criterion, the *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 type.
- *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 groups.
- *Group-theoretic reformulation of the Lehmer conjecture;* by the work of Breuillard et al., the Lehmer conjecture in algebraic number theory is equivalent to a problem about growth of certain matrix groups.
- *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 book. The goal of this book is to explain the basic terminology of geometric group theory, the standard proof techniques, and how these concepts can be applied to obtain the results listed above.

As the main characters in geometric group theory are groups, we will start by reviewing concepts and examples from group theory and by introducing constructions that allow to generate interesting groups (Chapter 2). Readers familiar with group theory and the standard examples of groups can happily skip this chapter.

Then we will introduce one of the main combinatorial objects in geometric group theory, Cayley graphs, and review basic notions concerning actions of groups (Chapter 3–4). A first taste of the power of geometric group theory is the geometric characterisation of free groups via actions on trees.

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 (Chapter 5).

After these preparations, we will enjoy the quasi-geometry of groups, including

- growth types (Chapter 6),
- hyperbolicity (Chapter 7),
- geometry at infinity (Chapter 8),
- amenability (Chapter 9).

Basics on fundamental groups, group (co)homology, and elementary properties of the hyperbolic plane, are collected in the appendices (Appendix A). As the proof of the pudding is in the eating, Appendix A.4 contains a list of programming tasks related to geometric group theory.

Literature. Geometric group theory is a vast, rapidly growing area of mathematics; therefore, not all aspects can be covered in this book. The selection of topics is biased by my own preferences, but I hope that this book will prepare and encourage the reader to discover more of geometric group theory. The standard resources for geometric group theory are:

- *Topics in Geometric group theory* by de la Harpe [77] (one of the first collections of results and examples in geometric group theory),
- *Geometric Group Theory* by Druţu and M. Kapovich, with an appendix by Nica [53] (the latest compendium on geometric group theory for advanced students and researchers),
- *Office Hours with a Geometric Group Theorist* edited by Clay and Margalit [41] (a recent collection of topics and examples with a focus on intuition),
- *Metric spaces of non-positive curvature* by Bridson and Haefliger [31] (a compendium on non-positive curvature and its relations with geometric group theory; parts of Chapter 7 and Chapter 8 follow this source),
- *Trees* by Serre [159] (the standard source for Bass-Serre theory).

A short and comprehensible introduction into curvature in classical Riemannian geometry is given in the book *Riemannian manifolds. An introduction to curvature* by Lee [96]. Background material on fundamental groups and covering theory can be found in the book *Algebraic Topology: An Introduction* by Massey [115].

Furthermore, I recommend to look at the overview articles by Bridson on geometric and combinatorial group theory [29, 30]. The original reference for modern large scale geometry of groups is the landmark paper *Hyperbolic groups* [74] by Gromov.

One sentence on notation. The natural numbers $\mathbb{N}$ contain 0.

Part I

Groups

2

Generating groups

As the main characters in geometric group theory are groups, we start by reviewing concepts and examples from group theory. In particular, we will present basic construction principles that allow to generate interesting examples of groups. This includes the description of groups in terms of generators and relations and the iterative construction of groups via semi-direct products, amalgamated free products, and HNN-extensions.

Readers familiar with (infinite) group theory and standard examples of groups can happily skip this chapter.

Overview of this chapter

2.1 Review of the category of groups

2.1.1 Abstract groups: axioms

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 [94, 150]. 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.,
$$\forall_{g \in G} \;\; e \cdot g = g = g \cdot e.$$
(This property uniquely determines the neutral element.)
- *Existence of inverses.* For every $g \in G$ there exists an *inverse element* $g^{-1} \in G$ with respect to "$\cdot$", i.e.,
$$g \cdot g^{-1} = e = g^{-1} \cdot g.$$
(This property uniquely determines the inverse element of 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 $H \times H \subset G \times G$. The *index* $[G : H]$ of a subgroup $H \subset G$ is the cardinality of the set $\{g \cdot H \mid g \in G\}$; here, we use the coset notation $g \cdot H := \{g \cdot h \mid h \in H\}$.

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 $\mathbb{N} = \{0, 1, \dots\}$ 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, H be 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 $g_1, g_2 \in G$. (Every group homomorphism maps the neutral element to the neutral element and inverses to inverses).
- 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, all trivial groups are (canonically) isomorphic. Hence, we usually speak of "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 (Exercise 2.E.3).
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 *symmetric group over* X. If $n \in \mathbb{N}$, then we abbreviate $S_n := S_{\{1,\dots,n\}}$. If $|X| \geq 3$, 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. Then G is isomorphic to a subgroup of S_G: 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 group of a group)**.** 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 (basic notions for metric spaces are recalled in Chapter 5.1). The set $\mathrm{Isom}(X)$ 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 2.2.20, Exercise 2.E.8).

Example 2.1.12 (Matrix groups). Let k be a commutative ring with unit, and let V be a k-module. Then the set $\mathrm{Aut}(V)$ of all k-linear isomorphisms $V \longrightarrow V$ forms a group with respect to composition. In particular, the set $\mathrm{GL}(n, k)$ of invertible $n \times n$-matrices over k is a group (with respect to matrix multiplication) for every $n \in \mathbb{N}$. Similarly, also $\mathrm{SL}(n, k)$, the subgroup of invertible matrices of determinant 1, is a group.

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

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

of field automorphisms of L fixing K is a group with respect to composition, the *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

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

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

In more conceptual language, these examples are all instances of the following general principle: If X is an object in a category C, then the set $\mathrm{Aut}_C(X)$ of C-isomorphisms of type $X \longrightarrow X$ is a group with respect to composition in C. We will now explain this in more detail:

Definition 2.1.15 (Category). A *category* C consists of the following components:

- A class $\mathrm{Ob}(C)$; the elements of $\mathrm{Ob}(C)$ are *objects of* C. (Classes are a generalisation of sets, allowing, e.g., for the definition of the class of all sets [164]).
- A set $\mathrm{Mor}_C(X, Y)$ for each choice of objects $X, Y \in \mathrm{Ob}(C)$; elements of $\mathrm{Mor}_C(X, Y)$ are called *morphisms from X to Y*. (We implicitly assume that morphism sets between different pairs of objects are disjoint.)
- For all objects $X, Y, Z \in \mathrm{Ob}(C)$ a composition

$$\begin{aligned} \circ\colon \mathrm{Mor}_C(Y, Z) \times \mathrm{Mor}_C(X, Y) &\longrightarrow \mathrm{Mor}_C(X, Z) \\ (g, f) &\longmapsto g \circ f \end{aligned}$$

  of morphisms.

These data have to satisfy the following conditions:

- For each object X in C there is a morphism $\mathrm{id}_X \in \mathrm{Mor}_C(X, X)$ with the following property: For all $Y \in \mathrm{Ob}(C)$ and all $f \in \mathrm{Mor}_C(X, Y)$ and $g \in \mathrm{Mor}_C(Y, X)$ we have
$$f \circ \mathrm{id}_X = f \qquad \text{and} \qquad \mathrm{id}_X \circ g = g.$$
(The morphism id_X is uniquely determined by this property; it is the *identity morphism of X in C*).
- Morphism composition is associative, i.e., for all $W, X, Y, Z \in \mathrm{Ob}(C)$ and all $f \in \mathrm{Mor}_C(W, X)$, $g \in \mathrm{Mor}_C(X, Y)$ and $h \in \mathrm{Mor}_C(Y, Z)$ we have
$$h \circ (g \circ f) = (h \circ g) \circ f.$$

Caveat 2.1.16. The concept of morphisms and compositions is modelled on the example of maps between sets and ordinary composition of maps. However, in general, morphisms in categories need not be given as maps between sets and composition need not be composition of maps!

The notion of categories contains all the ingredients necessary to talk about isomorphisms and automorphisms:

Definition 2.1.17 (Isomorphism)**.** Let C be a category. Objects $X, Y \in \mathrm{Ob}(C)$ are *isomorphic in C* if there exist $f \in \mathrm{Mor}_C(X, Y)$ and $g \in \mathrm{Mor}_C(Y, X)$ with
$$g \circ f = \mathrm{id}_X \qquad \text{and} \qquad f \circ g = \mathrm{id}_Y .$$
In this case, f and g are *isomorphisms in C* and we write $X \cong_C Y$ (or $X \cong Y$ if the category is clear from the context).

Definition 2.1.18 (Automorphism group)**.** Let C be a category and let X be an object of C. Then the set $\mathrm{Aut}_C(X)$ of all isomorphisms $X \longrightarrow X$ in C is a group with respect to composition in C (Proposition 2.1.19), the *automorphism group of X in C*.

Proposition 2.1.19 (Automorphism groups in categories)**.**
1. *Let C be a category and let $X \in \mathrm{Ob}(C)$. Then $\mathrm{Aut}_C(X)$ is a group.*
2. *Let G be a group. Then there exists a category C and an object X in C such that $G \cong \mathrm{Aut}_C(X)$.*

Proof. Ad 1. Because the composition of morphisms in C is associative, composition in $\mathrm{Aut}_C(X)$ is associative. The identity morphism id_X is an isomorphism $X \longrightarrow X$ (being its own inverse) and, by definition, id_X is the neutral element with respect to composition. Moreover, the existence of inverses is guaranteed by the definition of isomorphisms in categories.

Ad 2. We consider the category C that contains only a single object X. We set $\mathrm{Mor}_C(X, X) := G$ and we define the composition in C via the composition in G by

$$\circ\colon \operatorname{Mor}_C(X,X) \times \operatorname{Mor}_C(X,X) \longrightarrow \operatorname{Mor}_C(X,X)$$
$$(g,h) \longmapsto g \cdot h.$$

A straightforward computation shows that C indeed is a category and that $\operatorname{Aut}_C(X)$ is G. □

We will now illustrate these terms in more concrete examples:

Example 2.1.20 (Set theory)**.** The category Set of sets consists of:

- Objects: Let Ob(Set) be the class(!) of all sets.
- Morphisms: For sets X and Y, we let $\operatorname{Mor}_{\mathsf{Set}}(X,Y)$ be the set of all set-theoretic maps $X \longrightarrow Y$.
- Compositions are ordinary compositions of maps: For sets X, Y, Z we define the composition $\operatorname{Mor}_{\mathsf{Set}}(Y,Z) \times \operatorname{Mor}_{\mathsf{Set}}(X,Y) \longrightarrow \operatorname{Mor}_{\mathsf{Set}}(X,Z)$ to be ordinary composition of maps.

It is clear that this composition is associative. If X is a set, then the ordinary identity map

$$X \longrightarrow X$$
$$x \longmapsto x$$

is the identity morphism of X in Set. Objects in Set are isomorphic if and only if they have the same cardinality and for all sets X the symmetric group S_X coincides with $\operatorname{Aut}_{\mathsf{Set}}(X)$.

Example 2.1.21 (Algebra)**.** The category Group of groups consists of:

- Objects: Let Ob(Group) be the class of all groups.
- Morphisms: For groups G and H we let $\operatorname{Mor}_{\mathsf{Group}}(G,H)$ be the set of all group homomorphisms.
- Compositions: As compositions we choose ordinary composition of maps.

Analogously, one also obtains the category Ab of Abelian groups, the category $\mathsf{Vect}_{\mathbb{R}}$ of $\mathbb{R}$-vector spaces, the category $_R\mathsf{Mod}$ of left R-modules over a ring R, ... Objects in Group, Ab, $\mathsf{Vect}_{\mathbb{R}}$, $_R\mathsf{Mod}$, ... are isomorphic in the sense of category theory if and only if they are isomorphic in the algebraic sense. Moreover, both the category theoretic and the algebraic point of view result in the same automorphism groups.

Example 2.1.22 (Geometry of isometric embeddings)**.** The category $\mathsf{Met}_{\mathsf{isom}}$ of metric spaces and isometric embeddings consists of:

- Objects: Let $\operatorname{Ob}(\mathsf{Met}_{\mathsf{isom}})$ be the class of all metric spaces.
- Morphisms in $\mathsf{Met}_{\mathsf{isom}}$ are isometric embeddings (i.e., distance preserving maps) of metric spaces.
- The compositions are given by ordinary composition of maps.

Then objects in $\mathsf{Met}_{\mathsf{isom}}$ are isomorphic if and only if they are isometric and automorphism groups in $\mathsf{Met}_{\mathsf{isom}}$ are nothing but isometry groups of metric spaces.

Example 2.1.23 (Topology)**.** The category Top of topological spaces consists of:

- Objects: Let Ob(Top) be the class of all topological spaces.
- Morphisms in Top are continuous maps.
- The compositions are given by ordinary composition of maps.

Isomorphisms in Top are precisely the homeomorphisms; automorphism groups in Top are the groups of self-homeomorphisms of topological spaces.

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, 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 2.1.24 (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 $N \triangleleft G$.

Example 2.1.25 (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 2.1.26)).

Proposition 2.1.26 (Quotient group). *Let G be a group, and let N be a subgroup.*

1. *Let $G/N := \{g \cdot N \mid g \in G\}$. 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 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 every group H and every 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 $\overline{\varphi} \circ \pi = \varphi$:

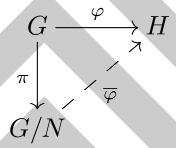

Proof. Ad 1. Suppose that N is normal in G. In this case, the composition map is well-defined (in the sense that the definition does not depend on the choice of the representatives of cosets): Let g_1, g_2, $\overline{g}_1$, $\overline{g}_2 \in G$ with

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

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: 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 2.1.27 (Quotient groups).

- Let $n \in \mathbb{Z}$. Then composition in the quotient group $\mathbb{Z}/n\mathbb{Z}$ is nothing but addition modulo n. If $n \neq 0$, then $\mathbb{Z}/n\mathbb{Z}$ is a *cyclic group of order* n; if $n = 0$, then $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}$ is an *infinite cyclic group*. We will also abbreviate $\mathbb{Z}/n := \mathbb{Z}/n\mathbb{Z}$.
- 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 $\sigma := (1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1)$ is isomorphic to $\mathbb{Z}/2$.

Example 2.1.28 (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 2.E.6), and the quotient

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

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 [173]. Curiously, one has for all $n \in \mathbb{N}$ that [119]

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

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.

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.E.4).

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 containing S is *freely generated by* S if F has the following universal property: For every group G and every map $\varphi\colon S \longrightarrow G$ there is a unique group homomorphism $\overline{\varphi}\colon F \longrightarrow G$ extending φ:

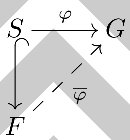

A group is *free* if it contains a free generating set.

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 $\mathbb{Z}/2$ and $\mathbb{Z}^2$ are *not* free (Exercise 2.E.11).

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.

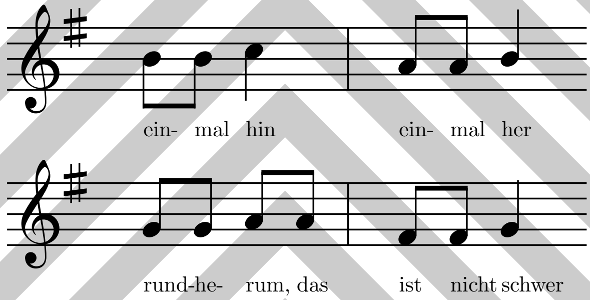

Figure 2.1.: Uniqueness of universal objects; *Brüderchen, komm tanz mit mir* (German folk song: *Little brother dance with me*). Literal translation of the song text displayed above: once to, once from, all around, that is not hard.

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. *Einmal hin:* 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 $\overline{\varphi}' \circ \varphi = \varphi'$.
   Einmal her: Analogously, there is a group homomorphism $\overline{\varphi}\colon F' \longrightarrow F$ satisfying $\overline{\varphi} \circ \varphi' = \varphi$:

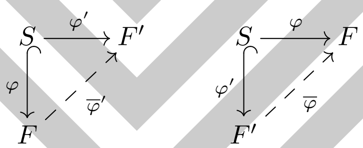

2. *Rundherum, das ist nicht schwer:* We now show that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$ and $\overline{\varphi}' \circ \overline{\varphi} = \mathrm{id}_{F'}$ and hence that φ and φ' are isomorphisms: The composition $\overline{\varphi} \circ \overline{\varphi}'\colon F \longrightarrow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

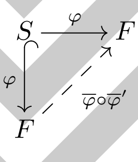

commutative. Moreover, also id_F is a group homomorphism fitting into this diagram. Because F is freely generated by S, the uniqueness part of the universal property thus tells us that these two homomorphisms have to coincide, i.e., that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$. Analogously, one shows that $\overline{\varphi}' \circ \overline{\varphi} = \mathrm{id}_{F'}$.

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. □

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

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

where $\widehat{S} := \{\widehat{s} \mid s \in S\}$ is a disjoint copy of S; i.e., $\widehat{\cdot} \colon S \longrightarrow \widehat{S}$ is a bijection and $S \cap \widehat{S} = \emptyset$. For every element s in S the element $\widehat{s}$ will play the role of the inverse of s in the group that we are about to 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

  $$\forall_{x,y \in A^*} \ \forall_{s \in S} \ \ x s \widehat{s} y \sim xy,$$
  $$\forall_{x,y \in A^*} \ \forall_{s \in S} \ \ x \widehat{s} s y \sim xy;$$

  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 $\cdot \colon F(S) \times F(S) \longrightarrow F(S)$ via

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

  for all x, $y \in A^*$ (because the generating cancellations in each of the factors map to generating cancellations of the concatenation).

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

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

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 $\sim$). Therefore, also

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

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

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

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 $\overline{\varphi} \circ i = \varphi$. Because $i(S)$ generates $F(S)$, there is no other 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

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

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 a 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 Cayley graphs (Chapter 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 every free generating set (Corollary 2.2.8); secondly, 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 (mapping to $\mathbb{Z}/2$) together with a counting argument (Exercise 2.E.12). 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 [115, Chapter VI.8] or via actions on trees (Chapter 4.2.3).

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 (Normal generation)**.**

- 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 set 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 $(S \cup S^{-1})^*$. If we want to emphasise the formality of inverses, we will also sometimes use words in $(S \cup \widehat{S})^*$ instead of $(S \cup S^{-1})^*$.

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

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

is said to be *generated by S with the relations R*.

If G is a group with $G \cong \langle S \,|\, R \rangle$, then the pair (S, R) is a *presentation of G*; by abuse of notation we also use the symbol $\langle S \,|\, R \rangle$ to denote this presentation.

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 $\langle S \,|\, R \rangle$ is a group in which the relations R hold as non-trivially as possible:

Proposition 2.2.18 (Universal property of generators and relations)**.** *Let S be a set and let $R \subset (S \cup S^{-1})^*$. The group $\langle S \,|\, R \rangle$ generated by S with relations R together with the canonical map $\pi \colon S \longrightarrow F(S)/\langle R \rangle^{\triangleleft}_{F(S)} = \langle S \,|\, R \rangle$ has the following universal property: For every group G and every 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 $\overline{\varphi} \colon \langle S \,|\, R \rangle \longrightarrow G$ 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, $\langle S \,|\, R \rangle$ (together with π) is determined uniquely (up to canonical isomorphism) by this universal property.

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 2.1.26) (Exercise 2.E.15). □

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 derived from the universal property (Exercise 2.E.14).

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*:

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

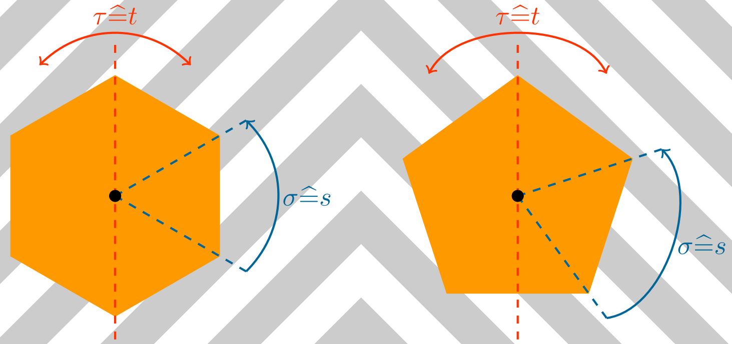

Figure 2.2.: Generators of the dihedral groups D_6 and D_5

This can be seen as follows: Let $\sigma \in \operatorname{Isom}(X_n)$ be rotation about $2\pi/n$ around the centre of the regular n-gon X_n and let τ be reflection along one of the diameters passing through one of the vertices (Figure 2.2). Then a straightforward calculation shows that

$$\sigma^n = \mathrm{id}_{X_n}, \qquad \tau^2 = \mathrm{id}_{X_n}, \qquad \tau \circ \sigma \circ \tau^{-1} = \sigma^{-1}.$$

Thus, the universal property of generators and relations (Proposition 2.2.18) provides us with a well-defined group homomorphism

$$\overline{\varphi}\colon D_n \longrightarrow \operatorname{Isom}(X_n)$$

with $\overline{\varphi}(\overline{s}) = \sigma$ and $\overline{\varphi}(\overline{t}) = \tau$, where $\overline{s}, \overline{t} \in D_n$ denote the elements of D_n represented by s and t, respectively.

In order to see that $\overline{\varphi}$ is an isomorphism, we construct the inverse homomorphism; however, for this direction, the universal property of generators and relations is *not* applicable – therefore, we have to construct the inverse by other means: Using the fact that isometries of X_n map (neighbouring) vertices to (neighbouring) vertices, we deduce that $\operatorname{Isom}(X_n)$ contains exactly $2 \cdot n$ elements, namely,

$$\mathrm{id}_{X_n}, \sigma, \dots, \sigma^{n-1}, \tau, \tau \circ \sigma, \dots, \tau \circ \sigma^{n-1}.$$

An elementary calculation then shows that

$$\begin{aligned} \psi\colon \operatorname{Isom}(X_n) &\longrightarrow D_n \\ \sigma^k &\longmapsto \overline{s}^k \\ \tau \circ \sigma^k &\longmapsto \overline{t} \cdot \overline{s}^k \end{aligned}$$

is a well-defined group homomorphism that is the inverse of $\overline{\varphi}$.

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

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

(Exercise 2.E.20). 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 [38].

Thompson's group F has many peculiar properties. For example, the commutator subgroup $[F, F]$, i.e., the subgroup of F that is generated by the set $\{[g, h] \mid g, h \in F\}$ is an example of an infinite simple group [38]. Moreover, it can be shown that F does not contain subgroups isomorphic to F_2 [38]. 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 [155]: "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 example, $\mathrm{BS}(1, 1) \cong \mathbb{Z}^2$ (Exercise 2.E.21). The family of Baumslag-Solitar groups contains many intriguing examples of groups. For instance, the group $\mathrm{BS}(2, 3)$ 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 [16], 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.

Further examples of prominent classes of group presentations are discussed in the exercises (Exercise 2.E.19ff).

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

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

is trivial: 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 or not 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 [150, 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 or not, is undecidable. In contrast, we will see in Chapter 7.4 that for certain geometric classes of groups the word problem is solvable.

The undecidability of the triviality problem and 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 [114], and there are far-reaching consequences for the global shape of moduli spaces [174].

2.2.4 Finitely presented groups

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 set S and a finite set $R \subset (S \cup S^{-1})^*$ 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.

[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.

Example 2.2.26 (Geometric finite presentation). If X is a path-connected CW-complex with finite 2-skeleton, then the fundamental group $\pi_1(X)$ of X is finitely presented (Exercise 2.E.25). For example, this implies that compact connected manifolds have finitely presented fundamental group. Conversely, every finitely presented group is the fundamental group of a finite CW-complex (Outlook 3.2.5) and of a closed manifold of dimension at least 4.

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

Example 2.2.27 (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 [15] (Exercise 2.E.27). This group is an example of a *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 often 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 2.2.29):

Theorem 2.2.28 (Uncountably many finitely generated groups). *There exist uncountably many isomorphism classes of groups generated by two elements.*

Before sketching Hall's proof [76, Theorem 7][77, Chapter III.C] of this theorem, we discuss an important consequence:

Corollary 2.2.29. *There are uncountably many isomorphism classes of 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 2.2.28. □

The proof of Theorem 2.2.28 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 2.2.30).
2. We then show that G even has uncountably many quotient groups that are pairwise non-isomorphic (Proposition 2.2.31).

Proposition 2.2.30 (Uncountably many normal subgroups). *There exists a group generated by two elements with uncountably many normal subgroups.*

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 under conjugation by all other group elements) isomorphic to the big additive group $\bigoplus_{\mathbb{N}} \mathbb{Z}$. 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 $\mathbb{N}$), and all these subgroups of C are normal in G because C is central in G.

An example of such a group is $G := \langle s, t \mid R \rangle$, where

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

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, using the so-called calculus of commutators, it can be shown that C contains the additive group $\bigoplus_{\mathbb{N}} \mathbb{Z}$ [76, p. 434f][110, Corollary 5.12]. Alternatively, one can give an explicit construction of such a group (Exercise 2.E.35). □

Proposition 2.2.31 (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 $G \longrightarrow Q$ (because every such homomorphism is uniquely determined by its values on a finite generating set of G); 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. □

Outlook 2.2.32 (Non-constructive existence proofs). 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) [8].

2.3 New groups out of old

In many categories, there are ways to construct objects out of given components; examples of such constructions are products and sums/pushouts (or, more generally, limits and colimits). In the world of groups, these correspond to direct products and (amalgamated) free products. There are two views on such constructions: through universal properties and through concrete construction recipes.

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., (amalgamated) free products and HNN-extensions.

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 $\prod_{i \in I} G_i$ 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 with families of homomorphisms to the 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 an injective group homomorphism, π is a surjective group homomorphism, 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 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 position 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, $i\colon N \longrightarrow N \rtimes_\varphi Q$ 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

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

does *not* split because there is no non-trivial homomorphism from the torsion group $\mathbb{Z}/2$ to $\mathbb{Z}$. One way to classify group extensions (with Abelian kernel) is to consider group cohomology [34, Chapter IV][101, Chapter 1.4.4].

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

- If N and Q are groups and $\varphi\colon Q \longrightarrow \mathrm{Aut}(N)$ 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

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

where $\varphi\colon \mathbb{Z}/2 \longrightarrow \operatorname{Aut}\mathbb{Z}/n$ 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 $\mathbb{Z}/2$ by $\mathbb{Z}$ with respect to multiplication by -1 (Exercise 2.E.31).

- Semi-direct products of the type $\mathbb{Z}^n \rtimes_\varphi \mathbb{Z}$ lead to interesting examples of groups provided the automorphism $\varphi(1) \in \mathrm{GL}(n, \mathbb{Z}) \subset \mathrm{GL}(n, \mathbb{R})$ is chosen suitably, e.g., if $\varphi(1)$ has interesting eigenvalues (Exercise 6.E.18).
- 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 $G \wr H$.
- Similarly, one can define lamplighter and wreath product groups using $\bigoplus_H G$ instead of $\prod_H G$ (Exercise 2.E.34).

2.3.2 Free products and amalgamated free 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 (a special type of colimits):

Definition 2.3.6 (Pushout of groups, free product (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 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 a *pushout of G_1 and G_2 over A (with respect to α_1 and α_2)* if the following universal property is satisfied:

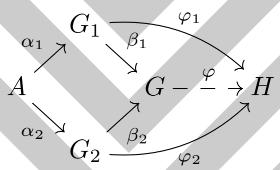

For every group H and all group homomorphisms $\varphi_1\colon G_1 \longrightarrow H$ and $\varphi_2\colon G_2 \longrightarrow H$ with $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$ there is exactly one homomorphism $\varphi\colon G \longrightarrow H$ with $\varphi \circ \beta_1 = \varphi_1$ and $\varphi \circ \beta_2 = \varphi_2$. Such a pushout is denoted by $G_1 *_A G_2$ (see Theorem 2.3.9 for existence and uniqueness).

Two special cases deserve their own names:

- 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*.
- If α_1 and α_2 both are injective, then the pushout group $G_1 *_A G_2$ is an *amalgamated free product of G_1 and G_2 over A (with respect to α_1 and α_2)*.

Caveat 2.3.7. In the situation of the above definition, in general, pushout groups and amalgamated free products do depend on the glueing 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 (Pushout groups, (amalgamated) free 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 $\mathbb{Z}/2 * \mathbb{Z}/2$; for instance, reflection at 0 and reflection at 1/2 provide generators of D_∞ corresponding to the obvious generators of $\mathbb{Z}/2 * \mathbb{Z}/2$ (Exercise 2.E.31).
- The matrix group $\mathrm{SL}(2,\mathbb{Z})$ is isomorphic to the amalgamated free product $\mathbb{Z}/6 *_{\mathbb{Z}/2} \mathbb{Z}/4$ [159, Example I.4.2] (Outlook 4.4.3).
- Pushout groups occur naturally in topology: By the theorem of Seifert and van Kampen, the fundamental group of a pointed space glued together out of two components is a pushout 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) [115, Chapter IV] (see Figure 2.3). A quick introduction to fundamental groups is given in Appendix A.1.

Theorem 2.3.9 (Pushout groups: uniqueness and construction). *All pushout groups exist and are unique up to canonical isomorphism.*

In particular, all amalgamated free products and all free products of groups 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 pushout groups: The idea is to use generators and relations to enforce the desired 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. Let

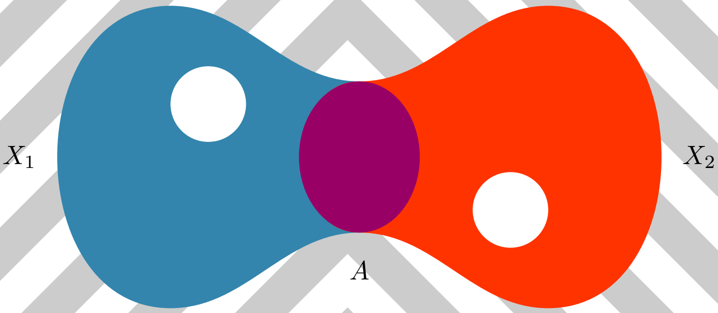

$$\pi_1(X_1 \cup_A X_2) \cong \pi_1(X_1) *_{\pi_1(A)} \pi_1(X_2)$$

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

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

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 pushout group of G_1 and G_2 over A: Let H be a group and let $\varphi_1 \colon G_1 \longrightarrow H$, $\varphi_2 \colon G_2 \longrightarrow H$ 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

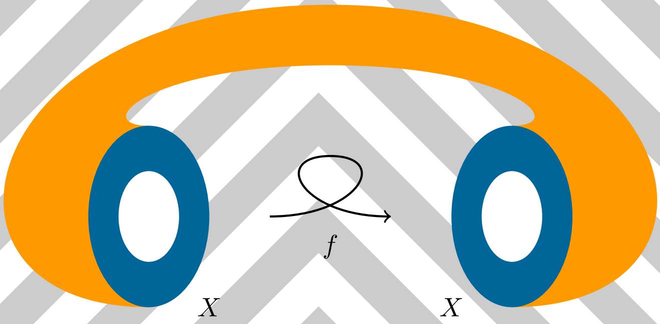

$\pi_1(\text{mapping torus of } f) \cong \pi_1(X)*_{\pi_1(f)}$

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

corresponding to this map provided by the universal property of generators and relations.

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 every 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 every presentation of the summands; this produces more efficient presentations of the amalgamated free product.

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

Definition 2.3.10 (HNN-extension)**.** Let G be a group, let $A, B \subset G$ be 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

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

One also says that t is the *stable letter* of this HNN-extension.

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 [107, Chapter IV][150, Chapter 12]. HNN-extensions are named after G. Higman, B.H. Neumann, and H. Neumann who were the first to systematically study such groups (Remark 2.3.12). Topologically, HNN-extensions arise naturally as fundamental groups of certain self-glueings, e.g., of mapping tori of maps that are injective on the level of fundamental groups [107, p. 180] (see Figure 2.4).

Outlook 2.3.11 (Amalgamated free products and HNN-extensions as building blocks)**.** The class of (non-trivial) amalgamated free 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 [166] (Theorem 8.2.14), and they are the starting point of Bass-Serre theory [159] (Outlook 4.2.7), which is concerned with actions of groups on trees (Outlook 4.2.7). Moreover, free groups, free products, amalgamated free products, and HNN-extensions can be understood in very concrete terms via suitable normal forms (Outlook 3.3.8).

Remark 2.3.12 (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:

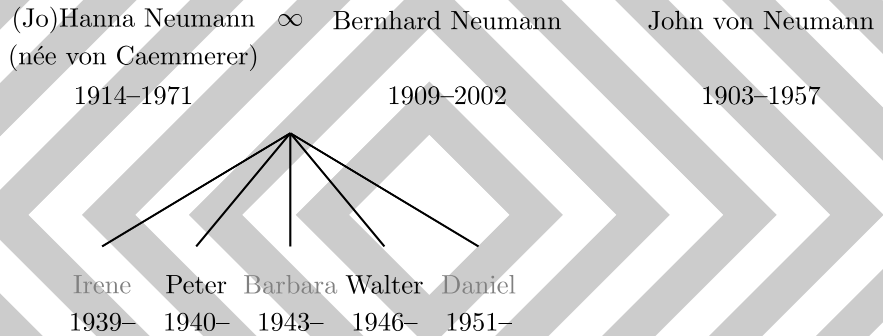

The contributions to geometric group theory of the (von) Neumanns are too numerous to be listed here [10, 9, 140]; for the topics in this book, the most important ones are:

- 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 (Outlook 4.2.13).
- There is a joint article by Bernhard, Hanna, and Peter Neumann [129].
- There is/was the von Neumann conjecture on the relation between non-amenability and free subgroups (Remark 9.1.12).

2.E Exercises

Basic group theory

Quick check 2.E.1 (Subgroups*). Let G be a group and let H and K be subgroups of G.

1. Is $H \cap K$ a subgroup of G ?
2. Is $H \cup K$ a subgroup of G ?

Quick check 2.E.2 (Squares of groups*).

1. Are the additive groups $\mathbb{Z}$ and $\mathbb{Z}^2$ isomorphic?
2. Are the additive groups $\mathbb{R}$ and $\mathbb{R}^2$ isomorphic?

Exercise 2.E.3 (Kernels and injectivity*). Let $\varphi\colon G \longrightarrow H$ be a group homomorphism.

1. Show that φ is injective if and only if $\ker \varphi$ is trivial.
2. Let $\overline{\varphi}\colon G/\ker\varphi \longrightarrow H$ be the group homomorphism from the quotient group induced by φ. Show that $\overline{\varphi}$ is injective.

Exercise 2.E.4 (Finitely generated groups**).

1. Is the additive group $\mathbb{Q}$ finitely generated?
2. Is the symmetric group S_X of an infinite set X finitely generated?

Exercise 2.E.5 (The normal subgroup trick**). Let G be a group.

1. Let $H, K \subset G$ be subgroups of finite index. Show that also $H \cap K$ has finite index in G.
2. Let $H \subset G$ be a subgroup and $S \subset G$ be a set of representatives of $\{g \cdot H \mid g \in G\}$. Show that

$$\bigcap_{g \in G} g \cdot H \cdot g^{-1} = \bigcap_{g \in S} g \cdot H \cdot g^{-1}.$$

3. Let $H \subset G$ be a subgroup of finite index. Show that there exists a normal subgroup $N \subset G$ of finite index with $N \subset H$.
   Hints. Consider $\bigcap_{g \in G} g \cdot H \cdot g^{-1}$ ...

Exercise 2.E.6 (Outer automorphism groups*).

1. Let G be a group. Show that the set $\operatorname{Inn}(G)$ of inner automorphisms of G is a normal subgroup of $\operatorname{Aut}(G)$.
2. Determine $\operatorname{Out}(\mathbb{Z})$.
3. Determine $\operatorname{Out}(\mathbb{Z}/2016)$ and $\operatorname{Out}(\mathbb{Z}/2017)$.

Exercise 2.E.7 (Galois groups***).

1. Find suitable categories that allow to interpret Galois groups as automorphism groups.
2. Find suitable categories that allow to interpret deck transformation groups as automorphism groups.

Basic isometry groups

Exercise 2.E.8 (Isometry group of the unit square**). Let $Q := [0,1] \times [0,1]$ be the unit square in $\mathbb{R}^2$ (with the Euclidean metric).

1. Give an algebraic description of the isometry group I of Q (e.g., by writing down the multiplication table).
2. Is there a group having the same number of elements as I that is not isomorphic to I ?

Exercise 2.E.9 (More isometry groups**). Is there for every $n \in \mathbb{N}_{>0}$ a subset $X_n \subset \mathbb{R}^2$ such that the isometry group of X_n is isomorphic to $\mathbb{Z}/n$?

Exercise 2.E.10 (Even more isometry groups**). Is there for every group G an $n \in \mathbb{N}$ and a subset $X \subset \mathbb{R}^n$ such that the isometry group of X is isomorphic to G ?

Free groups

Exercise 2.E.11 (Unfree groups*). Use the universal property of free groups to prove the following:

1. The additive group $\mathbb{Z}/2017$ is *not* free.
2. The additive group $\mathbb{Z}^2$ is *not* free.

Exercise 2.E.12 (Rank of free groups**).

1. Let S be a set and let F be the free group generated by S. Prove that if $S' \subset F$ is a generating set of F, then $|S'| \geq |S|$.
   Hints. If S is finite, one can apply the universal property of free groups to homomorphisms to $\mathbb{Z}/2$ and a counting argument. If S is infinite, one can use a cardinality argument or pass to the vector space $F_{\text{ab}} \otimes_{\mathbb{Z}} \mathbb{Q}$, where F_{ab} denotes the abelianisation of F (Exercise 2.E.18).
2. Conclude that all free generating sets of a free group have the same cardinality.
3. Show that the free group generated by two elements contains a subgroup that cannot be generated by two elements.
   Hints. Map surjectively to a big finite symmetric group and find a subgroup (e.g., Abelian) of this symmetric group that cannot be generated by two elements.

Exercise 2.E.13 (Mapping class groups***).

1. Look up the term *mapping class group* (e.g., of manifolds, topological spaces). Which formal similarity is there between this definition and the definition of outer automorphism groups of groups?
2. Let F be a free group of rank 2. How are $\operatorname{Out}(F)$ and the figure eight $S^1 \vee S^1$ (Figure 2.5) related?

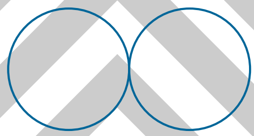

Figure 2.5.: The figure eight, $S^1 \vee S^1$

Generators and relations

Quick check 2.E.14 (Generators and relations, examples*).

1. Is the group $\langle x, y \mid xyx^{-1}y^{-1}\rangle$ isomorphic to $\mathbb{Z}^2$?
2. Are the groups $\langle s, t \mid t^2, tst^{-1} = s^{-1}\rangle$ and $\langle a, b \mid a^2, b^2\rangle$ isomorphic?
3. Is the group $\langle x, y \mid xy^{2014}x = yx^{2015}\rangle$ trivial?
4. Is the group $\langle x, y \mid xyx = yxy\rangle$ trivial?

Exercise 2.E.15 (Universal property for generators and relations*). Let S be a set and let $R \subset (S \cup S^{-1})^*$.

1. Show that the group $\langle S \mid R\rangle$ generated by S with the 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 every group H and every map $\varphi\colon S \longrightarrow H$ satisfying
$$\forall_{r\in R} \;\; \varphi^*(r) = e$$
there is exactly one group homomorphism $\overline{\varphi}\colon \langle S \mid R\rangle \longrightarrow H$ with
$$\overline{\varphi} \circ \pi = \varphi;$$
here, $\varphi^*\colon (S \cup S^{-1})^* \longrightarrow H$ is defined inductively by $\varphi^*(\varepsilon) = e$ and
$$\forall_{s\in S} \;\; \forall_{x\in(S\cup S^{-1})^*} \;\; \varphi^*(sx) = \varphi(s) \cdot \varphi^*(x)$$
$$\forall_{s\in S} \;\; \forall_{x\in(S\cup S^{-1})^*} \;\; \varphi^*(s^{-1}x) = \bigl(\varphi(s)\bigr)^{-1} \cdot \varphi^*(x).$$
2. Prove that up to canonical isomorphism there is exactly one group that has this universal property.

Exercise 2.E.16 (Finite normal generation of kernels**). Let $\varphi\colon G \longrightarrow H$ be a surjective group homomorphism, where G is finitely generated and H is finitely presented. Show that then $\ker\varphi$ is finitely normally generated, i.e., there is a finite set $N \subset G$ with

$$\ker\varphi = \langle N \rangle^{\triangleleft}_{G}.$$

Hints. First show that the property of finite presentability is independent of the chosen finite generating set.

Exercise 2.E.17 (Positive relations** [77, Exercise V.12]). A group presentation $\langle S \mid R \rangle$ is *positive* if $R \subset S^*$, i.e., if no negative exponents occur in any of the relations. Show that for every presentation $\langle S \mid R \rangle$ there is a positive presentation $\langle S' \mid R' \rangle$ with $\langle S' \mid R' \rangle \cong \langle S \mid R \rangle$ and

$$|S'| \leq |S| + 1 \qquad \text{and} \qquad |R'| \leq |R| + 1.$$

Exercise 2.E.18 (Abelianisation**). Let G be a group and let $[G,G]$ be its *commutator subgroup*, i.e., the subgroup generated by $\{[g,h] \mid g,h \in G\}$. The *commutator* of $g,h \in G$ is defined by $[g,h] := g \cdot h \cdot g^{-1} \cdot h^{-1}$. The quotient group

$$G_{\text{ab}} := G/[G,G]$$

is the *abelianisation of* G.

1. Prove that $[G,G]$ indeed is a normal subgroup of G and that the quotient group G_{ab} is Abelian.
2. Prove that abelianisation enjoys the following universal property: For every Abelian group H and every group homomorphism $\varphi\colon G \longrightarrow H$ there exists exactly one group homomorphism $\overline{\varphi}\colon G_{\text{ab}} \longrightarrow H$ satisfying $\overline{\varphi} \circ \pi = \varphi$, where $\pi\colon G \longrightarrow G_{\text{ab}}$ denotes the canonical projection.
3. How can this construction be turned into a functor $\cdot_{\text{ab}}\colon \mathsf{Group} \longrightarrow \mathsf{Ab}$?
4. Determine F_{ab} for all free groups F, using appropriate universal properties.
5. Let $G = \langle S \mid R \rangle$. Show that there is a canonical isomorphism

$$G_{\text{ab}} \cong \bigl\langle S \bigm| R \cup \{st = ts \mid s,t \in S\} \bigr\rangle.$$

Exercise 2.E.19 (The infinite dihedral group**). Let

$$D_\infty := \langle s,t \mid t^2, tst^{-1} = s^{-1} \rangle$$

be the *infinite dihedral group*. We consider $\mathbb{Z} \subset \mathbb{R}$ with the metric induced by the standard metric on $\mathbb{R}$. Show that $D_\infty \cong \operatorname{Isom}(\mathbb{Z})$.

Exercise 2.E.20 (Thompson's group F **). Let F denote Thompson's group

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

1. Show that
$$F \cong \langle a, b \mid [ab^{-1}, a^{-1}ba], [ab^{-1}, a^{-2}ba^2] \rangle$$
(in particular, F is finitely presented).
2. Look up in the literature how *PL-homeomorphisms* of $[0,1]$ are defined and how they are related to Thompson's group F. A graphical representation of an example of such a PL-homeomorphism is depicted in Figure 2.6.

Figure 2.6.: An example of a PL-homeomorphism of $[0,1]$

Exercise 2.E.21 (Baumslag-Solitar groups**).
1. Show that $\mathrm{BS}(1,1) \cong \mathbb{Z}^2$ (e.g., by comparing appropriate universal properties).
2. Let $m, n \in \mathbb{N}_{>0}$. Prove that $\mathrm{BS}(m,n)$ is infinite by studying the following matrices in $\mathrm{GL}(2, \mathbb{Q})$:
$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \qquad \begin{pmatrix} \frac{n}{m} & 0 \\ 0 & 1 \end{pmatrix}$$
3. Let $m, n \in \mathbb{N}_{>0}$. Show that $\mathrm{BS}(m,n)$ is *not* cyclic.
4. Show that
$$\begin{aligned} \varphi \colon \mathrm{BS}(2,3) &\longrightarrow \mathrm{BS}(2,3) \\ a &\longmapsto a^2 \\ b &\longmapsto b \end{aligned}$$
describes a well-defined surjective group homomorphism and that the commutator $[bab^{-1}, a]$ represents an element of $\ker \varphi$.

Exercise 2.E.22 (A normal form for $\mathrm{BS}(1,2)$ **). We consider the defining presentation $\langle a, b \mid bab^{-1} = a^2 \rangle$ of $\mathrm{BS}(1,2)$; let $\overline{a}, \overline{b} \in \mathrm{BS}(1,2)$ be the group elements corresponding to a and b, respectively.
1. Show that every element of $\mathrm{BS}(1,2)$ can be written in the form $\overline{b}^{-j} \cdot \overline{a}^k \cdot \overline{b}^\ell$ with $j, \ell \in \mathbb{N}$ and $k \in \mathbb{Z}$ and the following additional condition: The exponent k is odd or $j \cdot \ell = 0$.

2. Show that the exponents in the first part are unique.
Hints. Consider the matrices specified in Exercise 2.E.21 and elementary number theory.

Exercise 2.E.23 (Surface groups**)**.** For $n \in \mathbb{N}$ we define

$$G_n := \Big\langle a_1, \dots, a_n, b_1, \dots, b_n \;\Big|\; \prod_{j=1}^{n} [a_j, b_j] \Big\rangle.$$

Then by the Seifert and van Kampen theorem G_n is isomorphic to the fundamental group of an oriented closed connected surface of genus n (Figure 2.7).

1. Prove that for all $n, m \in \mathbb{N}$ we have $G_n \cong G_m$ if and only if $n = m$.
Hints. Abelianisation (Exercise 2.E.18) might help.
2. For which $n \in \mathbb{N}$ is the group G_n Abelian?

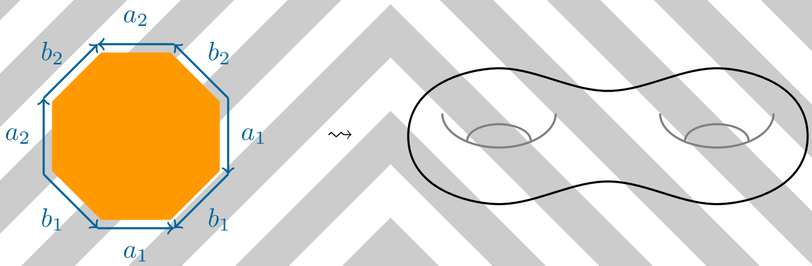

Figure 2.7.: The fundamental group of an oriented closed connected surface of genus 2

Exercise 2.E.24 (Coxeter groups**)**.** A *Coxeter group* is a group W such that there exist $n \in \mathbb{N}$ and a symmetric matrix $m \in M_{n\times n}(\mathbb{Z} \cup \{\infty\})$ with

$$W \cong \big\langle s_1, \dots, s_n \;\big|\; \{(s_j s_k)^{m_{jk}} \mid j, k \in \{1, \dots, n\}\} \big\rangle,$$

where the *Coxeter matrix* m satisfies the following: For all $j \in \{1, \dots, n\}$ we have $m_{jj} = 1$ and for all $j, k \in \{1, \dots, n\}$ we have $m_{kj} = m_{jk} \geq 2$; if $m_{jk} = \infty$, then the relation $(s_j s_k)^{m_{jk}} = e$ is viewed as empty condition.

1. Let $j, k \in \{1, \dots, n\}$ with $j \neq k$ and $m_{jk} = 2$. Show that the corresponding elements $\overline{s}_j$ and $\overline{s}_k$ commute in W.
2. Show that $\big\langle s_1, s_2, s_3 \;\big|\; (s_1 s_2)^2, (s_1 s_3)^2, (s_2 s_3)^2 \big\rangle \cong (\mathbb{Z}/2)^3$.
3. How can the isometry group of a regular n-gon be viewed as a Coxeter group?

More information on the rich geometry of Coxeter groups can be found in the book by Davis [45].

Exercise 2.E.25 (Geometric finite presentation***)**.** Let X be a path-connected CW-complex with finite 2-skeleton. Prove that the fundamental group of X is finitely presented. How can one read off a finite presentation from the 2-skeleton?

Hints. It can be helpful to use the fact every path-connected CW-complex is homotopy equivalent to a CW-complex with a single 0-cell. Then the theorem of Seifert and van Kampen will produce a finite presentation.

Exercise 2.E.26 (Braid groups**)**.** For $n \in \mathbb{N}$ the *braid group on n strands* is defined by

$$B_n := \bigl\langle s_1, \dots, s_{n-1} \bigm| \{s_j s_{j+1} s_j = s_{j+1} s_j s_{j+1} \mid j \in \{1, \dots, n-2\}\} \\ \{s_j s_k = s_k s_j \mid j, k \in \{1, \dots, n-1\}, |j-k| \geq 2\} \bigr\rangle.$$

1. Show that

$$\begin{aligned} B_n &\longrightarrow \mathbb{Z} \\ s_j &\longmapsto 1 \end{aligned}$$

   defines a well-defined group homomorphism. For which $n \in \mathbb{N}$ is this homomorphism surjective?
2. Show that

$$\begin{aligned} B_n &\longrightarrow S_n \\ s_j &\longmapsto (j, j+1) \end{aligned}$$

   defines a well-defined surjective homomorphism onto the symmetric group S_n.
3. Geometrically, the group B_n can be described as follows [88]:

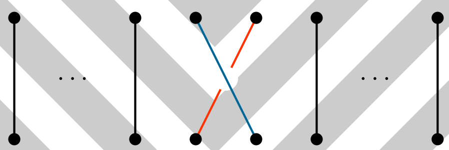

Figure 2.8.: Geometric generator s_j of the braid group B_n

   Elements of B_n are isotopy classes of n-braids. Here, an *n-braid* is a sequence $(\alpha_1, \dots, \alpha_n)$ of paths $[0,1] \longrightarrow \mathbb{R}^3$ with the following properties:
   - For each $j \in \{1, \dots, n\}$, the last coordinate of α_j is strictly increasing.
   - For each $j \in \{1, \dots, n\}$ we have $\alpha_j(0) = (j, 0, 0)$.

- There exists a permutation $\pi \in S_n$ such that for all $j \in \{1, \dots, n\}$ we have $\alpha_j(1) = \bigl(\pi(j), 0, 1\bigr)$.

An (ambient) *isotopy* between two n-braids α and β is a continuous map $F \colon \mathbb{R}^3 \times [0,1] \longrightarrow \mathbb{R}^3$ with the following properties: For all $t \in [0,1]$, the map $F(\,\cdot\,, t)$ is a self-homeomorphism of $\mathbb{R}^3$ that is the identity on $\mathbb{R}^2 \times \{0,1\}$ and that maps α to an n-braid, and furthermore we have $F(\,\cdot\,, 0) = \mathrm{id}_{\mathbb{R}^3}$, and $F(\,\cdot\,, 1)$ maps α to β. In this terminology, vertical concatenation (and rescaling) of braids corresponds to composition in B_n and the braid depicted in Figure 2.8 corresponds to the generator s_j of B_n.

Draw the geometric braid relations corresponding to the defining algebraic relations in the braid group B_n.

Exercise 2.E.27 (A finitely generated group that is not finitely presented**). We consider the group

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

The goal of this exercise is to prove that G is *not* finitely presentable.

1. Show that $G \cong \bigl\langle s, t \bigm| \{[s, t^n s t^{-n}] \mid n \in \mathbb{N}_{>0}\} \bigr\rangle$.
2. For $N \in \mathbb{N}_{>0}$ let $G_N := \bigl\langle s, t \bigm| \{[s, t^n s t^{-n}] \mid n \in \{1, \dots, N\}\} \bigr\rangle$. Show that the homomorphism $\pi_N \colon G_N \longrightarrow G_{N+1}$ given by the identity on $\{s, t\}$ is surjective but *not* injective.
   Hints. Use the universal property of generators and relations and try to map s to the transposition $(1\ 2) \in S_{2\cdot N+3}$ and t to the permutation $(1 \mapsto 3, 2 \mapsto 4, 3 \mapsto 5, \dots) \in S_{2\cdot N+3}$.
3. Use the second part to conclude that G is *not* finitely presentable.

Exercise 2.E.28 (Normal forms**). Let G be a group and let $S \subset G$ be a generating set. A *normal form for G over S* is a split of the canonical projection $(S \cup \widehat{S})^* \longrightarrow G$. We then say that G admits a *regular normal form* if G has a finite generating set and a normal form $N \colon G \longrightarrow (S \cup \widehat{S})^*$ for which the language $N(G) \subset (S \cup \widehat{S})^*$ is regular [40, Chapter 1][2]. Similarly, groups with context-free normal form are defined.

1. Give regular normal forms for $\mathbb{Z}$, $\mathbb{Z}/2017$, and $\mathbb{Z}^2$.
2. Show that the existence of a regular normal form is independent of the chosen finite generating set.
3. Use the pumping lemma for regular languages to show that every finitely generated infinite group with regular normal form contains an element of infinite order [117].
4. Use the pumping lemma for context-free languages to show that every finitely generated infinite group with context-free normal form contains an element of infinite order.

Exercise 2.E.29 (Random groups***). Look up in the literature how random finitely presented groups can be defined. There are several popular models; choose one of these models and describe it in detail.

New groups out of old

Quick check 2.E.30 (Special pushouts*). Let A and G be groups.

1. Does $G * A$ contain a subgroup that is isomorphic to G ?
2. What is the pushout group $G *_A 1$ with respect to a given homomorphism $A \longrightarrow G$ and the trivial homomorphism $A \longrightarrow 1$?
3. What is the pushout group $G *_A A$ with respect to a given homomorphism $A \longrightarrow G$ and id_A ?

Exercise 2.E.31 (The infinite dihedral group strikes back**). Let D_∞ be the infinite dihedral group (Exercise 2.E.19).

1. Show that $D_\infty \cong \mathbb{Z}/2 * \mathbb{Z}/2$.
2. Show that $\mathrm{Isom}(\mathbb{Z}) \cong \mathbb{Z} \rtimes_\varphi \mathbb{Z}/2$, where $\varphi\colon \mathbb{Z}/2 \longrightarrow \mathrm{Aut}(\mathbb{Z})$ is given by multiplication by -1.

Exercise 2.E.32 (Heisenberg group**). Let H be the *Heisenberg group*, i.e.,

$$H = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x,\ y,\ z \in \mathbb{Z} \right\} \subset \mathrm{SL}(3, \mathbb{Z}).$$

1. The Heisenberg group is an extension of $\mathbb{Z}^2$ by $\mathbb{Z}$:
$$1 \longrightarrow \mathbb{Z} \xrightarrow{\;i\;} H \xrightarrow{\;\pi\;} \mathbb{Z}^2 \longrightarrow 1;$$
here, $i\colon \mathbb{Z} \longrightarrow H$ and $\pi\colon H \longrightarrow \mathbb{Z}^2$ are defined as follows:
$$i\colon z \longmapsto \begin{pmatrix} 1 & 0 & z \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad \pi\colon \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \longmapsto (x, y).$$
Show that this group extension does *not* split.
2. Show that
$$\langle x, y, z \mid [x, z],\ [y, z],\ [x, y] = z \rangle$$
is a presentation of the Heisenberg group.

Exercise 2.E.33 (Equivalence of extensions**).

1. Look up how the notion of *equivalence* is defined for group extensions.
2. Give three pairwise non-equivalent group extensions of the following type:
$$1 \longrightarrow \mathbb{Z} \xrightarrow{\;?\;} \,?\, \xrightarrow{\;?\;} \mathbb{Z}/3 \longrightarrow 1.$$

Exercise 2.E.34 (Restricted lamplighter groups**). For a family $(G_i)_{i \in I}$ of groups, we write $\bigoplus_{i \in I} G_i$ for the subgroup of $\prod_{i \in I} G_i$ of all sequences that are trivial almost everywhere (i.e., in each sequence, only a finite number of entries is non-trivial). Let G be a non-trivial finitely generated group and let H be a finitely generated group.

1. Show that $\bigoplus_H G$ is *not* finitely generated.
2. Show that $(\bigoplus_H G) \rtimes_\varphi H$ is finitely generated, where φ denotes the shift action of H on $\bigoplus_H G$.

Exercise 2.E.35 (A finitely generated group with big centre**). The goal of this exercise is to construct a finitely generated group that contains a central subgroup isomorphic to $\bigoplus_{\mathbb{N}} \mathbb{Z}$. Let A and B be free $\mathbb{Z}$-modules with bases $(a_n)_{n\in\mathbb{N}_{>0}}$ and $(b_n)_{n\in\mathbb{Z}}$ (we view A and B as Abelian groups) and let $f\colon B \times B \longrightarrow A$ be the unique $\mathbb{Z}$-bilinear map with

$$f(b_m, b_n) = \begin{cases} a_{n-m} & \text{if } n > m \\ 0 & \text{if } n \leq m \end{cases}$$

for all $n, m \in \mathbb{Z}$. We then define the composition

$$\begin{aligned} H \times H &\longrightarrow H \\ \bigl((a,b),(a',b')\bigr) &\longmapsto \bigl(a + a' + f(b,b'), b + b'\bigr) \end{aligned}$$

on the set $H := A \times B$. Moreover, we consider the map

$$\begin{aligned} T\colon H &\longrightarrow H \\ \Bigl(a, \sum_{n\in\mathbb{Z}} \beta_n \cdot b_n\Bigr) &\longmapsto \Bigl(a, \sum_{n\in\mathbb{Z}} \beta_n \cdot b_{n+1}\Bigr) \end{aligned}$$

and set $G := H \rtimes_{(1\mapsto T)} \mathbb{Z}$.

1. Prove that H is a group with respect to the composition defined above.
2. Show that $A \times \{0\}$ is a central subgroup of H.
3. Compute the commutators $\bigl[(0,b_0),(0,b_n)\bigr]$ in H for all $n \in \mathbb{N}_{>0}$.
4. Conclude that $\{(0,b_n) \mid n \in \mathbb{Z}\}$ is a generating set of H.
5. Prove that T is a group automorphism of H.
6. Show that G is generated by $\sigma := ((0,b_0),0)$ and $\tau := ((0,0),1)$.
7. Show that $A \times \{0\} \times \{0\}$ is a central subgroup of G that is isomorphic to $\bigoplus_{\mathbb{N}} \mathbb{Z}$.
8. Show that $s \mapsto \sigma, t \mapsto \tau$ induces a well-defined epimorphism

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

   Conclude that the subgroup generated by $\{[s,t^n s t^{-n}] \mid n \in \mathbb{Z}\}$ contains a subgroup isomorphic to $\bigoplus_{\mathbb{N}} \mathbb{Z}$.

Quick check 2.E.36 (Baumslag-Solitar groups and HNN-extensions*). How can Baumslag-Solitar groups be viewed as HNN-extensions?

Exercise 2.E.37 (Ascending HNN-extensions*). An *ascending HNN-extension* is an HNN-extension of the form $G*_\vartheta$, where G is a group and $\vartheta\colon G \longrightarrow G$ is an injective homomorphism.

1. What happens if $\vartheta = \mathrm{id}_G$?
2. Let $\vartheta \in \mathrm{Aut}(G)$. Prove that $G*_\vartheta \cong G \rtimes_{1\mapsto\vartheta} \mathbb{Z}$.

Exercise 2.E.38 (Free wreath product**).

1. Let I be a set and let $(G_i)_{i\in I}$ be a family of groups. Formulate the universal property of the corresponding *free product group* $\bigstar_{i\in I} G_i$ as a suitable colimit universal property.
2. Indicate how such general free products can be constructed and how free groups can be viewed as such free products.
3. Let G and H be groups. Then H acts on the free product $\bigstar_H G$ by shifting the summands. We call the corresponding semi-direct product

$$G \wr_* H := \Bigl(\bigstar_H G\Bigr) \rtimes H$$

the *free wreath product* of G and H. Show that the obvious homomorphisms $H \longrightarrow G \wr_* H$ and $G \longrightarrow G \wr_* H$ induce an isomorphism

$$G * H \cong G \wr_* H.$$

Part II

Groups $\to$ Geometry

3

Cayley graphs

A fundamental question 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 already has a rudimentary geometric flavour and is discussed in this chapter.
2. As second step, one adds a metric structure to Cayley graphs via word metrics. We will study this step in Chapter 5.

We start by reviewing some basic notation from graph theory (Chapter 3.1). We will then introduce Cayley graphs and discuss basic examples of Cayley graphs (Chapter 3.2); in particular, we will 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 the corresponding generating set is free (Chapter 3.3).

Overview of this chapter

3.1 Review of graph notation

We start by reviewing basic terminology from graph theory; more information can be found in the literature [49, 79, 43]. In the following, we will always consider undirected, simple graphs without loops:

Definition 3.1.1 (Graph)**.** A *graph* is a pair $X = (V, E)$ 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* X.

In other words, graphs are a different point of view on (symmetric) 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 structure).

Definition 3.1.2 (Adjacent, neighbour, degree)**.** 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 $X_1 := (V, E)$ can be illustrated as in Figure 3.1; however, differently looking pictures can in fact represent the same graph (a graph is a combinatorial object!). In X_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):

$$\begin{aligned} X_2 &:= \bigl(\{1, \dots, 5\}, \{\{j, k\} \mid j, k \in \{1, \dots, 5\}, \ j \neq k\}\bigr), \\ X_3 &:= \bigl(\{1, \dots, 9\}, \\ &\qquad \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{2, 6\}, \{3, 7\}, \{4, 8\}, \{8, 9\}\}\bigr). \end{aligned}$$

The graph X_2 is a complete graph: all vertices are neighbours of each other.

Example 3.1.4 (Complete graphs)**.** Let $n, m \in \mathbb{N}$. The graph

$$K_n := \bigl(\{1, \dots, n\}, \{\{j, k\} \mid j, k \in \{1, \dots, n\}, j \neq k\}\bigr)$$

is "the" *complete graph* on n vertices. If $n > 0$, then K_n has exactly n vertices (each of degree $n - 1$) and $1/2 \cdot n \cdot (n - 1)$ edges. The graph

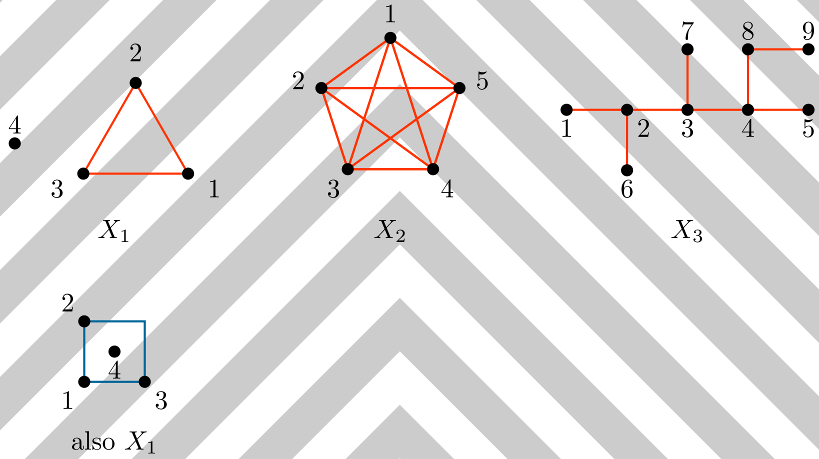

Figure 3.1.: Some graphs

$$K_{n,m} := \bigl(\{(1,0),\dots,(n,0),(1,1),\dots,(m,1)\},\\ \{\{(j,0),(k,1)\} \mid j \in \{1,\dots,n\}, k \in \{1,\dots,m\}\}\bigr)$$

is "the" *complete bipartite graph* graph. If $n, m > 0$, then $K_{n,m}$ has exactly $n + m$ vertices and $n \cdot m$ edges.

Definition 3.1.5 (Graph isomorphisms)**.** Let $X = (V, E)$ and $X' = (V', E')$ be graphs. The graphs X and X' are *isomorphic*, if there is a *graph isomorphism* between X and X', i.e., a bijection $f : V \longrightarrow V'$ such that for all $v, w \in V$ we have $\{v, w\} \in E$ if and only if $\bigl\{f(v), f(w)\bigr\} \in E'$. Thus, 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 [91].

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

Definition 3.1.6 (Paths, cycles)**.** Let $X = (V, E)$ be a graph.

- Let $n \in \mathbb{N} \cup \{\infty\}$. A *path in X 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*.

[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.

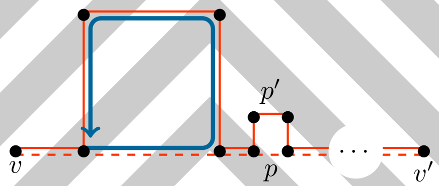

Figure 3.2.: Constructing a cycle (blue) out of two different paths.

- The graph X is called *connected* if every two of its vertices can be connected by a path in X.
- Let $n \in \mathbb{N}_{>2}$. A *cycle in X of length n* is a path $v_0, \dots, v_{n-1}$ in X with $\{v_{n-1}, v_n\} \in E$.

Example 3.1.7. In Example 3.1.3, the graphs X_2 and X_3 are connected, but X_1 is not connected (e.g., in X_1 there is no path connecting the vertex 4 to vertex 1). The sequence $1, 2, 3$ is a path in X_3, but $7, 8, 9$ and $2, 3, 2$ are no paths in X_3. In X_1, the sequence $1, 2, 3$ is a cycle.

Definition 3.1.8 (Tree)**.** A *tree* is a connected graph that does not contain any cycles. A graph that does not contain any cycles is a *forest*; so, a tree is the same as a connected forest.

Example 3.1.9 (Trees)**.** The graph X_3 in Example 3.1.3 is a tree, while X_1 and X_2 are not.

Proposition 3.1.10 (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 X be a graph such that every pair of vertices can be connected by exactly one path in X; in particular, X is connected. *Assume* for a contradiction that X 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, X is a tree.

Conversely, let X be a tree; in particular, X is connected, and every two vertices can be connected by a path in X. *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 X (see Figure 3.2), contradicting the fact that X is a tree. Hence, every two vertices of X can by connected by exactly one path in X. □

An alternative characterisation of finite trees is given in Exercise 3.E.4.

Trees can be viewed as basic ingredients of graphs: every connected graph contains a spanning tree (Exercise 3.E.6).

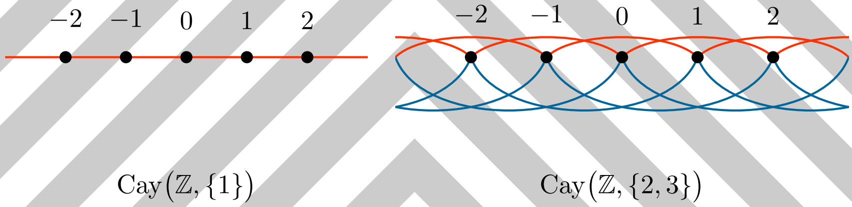

Figure 3.3.: Cayley graphs of the additive group $\mathbb{Z}$

Definition 3.1.11 (Spanning tree)**.** A *spanning tree* of a graph X is a subgraph of X that is a tree and contains all vertices of X. A *subgraph* of a graph (V, E) is a graph (V', E') with $V' \subset V$ and $E' \subset E$.

For example, in algebraic topology, spanning trees are used to calculate the fundamental group of connected 1-dimensional complexes. Moreover, we will use an equivariant version of spanning trees in Chapter 4.2.1 in order to characterise free groups.

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 right multiplication by an (inverse of an) element of the generating set in question. By definition, the Cayley graph with respect to a generating set S coincides with the Cayley graphs for S^{-1} and for $S \cup S^{-1}$.

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 depicted in Figure 3.3. When looking at these two graphs "from far away" they seem to have the

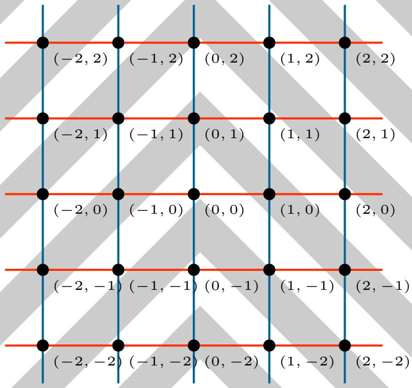

Figure 3.4.: The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$

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 introduce and study thoroughly in later chapters (Chapters 5–9.).

- 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.
- The Cayley graph of the cyclic group $\mathbb{Z}/6$ 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 set $\{\tau, \sigma\}$ is depicted in Figure 3.5.
  The Cayley graph $\operatorname{Cay}(S_3, S_3)$ is a complete graph on six vertices; similarly, $\operatorname{Cay}(\mathbb{Z}/6, \mathbb{Z}/6)$ 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. The question which groups admit isomorphic Cayley graphs is discussed in more detail in Outlook 3.2.4.
- The Cayley graph of a free group with respect to a free generating set is a tree (see Theorem 3.3.1 below).

Further examples of Cayley graphs are subject of various exercises (Chapter 3.E)

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.

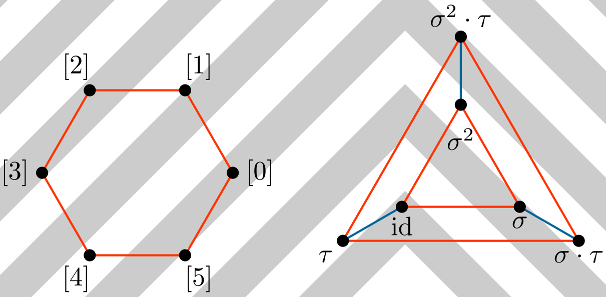

$\mathrm{Cay}\bigl(\mathbb{Z}/6, \{[1]\}\bigr)$ $\mathrm{Cay}\bigl(S_3, \{\tau, \sigma\}\bigr)$ $\mathrm{Cay}(S_3, S_3) \cong \mathrm{Cay}(\mathbb{Z}/6, \mathbb{Z}/6)$

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

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.

Already in this basic setup we can witness an interesting geometric phenomenon: rigidity.

Outlook 3.2.4 (Isomorphism rigidity of Cayley graphs)**.** It is natural to consider the combinatorial problem of which finitely generated groups *admit isomorphic Cayley graphs*, i.e., for which finitely generated groups G and H there exist finite generating sets $S \subset G$ and $T \subset H$ such that the graphs $\mathrm{Cay}(G, S)$ and $\mathrm{Cay}(H, T)$ are isomorphic. This question is related to the problem of determining for which groups isomorphisms/automorphisms of Cayley graphs are *affine*, i.e., given (up to translation by a fixed group element) by a group isomorphism. Both of these questions ask for rigidity properties of Cayley graphs, namely, how much of the algebraic structure is rigid enough to be visible in the combinatorics of all Cayley graphs of a given group.

These questions are well studied for finite groups [97]. For infinite groups, the following is known:

- Cayley graphs of finitely generated Abelian groups are rigid in the sense that automorphisms are affine on the free part and that these Cayley graphs remember the rank and the size of the torsion part [102].
- Moreover, it is known that automorphisms of Cayley graphs of finitely generated torsion-free nilpotent groups are affine and that Cayley graphs of finitely generated nilpotent groups remember the group modulo torsion [177].

- Finitely generated free groups admit isomorphic Cayley graphs if and only if they have the same rank. The proof is probabilistic; more precisely, it is based on the relation between the expected degree of random spanning forests and the first L^2-Betti number [108].

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).

Outlook 3.2.5 (Presentation complex, classifying space)**.** There are higher dimensional analogues of group presentations and Cayley graphs in topology:

Associated with a presentation of a group, there is the *presentation complex* [31, Chapter I.8A], which is a two-dimensional object. Roughly speaking, the presentation complex is the two-dimensional CW-complex given by

- taking a point,
- attaching a circle for every generator,
- and attaching a disk for every relation (in such a way that the boundary of the disk represents the word of the relation in the fundamental group of the glued circles).

By the Seifert and van Kampen theorem, the fundamental group of the presentation complex coincides with the given group. The presentation complex is finite/compact if and only if the underlying presentation is finite.

For example, the presentation complex associated with the presentation $\langle x, y \mid [x, y] \rangle$ is the torus and the presentation complex associated with the presentation $\langle x \mid x^2 \rangle$ is the projective plane $\mathbb{R}P^2$ (Figure 3.6).

More generally, every group admits a *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 [81, Chapter I.B]; one way to construct classifying spaces is to start with a presentation complex and then to add higher dimensional cells that kill the higher homotopy groups. These spaces are unique up to homotopy equivalence and allow to model group theory (both groups and homomorphisms) in topology. Classifying spaces play an important role in the study of group cohomology [34, 101] (Appendix A.2). Hence, classifying spaces (and their (co)homology) can be viewed as higher dimensional versions of group presentations.

For example, the torus is a classifying space for $\mathbb{Z}^2$ and the infinite-dimensional projective space $\mathbb{R}P^\infty$ is a classifying space for $\mathbb{Z}/2$.

How is all this related to Cayley graphs? The one-dimensional part (i.e., the 1-skeleton) of the universal covering of the presentation complex of a presentation $\langle S \mid R \rangle$ almost is the Cayley graph $\mathrm{Cay}(\langle S \mid R \rangle, S)$ (in case of generators of order 2 some modifications might be necessary) [45, Chapter 2.2]. We will return to this point of view in Outlook 4.1.21.

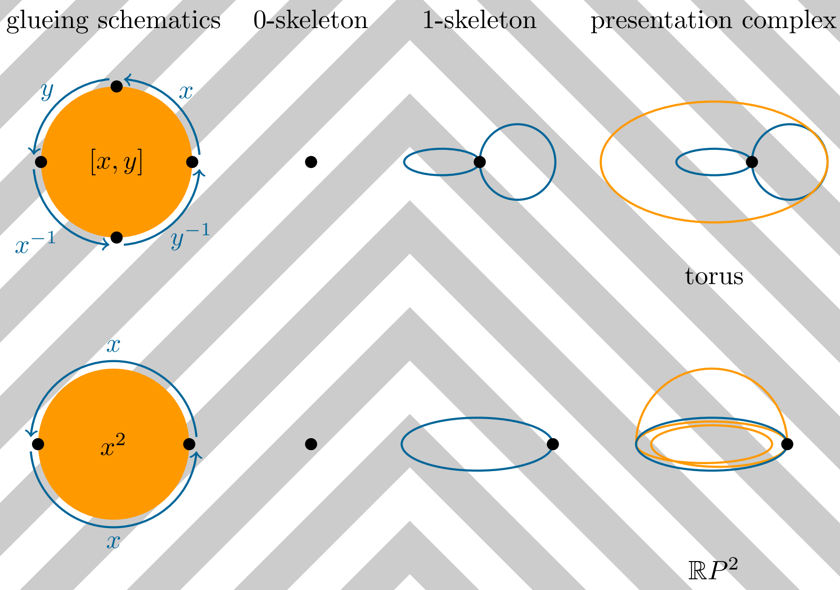

Figure 3.6.: Examples of presentation complexes

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 $\operatorname{Cay}(\mathbb{Z}/2, [1])$ consists of two vertices joined by an edge; clearly, this graph is a tree, but the group $\mathbb{Z}/2$ is not free.
- The Cayley graph $\operatorname{Cay}(\mathbb{Z}, \{-1, 1\})$ coincides with $\operatorname{Cay}(\mathbb{Z}, \{1\})$, 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 types of things that can go wrong:

Theorem 3.3.3 (Cayley trees and free groups)**.** *Let G be a group, 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 (Chapter 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 the cancellation 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 $(S \cup \widehat{S})^*$ be the set of words over S and formal inverses of elements of S.

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

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

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

Proposition 3.3.5 (Free groups via reduced words)**.** *Let S be a set.*

1. *The set $F_{\mathrm{red}}(S)$ of reduced words over $S \cup \widehat{S}$ forms a group with respect to the composition $F_{\mathrm{red}}(S) \times F_{\mathrm{red}}(S) \longrightarrow F_{\mathrm{red}}(S)$ 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 in $F_{\mathrm{red}}(S)$ (with $s_1, \dots, s_m \in S \cup \widehat{S}$), and

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

   In other words, the composition of reduced words is given by first concatenating the words and then reducing maximally at the concatenation position.
2. *The group $F_{\mathrm{red}}(S)$ is freely generated by S.*

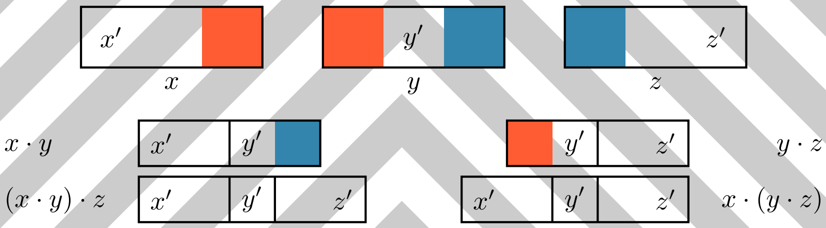

Figure 3.7.: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do not interfere

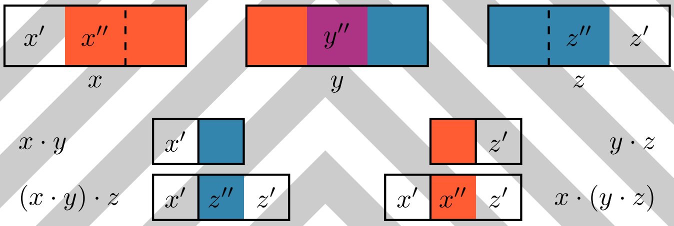

Figure 3.8.: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do interfere

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 hat 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 involving lots of indices, we explain the argument graphically (Figures 3.7 and 3.8): Let $x, y, z \in F_{\text{red}}(S)$; 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; 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 $F_{\text{red}}(S)$ 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

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

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

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

Corollary 3.3.6 (Normal form for free groups)**.** *Let S be a set. Every element of the free group $F(S) = (S \cup \widehat{S})^* / \sim$ can be represented by exactly one reduced word over $S \cup \widehat{S}$.* □

Corollary 3.3.7 (Word problem for free groups)**.** *The word problem in free groups with respect to free generating sets is solvable.*

Proof. Let F be a free group with free generating set S. If $w \in (S \cup \widehat{S})^*$, then we inductively reduce the word w until we reach a reduced word w'. Then the words w and w' represent the same element of F. Arguing as in the proof of the second part of Proposition 3.3.5 via the canonical isomorphism $F_{\text{red}}(S) \cong F$, we now only need to check whether w' is the empty word or not to determine whether the group element w is trivial or not. □

Outlook 3.3.8 (Reduced words in free products etc.)**.** Using the same method of proof as in Proposition 3.3.5, 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 $g_{j+1} \in G_1 \setminus \{e\}$.

Such reduced words can be composed by "concatenation and then maximal reduction at the concatenation position". The resulting group is the free product of G_1 and G_2 (all of this is not hard to check).

One can also describe amalgamated free products and HNN-extensions by suitable classes of reduced words [159, Chapter I][150, Chapter 11] (however, these generalisations are slightly more involved because more bookkeeping is needed and more ambiguities occur):

1. *Amalgamated free products.* Let A, G_1, G_2 be groups, let $\alpha_1 \colon A \longrightarrow G_1$, $\alpha_2 \colon A \longrightarrow G_2$ be injective group homomorphisms. Let $n \in \mathbb{N}$ and let $g_0, \dots, g_n \in G_1$, $h_0, \dots, h_n \in G_2$ with

$$\forall_{j \in \{1,\dots,n\}} \ g_j \notin \operatorname{im} \alpha_1 \qquad \text{and} \qquad \forall_{k \in \{0,\dots,n-1\}} \ h_k \notin \operatorname{im} \alpha_2.$$

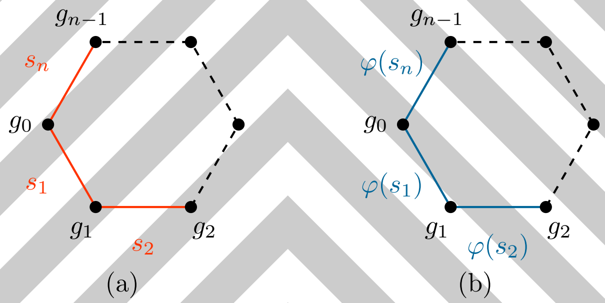

Figure 3.9.: Cycles lead to reduced words, and vice versa

Then the corresponding product $g_0 \cdot h_0 \cdot \dots \cdot g_n \cdot h_n$ in $G_1 *_A G_2$ is non-trivial. Moreover, every non-trivial element of $G_1 *_A G_2$ can be written in this so-called reduced form.

2. *HNN-Extensions.* Let G and A be groups, let $\vartheta \colon A \longrightarrow G$ be an injective group homomorphism, and let t denote the stable letter of $G*_\vartheta$. Let $n \in \mathbb{N}$, $m_1, \dots, m_n \in \mathbb{Z} \setminus \{0\}$, and $g_0, \dots, g_n \in G$ with

$$\forall_{j \in \{1,\dots,n\}} \;\; m_j < 0 \Longrightarrow g_j \notin A$$
$$\forall_{j \in \{1,\dots,n\}} \;\; m_j > 0 \Longrightarrow g_j \notin \vartheta(A).$$

Then $g_0 \cdot t^{m_1} \cdot g_1 \cdot \dots \cdot g_{n-1} \cdot t^{m_n} \cdot g_n$ is non-trivial in $G*_\vartheta$. Moreover, every element of $G*_\vartheta$ can be written in this so-called reduced form.

3.3.2 Free groups $\rightarrow$ 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 $F_{\mathrm{red}}(S)$ via an isomorphism that is the identity on S; without loss of generality we can therefore assume that F is $F_{\mathrm{red}}(S)$.

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

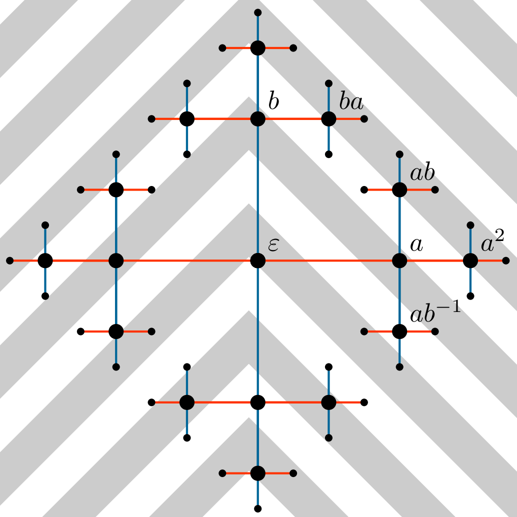

Figure 3.10.: Cayley graph of the free group of rank 2 with respect to a free generating set $\{a, b\}$

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

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 $F_{\text{red}}(S)$ via an isomorphism that is the identity on S.

Because $F_{\text{red}}(S)$ is freely generated by S, the universal property of free groups provides us with a group homomorphism $\varphi\colon F_{\text{red}}(S) \longrightarrow G$ 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 $s_1 \dots s_n \in F_{\text{red}}(S) \setminus \{\varepsilon\}$

with $s_1, \dots, s_n \in S \cup \widehat{S}$ 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 $s \cdot t \neq e$ holds in G 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

  $$g_{j+1} := g_j \cdot s_{j+1}$$

  for all $j \in \{0, \dots, n-2\}$ (Figure 3.9 (b)). The sequence $g_0, \dots, g_{n-1}$ is a cycle in $\mathrm{Cay}(G, S)$ because by minimality of the word $s_1 \dots s_n$, the elements $g_0, \dots, g_{n-1}$ are all distinct; moreover, $\mathrm{Cay}(G, S)$ 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 the hypothesis that $\mathrm{Cay}(G, S)$ is a tree.

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

3.E Exercises

Basic graph theory

Quick check 3.E.1 (Isomorphic graphs?*)**.** Which of the graphs in Figure 3.11 are isomorphic?

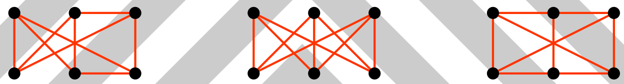

Figure 3.11.: Three 3-regular graphs

Quick check 3.E.2 (Index*)**.** Show that the index of this book is not a forest.
Hints. Look for a silly cycle.

Exercise 3.E.3 (Regular graphs of degree 2 *)**.** Classify (up to isomorphism) all connected graphs where each vertex has degree 2.

Exercise 3.E.4 (Characterisation of finite trees*)**.** Let $X = (V, E)$ be a finite connected graph with $V \neq \emptyset$. Show that X is a tree if and only if

$$|E| = |V| - 1.$$

Exercise 3.E.5 (Locally finite trees**)**.** Let T be a tree with infinitely many vertices that all have finite degree. Show that T contains an infinite path.

Exercise 3.E.6 (Spanning trees**)**.** Use Zorn's lemma to prove that every connected graph contains a spanning tree.

Exercise 3.E.7 (Infinite paths**)**.** Let X be a connected graph with infinitely many vertices.

1. Does X necessarily contain an infinite path?
2. Show that X contains an infinite path, if every vertex has finite degree.
   Hints. Consider a spanning tree of X and apply Exercise 3.E.5.

Exercise 3.E.8 (Marriage theorem***)**.** In the following, we will prove *Hall's marriage theorem*: Let W, M be non-empty sets and let $F \colon W \longrightarrow P^{\text{fin}}(M)$ be a map satisfying the *marriage condition*

$$\forall_{V \in P^{\text{fin}}(W)} \;\; F(V) := \Bigl| \bigcup_{w \in V} F(w) \Bigr| \geq |V|;$$

here, $P^{\mathrm{fin}}(W)$ denotes the set of all finite subsets of W. Then there exists a (W, M, F)-*marriage*, i.e., an injective map $\mu\colon W \longrightarrow M$ with

$$\forall_{w \in W} \;\; \mu(w) \in F(w).$$

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.

1. Formulate this theorem in terms of graph theory.
2. Prove the theorem in the case that W and M are finite.
   Hints. Proceed by induction over $|W|$. In the induction step, distinguish between two cases, depending on whether $|F(V)| \geq |V| + 1$ holds for all $V \in P^{\mathrm{fin}}(W)$ or not.
3. Prove the theorem in the general case.
   Hints. Use Zorn's lemma and the finite case. Zorn's lemma can, for example, be applied to the partially ordered set of all extendable partial marriages:
   - Let $W' \subset W$ be finite. Then an *extendable W'-marriage* is a $(W', M, F|_{W'})$-marriage μ' such that for all finite $W'' \subset W$ with $W' \subset W''$ there exists a $(W'', M, F|_{W''})$-marriage that extends μ'.
   - For a general (not necessarily finite) subset $W' \subset W$ an *extendable W'-marriage* is a $(W', M, F_{W'})$-marriage whose restriction to every finite subset $W'' \subset W'$ is an extendable W''-marriage.

Cayley graphs

Quick check 3.E.9 (Cayley graphs with few edges?*)**.**

1. Does there exist a group that has a Cayley graph with exactly 2016 vertices and exactly 2017 edges?
2. Does there exist a group that has a Cayley graph with exactly 2016 vertices and exactly 2016 edges?

Quick check 3.E.10 (Cycles in Cayley graphs*)**.** Let G be an Abelian group that is not cyclic.

1. Does G admit a Cayley graph with cycles of length 3 ?
2. Does every Cayley graph of G contain a cycle of length 4 ?

Exercise 3.E.11 (Characterising infinite groups**)**.** Let G be a group and let $S \subset G$ be a generating set of G. Show that G is infinite if and only if $\mathrm{Cay}(G, S)$ contains an infinite path.

Hints. If S is finite, then Exercise 3.E.7 will help. If S is infinite, then there is enough space for a naive inductive construction.

Exercise 3.E.12 (Klein bottle**). Let

$$\begin{aligned} \varphi\colon \mathbb{Z} &\longrightarrow \operatorname{Aut}(\mathbb{Z}) \\ n &\longmapsto (-1)^n \cdot \mathrm{id}_{\mathbb{Z}} \end{aligned}$$

and $G := \mathbb{Z} \rtimes_\varphi \mathbb{Z}$.

1. Show that G is *not* isomorphic to $\mathbb{Z}^2$.
2. Sketch a Cayley graph of G with respect to some finite generating set of your choice.
3. Prove that G is isomorphic to the fundamental group of the Klein bottle (Figure 3.12).

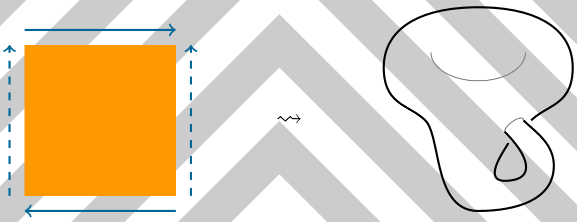

Figure 3.12.: The Klein bottle

Exercise 3.E.13 (Petersen graph**). Show that there is no group that has a finite generating set such that the associated Cayley graph is isomorphic to the Petersen graph (Figure 3.13).
Hints. It might be useful to first show that the Petersen graph is no Cayley graph of $\mathbb{Z}/10$ and no Cayley graph of D_5.

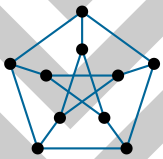

Figure 3.13.: The Petersen graph

Exercise 3.E.14 (Cayley graphs of free products*).

1. Sketch the Cayley graph of $\mathbb{Z}/2 * \mathbb{Z}/5$ with respect to some finite generating set of your choice.
   Hints. Use the description of free products in terms of reduced words (Outlook 3.3.8).

2. Sketch the Cayley graph of D_∞ with respect to some finite generating set of your choice.

Exercise 3.E.15 ((Im)Possible Cayley graphs?**).
1. Is there a finitely generated group G with a finite generating set S such that the corresponding Cayley graph $\operatorname{Cay}(G,S)$ is a a tree all of whose vertices have degree 3 (Figure 3.14.(a))?
2. Is there a finitely generated group G with a finite generating set S such that the corresponding Cayley graph $\operatorname{Cay}(G,S)$ is isomorphic to the graph in Figure 3.14.(b)?

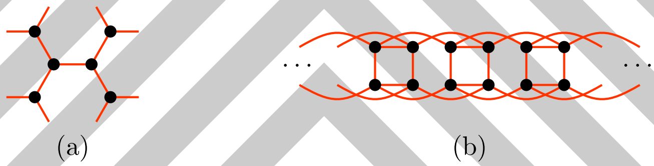

Figure 3.14.: (Im)Possible Cayley graphs?

Exercise 3.E.16 (Cayley graph of $\operatorname{BS}(1,2)$ **). Sketch the Cayley graph of the Baumslag-Solitar group $\operatorname{BS}(1,2) = \langle a,b \mid bab^{-1} = a^2\rangle$ with respect to the generating set $\{a,b\}$.
Hints. Use the normal form from Exercise 2.E.22 or a suitable description of HNN-extensions in terms of reduced words.

Free groups via reduced words

Quick check 3.E.17 (Powers in free groups*).
1. Which elements $g \in F_2$ satisfy $g^2 = e$?
2. Which elements $g,h \in F_2$ satisfy $g^2 = h^2$?
3. Are there elements $g,h \in F_2$ with $g^{2017} = h^{2018}$?

Exercise 3.E.18 (Trivial inner automorphisms of free groups*). Let S be a set.
1. Let $s \in S$ and let $g \in F(S)$ with $s = g \cdot s \cdot g^{-1}$. Show that there exists a $k \in \mathbb{Z}$ with $g = s^k$.
2. Conclude: If $|S| \geq 2$ and $g \in F(S)$ satisfies
$$\forall_{x \in F(S)} \;\; g \cdot x \cdot g^{-1} = x,$$
then $g = e$.

Exercise 3.E.19 (A large spanning tree**). Find a finite generating set S of the free group F of rank 2 such that $\operatorname{Cay}(F,S)$ contains a spanning tree that is regular of degree 6.

Exercise 3.E.20 (Free Abelian subgroups of $\operatorname{Out}(F_n)$**). Let $n \in \mathbb{N}_{\geq 2}$ and let F_n be the free group of rank n.

1. Show that $\operatorname{Aut}(F_n)$ contains a subgroup that is isomorphic to $\mathbb{Z}^{2 \cdot n-2}$.
2. Show that $\operatorname{Out}(F_n)$ contains a subgroup that is isomorphic to $\mathbb{Z}^{2 \cdot n-3}$.

Hints. Let $\{x_1, \dots, x_n\}$ be a free generating set of F_n. Consider the building blocks "$x_j \mapsto x_j \cdot x_1$" and "$x_j \mapsto x_1 \cdot x_j$" ... Exercise 3.E.18 might help to prove that the constructed subgroup indeed is free Abelian of the correct rank.

Isomorphisms of Cayley graphs

Exercise 3.E.21 (Isomorphic Cayley graphs**).

1. Show that there exist finite generating sets S of $\mathbb{Z}$ and T of D_∞ with $\operatorname{Cay}(\mathbb{Z}, S) \cong \operatorname{Cay}(D_\infty, T)$.
2. Show that there exist finite generating sets S of $\mathbb{Z} \times \mathbb{Z}/2$ and T of D_∞ with $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2, S) \cong \operatorname{Cay}(D_\infty, T)$.

Exercise 3.E.22 (Isomorphic Cayley graphs?!***). Cay and Ley obtained their bachelor degree in "Evaluation of children's behaviour and transport logistics (EU-directive *St. Nicholas*)" and are now discussing how to arrange the reindeers in front of their sleighs:

Cay Ha, I patented all arrangements of reindeers and cords between them that look like Cayley graphs of $\mathbb{Z} \times \mathbb{Z}/2$ (with respect to finite generating sets)!

Ley Why would I care? My sleigh is faster anyway – I use a Cayley graph of $\mathbb{Z}$.

Cay Uh-oh, that will be expensive for you; you'll have to pay serious licence fees to me: I'll easily find a finite generating set S of $\mathbb{Z} \times \mathbb{Z}/2$ such that your sluggish sleigh arrangement is $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2, S)$.

Ley Are you out of your mind? Despite of $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2$ being quasi-isometric there are no finite generating sets S and S' of $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2$ respectively such that $\operatorname{Cay}(\mathbb{Z}, S)$ and $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2, S')$ are isomorphic.

Are you able to help?

Hints. Look at the automorphism given by inversion.

Exercise 3.E.23 (Free products with isomorphic Cayley graphs**). Let $n \in \mathbb{N}_{>0}$ and let

$$G_n := \langle a_1, \dots, a_n \mid a_1^n, \dots, a_n^n \rangle \cong (\mathbb{Z}/n)^{*n}.$$

Show that

$$\operatorname{Cay}\bigl(G_n, \{a_1, \dots, a_n\}\bigr) \cong \operatorname{Cay}\bigl(F(a_1, \dots, a_{n-1}), \{a_1, \dots, a_{n-1}, a_1 \cdot \dots \cdot a_{n-1}\}\bigr).$$

Exercise 3.E.24 (Automorphisms of Cayley trees**). Let F be a free group of rank 2 and let $S \subset F$ be a free generating set. Show that the Cayley graph $\operatorname{Cay}(F, S)$ admits uncountably many graph automorphisms.

Exercise 3.E.25 (Cayley graphs of free groups$^\infty$*). Let F and F' be finitely generated free groups. Show that the following are equivalent:

1. The free groups F and F' have the same rank.
2. There exist finite generating sets $S \subset F$ and $S' \subset F'$ such that the graphs $\operatorname{Cay}(F, S)$ and $\operatorname{Cay}(F', S')$ are isomorphic.

Hints. There is a probabilistic proof of this fact (Outlook 3.2.4). However, no elementary, geometric, proof is known.

Chromatic number of groups$^+$

Cayley graphs of groups allow to apply invariants from graph theory to groups, for example the chromatic number. The following exercises will discuss some basic properties and problems related to chromatic numbers of groups, as introduced by Babai [11, 165].

Definition 3.E.1 (Chromatic number). Let $X = (V, E)$ be a graph. Let C be a set. A *colouring of X by C* is a map $c \colon V \longrightarrow C$ satisfying

$$\forall_{\{v,w\} \in E} \;\; c(v) \neq c(w).$$

The *chromatic number* $\operatorname{ch}(X)$ of X is the smallest $n \in \mathbb{N}$ such that X admits a colouring by $\{1, \dots, n\}$.

Quick check 3.E.26 (Chromatic number of small graphs*). Let $n, m \in \mathbb{N}$.

1. Is $\operatorname{ch}(K_n) = n$?
2. Is $\operatorname{ch}(K_{n,m}) = \min(n, m)$?

Definition 3.E.2 (Chromatic number of a group). Let G be a group. Then the *chromatic number* $\operatorname{ch}(G)$ *of G* is defined by

$$\operatorname{ch}(G) := \inf\bigl\{\operatorname{ch}(\operatorname{Cay}(G, S)) \bigm| S \subset G \text{ generates } G\bigr\} \in \mathbb{N} \cup \{\infty\}.$$

Quick check 3.E.27 (High chromatic numbers?*).

1. Does S_6 contain a generating set S with $\operatorname{ch}(\operatorname{Cay}(S_6,)) \geq 2017$?
2. Does $\mathbb{Z}$ contain a generating set S with $\operatorname{ch}(\operatorname{Cay}(\mathbb{Z}, S)) \geq 2017$?

Exercise 3.E.28 (Chromatic number of small groups*). Let $n \in \mathbb{N}$.

1. Determine $\operatorname{ch}(\mathbb{Z}/n)$.
2. Determine $\operatorname{ch}(\mathbb{Z}^n)$.

Exercise 3.E.29 (Groups with chromatic number 2 ** [11]). Let G be a group.

1. Let $N \subset G$ be a normal subgroup. Show that $\operatorname{ch} G \leq \operatorname{ch}(G/N)$.
2. Show that $\operatorname{ch}(G) = 2$ if and only if G contains a subgroup of index 2.

Exercise 3.E.30 (Chromatic number of finitely generated groups***). Let G be a finitely generated group.

1. Show that

$$\operatorname{ch} G = \min\bigl\{\operatorname{ch}(\operatorname{Cay}(G,S)) \bigm| S \subset G \text{ generates } G \text{ and } S \text{ is finite}\bigr\}.$$

2. Let $S \subset G$ be a finite generating set of G. Show that

$$\operatorname{ch}\bigl(\operatorname{Cay}(G,S)\bigr) = \max\bigl\{\operatorname{ch} X \bigm| X \text{ is a finite subgraph of } \operatorname{Cay}(G,S)\bigr\}.$$

Exercise 3.E.31 (Groups with large chromatic number?!$^\infty$*). Let $n \in \mathbb{N}$. Does there exist a finitely generated group G with $\operatorname{ch} G \geq n$?!

Hints. This is an open problem. Reasonable answers exist for the expected chromatic number of random generating sets [4].

4 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 group actions, which can be viewed as a generalisation of seeing groups as symmetry groups. We start by recalling some basic concepts about group actions (Chapter 4.1).

As we have seen, free groups can be characterised combinatorially as the groups admitting trees as Cayley graphs (Chapter 3.3). In Chapter 4.2, we will prove that this characterisation can be promoted 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! (Chapter 4.2.3)

Another group action tool that helps to recognise that a given group is free is the ping-pong lemma (Chapter 4.3); this is particularly useful when proving that certain matrix groups are free – which also is a purely algebraic statement (Chapter 4.4).

Overview of this chapter

4.1 Review of group actions

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

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 $G \longrightarrow \operatorname{Aut}_C(X)$. 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*:

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

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

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

  In other words: group actions are a concept generalising automorphism and symmetry groups.
- Let G be a group and let X be a set. If $\varrho \colon G \longrightarrow \operatorname{Aut}_{\mathrm{Set}}(X)$ 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 $\varrho \colon G \longrightarrow \operatorname{Aut}_{\mathrm{Set}}(X)$ 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

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

  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); in these cases, automorphisms are equivalence classes of maps of sets and composition of morphisms is performed 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. Further introductory material on group actions and symmetry can be found in Armstrong's book [6].

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

$$\begin{aligned} G &\longrightarrow S_G = \operatorname{Aut}_{\mathsf{Set}}(G) \\ g &\longmapsto (h \mapsto g \cdot h) \end{aligned}$$

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

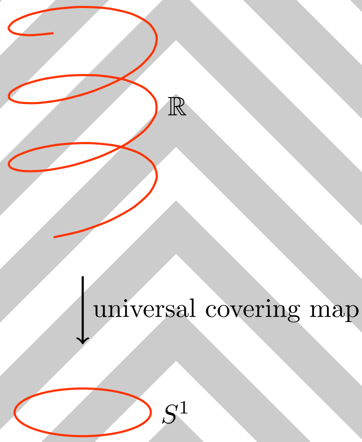

Figure 4.1.: Universal covering of S^1

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

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

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. Moreover, the centre of the square is fixed by *all* isometries of the square.

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 universal covering space $\widetilde{X}$, a path-connected space covering X that has trivial fundamental group [115, Chapter V].

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) [115, 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 the free group 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 glueing 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

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

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 determine 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 $\mathrm{Cay}(G, S)$. □

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).

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 [115, Chapter V]. It is worthwhile to check this assertion in the cases of Example 4.1.7.

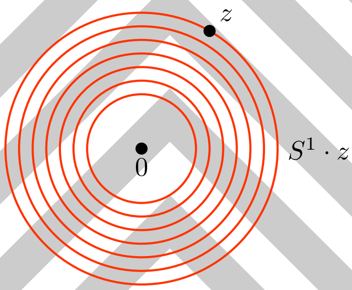

Figure 4.2.: Orbits of the rotation action of S^1 on $\mathbb{C}$

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 $Q = [0,1] \times [0,1]$ 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)**.** *Every 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. This can be shown by looking at a "minimal" orbit of a vertex and paths between vertices of this orbit (Exercise 4.E.7). □

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

1. *If $x \in X$, then the map*

$$A_x \colon G/G_x \longrightarrow G \cdot x$$
$$g \cdot G_x \longmapsto g \cdot x$$

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 values on cosets do 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 also 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, $g_1 \cdot G_x = g_1 \cdot (g_1^{-1} \cdot g_2) \cdot G_x = g_2 \cdot G_x$. 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 left hand side: We know $|G/G_x| \cdot |G_x| = |G|$ because every coset of G_x in G has the same size as G_x; therefore, using the first part, we obtain

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

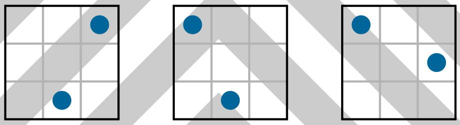

Figure 4.3.: The same punchcard seen from different sides/angles

4.1.3 Application: Counting via group actions

Group actions can be used to solve counting problems. A standard example from algebra is the proof of the Sylow theorems in finite group theory: in this proof, the conjugation action of a group on the set of certain subgroups is considered [6, Chapter 20]. Moreover, group actions also provide a convenient means to organise the proof of normal form theorems for amalgamated free products and HNN-extensions [150, Chapter 11]

Another class of examples arises in combinatorics:

Example 4.1.18 (Counting punchcards [3])**.** 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 $D_4 = \langle s,t \,|\, t^2, s^4, tst^{-1} = s^{-1}\rangle$ (Example 2.2.20) the number of configurations fixed by this element. Taking into account 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
$\overline{t}$, $\overline{s}^2 \cdot \overline{t}$	$3+3=6$
$\overline{s} \cdot \overline{t}$, $\overline{t} \cdot \overline{s}$	$3+3=6$

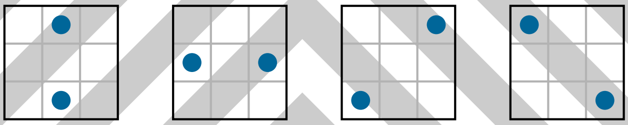

Figure 4.4.: The punchcards fixed by rotation around π

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 Proposition 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.

4.1.4 Transitive actions

Transitive actions on "connected spaces" yield generating sets through "close neighbours". A first instance of this general principle is Proposition 4.1.20. A metric version of this principle is the Švarc-Milnor lemma (Chapter 5.4).

Definition 4.1.19 (Transitive action on a set)**.** A group action on a set is *transitive* if it has at most one orbit.

For example, if G is a group and $S \subset G$ is a generating set, then the left translation action of G on the vertices of $\operatorname{Cay}(G, S)$ is transitive (and free). In fact, this property can be used to characterise Cayley graphs in terms of actions on graphs:

Proposition 4.1.20 (Actions on graphs yield Cayley graphs)**.** *Let G be a group and let G act on a connected graph $X = (V, E)$ by graph automorphisms. If this action is free and transitive on V of vertices of X and if $x \in V$, then the set*

$$S := \bigl\{ s \in G \bigm| \{x, s \cdot x\} \in E \bigr\}$$

generates G and the Cayley graph $\operatorname{Cay}(G, S)$ is isomorphic to X.

Proof. As first step, we use the action on the vertices to identify V with the acting group G: Because the G-action on V is free and transitive, the map

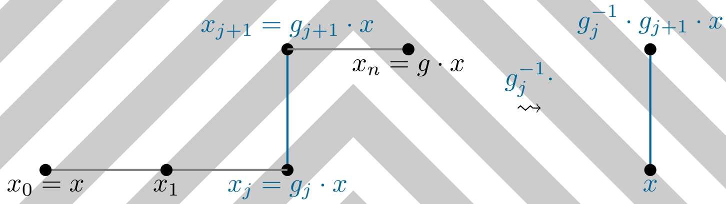

Figure 4.5.: From paths to words, using a transitive action

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

is bijective (Proposition 4.1.17). As second step, we show that the set S indeed is a generating set of G: Let $g \in G$. Because the graph X is connected, there is a path $x_0 = x, x_1, \dots, x_n = g \cdot x$ in X joining x and $g \cdot x$. We now only need to translate this path into steps in the group G, using the action of G on X (Figure 4.5). For all $j \in \{0, \dots, n-1\}$ we let

$$g_j := \varphi^{-1}(x_j) \in G$$

and

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

Why is $s_j \in S$? Because $x_0, \dots, x_n$ is a path in X, we know that $\{x_j, x_{j+1}\}$ is an edge of X. As G acts by graph automorphisms, also $\{g_j^{-1} \cdot x_j, g_j^{-1} \cdot x_{j+1}\}$ is an edge of X. By construction,

$$g_j^{-1} \cdot x_j = \bigl(\varphi^{-1}(x_j)\bigr)^{-1} \cdot x_j = \bigl(\varphi^{-1}(x_j)\bigr)^{-1} \cdot \varphi^{-1}(x_j) \cdot x = x$$

and

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

Thus, $s_j \in S$. On the other hand, by definition,

$$\begin{aligned} g = g_n &= g_0 \cdot g_0^{-1} \cdot g_1 \cdot \dots \cdot g_{n-1} \cdot g_{n-1}^{-1} \cdot g_n \\ &= e \cdot s_0 \cdot s_1 \cdot \dots \cdot s_{n-1}. \end{aligned}$$

Therefore, S is a generating set of G.

It remains to prove that $\mathrm{Cay}(G, S)$ is isomorphic to the given graph X. To this end, we prove that φ induces such a graph isomorphism: We already know that φ is bijective on the vertices. What about the edges? Let $g, h \in G$. Then

$$\bigl\{\varphi(g), \varphi(h)\bigr\} = \{g \cdot x, h \cdot x\}.$$

Because G acts by graph automorphisms on X, this set is an edge of X if and only if

$$\{x, g^{-1} \cdot hx\} = \{g^{-1} \cdot (g \cdot x), g^{-1} \cdot (h \cdot x)\}$$

is an edge of X. By construction of S, this is equivalent to $g^{-1} \cdot h \in S$, whence equivalent to $\{g, h\}$ being an edge of $\mathrm{Cay}(G, S)$. Hence, $\mathrm{Cay}(G, S)$ is isomorphic to the graph X. □

Outlook 4.1.21 (Cayley complex)**.** Let G be a group. By Proposition 4.1.20, a Cayley graph of G is nothing but a connected graph with a G-action whose induced action on the vertices is free and transitive. Cayley complexes are a two-dimensional version of this concept: A *Cayley complex* of G is a simply connected two-dimensional CW-complex with a cellular action by G such that the induced action on the vertices is free and transitive. The condition of being simply connected is a higher version of connectedness from algebraic topology (Definition A.1.3) and CW-complexes are a topological higher-dimensional generalisation of graphs.

If $\langle S \mid R \rangle$ is a presentation of G, then the universal covering of the presentation complex of $\langle S \mid R \rangle$ is a Cayley complex of G (Outlook 3.2.5). Moreover, for every generating set S of G there is a Cayley complex of G such that its 1-skeleton corresponds (almost) to $\mathrm{Cay}(G, S)$ [45, Chapter 2.2].

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 a set $S \subset F$; then the Cayley graph $\mathrm{Cay}(F, S)$ is a tree by Theorem 3.3.1. We consider the left translation action of F on $\mathrm{Cay}(F, S)$.

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 elements of order 2; therefore, the left translation action of F on $\mathrm{Cay}(F, S)$ is free by Proposition 4.1.10. □

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

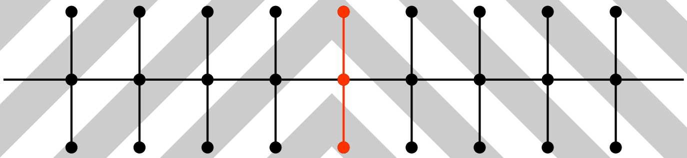

Figure 4.6.: A spanning tree (red) for a shift action of $\mathbb{Z}$

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 (equivariant) spanning trees, which we will discuss in the following section.

4.2.1 Spanning trees for group actions

Spanning trees for group actions are a natural generalisation of spanning trees of graphs:

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

Example 4.2.3 (Spanning trees)**.** We consider the action of $\mathbb{Z}$ by "horizontal" shifting on the (infinite) tree depicted in Figure 4.6. 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.*

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 of 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 every totally ordered chain of T_G has an upper bound in T_G (namely the "union" over all trees in this chain). By Zorn's lemma, 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

Figure 4.7.: Contracting the spanning tree and all its translates to vertices (red), in the situation of Example 4.2.3

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. Because the path p' is shorter than the path p, 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 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. □

4.2.2 Reconstructing a Cayley tree

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 4.7 shows this in the situation of Example 4.2.3); here, $g \cdot T'$ denotes the subgraph of T obtained by translating T' by g. This idea of contracting T' can be made precise and concludes the proof with an application of Proposition 4.1.20 (Remark 4.2.5). However, we prefer to proceed directly in the original tree T:

As in Proposition 4.1.20, the candidate for a generating set comes 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

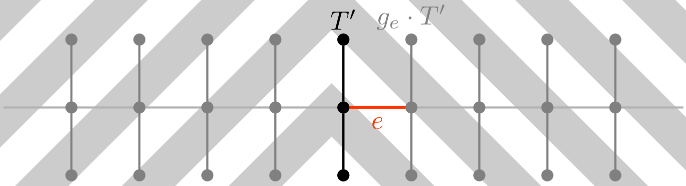

Figure 4.8.: An essential edge (red) for the shift action of $\mathbb{Z}$ (Example 4.2.3)

to T' (then the other vertex cannot belong to T' as well, by uniqueness of paths in trees (Proposition 3.1.10)).

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 $g_e \cdot T'$. 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:

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 4.9).

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

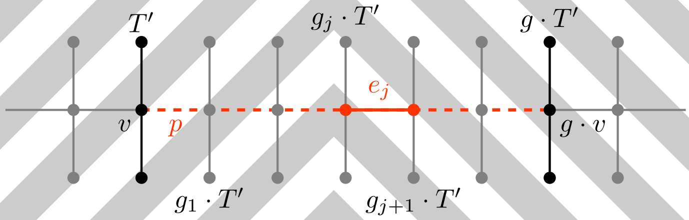

Figure 4.9.: The set $\widetilde{S}$ generates G

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

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

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

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 $\operatorname{Cay}(G, \widetilde{S})$).

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 $\operatorname{Cay}(G, S)$ 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 $\operatorname{Cay}(G, S) = \operatorname{Cay}(G, \widetilde{S})$. 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 4.10). 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 the hypothesis that T is a tree. □

Remark 4.2.5 (Contracting a spanning tree)**.** The contraction construction idea in the previous proof can also be implemented as follows: Let $X = (V, E)$ be the following graph:

- We set
$$V := \{g \cdot T' \mid g \in G\}$$
- and

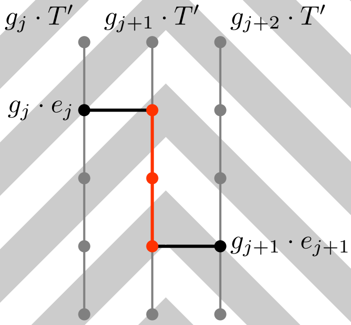

Figure 4.10.: Cycles in $\operatorname{Cay}(G, \widetilde{S})$ lead to cycles in T by connecting translates of essential edges in the corresponding translates of T' (red path)

$$E := \bigl\{\{g \cdot T', h \cdot T'\} \bigm| g, h \in G \text{ and } g \cdot T' \text{ and } h \cdot T' \text{ are adjacent in } T\bigr\}.$$

For $g, h \in G$ we call $g \cdot T'$ and $h \cdot T'$ *adjacent in* T, if $g \cdot T' \neq h \cdot T'$ and there exist vertices $v \in g \cdot T'$ and $w \in h \cdot T'$ such that $\{v, w\}$ is an edge of T.

One then establishes the following facts (Exercise 4.E.11):

1. For all $g, h \in G$ with $g \neq h$, the copies $g \cdot T'$ and $h \cdot T'$ have no common vertex. If $g \cdot T'$ and $h \cdot T'$ are adjacent, then there is a unique connecting edge between $g \cdot T'$ and $h \cdot T'$.
2. The graph X is a tree.
3. The graph X admits a G-action that is free and transitive on the set V of vertices.

Therefore, Proposition 4.1.20 (applied to the action of G on X) shows that G admits a generating set $\widetilde{S}$ such that the Cayley graph $\operatorname{Cay}(G, \widetilde{S})$ is a tree (namely, the tree X). Because G acts freely on the tree T, the group G cannot contain any elements of order 2 (Proposition 4.1.16). Therefore, one can select a generating set $S \subset \widetilde{S}$ that satisfies the hypotheses of Theorem 3.3.3.

Remark 4.2.6 (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 [115, Chapter VI].

Outlook 4.2.7 (Bass-Serre theory)**.** We characterised free groups as those groups that admit free actions on trees. What happens if we relax the freeness condition for actions on trees? The ultimate result regarding actions on trees is given by Bass-Serre theory in terms of so-called graphs of groups and their fundamental groups [159, Chapter 5][66, Chapter 6]. Roughly speaking, if a group G admits an action on a tree (without inversions), then G can be decomposed into groups, where the combinatorics of this decomposition is related to the orbit structure and the corresponding stabilisers of the G-action on the tree.

The simplest cases of such decompositions are amalgamated free products and HNN-extensions. In other words, free products, amalgamated free products, and HNN-extensions also admit characterisations in terms of actions on trees with suitable orbit structures and stabilisers.

4.2.3 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 subgroups in many situations that are algebraically rather inaccessible:

Corollary 4.2.8 (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. □

Example 4.2.9. Free groups do *not* contain subgroups that are isomorphic to $\mathbb{Z}^2$: Let F be a free group and let $H \subset F$ be a subgroup. Then H is free (by the Nielsen-Schreier theorem, Corollary 4.2.8). Because $\mathbb{Z}^2$ is *not* free (Exercise 2.E.11), we obtain that $H \not\cong \mathbb{Z}^2$. We will see a vast, geometric, generalisation of this fact in Corollary 7.5.15.

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 4.2.10 (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 $T := \operatorname{Cay}(F, S)$; 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 therefore T' has exactly $k-1$ edges (Exercise 3.E.4). 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. □

Remark 4.2.11 (Topological proof of the Nielsen-Schreier theorem). A topological version of the proof of the Nielsen-Schreier theorem can be given via covering theory [115, 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 $\overline{X}$ inherits the structure of a one-dimensional CW-complex. On the other hand, every 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 4.2.12. *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. Using a surjective homomorphism $F \longrightarrow \mathbb{Z}$, one can construct subgroups of finite index in F. We can then apply the quantitative version of the Nielsen-Schreier theorem to obtain free subgroups of large rank (Exercise 4.E.12). □

Outlook 4.2.13 (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 (Remark 2.3.12) 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 [62] and I. Mineyev [121, 122], in 2011. A short, rather elementary, proof was provided by W. Dicks [47].

Corollary 4.2.14. *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 $[F(S) : H'] = [G : H]$.

By Corollary 4.2.10, 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 4.2.14 via the Švarc-Milnor lemma (Corollary 5.4.5).

Corollary 4.2.15 (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 4.11 for the free product $\mathbb{Z}/2 * \mathbb{Z}/3$). Using the description of the free product $G * H$ in terms of reduced words (Outlook 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

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

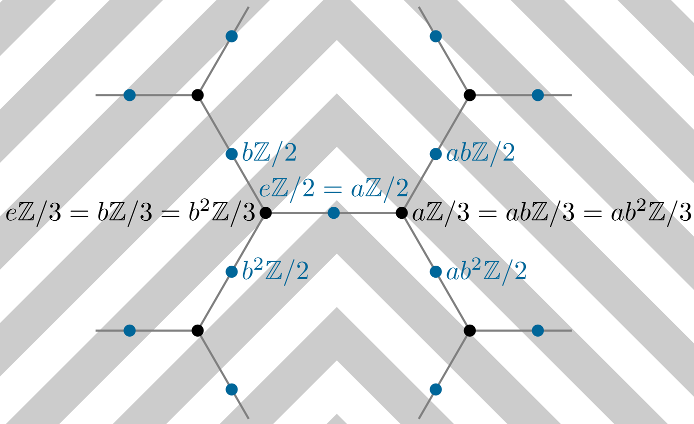

Figure 4.11.: The tree for the free product $\mathbb{Z}/2 * \mathbb{Z}/3 \cong \langle a, b \,|\, a^2, b^3 \rangle$

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, every 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 [159].

4.3 The ping-pong lemma

The following sufficient criterion for freeness via suitable actions is due to F. Klein; it should be noted that there are lots of variations of this principle in the literature that all go by the name of ping-pong lemma.

Theorem 4.3.1 (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 contained in A and such that for all $n \in \mathbb{Z} \setminus \{0\}$ we have*

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

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 $F_{\mathrm{red}}(\{\alpha, \beta\})$

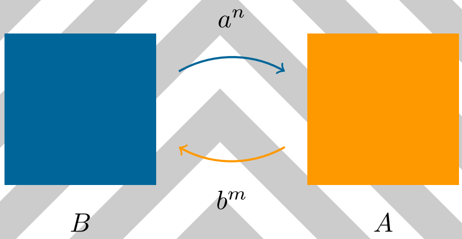

Figure 4.12.: The ping-pong lemma

there is a group homomorphism $\varphi\colon F_{\text{red}}(\{\alpha,\beta\}) \longrightarrow G$ 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 4.12)
$$\begin{aligned} B = e\cdot B = \varphi(w)\cdot B & \\ = a^{n_0}\cdot b^{m_1}\cdot a^{n_1}\cdot\dots\cdot b^{m_k}\cdot a^{n_k}\cdot B & \\ \subset a^{n_0}\cdot b^{m_1}\cdot a^{n_1}\cdot\dots\cdot b^{m_k}\cdot A & \qquad \text{ping!} \\ \subset a^{n_0}\cdot b^{m_1}\cdot a^{n_1}\cdot\dots\cdot a^{n_{k-1}}\cdot B & \qquad \text{pong!} \\ \subset \cdots & \qquad \vdots \\ \subset a^{n_0}\cdot B & \\ \subset A, & \end{aligned}$$
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 is impossible.

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

Outlook 4.3.2 (Ping-pong lemma for free products)**.** Similarly, using the description of free products in terms of reduced words (Outlook 3.3.8), one can show the following [77, 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 contained 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 $G \cong G_1 * G_2$.

The ping-pong lemma is a standard tool to establish that certain matrix groups are free (Chapter 4.4). Further examples are given in de la Harpe's book [77, 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 (Exercise 4.E.16).

4.4 Free subgroups of matrix groups

Via the ping-pong lemma we can establish that certain matrix groups are free. We will illustrate this first in a simple example in $\mathrm{SL}(2,\mathbb{Z})$ (Chapter 4.4.1), which has applications in graph theory (Chapter 4.4.2); finally, we will briefly discuss the Tits alternative (Chapter 4.4.3).

4.4.1 Application: The group $\mathrm{SL}(2,\mathbb{Z})$ is virtually free

As a first example, we consider the case of the modular group:

Example 4.4.1 (A free subgroup of $\mathrm{SL}(2,\mathbb{Z})$)**.** Let a, $b \in \mathrm{SL}(2,\mathbb{Z})$ be given by

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

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 consider the subsets

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

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 free subgroup of rank 2 has index 12 in $\mathrm{SL}(2,\mathbb{Z})$:

Proposition 4.4.2 ($\mathrm{SL}(2,\mathbb{Z})$ is virtually free)**.** *Let*

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

Then $F := \langle \{a,b\} \rangle_{\mathrm{SL}(2,\mathbb{Z})}$ has finite index in $\mathrm{SL}(2,\mathbb{Z})$. More specifically: We consider the subgroups

$$\begin{aligned} F' &:= \left\{ \begin{pmatrix} 4m+1 & 2r \\ 2s & 4n+1 \end{pmatrix} \;\middle|\; m,n,r,s \in \mathbb{Z},\ \det \begin{pmatrix} 4m+1 & 2r \\ 2s & 4n+1 \end{pmatrix} = 1 \right\} \\ G &:= \left\{ \begin{pmatrix} 2m+1 & 2r \\ 2s & 2n+1 \end{pmatrix} \;\middle|\; m,n,r,s \in \mathbb{Z},\ \det \begin{pmatrix} 2m+1 & 2r \\ 2s & 2n+1 \end{pmatrix} = 1 \right\} \end{aligned}$$

of $\mathrm{SL}(2,\mathbb{Z})$.

1. *Then* $[G : F'] = 2$ *and* $\bigl[\mathrm{SL}(2,\mathbb{Z}) : G\bigr] = 6$. *In particular,*

$$\bigl[\mathrm{SL}(2,\mathbb{Z}) : F'\bigr] = 12.$$

2. *Moreover,* $F = F'$.

Proof. Ad 1. We first show that $[G : F'] = 2$: Let

$$x := \begin{pmatrix} 2m+1 & 2r \\ 2s & 2n+1 \end{pmatrix} \in G.$$

Considering the determinant condition modulo 4 shows that either both m and n are odd, or they are both even. If m and n are even, then clearly $x \in F'$. In the case that m and n are odd, a straightforward calculation shows that $(-E_2) \cdot x \in F'$, where E_2 denotes the unit 2×2-matrix. Thus,

$$\{g \cdot F' \mid g \in G\} = F' \sqcup (-E_2) \cdot F',$$

which proves that $[G : F'] = 2$.

We now show that $\bigl[\mathrm{SL}(2,\mathbb{Z}) : G\bigr] = 6$: By definition, G is the kernel of the homomorphism $\mathrm{SL}(2,\mathbb{Z}) \longrightarrow \mathrm{SL}(2,\mathbb{Z}/2)$ given by reduction modulo 2. Therefore,

$$\bigl[\mathrm{SL}(2,\mathbb{Z}) : G\bigr] = \bigl|\mathrm{SL}(2,\mathbb{Z}/2)\bigr|.$$

A simple counting argument shows that $\bigl|\mathrm{SL}(2,\mathbb{Z}/2)\bigr| = 6$ (Exercise 4.E.18).

Elementary group theory shows that the index is multiplicative with respect to intermediate groups and so $\bigl[\mathrm{SL}(2,\mathbb{Z}) : F'\bigr] = 6 \cdot 2 = 12$.

Ad 2. A straightforward induction shows that $F \subset F'$. Why does also the converse inclusion $F' \subset F$ hold? Let

$$x = \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix} \in F'.$$

We show by induction over $\min(|x_{11}|, |x_{12}|)$ that $x \in F$, using the following arguments:

- *Base case.* If $x_{12} = 0$, then the determinant 1 condition implies that $x_{11} = 1 = x_{22}$, and so $x \in \langle b \rangle_{\mathrm{SL}(2,\mathbb{Z})} \subset F$.
- *Induction step I.* If $|x_{12}| \geq |x_{11}|$, then we proceed as follows: We use integer division to find $k \in \mathbb{Z}$ and $R \in \{0, \dots, |2x_{11}| - 1\}$ with

  $$x_{12} + |x_{11}| = -k \cdot 2x_{11} + R.$$

  We then consider the matrix

  $$x' := x \cdot a^k = \begin{pmatrix} x_{11} & x_{12} + 2k \cdot x_{11} \\ x_{21} & x_{22} + 2k \cdot x_{21} \end{pmatrix}.$$

  By construction, we have $x' \in x \cdot F$ and

  $$\begin{aligned} |x'_{12}| &= |x_{12} + 2k \cdot x_{11}| = \bigl|-|x_{11}| - 2k \cdot x_{11} + R + 2k \cdot x_{11}\bigr| \\ &\leq |x_{11}| = |x'_{11}|. \end{aligned}$$

  Moreover, parity shows that $|x'_{12}| \neq |x'_{11}|$ and so $|x'_{12}| < |x'_{11}|$. In particular, $\min(|x'_{11}|, |x'_{12}|) = |x'_{12}| < |x'_{11}| = |x_{11}| = \min(|x_{11}|, |x_{12}|)$.
- *Induction step II.* Similarly, if $|x_{12}| < |x_{11}|$, then we can find $k \in \mathbb{Z}$ and $R \in \{0, \dots, |2x_{12}| - 1\}$ with $x_{11} + |x_{12}| = -k \cdot 2x_{12} + R$. We consider

the matrix

$$x' := x \cdot b^k \in x \cdot F$$

and obtain $\min(|x'_{11}|, |x'_{12}|) < \min(|x_{11}|, |x_{12}|)$, similar to the previous case.

Therefore, inductively, we obtain that $x \in F$. □

The fact that free groups can be embedded into $\mathrm{SL}(2, \mathbb{Z})$ also has other interesting group-theoretic consequences for free groups; for example, finitely generated free groups can be approximated in a reasonable way by finite groups (Exercise 4.E.26).

Outlook 4.4.3 ($\mathrm{SL}(2,\mathbb{Z})$ as amalgamated free product)**.** The discussion above can be extended to prove the following fact [159, Example I.4.2]: Let

$$G_1 := \left\langle \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix} \right\rangle_{\mathrm{SL}(2,\mathbb{Z})}, \quad G_2 := \left\langle \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \right\rangle_{\mathrm{SL}(2,\mathbb{Z})}, \quad A := G_1 \cap G_2.$$

Then $G_1 \cong \mathbb{Z}/6$, $G_2 \cong \mathbb{Z}/4$, and $A \cong \mathbb{Z}/2$ and the inclusions of A into G_1 and G_2 (and into $\mathrm{SL}(2,\mathbb{Z})$) induce an isomorphism

$$\mathbb{Z}/6 *_{\mathbb{Z}/2} \mathbb{Z}/4 \cong G_1 *_A G_2 \cong \mathrm{SL}(2,\mathbb{Z}).$$

4.4.2 Application: Regular graphs of large girth

We will now discuss a graph-theoretic application of Example 4.4.1, namely the construction of regular graphs with few vertices and large girth.

Definition 4.4.4 (Girth)**.** The *girth* $g(X)$ of a graph X is the length of a shortest cycle in X. By definition, forests have infinite girth.

Example 4.4.5 (Girth of basic graphs)**.** If $n \in \mathbb{N}_{\geq 3}$, then the complete graph K_n satisfies $g(K_n) = 3$ and $g(\mathrm{Cay}(\mathbb{Z}/n, \{[1]\})) = n$.

It is a classical construction problem from graph theory to find graphs of large girth that satisfy additional constraints. A prominent example is the probabilistic proof [20] of the existence of finite graphs of large girth and large chromatic number (Definition 3.E.1), which shows that colouring graphs indeed is a global problem. A first, constructive, step in this direction is Mycielski's iterated graph construction [127] (Exercise 4.E.23). Another construction problem of this type is to exhibit regular graphs of large girth with "few" vertices.

Margulis [113] solved this problem, using Cayley graphs; for simplicity, we only treat the case of 4-regular graphs.

Theorem 4.4.6 (Regular graphs of large girth)**.** *Let $N \in \mathbb{N}_{\geq 5}$. Then there exists a graph X_N with the following properties:*

- *The graph X_N is 4-regular.*
- *The graph X_N has at most N^3 vertices.*
- *The graph X_N satisfies*

$$g(X_N) \geq 2 \cdot \log_d \frac{N}{2} - 1,$$

where $d := 1 + \sqrt{2}$.

It should be noted that this result is asymptotically optimal in the sense that the girth grows at most logarithmically in the vertices (Exercise 4.E.22). Moreover, the proof by Margulis is constructive.

Proof. We construct the desired graphs explicitly: Let

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

be the matrices in $\mathrm{SL}(2,\mathbb{Z})$ of Example 4.4.1. Then $F := \langle \{a,b\} \rangle_{\mathrm{SL}(2,\mathbb{Z})}$ is free of rank 2. For $N \in \mathbb{N}_{\geq 5}$ we consider the homomorphism

$$\varphi_N \colon \mathrm{SL}(2,\mathbb{Z}) \longrightarrow \mathrm{SL}(2,\mathbb{Z}/N)$$

given by reduction modulo N and we set

$$a_N := \varphi_N(a) \qquad \text{and} \qquad b_N := \varphi_N(b).$$

Moreover, we define

$$\begin{aligned} G_N &:= \langle a_N, b_N \rangle_{\mathrm{SL}(2,\mathbb{Z}/N)} \subset \mathrm{SL}(2,\mathbb{Z}/N), \\ X_N &:= \mathrm{Cay}(G_N, \{a_N, b_N\}). \end{aligned}$$

In the following, we will show that this graph X_N has the claimed properties: By construction, X_N is 4-regular and X_N has at most

$$\bigl|\mathrm{SL}(2,\mathbb{Z}/N)\bigr| \leq N^3$$

vertices. Therefore, it remains to prove the lower girth bound: To this end, we consider two different paths in X_N having the same start and endpoints, i.e., we consider reduced words $w, v \in F_{\mathrm{red}}(\alpha,\beta)$ such that $w \neq v$ and

$$\varphi_N(\overline{w}) = \varphi_N(\overline{v}) \quad \text{in } G_N;$$

here, $\alpha \neq \beta$, and $\overline{w}, \overline{v} \in F$ denote the images of w and v respectively under the canonical homomorphism $F_{\mathrm{red}}(\alpha,\beta) \longrightarrow F$ given by $\alpha \mapsto a, \beta \mapsto b$. In other words, we evaluate w and v on a, b and on a_N, b_N. By definition of $g(X_N)$, we may assume that the lengths $m, n \in \mathbb{N}$ of w and v respectively satisfy

$$g(X_N) = m + n \qquad \text{and} \qquad \max(m, n) \leq \frac{m+n+1}{2}.$$

Because F is free and $w \neq v$, we obtain $\overline{w} \neq \overline{v}$. Let

$$c := \overline{w} - \overline{v} \in M_{2\times 2}(\mathbb{Z}).$$

Then $c \neq 0$, but the reduction c_N of c in $M_{2\times 2}(\mathbb{Z}/N)$ is $\varphi_N(\overline{w}) - \varphi_N(\overline{v}) = 0$. Therefore, all entries of c are divisible by N, i.e., there exists $c' \in M_{2\times 2}(\mathbb{Z})$ with

$$c = N \cdot c'.$$

In particular, we obtain the following estimates for the operator norms (where we consider the action of $\mathrm{SL}(2, \mathbb{Z})$ on $\mathbb{R}^2$ by matrix multiplication)

$$\|\overline{w}\| + \|\overline{v}\| \geq \|c\| = N \cdot \|c'\| \geq N \cdot \max\bigl(|c'_{11}|, |c'_{12}|, |c'_{21}|, |c'_{22}|\bigr) \geq N;$$

the last inequality follows because the entries of c' are integral and $c' \neq 0$. On the other hand, a straightforward calculation shows

$$\|a\| \leq 1 + \sqrt{2} = d \qquad \text{and} \qquad \|b\| \leq d.$$

Therefore, we obtain

$$N \leq \|c\| \leq d^m + d^n \leq 2 \cdot d^{\frac{m+n+1}{2}} = 2 \cdot d^{\frac{g(X_N)+1}{2}},$$

which gives the desired lower bound for $g(X_N)$. □

4.4.3 Application: The Tits alternative

In contrast to $\mathrm{SL}(2, \mathbb{Z})$, for $n \in \mathbb{N}_{\geq 3}$, the groups $\mathrm{SL}(n, \mathbb{Z})$ do *not* contain a free group of finite index (Exercise 4.E.19). But non-Abelian free groups appear frequently as building blocks in linear groups; more precisely, J. Tits [172] discovered the following:

Theorem 4.4.7 (Tits alternative). *For all fields K and all $n \in \mathbb{N}_{\geq 1}$ the following holds: If G is a finitely generated 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.*

Solvable groups are discussed in more detail in Chapter 6.3.1; the definition of solvability is recalled in Definition 6.3.3.

In the following, we will sketch the main steps of the proof of the Tits alternative. As we will see below, a complete proof requires more machinery and background in linear algebraic groups and number theory; a detailed proof can be found in the book by Druţu and Kapovich [53]. The proof of the Tits alternative consists of the following components:

- The two alternatives exclude each other (Exercise 4.E.20).
- Recognition of free groups via the ping-pong lemma.
- Eigenvalue analysis to set up the ping-pong lemma.

In Example 4.4.1 we have seen one way to find a free subgroup of rank 2 in $\mathrm{SL}(2,\mathbb{Z})$. However, this example of free linear groups does not generalise well to larger classes of groups. Therefore, we consider a slightly different type of examples:

Example 4.4.8 (Another free linear group). Let $\lambda \in \mathbb{C}$. We consider the matrices

$$a := \begin{pmatrix} \lambda & 0 \\ 0 & 1 \end{pmatrix}, \qquad b := c \cdot a \cdot c^{-1}, \qquad c := \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

in $\mathrm{GL}(2,\mathbb{C})$ and the action of $\mathrm{GL}(2,\mathbb{C})$ on $\mathbb{C}^2$ by matrix multiplication. A straightforward calculation then shows that the subsets

$$\begin{aligned} A &:= \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{C}^2 \;\middle|\; \frac{|x|}{|y|} \in (1-\varepsilon, 1+\varepsilon) \right\}, \\ B &:= c^{-1} \cdot A \end{aligned}$$

satisfy the condition of the ping-pong lemma (Theorem 4.3.1) provided that $|\lambda|$ is large enough and $\varepsilon \in \mathbb{R}_{>0}$ is small enough (Exercise 4.E.21). Hence, in this case the subgroup $\langle a, b \rangle_{\mathrm{GL}(2,\mathbb{C})}$ is free of rank 2.

Of course, we can argue similarly if a has eigenvalues λ_1 and λ_2 with big ratio $|\lambda_1|/|\lambda_2|$.

In this example, the attracting/repelling nature of eigenspaces is essential; a convenient way to describe this phenomenon is to pass to the corresponding action on projective space and to formulate the attraction/repelling properties for the *points* in projective space associated with the one-dimensional eigenspaces of a (and b).

This example suggests that the structure of the spectrum of the matrices in question plays a central role in finding free subgroups in matrix groups.

As an experiment, let us try to prove the Tits alternative "by hand" for finitely generated subgroups G of $\mathrm{GL}(2,\mathbb{C})$. If G does not contain a solvable subgroup of finite index, then as suggested above, we will try to find a matrix in G with a unique largest eigenvalue and a suitable conjugate of this matrix. Looking at the possible Jordan normal forms shows that a priori it is not entirely clear that we will find a matrix a in G with eigenvalues of different norms, and then that we will find a conjugate b of a such that the eigenspaces of a and b are not related by inclusion. Therefore, even in the simple-looking case of $\mathrm{GL}(2,\mathbb{C})$ input from the theory of linear algebraic groups and of normed fields enters.

Sketch of proof of the Tits alternative over the field $\mathbb{C}$. Let $d \in \mathbb{N}$ and let $G \subset \mathrm{GL}(d,\mathbb{C})$ be a finitely generated group that does not contain a solvable subgroup of finite index. We now indicate how to find a free subgroup in G of rank 2, using the theory of linear algebraic groups and of normed fields:

1. *Finding diagonalisable elements in G.* Dividing out the solvable radical, one can show that it is sufficient to consider the case that the Zariski closure $\overline{G}$ of G is a semi-simple linear algebraic group.
   The structure theory of semi-simple linear algebraic groups implies:
   - The group $\overline{G}$ is perfect, i.e., $\overline{G} = [\overline{G}, \overline{G}]$; in particular, we may assume that G is a subgroup of $\mathrm{SL}(d, \mathbb{C})$.
   - The elements of finite order in $\overline{G}$ are *not* dense in $\overline{G}$.
   - The diagonalisable elements of $\overline{G}$ contain a dense open subset of $\overline{G}$.

   Therefore, we will find a diagonalisable matrix $a \in G$ of infinite order; in particular, we have $\det a = 1$ (so, a has at least two different eigenvalues) and one of the eigenvalues of a is *not* a root of unity.
2. *Finding a unique large eigenvalue.* As far as we know so far, all of the eigenvalues of a might lie on the unit circle in $\mathbb{C}$. The beautiful idea of Tits is to change the point of view and to consider other normed fields: Let $S \subset G$ be a finite generating set, let $B \subset \mathbb{C}^d$ be an eigenbasis of a, and let k be the field extension of $\mathbb{Q}$ generated by the matrix entries of S and a with respect to B. We can then view $\overline{G}$ as an algebraic subgroup of $\mathrm{GL}(d, k)$.
   By the first step, there is an eigenvalue $\lambda \in k^\times$ of a that is not a root of unity. Because k is a finitely generated extension of $\mathbb{Q}$, there exists an extension of k to a locally compact field k' with absolute value $|\cdot|'$ that satisfies
   $$|\lambda|' \neq 1.$$
   Passing to a^{-1} if necessary we hence may assume that $|\lambda|' > 1$.
   Let μ be a $|\cdot|'$-maximal eigenvalue of a. Passing to suitable exterior powers, we may assume that the eigenspace of μ is one-dimensional and that G acts absolutely irreducibly on ${k'}^d$. By now, we did not only change the field of definition but possibly also the linear representation of our group G !
3. *Finding a unique maximal and a unique minimal eigenvalue.* A careful analysis of conjugates/commutators and Jordan form calculations show (using the element a from the previous step) that the set of elements of G that have a unique maximal and a unique minimal eigenvalue (with respect to $|\cdot|'$) is Zariski dense in G. Therefore, we can find such an element a' that in addition is also diagonalisable over the algebraic closure of k'. In view of absolute irreducibility of the G-action, we can pass to a finite extension of k' so that a' is diagonalisable and the G-action still is absolutely irreducible. Using this irreducibility, one can find a suitable conjugate b' so that the ping-pong lemma can be applied to (large powers of) a' and b' in a similar way as in Example 4.4.8. □

A quantitative version of the Tits alternative was recently established by Breuillard (see Chapter 6.4.3).

4.E Exercises

General actions

Quick check 4.E.1 (Fixed sets of continuous actions*). We consider a group action of a group G on a topological space X by homeomorphisms. Let $H \subset G$ be a subset of G.

1. Is the fixed set X^H always open?
2. Is the fixed set X^H always closed?

Quick check 4.E.2 (Actions on the real line?*).

1. Is there a free isometric action of $\mathbb{Z}/2$ on $\mathbb{R}$?
2. Is there a free isometric action of $\mathbb{Z}/2$ on $\mathbb{R} \setminus \{0\}$?
3. Is there a free isometric action of $\mathbb{Z}/3$ on $\mathbb{R}$?
4. Is there a free isometric action of $\mathbb{Z}^2$ on $\mathbb{R}$?

Exercise 4.E.3 (Fixed points of matrix groups*). Let $n \in \mathbb{N}$, let $a \in \operatorname{GL}(n, \mathbb{R})$, and let $G := \langle a \rangle_{\operatorname{GL}(n,\mathbb{C})}$.

1. Suppose that G acts freely by matrix multiplication on $\mathbb{R}^n \setminus \{0\}$. Show that G then also acts freely by matrix multiplication on $\mathbb{C}^n \setminus \{0\}$.
2. Let $n \geq 2$. Give an example of a non-trivial element $a \in \operatorname{SL}(n, \mathbb{R})$ such that G acts freely on $\mathbb{R}^n \setminus \{0\}$.

Exercise 4.E.4 (Conjugation of permutations**). Determine the fixed sets and stabiliser groups of the conjugation action of S^3 on itself.

Exercise 4.E.5 (Rubik's cube**). Model playing with Rubik's cube by a suitable group action. Give some examples of interesting group elements in this group.

Exercise 4.E.6 (Essentially free actions**). Let G be a countable group acting on a probability space (X, μ) by measure-preserving measurable isomorphisms. Such an action is *essentially free* if for all $g \in G \setminus \{e\}$ we have $\mu(X^g) = 0$. We consider the shift action of $\mathbb{Z}$ on the infinite product $X := \bigotimes_{\mathbb{Z}} \bigl(\{0,1\}, 1/2\delta_0 + 1/2\delta_1\bigr)$. I.e., X is the probability space modelling a bi-infinite sequence of independent coin tosses with a fair two-sided coin.

1. Show that this action is *not* free.
2. Show that this action is essentially free.

Actions on graphs

Exercise 4.E.7 (Actions of finite groups on trees**). Prove (without using the characterisation of free groups in terms of free actions on trees) that every

action of a finite group on a non-empty tree has a global fixed point (i.e., a vertex or an edge on which all group elements act trivially).
Hints. Consider a "minimal" orbit of a vertex and paths between vertices of this orbit.

Quick check 4.E.8 (Free actions on graphs?*). Which groups admit free actions on the graphs in Figure 4.13 ?

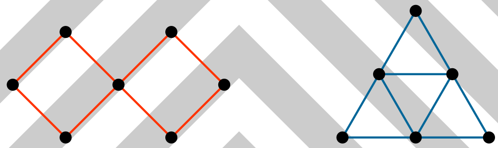

Figure 4.13.: Do these graphs admit interesting free actions?

Quick check 4.E.9 (Actions on trees*).
1. Is every action of a free group on a tree free?
2. Suppose that a free group acts freely on a graph X. Is X then a tree?

Quick check 4.E.10 (Spanning trees for group actions*).
1. Sketch a spanning tree for the action of the group $\mathbb{Z}$ on the Cayley graph $\operatorname{Cay}(F(\{a,b\}),\{a,b\})$, where the action is given by left translation by the powers of a. Are all spanning trees of this action isomorphic?
2. The group $\mathbb{Z}/4$ acts on $\operatorname{Cay}(\mathbb{Z}^2,\{(1,0),(0,1)\})$ via rotation by π; i.e., the generator $[1] \in \mathbb{Z}/4$ acts by rotation by π around 0. Sketch a spanning tree for this action. Are all spanning trees of this action isomorphic?

Exercise 4.E.11 (Contracting a spanning tree**). Fill in the details in Remark 4.2.5 to complete this version of the proof of Theorem 4.2.1.

Exercise 4.E.12 (Subgroups of large rank**). Let F be a free group of rank at least 2. Prove that for every $n \in \mathbb{N}$ there is a free subgroup of F that has rank at least n.

Exercise 4.E.13 (Rank gradient**). The *rank* $\operatorname{rk} G$ of a group G is the minimal cardinality of a generating set of G. The *rank gradient* $\operatorname{rg} G$ of a finitely generated group G is defined by

$$\operatorname{rg} G := \inf_{H \in S(G)} \frac{\operatorname{rk} H}{[G : H]},$$

where $S(G)$ denotes the set of all finite index subgroups of G.
1. Determine $\operatorname{rg} \mathbb{Z}^d$ for all $d \in \mathbb{N}$.
2. Determine the rank gradient of finitely generated free groups.

Exercise 4.E.14 (Characterisation of finite cyclic groups**)**.**
1. Find a class C of graphs with the following property: A group is finite cyclic (i.e., generated by an element of finite order) if and only if it admits a free action on some graph in C.
2. Is there such a class of graphs for the set $\{\mathbb{Z}/n \times \mathbb{Z}/n \mid n \in \mathbb{N}_{>0}\}$ of isomorphism types of groups?

Ping-pong

Quick check 4.E.15 (Ping-pong?*)**.** Let G be a group that is generated by the elements a, b. Suppose that there exists an action of G on a set $X = A \sqcup B$, where A and B are non-empty and

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

1. Is G free of rank 2 if a and b have infinite order?
2. Is $G \cong \mathbb{Z}/2 * \mathbb{Z}/2$ if a and b have order 2 ?

Exercise 4.E.16 (Free subgroups of the homeomorphism group of $\mathbb{R}$ ** [77, Example II.29])**.** We consider the homeomorphism

$$\begin{aligned} f \colon [0,1] &\longrightarrow [0,1] \\ t &\longmapsto \begin{cases} 4 \cdot t & \text{if } t \in [0, 1/5] \\ \frac{4}{5} + \frac{1}{4} \cdot \left(t - \frac{1}{5}\right) & \text{if } t \in [1/5, 1] \end{cases} \end{aligned}$$

and the maps

$$\begin{aligned} a \colon \mathbb{R} &\longrightarrow \mathbb{R} \\ t &\longmapsto \lfloor t \rfloor + f(\{t\}) \end{aligned}$$

(where $\lfloor \cdot \rfloor$ denotes the lower integral part and $\{\cdot\} := \mathrm{id} - \lfloor \cdot \rfloor$ denotes the fractional part) and

$$b := c \cdot a \cdot c^{-1},$$

where $c \colon t \mapsto t - 1/2$.
1. Show that a and b are self-homeomorphisms of $\mathbb{R}$.
2. Show that a and b generate a free group of rank 2 in the self-homeomorphism group of $\mathbb{R}$.
Hints. Consider the sets $f^n([1/5, 1])$ and $f^{-n}([0, 4/5])$ for $n \in \mathbb{N}_{>0}$ as well as

$$\bigcup_{k \in \mathbb{Z}} \Big[k - \frac{1}{5}, k + \frac{1}{5}\Big] \qquad \text{and} \qquad \bigcup_{k \in \mathbb{Z}} \Big[k + \frac{1}{2} - \frac{1}{5}, k + \frac{1}{2} + \frac{1}{5}\Big].$$

Exercise 4.E.17 (Free rotation group***). Show that the special orthogonal group SO(3) contains a free subgroup of rank 2.
Hints. Consider the matrices

$$\begin{pmatrix} \frac{3}{5} & \frac{4}{5} & 0 \\ -\frac{4}{5} & \frac{3}{5} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{3}{5} & -\frac{4}{5} \\ 0 & \frac{4}{5} & \frac{3}{5} \end{pmatrix}$$

and divisibility by 5 in $\mathbb{Z}^3 \subset \mathbb{R}^3$ [170, Chapter 2.2].

Special linear groups

Exercise 4.E.18 (Small linear group*). Show that $\mathrm{SL}(2, \mathbb{Z}/2)$ contains exactly six elements.

Exercise 4.E.19 ($\mathrm{SL}(n, \mathbb{Z})$ is not virtually free**). Let $n \in \mathbb{N}_{\geq 3}$. Show that $\mathrm{SL}(n, \mathbb{Z})$ does *not* contain a free group of finite index.
Hints. Try to find a subgroup of $\mathrm{SL}(n, \mathbb{Z})$ that is isomorphic to $\mathbb{Z}^2$.

Exercise 4.E.20 (Free subgroups imply non-solvability**). Let G be a group that contains a free subgroup of rank 2.
1. Let $G' \subset G$ be a subgroup of finite index. Show that then G' contains a free subgroup of rank 2.
2. Show that the commutator subgroup $[G, G]$ of G contains a free subgroup of rank 2.
3. In particular, deduce that groups that contain a free subgroup of rank 2 do *not* contain a solvable subgroup of finite index.
4. Let $n \in \mathbb{N}_{\geq 2}$. Conclude that $\mathrm{SL}(n, \mathbb{Z})$ does *not* contain a solvable subgroup of finite index.

Exercise 4.E.21 (Eigenspace ping-pong**). We use the notation from Example 4.4.8.
1. Sketch the eigenspaces of the $\mathbb{R}$-versions of a and b in $\mathbb{R}^2$.
2. Find explicit bounds for $|\lambda|$ and ε so that the ping-pong lemma indeed can be applied to the situation specified in Example 4.4.8.

Girth of graphs

Exercise 4.E.22 (Moore bound**). Let $d \in \mathbb{N}_{\geq 2}$. Show that there exists a constant $C \in \mathbb{R}_{>0}$ such that: For all finite d-regular graphs $X = (V, E)$ we have

$$g(X) \leq 2 \cdot \log_{d-1} |V| + C.$$

Hints. Choose a vertex v and look at the subgraph (subtree!) of vertices that can be reached from v by paths of length smaller than $1/2 \cdot g(X)$.

Exercise 4.E.23 (The Mycielski graph**). Given a graph $X = (V, E)$, we can construct the associated *Mycielski graph* $\mu(X)$ as follows: The graph $\mu(X)$ has the vertex set

$$\big(V \times \{0,1\}\big) \sqcup \{u\}$$

(where $u \notin V \times \{0,1\}$) and the edge set

$$\begin{aligned} &\Big\{\{(v,0),(w,0)\} \;\Big|\; \{v,w\} \in E\Big\} \\ \cup &\Big\{\{(v,0),(w,1)\} \;\Big|\; \{v,w\} \in E\Big\} \\ \cup &\Big\{\{(v,1),u\} \;\Big|\; v \in V\Big\}. \end{aligned}$$

In other words, $\mu(X)$ consists of three layers (Figure 4.14): The first layer is X itself. The second layer is a copy of the vertices of X that are connected with the original vertices of X as indicated by the edges of X. The third layer is an additional vertex u that is connected with all the vertices in the second layer.

Prove the following inheritance properties of the Mycielski graph construction:

1. If X is a graph with $g(X) \geq 3$, then also $g(\mu(X)) \geq 3$.
2. We have $\operatorname{ch}(\mu(X)) \geq \operatorname{ch} X + 1$.

Hints. The chromatic number is introduced in Definition 3.E.1.

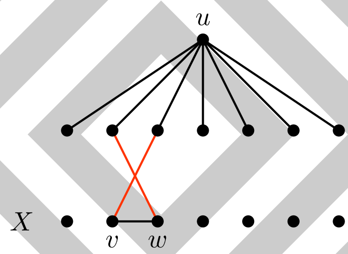

Figure 4.14.: The Mycielski graph construction, schematically

Exercise 4.E.24 (Graphs with few vertices and large girth – higher degree***). Generalise the techniques of the proof of Theorem 4.4.6 to produce regular graphs of degree $6, 8, \dots$ with high girth and few vertices.

Residually finite groups$^+$

It is a common theme of group theory to approximate groups by simpler groups. One instance of this is the class of residually finite groups:

Definition 4.E.1 (Residually finite group)**.** A group G is *residually finite* if it has the following property: For every $g \in G \setminus \{e\}$ there exists a finite group F and a homomorphism $\varphi\colon G \longrightarrow F$ with $\varphi(g) \neq e$.

Exercise 4.E.25 (Residually finite groups*)**.**

1. Prove that subgroups of residually finite groups are residually finite.
2. Prove that a group G is residually finite if and only if
$$\bigcap_{H \in N(G)} H = \{e\},$$
where $N(G)$ denotes the set of all finite index normal subgroups of G.
3. Prove that a group G is residually finite if and only if
$$\bigcap_{H \in S(G)} H = \{e\},$$
where $S(G)$ denotes the set of all finite index subgroups of G.
Hints. The normal subgroup trick (Exercise 2.E.5) might be useful.
4. Prove that a group G is residually finite if and only if the diagonal action of G on $\prod_{H \in N(G)} G/H$ is free.

Exercise 4.E.26 (Free groups are residually finite**)**.**

1. Let $n \in \mathbb{N}$. Prove that the matrix group $\mathrm{SL}(n, \mathbb{Z})$ is residually finite.
Hints. Consider for $N \in \mathbb{Z}$ the canonical homomorphism
$$\mathrm{SL}(n, \mathbb{Z}) \longrightarrow \mathrm{SL}(n, \mathbb{Z}/N)$$
given by reduction mod N.
2. Deduce that free groups of rank 2 are residually finite.
3. Conclude that all free groups of finite rank are residually finite.

Exercise 4.E.27 (Non-residually finite groups**)**.**

1. Show that the additive group $\mathbb{Q}$ is *not* residually finite.
2. Let $\sigma := (x \mapsto x + 1) \in S_{\mathbb{Z}}$ be the shift map and let $\tau \in S_{\mathbb{Z}}$ be the bijection swapping 1 and 2 (and fixing everything else). We then consider the finitely generated subgroup $G := \langle \sigma, \tau \rangle_{S_{\mathbb{Z}}}$ of $S_{\mathbb{Z}}$. Show that G is *not* residually finite.
Hints. Let F be a finite group and let $\varphi\colon G \longrightarrow F$ be a group homomorphism. Let $n \in \mathbb{N}_{\geq 5}$ with $n! > |F|$. Show that φ induces a well-defined group homomorphism $\overline{\varphi}\colon S_n \longrightarrow F$. What happens to the cycle $(1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1)$ under $\overline{\varphi}$ and φ ?

Exercise 4.E.28 (Residually finite groups are Hopfian**)**.** A group G is *Hopfian* if every surjective group homomorphism $G \longrightarrow G$ is an automorphism of G.

1. Show that all finitely generated residually finite groups are Hopfian.
Hints. If G is a finitely generated group and F is a finite group, then the set $\mathrm{Hom}(G, F)$ is finite. Hence, every surjective homomor-

phism $\varphi\colon G \longrightarrow G$ defines a permutation of $\operatorname{Hom}(G,F)$ by composition ...

2. Show that free groups of infinite rank are residually finite but not Hopfian.

Exercise 4.E.29 (Profinite completion**)**.** Let G be a group. The *profinite completion of* G is the group

$$\widehat{G} := \varprojlim_{H \in N(G)} G/H \subset \prod_{H \in N(G)} G/H,$$

where $N(G)$ denotes the set of all finite index normal subgroups of G, partially ordered by inclusion (the structure maps in this inverse system then are the associated canonical projection homomorphisms); more explicitly, $\widehat{G}$ is the group of all sequences in $\prod_{H \in N(G)} G/H$ whose entries are compatible with respect to the canonical projections. We can view $\widehat{G}$ as a topological group by taking the discrete topology on all finite quotients G/H and the product topology on $\prod_{H \in N(G)} G/H$.

1. Show that $\widehat{G}$ indeed is a topological group, i.e., that the inversion map $\widehat{G} \longrightarrow \widehat{G}$ and the composition map $\widehat{G} \times \widehat{G} \longrightarrow \widehat{G}$ are continuous.
2. Show that the image of the diagonal map $G \longrightarrow \widehat{G}$ is dense in $\widehat{G}$.
3. Show that the diagonal map $G \longrightarrow \widehat{G}$ is injective if and only if G is residually finite.

Exercise 4.E.30 (A topological characterisation of residual finiteness***)**.** Formulate and prove a characterisation of residual finiteness in terms of covering theory.

A far-reaching generalisation of approximating groups by finite pieces is the class of sofic groups [39].

The first Grigorchuk group$^+$

The (first) Grigorchuk group [69] is defined as follows, via tree automorphisms:

We first introduce the underlying tree. Let $\{0,1\}^*$ the set of finite sequences over $\{0,1\}$; we will denote the empty sequence by ε and we will concatenate such finite sequences by just writing them next to each other. Then the *rooted binary tree* (Figure 4.15) is the graph

$$T := \bigl(\{0,1\}^*, \{(w,wx) \mid w \in \{0,1\}^*, x \in \{0,1\}\}\bigr).$$

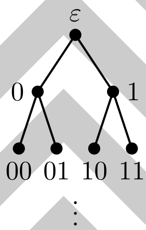

Figure 4.15.: The rooted binary tree T

Exercise 4.E.31 (The rooted binary tree*).

1. Show that T indeed is a tree.
2. Show that T has a single vertex of degree 2 (namely ε) and that all other vertices have degree 3.
3. Show that every graph automorphism of T fixes ε.
4. Show that the group $\operatorname{Aut}(T)$ of graph automorphisms of T is uncountable.

We will now introduce four special automorphisms a, b, c, d of T (which will then form the generating set of the Grigorchuk group). We define a by swapping the first subtrees of T:

$$\begin{aligned} a \colon \{0,1\}^* &\longrightarrow \{0,1\}^* \\ \varepsilon &\longmapsto \varepsilon \\ 0w &\longmapsto 1w \\ 1w &\longmapsto 0w \end{aligned}$$

Then we define b, c, d by mutual induction over the length of words by

$$\begin{aligned} b \colon \{0,1\}^* &\longrightarrow \{0,1\}^* & c \colon \{0,1\}^* &\longrightarrow \{0,1\}^* & d \colon \{0,1\}^* &\longrightarrow \{0,1\}^* \\ \varepsilon &\longmapsto \varepsilon & \varepsilon &\longmapsto \varepsilon & \varepsilon &\longmapsto \varepsilon \\ 0w &\longmapsto 0a(w) & 0w &\longmapsto 0a(w) & 0w &\longmapsto 0w \\ 1w &\longmapsto 1c(w) & 1w &\longmapsto 1d(w) & 1w &\longmapsto 0b(w). \end{aligned}$$

These definitions are illustrated in Figure 4.16.

Quick check 4.E.32 (Generators of the Grigorchuk group*).

1. What is $a(001101)$?
2. What is $b(10011011001)$?
3. What is $c(10011011001)$?
4. Why does the inductive definition of b, c, d work?
5. Why are a, b, c, d indeed graph automorphisms of T ?

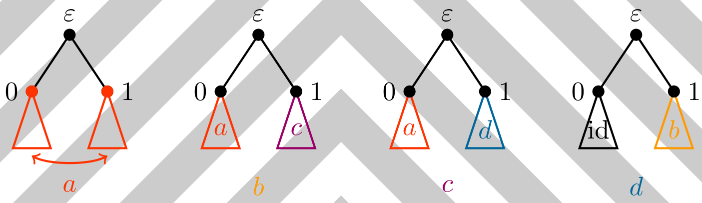

Figure 4.16.: The generators of the Grigorchuk group, schematically

Definition 4.E.2 (First Grigorchuk group)**.** The subgroup

$$\mathrm{Gri} := \langle a, b, c, d \rangle_{\mathrm{Aut}(T)}$$

of $\mathrm{Aut}(T)$ is the *(first) Grigorchuk group.*

Exercise 4.E.33 (Elementary relations in the Grigorchuk group*)**.**

1. Prove that $a, b, c, d \in \mathrm{Gri}$ have order 2.
2. Prove that

   $$b \cdot c = d = c \cdot b, \qquad d \cdot c = b = c \cdot d, \qquad d \cdot b = c = b \cdot d.$$

   holds in Gri.

By construction, the Grigorchuk group Gri acts on the rooted binary tree T. The subgroups that preserve the first levels of the tree are an important tool in studying the whole group Gri:

Definition 4.E.3 (Level stabiliser subgroups of the Grigorchuk group)**.** For $n \in \mathbb{N}$ we define

$$L_n := \bigl\{ g \in \mathrm{Gri} \bigm| \forall_{w \in \{0,1\}^n} \;\; g(w) = w \bigr\} \subset \mathrm{Gri}\,.$$

Exercise 4.E.34 (Basic properties of level stabilisers*)**.** Let $n \in \mathbb{N}$.

1. Show that L_n is a normal subgroup of Gri of finite index.
2. Show that $L_{n+1} \subset L_n$.
3. Show that $\bigcap_{n \in \mathbb{N}} L_n = \{e\}$ and conclude that Gri is residually finite.

Definition 4.E.4 (The child homomorphism)**.** Every $g \in L_1$ defines two automorphisms $g_0, g_1 \in \mathrm{Gri}$ by

$$\forall_{w \in \{0,1\}^*} \;\; g(0w) = 0 g_0(w)$$
$$\forall_{w \in \{0,1\}^*} \;\; g(1w) = 1 g_1(w).$$

We then set

$$\begin{aligned} \varphi \colon L_1 &\longrightarrow \mathrm{Gri} \times \mathrm{Gri} \\ g &\longmapsto (g_0, g_1) \end{aligned}$$

and we write $\varphi_0, \varphi_1 \colon L_1 \longrightarrow \mathrm{Gri}$ for the compositions with the projections on the two factors.

Exercise 4.E.35 (Basic properties of the child homomorphism*).

1. Let $g \in L_1$. Prove that indeed $g_0, g_1 \in \mathrm{Gri}$ holds.
2. Show that $\varphi \colon L_1 \longrightarrow \mathrm{Gri} \times \mathrm{Gri}$ is an injective group homomorphism.
3. Show that $\varphi_0(L_{n+1}) \subset L_n$ and $\varphi_1(L_{n+1}) \subset L_n$ for all $n \in \mathbb{N}$.
4. Prove that $\varphi_0, \varphi_1 \colon L_1 \longrightarrow \mathrm{Gri}$ are surjective.
   Hints. Look at $b, c, d, a \cdot b \cdot a, a \cdot c \cdot a, a \cdot d \cdot a$.
5. Conclude that the group Gri is infinite.

The Grigorchuk group exhibits the following self-similarity property: The group Gri is almost isomorphic to $\mathrm{Gri} \times \mathrm{Gri}$. More precisely, the finite index subgroup L_1 is isomorphic to a finite index subgroup of $\mathrm{Gri} \times \mathrm{Gri}$:

Exercise 4.E.36 (Weak self-similarity of the Grigorchuk group**). We now let $N := \langle b \rangle^{\triangleleft}_{\mathrm{Gri}}$ be the normal subgroup generated by b in Gri.

1. Show that the subgroup $\langle a, d \rangle_{\mathrm{Gri}}$ is finite (more precisely: this subgroup is isomorphic to D_4).
2. Show that L_1 is generated by $\{b, c, d, a \cdot b \cdot a, a \cdot c \cdot a, a \cdot d \cdot a\}$.
3. Show that N has finite index in Gri.
4. Show that $N \times \{e\} \subset \varphi(L_1)$ and $\{e\} \times N \subset \varphi(L_1)$.
5. Show that $N \times N \subset \varphi(L_1)$.
6. Conclude that $\varphi(L_1)$ has finite index in Gri.

Exercise 4.E.37 (The Grigorchuk group is a torsion group***).

1. Show that all elements of L_1 have finite order.
   Hints. First write elements of L_1 in terms of the generators a, b, c, d and then proceed by induction over the minimal number of generators needed to represent a given element (i.e., induction over the so-called word length).
2. Conclude that all elements of Gri have finite order.

Hence, the Grigorchuk group is a finitely generated infinite torsion group that is almost self-similar. The first Grigorchuk group Gri has further notable properties [77, Chapter VIII][69]: For example, it is known that Gri is *not* finitely presented, that all of its quotient groups are finite, and that it is of intermediate growth (Exercises 6.E.12 and 6.E.13).

5

Quasi-isometry

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 in the associated Cayley graph $\text{Cay}(G, S)$ induce a metric on G, the 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 geometry on a group independent of the choice of generating sets we pass to large scale geometry. Using the language of quasi-geometry, we arrive at such a notion for finitely generated groups – the quasi-isometry type, which is central to geometric group theory.

We start with some generalities on isometries, bilipschitz equivalences, and quasi-isometries. As next step, we will specialise to the case of finitely generated groups. The key to linking the geometry of groups to actual geometry is the Švarc-Milnor lemma (Chapter 5.4). Moreover, we will discuss Gromov's dynamic criterion for quasi-isometry (Chapter 5.5) and give an outlook on geometric properties of groups and quasi-isometry invariants (Chapter 5.6).

Overview of this chapter

5.1 Quasi-isometry types of metric spaces

In the following, we consider 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 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. A category theoretic framework will be explained in Remark 5.1.12.

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

Definition 5.1.1 (Metric space)**.** A *metric space* is a pair (X, d) consisting of a set X and 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:

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

Sometimes we will abuse notation and say that X is a metric space if the metric is clear from the context.

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

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

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

Also bilipschitz equivalences preserve 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).

- 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 f is a (c,b)*-quasi-isometric embedding*, i.e.,
$$\forall_{x,x' \in X} \quad \frac{1}{c} \cdot d_X(x,x') - b \leq d_Y\bigl(f(x), f(x')\bigr) \leq c \cdot d_X(x,x') + b.$$
- A map $f'\colon X \longrightarrow Y$ has *finite distance from* f if there is a $c \in \mathbb{R}_{\geq 0}$ with
$$\forall_{x \in X} \quad d_Y\bigl(f(x), f'(x)\bigr) \leq c.$$
- 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_X and $f \circ g$ has finite distance from id_Y.
- The 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)**.** Every isometry is a bilipschitz equivalence, and every bilipschitz equivalence is a quasi-isometry. In general, the converse does not hold:

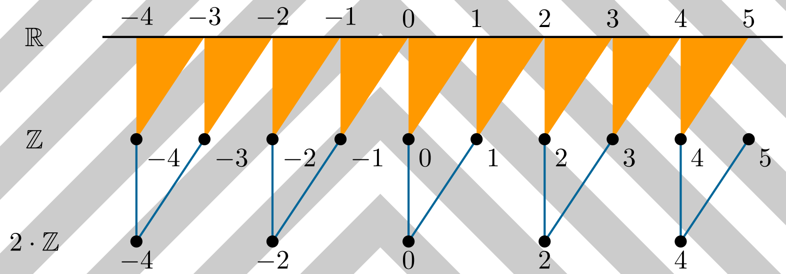

Figure 5.1.: The metric spaces $\mathbb{R}$, $\mathbb{Z}$, and $2 \cdot \mathbb{Z}$ are quasi-isometric

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).

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:

- in general, quasi-isometries are neither injective, nor surjective,
- in general, quasi-isometries are not continuous at all,
- in general, quasi-isometries do not have finite distance to an 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 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 5.E.24, Example 6.2.8).

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 $f\colon X \longrightarrow Y$ has* 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 $c \in \mathbb{R}_{>0}$ 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 find 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

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

By the axiom of choice, there exists 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: 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: 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

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

(The same argument shows that quasi-inverses of quasi-isometric embeddings are quasi-isometric embeddings). □

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

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

1. *Every map at finite distance of a quasi-isometric embedding is a quasi-isometric embedding.*
2. *Every 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 5.E.3 and Exercise 5.E.4). □

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 $\operatorname{Mor}_{\mathsf{QMet}'}(X,Y)$ 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

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

- 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).

- 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 [154].

5.2 Quasi-isometry types of groups

Every 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 $X = (V, E)$ be a connected graph. Then the map

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

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

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, word length). 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, $h \in G$. The distance $d_S(e, g)$ is also called *word length* of g with respect to S.

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.

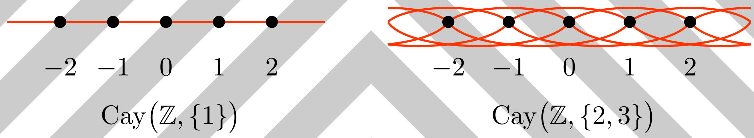

Figure 5.2.: Cayley graphs of $\mathbb{Z}$ with the same large scale geometry

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, every 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, the maximum

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

is finite. Let g, $h \in G$ and let $n := d_S(g, h)$. Then we can write $g^{-1} \cdot h$ as $s_1 \cdot \dots \cdot s_n$ 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. $\square$

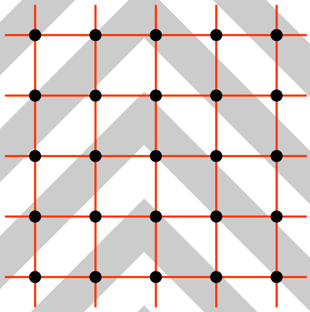

Figure 5.3.: The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ resembles the geometry of the Euclidean plane $\mathbb{R}^2$

Example 5.2.6 (Cayley graphs of $\mathbb{Z}$)**.** Two different Cayley graphs for the additive group $\mathbb{Z}$ with respect to finite generating sets are depicted in Figure 5.2.

For infinite generating sets the first part of the previous 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 5.2.7 (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 every) 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 every) 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 5.2.8 ($\mathbb{Z}^n \sim_{\mathrm{QI}} \mathbb{R}^n$)**.** 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, we will see soon 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 (Chapter 5.4).

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:

Proposition 5.2.9 (Quasi-isometry vs. bilipschitz equivalence). *Bijective quasi-isometries between finitely generated groups (with respect to the word metric of certain finite generating sets) are bilipschitz equivalences.*

Proof. The proof is based on the fact that the minimal non-trivial distance between two group elements is 1; one can then trade the additive constants in a bijective quasi-isometry for a contribution in the multiplicative constants (Exercise 5.E.5). □

However, not all infinite finitely generated groups that are quasi-isometric are bilipschitz equivalent. We will study in Chapter 9.4 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 5.2.10 (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 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 5.2.11 (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 of a finite generating set; because balls of finite radius with respect to word metrics of finite generating sets are finite (Remark 5.2.10), it follows that the group in question has to be finite.

In contrast, 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 $\mathbb{Z}$:

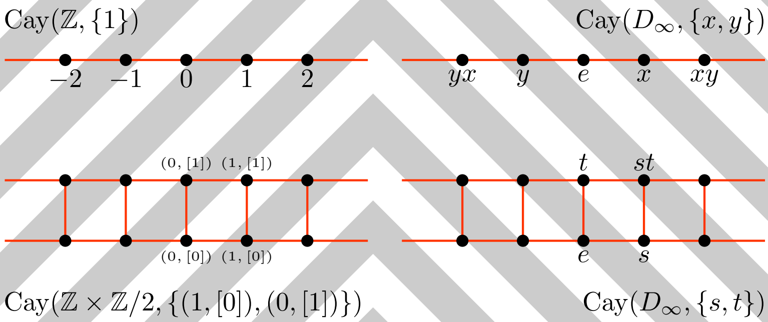

Figure 5.4.: The groups $\mathbb{Z}$, $\mathbb{Z} \times \mathbb{Z}/2$, and D_∞ are quasi-isometric

Example 5.2.12 (Some groups quasi-isometric to $\mathbb{Z}$)**.** The groups $\mathbb{Z}$, $\mathbb{Z} \times \mathbb{Z}/2$, 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 (Exercise 2.E.19, Exercise 2.E.31, Exercise 3.E.21):

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

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 $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2, \{(1, [0]), (0, [1])\})$; in particular, D_∞ and $\mathbb{Z} \times \mathbb{Z}/2$ 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 $\mathbb{Z} \times \mathbb{Z}/2$ are bilipschitz equivalent.

Caveat 5.2.13 (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 (Exercise 3.E.22). 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 [102]. More generally, a similar classification also applies to finitely generated nilpotent groups [177].

One can also show by elementary arguments that $\mathbb{Z}$ and $\mathbb{Z}^n$ are *not* quasi-isometric whenever $n \in \mathbb{N}_{\geq 2}$ (Exercise 5.E.24). More conceptual arguments will be given in Chapter 6.

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

Theorem 5.2.14 (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 several proofs of this result later when we have more tools available (Chapter 6.3.6, Corollary 7.5.8, Theorem 8.2.14).

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 Quasi-geodesics and quasi-geodesic spaces

In metric geometry, it is useful to require that the metric on the space in question is (quasi-)geodesic, i.e., that its metric can be realised (up to some uniform error) by paths. For example, this will be an important hypothesis in the Švarc-Milnor lemma.

5.3.1 (Quasi-)Geodesic spaces

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, the 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 5.E.9).
- The sphere S^2 with the standard round Riemannian metric is a geodesic metric space. The geodesics are parts of great circles on S^2. However, antipodal points can be joined by infinitely many different geodesics.
- The hyperbolic plane $\mathbb{H}^2$ is a geodesic metric space (Appendix A.3). In the Poincaré disk model, geodesics are parts of circles that intersect the boundary circle orthogonally.

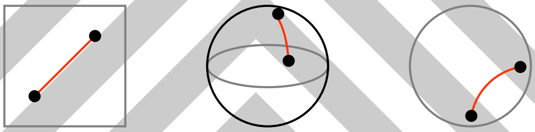

Figure 5.5.: Geodesic spaces and example geodesics

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

Finitely generated groups together with a word metric coming from a finite generating set 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 space)**.** Let (X, d) be a metric space, let $c \in \mathbb{R}_{>0}$, and let $b \in \mathbb{R}_{\geq 0}$.

- 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'.

Every 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 $X = (V, E)$ 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 X, and every path in the graph X 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 5.E.9).

5.3.2 Geodesification via geometric realisation of graphs

Sometimes it is more convenient to be able to argue via geodesics than via quasi-geodesics. Therefore, we explain how we can associate a geodesic space

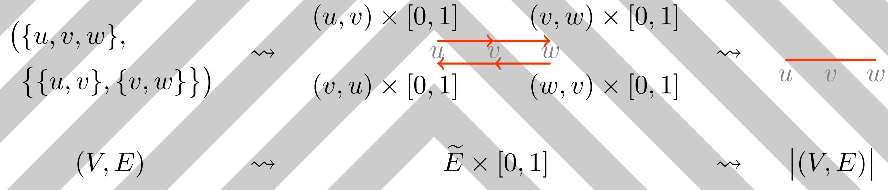

Figure 5.6.: Geometric realisation of graphs

with a connected graph and how quasi-geodesic spaces can be replaced by geodesic spaces via graphs:

Roughly speaking the geometric realisation of a graph is obtained by glueing 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 5.6):

Definition 5.3.6 (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

$$|X| := \widetilde{E} \times [0,1] \,/\!\sim .$$

Here,

$$\widetilde{E} := \bigl\{(u,v) \bigm| u, v \in V,\ \{u,v\} \in E\bigr\}$$

is the set of all directed edges (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 5.7)

- the elements coincide, i.e., $((u, v), t) = ((u', v'), t')$, or
- the elements describe the same vertex lying on both edges, i.e.,

Figure 5.7.: Parametrisations describing the same points in the geometric realisation

- $u = u'$ 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

$$d_{|X|}\bigl(\bigl[((u,v),t)\bigr], \bigl[((u'v'),t')\bigr]\bigr)$$
$$:= \begin{cases} |t - t'| & \text{if } (u,v) = (u',v') \\ |t - (1-t')| & \text{if } (u,v) = (v',u') \\ \min\bigl(t + d_X(u,u') + t' & \\ \quad , t + d_X(u,v') + 1 - t' & \\ \quad , 1 - t + d_X(v,u') + t' & \\ \quad , 1 - t + d_X(v,v') + 1 - t'\bigr) & \text{if } \{u,v\} \neq \{u',v'\} \end{cases}$$

for all $[((u,v),t)]$, $[((u',v'),t')] \in |X|$, where d_X denotes the metric on V induced from the graph structure (Definition 5.2.1).

Clearly, this construction can be extended to a functor from the category of graphs to the category of metric spaces. Hence, every action of a group on a graph induces a corresponding piecewise linear isometric action of the group on the geometric realisation of the given graph. 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 5.3.7 (Geometric realisations).
- 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 5.8).
- The geometric realisation of $\mathrm{Cay}(\mathbb{Z}, \{1\})$ is isometric to the real line $\mathbb{R}$ with the standard metric (Exercise 5.E.11).
- 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 $\mathbb{R}^2$ (Exercise 5.E.11).

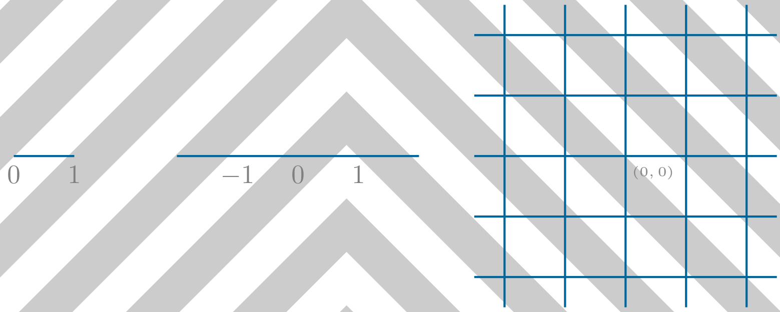

$|(\{0,1\},\{\{0,1\}\})|$ $|\mathrm{Cay}(\mathbb{Z},\{1\})|$ $|\mathrm{Cay}(\mathbb{Z}^2,\{(1,0),(0,1)\})|$

Figure 5.8.: Geometric realisations (Example 5.3.7)

Proposition 5.3.8 (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. This follows from straightforward calculations (Exercise 5.E.12). ☐

More generally, every quasi-geodesic space can be approximated by a geodesic space:

Proposition 5.3.9 (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.*

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 (this depends on the quasi-geodesity constants for X).

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 5.3.8). Hence, X and $|Y|$ are quasi-isometric. Moreover, $|Y|$ is a geodesic metric space by Proposition 5.3.8. We leave the details as an exercise (Exercise 5.E.12). ☐

However, 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.

5.4 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 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."

We start with a metric formulation of the Švarc-Milnor lemma for quasi-geodesic spaces; in a second step, we will deduce a more topological version, the version commonly used in applications.

Proposition 5.4.1 (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*

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

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: The argument follows the transitivity principle used in the proof of Proposition 4.1.20. 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

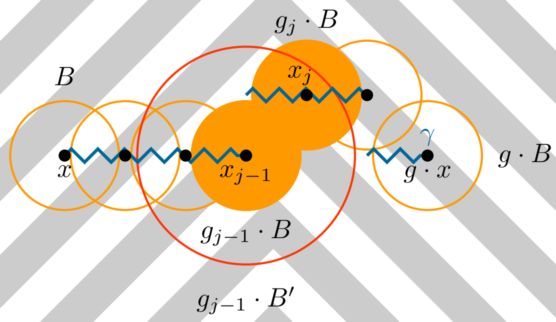

Figure 5.9.: Covering a quasi-geodesic by translates of B

quasi-geodesic (Figure 5.9): 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

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

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: As γ is a (c,b)-quasi-geodesic, we obtain

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

Therefore, $x_j \in B^{X,d}_{2\cdot b}(g_{j-1} \cdot B) = g_{j-1} \cdot B^{X,d}_{2\cdot b}(B) = g_{j-1} \cdot B'$ (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.

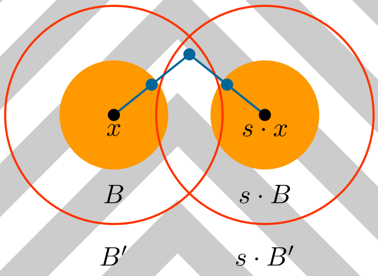

Figure 5.10.: If $s \in S$, then $d(x, s \cdot x) \leq 2 \cdot (\operatorname{diam} B + 2 \cdot b)$

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: 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

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

Conversely, we obtain a uniform upper bound of $d(\varphi(e), \varphi(g))$ in terms of the word length $d_S(e, g)$ 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 5.10) 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 $\operatorname{diam} B$ is assumed to be finite).

Because

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

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

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 obvious 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 5.4.2 (Švarc-Milnor lemma, topological formulation)**.** *Let G be a group acting by isometries on a (non-empty) proper and 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 5.4.3 (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 [115, Chapter V].
  - 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 $G \setminus X$ is compact with respect to the quotient topology.

  Example 5.4.4 (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 $H := \{z \in \mathbb{C} \mid \operatorname{Re} z > 0\}$ (Appendix A.3) is *not* cocompact (Exercise 5.E.20).

Proof of Corollary 5.4.2. Under the given assumptions, the metric space X is $(1,b)$-quasi-geodesic for every $b \in \mathbb{R}_{>0}$. In order to be able to apply the Švarc-Milnor lemma (Proposition 5.4.1), 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

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

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 5.4.1). □

5.4.1 Application: (Weak) commensurability

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

- Finite index subgroups of finitely generated groups are finitely generated.
- (Weakly) commensurable groups are quasi-isometric.
- Certain groups arising in geometric topology 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 5.4.5. *Finite index subgroups of finitely generated groups are finitely generated and quasi-isometric to the ambient group (via the inclusion map).*

Proof. Let G be a finitely generated group, and let $H \subset G$ 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 5.4.1): The space (G, d_S) is $(1,1)$-quasi-geodesic. Moreover, we let $B \subset G$ be a finite set of representatives of $H \setminus G$ (hence, the diameter of B is finite). Then $H \cdot B = G$, the set $B' := B_2^{G,d_S}(B)$ is finite, and so the set

$$T := \{h \in H \mid h \cdot B' \cap B' \neq \emptyset\}$$

is finite.

Therefore, H is finitely generated (by T) 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). □

Pursuing this line of thought leads to the notion of (weak) commensurability of groups:

Definition 5.4.6 ((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 5.E.16).

Corollary 5.4.7 (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 every group weakly commensurable to G is finitely generated and quasi-isometric to G.

Proof. Ad 1. In view of Corollary 5.4.5, 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 5.4.1). Therefore, we obtain $G \sim_{\mathrm{QI}} G/N$. □

Example 5.4.8 (Commensurability).

- 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 (Exercise 4.E.12), 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

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$$

is free of rank 2 (Example 4.4.1) and has index 12 in $\mathrm{SL}(2, \mathbb{Z})$ (Proposition 4.4.2). 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 5.4.9. Not all quasi-isometric groups are commensurable [77, 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 χ (an invariant from algebraic topology) of the corresponding classifying spaces is multiplicative under finite coverings [34]. Hence, commensurable groups G and G' (that admit sufficiently finite models of classifying spaces) satisfy

$$\chi(G) = 0 \iff \chi(G') = 0.$$

However, the inheritance properties of the Euler characteristic [34] yield

$$\begin{aligned}
\chi\bigl((F_3 \times F_3) * F_3\bigr) &= \chi(F_3) \cdot \chi(F_3) + \chi(F_3) - 1 \\
&= (1-3) \cdot (1-3) + (1-3) - 1 \\
&\neq 0 \\
&= (1-3) \cdot (1-3) + (1-4) - 1 \\
&= \chi(F_3) \cdot \chi(F_3) + \chi(F_4) - 1 \\
&= \chi\bigl((F_3 \times F_3) * F_4\bigr).
\end{aligned}$$

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 [77, III.18(xi)].

5.4.2 Application: Geometric structures on manifolds

As second example, we look at applications of the Švarc-Milnor lemma in algebraic topology and Riemannian geometry via fundamental groups; a concise introduction to Riemannian geometry is given by Lee's book [96].

Corollary 5.4.10 (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*

$$\begin{aligned}
\pi_1(M) &\longrightarrow \widetilde{M} \\
g &\longmapsto g \cdot x
\end{aligned}$$

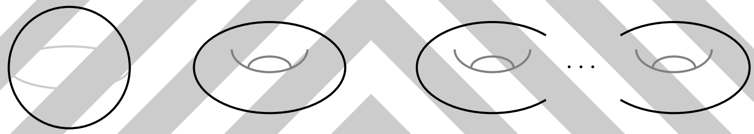

Figure 5.11.: Oriented closed connected surfaces

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 5.4.2) finishes the proof. □

We give a sample application of this consequence of the Švarc-Milnor lemma to Riemannian geometry:

Definition 5.4.11 (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 Riemannian manifold is called *hyperbolic* if its Riemannian universal covering is isometric to the hyperbolic space of the same dimension.

Being flat is the same as having vanishing sectional curvature and being hyperbolic is the same as having constant sectional curvature -1 [96, Chapter 11].

Example 5.4.12 (Surfaces)**.** Oriented closed connected surfaces are determined up to homeomorphism/diffeomorphism by their genus (i.e., the number of "handles", see Figure 5.11) [115, 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

$$\left\langle a_1, \dots, a_g, b_1, \dots, b_g \;\middle|\; \prod_{j=1}^{g} [a_j, b_j] \right\rangle$$

and one can show that these surfaces admit hyperbolic Riemannian metrics [18, Chapter B.1, B.3][167].

Corollary 5.4.13 ((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.*

Proof. This is a direct consequence of Corollary 5.4.10. □

Moreover, by the Bonnet-Myers theorem, closed connected Riemannian manifolds of positive sectional curvature have finite fundamental group [96, Theorem 11.7, Theorem 11.8].

So, classifying finitely generated groups up to quasi-isometry and studying the quasi-geometry of finitely generated groups gives insights into the geometry and topology of smooth/Riemannian manifolds.

5.5 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 5.5.1 (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. $G \cdot K = X = K \cdot H$.

[1] A *right action* of a group H on a set X is a map $X \times H \longrightarrow X$ 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 of H on X is the same as an antihomomorphism $H \longrightarrow S_X$.

2. The sets
$$F_G := \{g \in G \mid g \cdot K \cap K \neq \emptyset\},$$
$$F_H := \{h \in H \mid K \cdot h \cap K \neq \emptyset\}$$
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 $F_G(h) \subset G$ with $K \cdot h \subset F_G(h) \cdot K$.

Example 5.5.2 (Set-theoretic coupling for finite index subgroups)**.** 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 5.5.3 (Quasi-isometry and set-theoretic couplings)**.** *Let G and H be two finitely generated groups that admit a set-theoretic coupling. Then*
$$G \sim_{\text{QI}} H.$$

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 5.5.1. We prove $G \sim_{\text{QI}} H$ 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: 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, $K \cdot f(g)^{-1} \cdot f(g') \cap K \cdot F_H(S)^n \neq \emptyset$. 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. □

Outlook 5.5.4 (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 *cocycles* (which are algebraic objects); considering cocycles up to an appropriate equivalence relation ("being a coboundary") then gives rise to *cohomology groups* (Appendix A.2). In this way, aspects of group actions on a space can be translated into an algebraic theory. In particular, the characterisation of quasi-isometry of finitely generated groups through couplings leads to quasi-isometry invariance of certain (co)homological invariants [67, 161, 157, 98].

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.E.18).

The converse of Proposition 5.5.3 also holds: whenever two finitely generated groups are quasi-isometric, then there exists a coupling (even a topological coupling) between them:

Definition 5.5.5 (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 neighbour-

[2]There are several *different* notions of local compactness in the literature!

hood $K \subset X$ 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:

Theorem 5.5.6 (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.*

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 forms 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

$$\{g \in G \mid g \cdot K \cap K \neq \emptyset\} \qquad \text{and} \qquad \{h \in H \mid K \cdot h \cap K \neq \emptyset\}$$

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, this topological coupling is also a set-theoretic coupling for G and H.

Ad "$2 \Longrightarrow 3$". This was proved in Proposition 5.5.3.

Ad "$3 \Longrightarrow 1$". Suppose that the finitely generated groups G and H are quasi-isometric. We now explain 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

$$X := \Bigl\{f\colon G \longrightarrow H \times F \Bigm| f \text{ has } C\text{-dense image in } H \times F, \text{ and } \forall_{g,g' \in G}\ \frac{1}{C} \cdot d_S(g,g') \leq d_{T \times F}\bigl(f(g), f(g')\bigr) \leq C \cdot d_S(g,g')\Bigr\}$$

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 $g,\ g' \in G$ satisfy $f(g) = f(g')$, 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.E.5 from the fact that different elements of a finitely generated group have distance at least 1 with respect to every 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 consider the following left G-action and right H-action on X:

$$\begin{aligned} G \times X &\longrightarrow X \\ (g, f) &\longmapsto \bigl(x \mapsto f(g^{-1} \cdot x)\bigr) \end{aligned}$$

$$\begin{aligned} X \times H &\longrightarrow X \\ (f, h) &\longmapsto \bigl(x \mapsto f(x) \cdot (h, e)\bigr) \end{aligned}$$

By construction, 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 [89, 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 indeed proper and cocompact [156]. □

Outlook 5.5.7 (A dynamic criterion for bilipschitz equivalence)**.** It is also possible to formulate and prove a dynamic criterion for bilipschitz equivalence, using couplings of continuous actions on Cantor sets [118, Theorem 3.2][98].

5.5.1 Application: Comparing uniform lattices

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 5.5.9).

Definition 5.5.8 (Uniform lattice)**.** 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 inversion 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 exists an open neighbourhood U of the neutral element e in G such that $U \cap \Gamma = \{e\}$.

Corollary 5.5.9 (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 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. Moreover, these actions are cocompact and proper. Hence, the ambient group G serves a topological coupling for Γ and Λ. So, Γ and Λ are quasi-isometric by the dynamic criterion (Theorem 5.5.6). ☐

Therefore, quasi-isometry invariants can sometimes be used to prove that a given finitely generated group is *not* a uniform lattice in a specific locally compact topological group.

Example 5.5.10 (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 $\mathbb{R}^2$. In particular, 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.E.24).
- Let $H_\mathbb{R}$ be the *real Heisenberg group*, and let H be the *Heisenberg group*, i.e.,

$$H_\mathbb{R} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \, x, y, z \in \mathbb{R} \right\}, \quad H := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \, x, y, z \in \mathbb{Z} \right\}.$$

  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 5.E.21).

So, every 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 halfplane is not cocompact, Exercise 5.E.20); 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 $\mathrm{Isom}(\widetilde{M})$ [156, Theorem 2.35].

Outlook 5.5.11 (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 [63].

5.6 Quasi-isometry invariants

The central classification problem of geometric group theory is to classify 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.

5.6.1 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 is to construct suitable invariants. We start with the simplest case, namely set-valued quasi-isometry invariants:

Definition 5.6.1 (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 all finitely generated groups G, H with $G \sim_{\mathrm{QI}} H$ satisfy

$$I(G) = I(H).$$

Proposition 5.6.2 (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 5.6.3. 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 5.6.4 (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 5.2.11).
- *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 5.4.8).

In order to obtain interesting classification results we need further quasi-isometry invariants. In the following chapters, we will, for instance, study

- the growth of groups (Chapter 6),
- hyperbolicity (Chapter 7),
- ends of groups (i.e., geometry at infinity) (Chapter 8),
- and amenability (Chapter 9).

Caveat 5.6.5. If a quasi-isometry invariant has only a countable range of possible values, then it will *not* be a complete invariant: There exist uncountably many quasi-isometry classes of finitely generated groups. This fact is a quasi-geometric version of Theorem 2.2.28 and it can, for example, be proved by producing uncountably many different growth types of groups [69] or via small cancellation theory and the geometry of loops in Cayley graphs [23].

5.6.2 Geometric properties of groups and rigidity

In geometric group theory, it is common to use the following term:

Definition 5.6.6 (Geometric property of groups). Let P be a property of finitely generated groups (i.e., every 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 (i.e., if "having property P" is a quasi-isometry invariant).

Example 5.6.7 (Geometric properties).

- Being finite is a geometric property of groups (Example 5.2.11).
- 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. We list only the most basic instances, more complete lists can be found in the book of Druţu and Kapovich [53]:

- Being virtually[3] infinite cyclic is a geometric property (Chapter 6.3).
- More generally, for every $n \in \mathbb{N}$ the property of being virtually $\mathbb{Z}^n$ is geometric (Chapter 6.3).
- Being finitely generated and virtually free is a geometric property [178, 53].
- Being finitely generated and virtually nilpotent is a geometric property of groups (Chapter 6.3).
- Being finitely presented is a geometric property of groups [31, Proposition I.8.24] (Exercise 6.E.35).

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 a *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.

5.6.3 Functorial quasi-isometry invariants

A refined setup for quasi-isometry invariants is the formalisation of quasi-isometry invariants as functors between categories. Functors translate objects and morphisms between categories:

Definition 5.6.8 (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 $F(\operatorname{id}_X) = \operatorname{id}_{F(X)}$.
- 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 5.6.9 (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 the morphism $F(f) \in \operatorname{Mor}_D(F(X),F(Y))$ 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 can be viewed as a functor from the category of graphs to the category of metric spaces; group (co)homology is a functor from the category of groups to the category of graded modules (Appendix A.2).

Definition 5.6.10 (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.

Functorial quasi-isometries refine ordinary quasi-isometry invariants: If $F\colon \mathsf{QMet} \longrightarrow C$ is a functorial quasi-isometry invariant, then taking the isomorphism classes of values yields a set-valued quasi-isometry invariant (provided that the isomorphism classes of C form a set). However, the functor F contains more information: We do not only get isomorphic values on quasi-isometric objects, but we also get relations between the values of F in the presence of quasi-isometric embeddings (Example 6.2.9, Example 8.3.9).

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 quasi-hyperbolic spaces to the category of topological spaces (Chapter 8). Also growth types of finitely generated groups can be viewed as functorial quasi-isometry invariants from the subcategory of QMet generated by finitely generated groups to the (category assoicated with the) partially ordered set of growth types (Chapter 6).

Moreover, there is a general principle turning functors from algebraic topology into quasi-isometry invariants, based on the coarsening construction by Higson and Roe [83, 148, 132]; a more general and more conceptual approach was recently developed by Bunke and Engel [36]. For simplicity, we restrict ourselves to the domain category of uniformly discrete spaces of bounded geometry.

Definition 5.6.11 (UDBG space)**.** A metric space (X, d) 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

$$\forall_{x\in X} \ \ \bigl|B_r^{X,d}(x)\bigr| \leq K_r.$$

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 the maximal quasi-isometry invariant contained in the given functor:

Theorem 5.6.12 (Coarsening of functors)**.** *Let C be a category that is closed under direct limits, and let $F\colon \mathsf{Simp}_{\mathsf{h}}^{\mathsf{lf}} \longrightarrow C$ 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 (Figure 5.12)*

$$c \circ c_F = c_G.$$

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

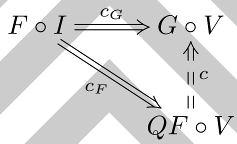

Figure 5.12.: Coarsening of functors: universal property

relating two given functors [109]. *Simplicial complexes* are combinatorial approximations of topological spaces, generalising the concept of a graph to higher dimensions [48, Chapter 8.1]; in the following, we will require some basic familiarity with simplicial complexes. 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; then the set of vertices is a UDBG space with respect to this metric. Moreover, we consider the corresponding canonical functors

$$I\colon \mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}} \longrightarrow \mathsf{Simp}^{\mathsf{lf}}_{\mathsf{h}}$$
$$V\colon \mathsf{Simp}^{\mathsf{ulf}}_{\mathsf{QI}} \longrightarrow \mathsf{UDBG}\,.$$

Sketch of proof of Theorem 5.6.12. 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 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: For $n \in \mathbb{N}$ 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 5.13).

- *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
$$QF(X) := \varinjlim_{r\to\infty} F\bigl(R_r(X)\bigr),$$
where the direct limit is taken by applying F to the above inclusions.

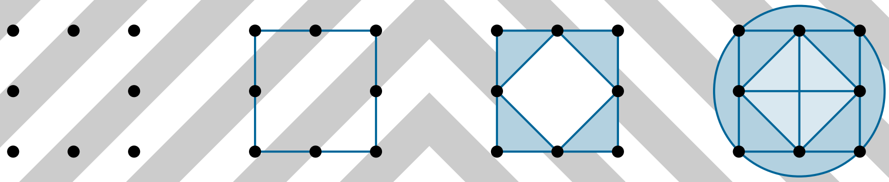

Figure 5.13.: Toy example for Rips complexes of a discrete metric space X; here, X has the metric induced from $\mathbb{R}^2$ and the large square has sides of length 2.

- *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 $s \in \mathbb{R}_{\geq 0}$ such that the quasi-isometric embedding $f\colon X \longrightarrow Y$ 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

$$F(X) \longrightarrow F\bigl(R_1(V(X))\bigr) \longrightarrow \varinjlim_{r \to \infty} F\bigl(R_r(X)\bigr) = QF(X).$$

by composing F applied to the inclusion with the structure morphism in the direct limit.

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

$$\varinjlim_{r \to \infty} G\bigl(V(R_r(X))\bigr) \cong G\bigl(V(X)\bigr).$$

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 every UDBG space is quasi-isometric to a space in the image of V, this natural transformation can be extended to UDBG. □

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 at a later stage in the Rips construction). In contrast, truly locally finite theories do have interesting coarsenings.

Outlook 5.6.13 (Uniformly finite homology)**.** For example, the coarsening of locally finite simplicial homology with uniformly bounded coefficients coincides with *uniformly finite homology*. Uniformly finite homology also has a concrete description in terms of combinatorial chains that satisfy geometric finiteness conditions [19, 175, 132] (Exercise 5.E.31ff). Applications of uniformly finite homology include the following:

- Uniformly finite homology allows to characterise amenable groups [19] (Chapter 9.2.4).
- Uniformly finite homology measures the difference between bilipschitz equivalences and quasi-isometries of UDBG spaces [175, 56] (Theorem 9.4.11).
- Uniformly finite homology can be used to show the existence of certain aperiodic tilings [19, 112].
- Uniformly finite homology is a tool that allows to study certain large scale notions of dimension [51].
- Uniformly finite homology has applications in the context of existence of positive scalar curvature metrics on smooth manifolds [19, 58].

5.E Exercises

Quasi-isometry and bilipschitz equivalence

Quick check 5.E.1 (Quasi-isometric embedding constants*). Let $c, c' \in \mathbb{R}_{>0}$ and $b, b' \in \mathbb{R}_{\geq 0}$.

1. Let $c' \geq c$ and $b' \geq b$. Is then every (c, b)-quasi-isometric embedding also a (c', b')-quasi-isometric embedding?
2. Let $c' \leq c$ and $b' \leq b$. Is then every (c, b)-quasi-isometric embedding also a (c', b')-quasi-isometric embedding?

Quick check 5.E.2 (Quasi-isometry of metric spaces*).

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

Exercise 5.E.3 (Maps close to quasi-isometric embeddings*).

1. Show that every map at finite distance of a quasi-isometric embedding is a quasi-isometric embedding.
2. Show that every map at finite distance of a quasi-isometry is a quasi-isometry.

Exercise 5.E.4 (Inheritance properties of quasi-isometric embeddings*). Let X, Y, Z be metric spaces and let $f, f' \colon X \longrightarrow Y$ be maps that have finite distance from each other.

1. Show that for all maps $g \colon Z \longrightarrow X$ the compositions $f \circ g$ and $f' \circ g$ have finite distance from each other.
2. Show that 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.

Conclude the following:

3. Compositions of quasi-isometric [bilipschitz] embeddings are quasi-isometric [bilipschitz] embeddings.
4. Compositions of quasi-isometries [bilipschitz equivalences] are quasi-isometries [bilipschitz equivalences].

Exercise 5.E.5 (Bijective quasi-isometries**).

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?

Exercise 5.E.6 (Quasi-isometry group**). Determine the quasi-isometry group of $\{n^3 \mid n \in \mathbb{Z}\}$ (with respect to the standard metric induced from $\mathbb{R}$).

Exercise 5.E.7 (Counting preimages**). Let X and Y be UDBG spaces (Definition 5.6.11) and let $f \colon X \longrightarrow Y$ be a quasi-isometry. Show that there are $c, C \in \mathbb{R}_{>0}$ with the following properties:

- The map $f \colon X \longrightarrow Y$ is a (c, c)-quasi-isometric embedding with c-dense image.
- For all finite sets $F \subset Y$ we have

$$\bigl|f^{-1}(B_c^Y(F))\bigr| \geq \frac{1}{C} \cdot |F| \qquad \text{and} \qquad \bigl|f^{-1}(F)\bigr| \leq C \cdot |F|.$$

Quasi-geodesic spaces

Quick check 5.E.8 (Quasi-isometry invariance of being (quasi-)geodesic*). Let X and Y be metric spaces and let $f \colon X \longrightarrow Y$ be a quasi-isometry.

1. Let X be geodesic. Is then also Y geodesic?
2. Let X be quasi-geodesic. Is then also Y quasi-geodesic?

Exercise 5.E.9 (Swiss cheese*). We consider $X := \mathbb{R}^2 \setminus \{0\}$ with the metric induced from the standard metric on $\mathbb{R}^2$.

1. Show that the space X is path-connected but not geodesic.
2. Show that for every $\varepsilon \in \mathbb{R}_{>0}$ the space X is $(1, \varepsilon)$-quasi-geodesic.

Illustrate your arguments with suitable pictures!

Exercise 5.E.10 (The maximum metric**). We consider the maximum metric d_∞ on $\mathbb{R}^2$.

1. Show that the space $(\mathbb{R}^2, d_\infty)$ is geodesic but not uniquely geodesic. (A space is called *uniquely geodesic* if every pair of points can be joined by a unique geodesic.)
2. Is the space $(\mathbb{R}^2 \setminus \{0\}, d_\infty)$ geodesic?

Illustrate your arguments with suitable pictures!

Exercise 5.E.11 (Geometric realisation of Cayley graphs**).

1. Show that $|\mathrm{Cay}(\mathbb{Z}, \{1\})|$ is isometric to the real line $\mathbb{R}$.
2. Show that $|\mathrm{Cay}(\mathbb{Z}^2, \{(1,0),(0,1)\})|$ is isometric to $\mathbb{R} \times \mathbb{Z} \cup \mathbb{Z} \times \mathbb{R} \subset \mathbb{R}^2$ with the metric induced from the ℓ^1-metric on $\mathbb{R}^2$.
3. Is the geometric realisation of the Cayley graph $\mathrm{Cay}(F, S)$ of a free group F of rank 2 with respect to a free generating set S of F isometric to a subset of $\mathbb{R}^2$ (with respect to the Euclidean metric)?
4. What is the relation between the geometric realisation of the Cayley graph $\mathrm{Cay}(\mathbb{Z}/2017, \{[1]\})$ and the circle S^1 ?

Exercise 5.E.12 (Quasi-geodesic $\rightsquigarrow$ geodesic**).

1. Show that the geometric realisation of a connected graph is a geodesic metric space.

2. Let $X = (V, E)$ be a connected graph. Show that the canonical map $V \hookrightarrow |X|$ is an isometric embedding that is a quasi-isometry (where we equip V with the metric induced from the graph structure on X).
3. Fill in the details of the proof of Proposition 5.3.9 to show that every quasi-geodesic metric space is quasi-isometric to a geodesic metric space.

Exercise 5.E.13 (Infinite geodesics in groups**)**.** Let G be a finitely generated group with finite generating set S. Show that G is infinite if and only if $|\mathrm{Cay}(G, S)|$ contains an infinite geodesic.
Hints. Exercise 3.E.11 might help.

Group actions and quasi-isometry

Quick check 5.E.14 (Švarc-Milnor lemma*)**.** For each of the following group actions name one of the conditions of the Švarc-Milnor lemma that is satisfied, and one that is not.

1. The action of $\mathrm{SL}(2, \mathbb{Z})$ on $\mathbb{R}^2$ given by matrix multiplication.
2. The action of $\mathbb{Z}$ on $X := \{(r^3, s) \mid r, s \in \mathbb{Z}\}$ (with respect to the metric induced from the Euclidean metric on $\mathbb{R}^2$) that is given by

$$\begin{aligned} \mathbb{Z} \times X &\longrightarrow X \\ (n, (r^3, s)) &\longmapsto (r^3, s + n). \end{aligned}$$

Exercise 5.E.15 (Švarc-Milnor lemma via quasi-isometric actions***)**.** Formulate and prove a truly quasi-geometric version of the Švarc-Milnor lemma, i.e., a version of the Švarc-Milnor lemma where the given group action is an action by quasi-isometries instead of isometries.

Exercise 5.E.16 (Commensurability*)**.**

1. Show that commensurability forms an equivalence relation on the class of all groups.
2. Show that weak commensurability forms an equivalence relation on the class of all groups.

Hints. Exercise 2.E.5 might help.

Quick check 5.E.17 (A set-theoretic coupling of $\mathbb{Z}$ with $\mathbb{Z}$?*)**.** Let $G := \mathbb{Z}$ and $H := \mathbb{Z}$. We consider the left action of G on $\mathbb{R}^2$ given by

$$\begin{aligned} G \times \mathbb{R}^2 &\longrightarrow \mathbb{R}^2 \\ (n, (x, y)) &\longmapsto (n + x, y) \end{aligned}$$

and the right action of H on $\mathbb{R}^2$ given by

$$\begin{aligned} \mathbb{R}^2 \times H &\longrightarrow \mathbb{R}^2 \\ \bigl((x,y),n\bigr) &\longmapsto (x, y+n). \end{aligned}$$

Is $\mathbb{R}^2$ together with these actions of G and H a set-theoretic coupling of G and H ?

Exercise 5.E.18 (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.

Exercise 5.E.19 (Measure equivalence***). Look up in the literature when two groups are called *measure equivalent*, and compare this definition with the dynamic criterion for quasi-isometry.

Exercise 5.E.20 (Möbius transformations***). The group $\mathrm{SL}(2,\mathbb{Z})$ acts on the upper halfplane $H := \bigl\{z \in \mathbb{C} \bigm| \mathrm{Re}(z) > 0\bigr\} \subset \mathbb{C}$ via Möbius transformations by

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

More details on this action can also be found in Appendix A.3.

1. Let $D := \bigl\{z \in H \bigm| |z| \geq 1 \text{ and } |\mathrm{Re}\, z| \leq 1/2\bigr\}$, see Figure 5.14. Show that
$$\mathrm{SL}(2,\mathbb{Z}) \cdot D = H.$$
Illustrate your arguments by suitable pictures!
2. Prove that this action is *not* cocompact with respect to the standard topology on $H \subset \mathbb{C}$.

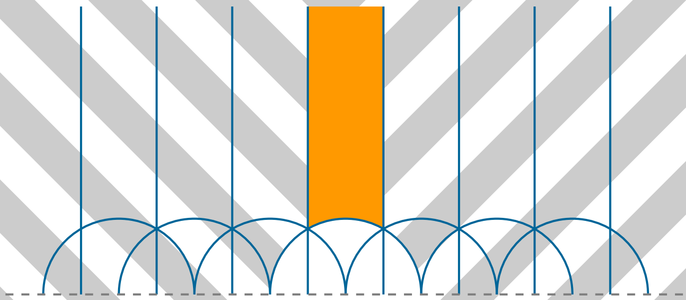

Figure 5.14.: The subset D in the upper halfplane and some geodesics

Exercise 5.E.21 (The Heisenberg group as a lattice**). Let $H_\mathbb{R}$ be the *real Heisenberg group* and let $H \subset H_\mathbb{R}$ be the Heisenberg group:

$$H_\mathbb{R} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \; x,\ y,\ z \in \mathbb{R} \right\}, \quad H := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \; x,\ y,\ z \in \mathbb{Z} \right\}.$$

We equip $H_\mathbb{R}$ with the topology given by convergence of all matrix coefficients.

1. Show that $H_\mathbb{R}$ is a locally compact topological group with respect to this topology.
2. Show that the subgroup H is a cocompact lattice in $H_\mathbb{R}$.

Quasi-isometry of groups

Quick check 5.E.22 (Homomorphisms and quasi-isometry*). Characterise all group homomorphisms between finitely generated groups that are quasi-isometries.

Exercise 5.E.23 (Quasi-dense subgroups*). Let G be a finitely generated group with finite generating set $S \subset G$ and let $H \subset G' \subset G$ be subgroups. Moreover, let $H \subset G'$ be quasi-dense with respect to d_S, i.e., there exists a $c \in \mathbb{R}_{>0}$ with

$$\forall_{g \in G'} \ \exists_{h \in H} \ d_S(g,h) \leq c.$$

Prove that then H has finite index in G'.

Exercise 5.E.24 (Groups not quasi-isometric to $\mathbb{Z}$ **). Let $n \in \mathbb{N}_{\geq 2}$.

1. Show that every 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 the cross $X := (\mathbb{Z} \times \{0\}) \cup (\{0\} \times \mathbb{Z})$ 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.

Exercise 5.E.25 ((Free) products and bilipschitz equivalence*). Let G, G' and H be finitely generated groups and suppose that G and G' are bilipschitz equivalent.

1. Are then also $G \times H$ and $G' \times H$ bilipschitz equivalent?
2. Are then also $G * H$ and $G' * H$ bilipschitz equivalent?

Exercise 5.E.26 ((Free) products and quasi-isometry**). Let G, G' and H be finitely generated groups and suppose that G and G' are quasi-isometric.

1. Are then also $G \times H$ and $G' \times H$ quasi-isometric?
2. Are then also $G * H$ and $G' * H$ quasi-isometric?

Quick check 5.E.27 (Geometric properties*).

1. Is the property of being generated by 2017 elements a geometric property of finitely generated groups?
2. Is the property of being isomorphic to a subgroup of $\mathbb{Z}^{2017}$ a geometric property of finitely generated groups?
3. Is the property of being infinite and torsion-free a geometric property of finitely generated groups?
4. Is the property of being a 2017-torsion group a geometric property of finitely generated groups?
5. Is the property of being a free product of two non-trivial groups a geometric property of finitely generated groups?

Quick check 5.E.28 (Quasi-isometries of groups*).

1. Let G and H be finitely generated groups and suppose that there is a quasi-isometric embedding $G \longrightarrow H$. Does this imply that there is a quasi-isometric embedding $H \longrightarrow G$?
2. For which finitely generated groups G are G and $\operatorname{Hom}(G, \mathbb{Z}/2)$ quasi-isometric?

Exercise 5.E.29 (Diameters**). Let F be the set of all generating sets of the symmetric group S_3. Determine the maximal diameter of S_3, i.e., determine the quantity

$$\max_{S \in F} \operatorname{diam} \operatorname{Cay}(S_3, S).$$

Exercise 5.E.30 (Coarse equivalence***).

1. Look up the terms *coarse embedding* and *coarse equivalence* in the literature.
2. Is every inclusion of a finitely generated subgroup of a finitely generated group a coarse embedding?
3. Is every coarse embedding between finitely generated groups a quasi-isometric embedding?
4. Show that finitely generated groups are coarsely equivalent if and only if they are quasi-isometric.

Uniformly finite homology$^+$

We will briefly describe the construction of uniformly finite homology via explicit geometric chains:

Let R be a commutative normed ring with unit (e.g., $\mathbb{R}$ or $\mathbb{Z}$), let (X, d) be a UDBG space, and let $n \in \mathbb{N}$. We then write $C_n^{\mathrm{uf}}(X; R)$ for the R-module of all bounded functions $c \colon X^{n+1} \longrightarrow R$ satisfying the following property: There is an $r \in \mathbb{R}_{>0}$ with

$$\forall_{x \in X^{n+1}} \quad \max\bigl\{d(x_j, x_k) \bigm| j, k \in \{0, \dots, n\}\bigr\} \geq r \Longrightarrow c(x) = 0.$$

The elements of $C_n^{\mathrm{uf}}(X;R)$ are called *uniformly finite chains in X with coefficients in R*; usually, such functions $c\colon X^{n+1} \longrightarrow R$ are denoted as formal sums of the form $\sum_{x\in X^{n+1}} c(x)\cdot x$.

Exercise 5.E.31 (Uniformly finite homology, boundary operator**)**.** Let X be a UDBG space, let $n \in \mathbb{N}_{>0}$ and let R be a normed ring. Show that

$$\begin{aligned}\partial_n\colon C_n^{\mathrm{uf}}(X;R) &\longrightarrow C_{n-1}^{\mathrm{uf}}(X;\mathbb{R})\\ \sum_{x\in X^{n+1}} c_x\cdot x &\longmapsto \sum_{x\in X^{n+1}}\sum_{j=0}^{n}(-1)^j\cdot c_x\cdot(x_0,\dots,x_{j-1},x_{j+1},\dots,x_n)\end{aligned}$$

describes a well-defined R-linear map that satisfies

$$\partial_n\circ\partial_{n+1}=0.$$

In addition, we define $\partial_0 := 0$. Hence, $C_*^{\mathrm{uf}}(X;\mathbb{R})$ is a chain complex with respect to ∂_*.

Exercise 5.E.32 (Uniformly finitely homology, induced chain map**)**.** Let X and Y be UDBG spaces, let R be a normed ring, and let $f\colon X\longrightarrow Y$ be a quasi-isometric embedding. Prove that for $n\in\mathbb{N}$ the expression

$$\begin{aligned}C_n^{\mathrm{uf}}(f;R)\colon C_n^{\mathrm{uf}}(X;R) &\longrightarrow C_n^{\mathrm{uf}}(Y;R)\\ \sum_{x\in X^{n+1}} c_x\cdot x &\longmapsto \sum_{x\in X^{n+1}} c_x\cdot\bigl(f(x_0),\dots,f(x_n)\bigr)\end{aligned}$$

describes a well-defined R-linear map that satisfies

$$\partial_{n+1}\circ C_{n+1}^{\mathrm{uf}}(f;R) = C_n^{\mathrm{uf}}(f;R)\circ\partial_{n+1}.$$

I.e., $C_*^{\mathrm{uf}}(f;R)\colon C_*^{\mathrm{uf}}(X;R)\longrightarrow C_*^{\mathrm{uf}}(Y;R)$ is a chain map.

Definition 5.E.1 (Uniformly finite homology)**.** Let R be a normed ring and let $n\in\mathbb{N}$. If X and Y are UDBG spaces and $f\colon X\longrightarrow Y$ is a quasi-isometric embedding, we define *uniformly finite homology of X* by

$$H_n^{\mathrm{uf}}(X;R) := \frac{\ker(\partial_n\colon C_n^{\mathrm{uf}}(X;R)\to C_{n-1}^{\mathrm{uf}}(X;R))}{\mathrm{im}(\partial_{n+1}\colon C_{n+1}^{\mathrm{uf}}(X;R)\to C_n^{\mathrm{uf}}(X;R))}$$

and the *map induced by f* by

$$\begin{aligned}H_n^{\mathrm{uf}}(f;R)\colon H_n^{\mathrm{uf}}(X;R) &\longrightarrow H_n^{\mathrm{uf}}(Y;R)\\ [c] &\longmapsto \bigl[C_n^{\mathrm{uf}}(f;R)(c)\bigr].\end{aligned}$$

Exercise 5.E.33 (Uniformly finite homology**)**.** Let R be a normed ring and let $n\in\mathbb{N}$.

1. Show that $H_n^{\mathrm{uf}}(\,\cdot\,;R)$ is well-defined (both on objects and on morphisms).
2. Show that $H_n^{\mathrm{uf}}(\,\cdot\,;R)\colon \mathsf{UDBG} \longrightarrow {}_R\mathsf{Mod}$ is a well-defined functor.
3. Conclude that quasi-isometries between UDBG spaces induce isomorphisms in uniformly finite homology.

In particular, we can define uniformly finite homology $H_*^{\mathrm{uf}}(G;R)$ for finitely generated groups G (by choosing a finite generating set and taking the associated word metric on G).

Exercise 5.E.34 (Uniformly finite homology of bounded spaces*)**.** Determine the uniformly finite homology of UDBG spaces of bounded diameter.

Alternative descriptions of uniformly finite homology can be obtained via simplicial topology and group homology:

Exercise 5.E.35 (Uniformly finite homology via coarsening***)**.** Show that uniformly finite homology is naturally isomorphic to the coarsening of locally finite simplicial homology with uniformly bounded coefficients.
Hints. How does the Rips complex construction translate into metric conditions on chains?

Exercise 5.E.36 (Uniformly finite homology of groups***)**.** Let R be a normed ring. Show that for finitely generated groups G, uniformly finite homology $H_*^{\mathrm{uf}}(G;R)$ with coefficients in R is naturally isomorphic to group homology $H_*(G;\ell^\infty(G;R))$ with ℓ^∞-coefficients.
Hints. A quick description of group homology as well as suitable references for group homology are given in Appendix A.2.

Part III

Geometry of groups

6

Growth types of groups

The first quasi-isometry invariant we discuss in detail is the growth type. Basically, we measure the "volume" 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 Gromov's polynomial growth theorem, all finitely generated groups of polynomial growth are virtually nilpotent! We will discuss this theorem in Chapter 6.3. In contrast, in Chapter 6.4 we will study some aspects of groups with exponential growth.

Overview of this chapter

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,

$$B_r^{G,S}(e) := B_r^{G,d_S}(e) = \{g \in G \mid d_S(g,e) \leq r\}$$

denotes the (closed) ball of radius r around e with respect to the word metric d_S on G.

This definition makes sense because balls for word metrics with respect to finite generating sets are finite (Remark 5.2.10).

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.
- 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):

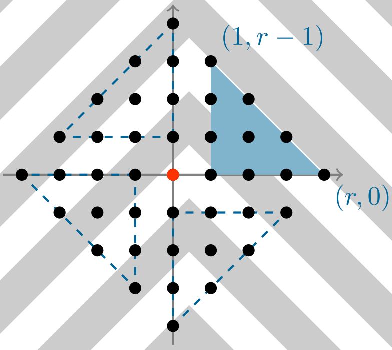

Figure 6.1.: The r-ball in $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ around $(0,0)$

$$\begin{aligned} \beta_{\mathbb{Z}^2,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longrightarrow 1 + 4 \cdot \sum_{j=1}^{r} (r+1-j) = 2 \cdot r^2 + 2 \cdot r + 1. \end{aligned}$$

- More generally, if $n \in \mathbb{N}$, then the growth functions of $\mathbb{Z}^n$ grow like a polynomial of degree n (Exercise 6.E.2).
- 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 6.E.6).
- 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 $r \in \mathbb{N}$.*

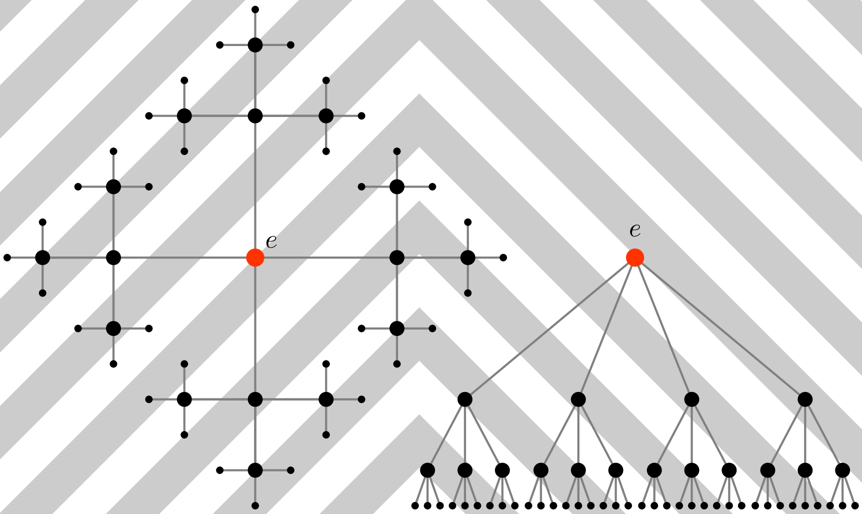

Figure 6.2.: The 3-ball in $\operatorname{Cay}\big(F(\{a,b\}),\{a,b\}\big)$ around e, drawn in two ways

3. A general upper bound. *For all $r \in \mathbb{N}$ we have*

$$\beta_{G,S}(r) \leq \beta_{F(S),S}(r).$$

Proof. Ad 1./2. This follows easily from the definition of the word metric d_S on G (Exercise 6.E.5).

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

$$\beta_{G,S}(r) = \big|B_r^{G,S}(e)\big| = \big|\varphi(B_r^{F(S),S}(e))\big| \leq \big|B_r^{F(S),S}(e)\big| = \beta_{F(S),S}(r)$$

for all $r \in \mathbb{N}$. The growth function $\beta_{F(S),S}$ is calculated in Example 6.1.2. $\square$

6.2 Growth types of groups

As we have seen, different finite generating sets can lead to different growth functions; however, one might suspect already that growth functions coming from different generating sets only differ by uniform multiplicative and additive error terms.

6.2.1 Growth types

We therefore introduce the following notion of equivalence for growth functions:

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

$$\forall_{r \in \mathbb{R}_{\geq 0}} \;\; f(r) \leq c \cdot g(c \cdot r + b) + b.$$

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 is *not* total (Exercise 6.E.4).

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

$$(x \mapsto x^a) \prec (x \mapsto x^{a'}) \iff a \leq a',$$

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 6.E.3). Moreover, there exist generalised growth functions f such that

$$\begin{aligned} f &\prec (x \mapsto a^x) \qquad \text{and} \qquad f \not\sim (x \mapsto a^x), \text{ and} \\ f &\succ (x \mapsto x^{a'}) \qquad \text{and} \qquad f \not\sim (x \mapsto x^{a'}) \end{aligned}$$

holds for all $a \in \mathbb{R}_{>1}$, $a' \in \mathbb{R}_{\geq 0}$ (Exercise 6.E.3).

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.

More explicitly, $\beta_{G,S}$ 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 will 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-isometries)**.** *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). ☐

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 Example 6.1.2), 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). From Example 6.1.2 we obtain:

- If $n \in \mathbb{N}$, then $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$ (Exercise 6.E.2).
- The Heisenberg group has the growth type of $(x \mapsto x^4)$ (Exercise 6.E.6).
- Non-Abelian free groups of finite rank have the growth type of the exponential function $(x \mapsto e^x)$.

The groups $\mathbb{Z}^n$ and the Heisenberg group hence have polynomial growth, while non-Abelian free groups have exponential growth.

Example 6.2.8 (Quasi-isometry classification of 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 from Example 6.2.2, Example 6.2.7, and Corollary 6.2.6. Hence, also for all $m, n \in \mathbb{N}$:

$$\mathbb{R}^m \sim_{\mathrm{QI}} \mathbb{R}^n \iff m = n.$$

More generally: If A is a finitely generated Abelian group, then by the structure theorem of finitely generated Abelian groups, there is a unique number $r \in \mathbb{N}$ and a finite Abelian group T (unique up to isomorphism) with

$$A \cong \mathbb{Z}^r \oplus T;$$

one then defines $\mathrm{rk}_{\mathbb{Z}} A := r$. Hence, combining the above observation with Corollary 5.4.5, we obtain for all finitely generated Abelian groups A and A' the equivalence

$$A \sim_{\mathrm{QI}} A' \iff \mathrm{rk}_{\mathbb{Z}} A = \mathrm{rk}_{\mathbb{Z}} A'.$$

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 [103].

Example 6.2.9 (Distinguishing quasi-isometry types of basic groups).

- We obtain for the Heisenberg group H that $H \not\sim_{\mathrm{QI}} \mathbb{Z}^3$ (Example 6.2.7), 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! Even worse, H cannot quasi-isometrically embed into $\mathbb{Z}^3$.
- Let F be a non-Abelian free group of finite rank, and let $n \in \mathbb{N}$. Because F grows exponentially, but $\mathbb{Z}^n$ and H have polynomial growth, we obtain

$$F \not\sim_{\mathrm{QI}} \mathbb{Z}^n \qquad \text{and} \qquad F \not\sim_{\mathrm{QI}} H.$$

Grigorchuk was the first to show that there indeed exist groups that have intermediate growth [69][77, Chapter VIII]:

Theorem 6.2.10 (Existence of groups of intermediate growth). *There exists a finitely generated group of intermediate growth.*

An example of such a group is the first Grigorchuk group (Definition 4.E.2), which can be described via automorphisms of trees or as an automatic group. A strategy to prove that this group has intermediate growth is outlined in Exercise 6.E.13. Furthermore, this group also has several other interesting properties [77, Chapter VIII]; for example, it is a finitely generated infinite torsion group (Exercise 4.E.37), and it is commensurable to the direct product with itself (Exercise 4.E.36).

Proposition 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

$$d_{S'}(h,e) \le d_T(h,e) \le r,$$

and so $B_r^{H,T}(e) \subset B_r^{G,S'}(e)$. In particular,

$$\beta_{H,T}(r) \le \beta_{G,S'}(r),$$

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 $\beta_{H,T} \prec \beta_{G,S}$. □

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 (Exercise 6.E.18).

Caveat 6.2.13 (Distorted subgroups)**.** 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 (Exercise 2.E.32) 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, each 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 [120]:

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 every) finite generating set) of the fundamental group $\pi_1(M)$ (which is finitely generated by Corollary 5.4.10). 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

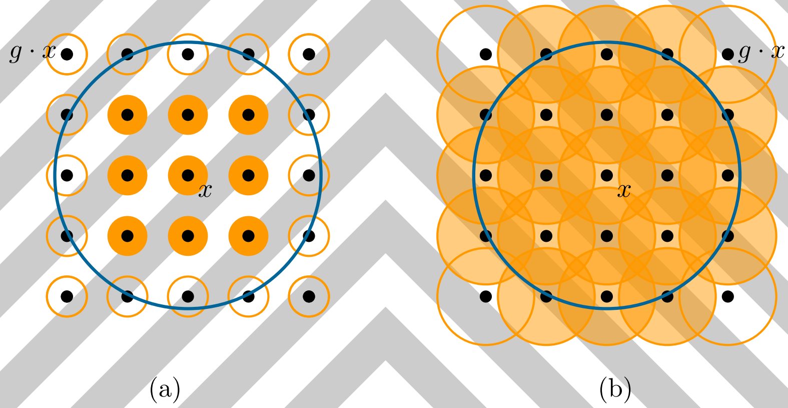

Figure 6.3.: Packing balls into balls, and covering balls by balls

Riemannian metric on $\widetilde{M}$*, and "*$\mathrm{vol}_{\widetilde{M}}$*" denotes the Riemannian volume with respect to the Riemannian metric on* $\widetilde{M}$*.*

Sketch of proof. By the Švarc-Milnor lemma (Corollary 5.4.10), the map

$$\begin{aligned} \varphi\colon \pi_1(M) &\longrightarrow \widetilde{M} \\ g &\longmapsto g \cdot x \end{aligned}$$

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 computation – using that φ is a quasi-isometry – then proves this claim.

More details are given in de la Harpe's book [77, 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*
$$\bigl(r \mapsto \operatorname{vol}_{\widetilde M} B_r^{\widetilde M}(x)\bigr) \succ \bigl(r \mapsto \operatorname{vol}_{\widetilde N} B_r^{\widetilde N}(y)\bigr)$$
holds for 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 *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 fundamental classes of the manifolds in question [50, Chapter VIII.4].

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 $\operatorname{im} \pi_1(p) = G$ [115, 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 $p \circ \overline{f} = f$ [115, Theorem V.5.1]:

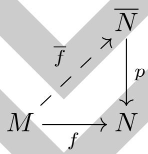

Looking at the induced diagram in singular homology with integral coefficients in top degree $n := \dim M = \dim N = \dim \overline{N}$ shows $H_n(\overline{N};\mathbb{Z}) \not\cong 0$:

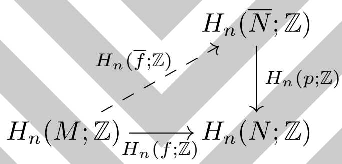

Namely, $H_n(M;\mathbb{Z}) \cong \mathbb{Z} \cong H_n(N;\mathbb{Z})$ and $H_n(\overline{f};\mathbb{Z})$ is multiplication by the mapping degree $\deg f$. In particular, $\overline{N}$ is compact [50, Corollary VIII.3.4]. On the other hand, $|\deg p|$ coincides with the number of sheets of the covering p [50, Proposition VIII.4.7], which in turn equals the index of $\operatorname{im} \pi_1(p)$ in $\pi_1(N)$ [115, p. 133]. Hence,

$$\left[\pi_1(N) : G\right] = \left[\pi_1(N) : \operatorname{im} \pi_1(p)\right] = |\deg p| < \infty,$$

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)**.** *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.*

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 [146, Chapter 3.4] (Proposition A.3.28 for $\mathbb{H}^2$). □

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 [100, 101].

6.3 Groups of polynomial growth

One of the milestones 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 [72] was subsequently simplified by van den Dries and Wilkie [52, 111]; alternative proofs have been given by Kleiner [90], Shalom and Tao [162], Ozawa [141], and Breuillard, Green, and Tao [26]. A complete proof is also given in the textbook by Druţu and Kapovich [53].

Theorem 6.3.1 (Gromov's polynomial growth theorem)**.** *Finitely generated groups have polynomial growth if and only if they are virtually nilpotent.*

In Chapter 6.3.1 and 6.3.2 we briefly discuss nilpotent groups and their growth properties; in Chapter 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.

Recall that 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 [94, 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 (Nilpotent/solvable groups)**.**

- 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 $n \in \mathbb{N}_{\geq 5}$ [94, 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 6.E.17).
- Free groups of rank at least 2 are *not* virtually solvable (Exercise 4.E.20).

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 normal in $C_{(j)}(G)$.*
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 (Exercise 6.E.14). □

Except for the last statement on centrality, the analogous statements also hold for the derived series instead of the lower central series (Exercise 6.E.15).

6.3.2 Growth of nilpotent groups

The growth type of finitely generated nilpotent (and hence of virtually nilpotent) groups can be expressed 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? The quotient group $C_{(j)}(G)/C_{(j+1)}(G)$ is Abelian (Proposition 6.3.5). Moreover, it can be shown that it is finitely generated (because G is finitely generated [179, Lemma 3.7][110, Theorem 5.4]). Therefore, $\operatorname{rk}_{\mathbb{Z}} C_{(j)}(G)/C_{(j+1)}(G)$ is well-defined.

The proof of Theorem 6.3.6 proceeds by induction over the nilpotence degree n and uses a suitable normal form of group elements in terms of the lower central series [14, 53]; the arguments from the computation of the growth rate of the Heisenberg group (Exercise 6.E.6) give a first impression of how this inductive proof works. We will refrain from going into the details for the general case.

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 6.E.18).

Wolf [179] and Milnor [120] 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.*

This theorem seems to be needed in all proofs of the polynomial growth theorem known so far.

6.3.3 Polynomial growth implies virtual nilpotence

We will now sketch Gromov's argument that finitely generated groups of polynomial growth are virtually nilpotent, mainly following the exposition by van den Dries and Wilkie [52]:

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 [72]:

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, Tao and Shalom, and Ozawa 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 [72]. Using the solution of Hilbert's fifth problem [125, 171], 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 $\mathrm{GL}(n, \mathbb{C})$ (Chapter 4.4.3)) allow to construct a surjective homomorphism from a finite index subgroup of G' to $\mathbb{Z}$ (see also Exercise 6.E.19).

A detailed proof is given in the paper by van den Dries and Wilkie [52]. Gromov's considerations of the sequence $(G, 1/n \cdot d_S)_{n \in \mathbb{N}}$ are a precursor of asymptotic cones [54, 53]. □

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 (Finding a subgroup of lower growth rate)**.** *Let $d \in \mathbb{N}$ and let G be a finitely generated group of polynomial growth of degree 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. This is proved in Exercise 6.E.20.

Ad 2. By the first part, we find a finite generating set $S \subset G$ that contains a finite generating set $T \subset K$ of K and that contains 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 [52, Lemma 2.1]:

Lemma 6.3.11. *Let*

$$1 \longrightarrow K \longrightarrow G \xrightarrow{\pi} \mathbb{Z} \longrightarrow 1$$

be an extension of groups, where G is finitely generated and K is virtually solvable. Then also G is virtually solvable.

Proof. Let $H \subset K$ be a solvable subgroup of K of finite index m. The intersection H' of all subgroups of index m in K is a solvable group (as subgroup of H) of finite index in K (Exercise 2.E.5) and all $\varphi \in \operatorname{Aut}(K)$ satisfy $\varphi(H') \subset H'$ (because automorphisms map subgroups of index m to subgroups of index m). Let $g \in G$ with $\pi(g) = 1 \in \mathbb{Z}$ and let

$$G' := \langle H' \cup \{g\} \rangle_G \subset G.$$

We now prove that G' is a solvable subgroup of G of finite index: The subgroup H' is normal in G' (because conjugation by g is an automorphism of K) and it follows that

$$G' \cap H = H'.$$

Therefore, G' fits into an extension $1 \longrightarrow H' \longrightarrow G' \longrightarrow \mathbb{Z} \longrightarrow 1$ and solvability of H' implies that G' is solvable. Moreover, a straightforward calculation shows that $[G : G'] = [K : H']$, which is finite. □

Let us continue with our previous considerations: Because G has polynomial growth, Theorem 6.3.8 lets us deduce that G indeed is virtually nilpotent, as desired. This finishes the sketch proof of Gromov's polynomial growth theorem.

6.3.4 Application: Virtual nilpotence is geometric

As a first application we show that being virtually nilpotent is a geometric property of finitely generated groups in the sense of Definition 5.6.6. In contrast, from the algebraic definition of virtually nilpotent groups it is not clear at all that this property is preserved under quasi-isometries.

Corollary 6.3.12. *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 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 the exponent d is a *natural number.*

Corollary 6.3.13 (Integrality of polynomial growth). *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. □

Corollary 6.3.14 (Integrality of polynomial growth of manifolds). *Let M be a closed connected Riemannian manifold whose Riemannian universal covering $\widetilde{M}$ has polynomial volume growth. Then there is a $d \in \mathbb{N}$ such that for all $x \in \widetilde{M}$ we have*

$$\bigl(r \mapsto \operatorname{vol}_{\widetilde{M}} B_r^{\widetilde{M}}(x)\bigr) \sim (n \mapsto n^d).$$

Proof. The volume growth of $\widetilde{M}$ coincides with the growth type of the fundamental group $\pi_1(M)$ (Proposition 6.2.14). Therefore, the previous corollary implies integrality of the growth exponent. □

6.3.6 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 in the following sense:

Corollary 6.3.15 (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 more elementary proofs of the quasi-isometry rigidity of $\mathbb{Z}$ in Chapter 7 (Corollary 7.5.8) and Chapter 8 (Exercise 8.E.11).

More generally, a similar argument yields quasi-isometry rigidity of higher-dimensional Abelian groups [32, Theorem 5.8]:

Corollary 6.3.16 (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 [31, p. 149f]. □

It turns out that it is also possible to prove quasi-isometry rigidity of $\mathbb{Z}^n$ without referring to the polynomial growth theorem [161, 44]. On the other hand, a full quasi-isometry classification of virtually nilpotent groups is out of reach.

6.3.7 Application: Expanding maps of manifolds

We conclude the discussion of polynomial growth with Gromov's geometric application [72, Geometric corollary on p. 55] of the polynomial growth theorem (Theorem 6.3.1) to infra-nil-endomorphisms:

Corollary 6.3.17. *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

$$\forall_{x,x' \in X} \quad x \neq x' \Longrightarrow d_Y\bigl(f(x), f(x')\bigr) > d_X(x, x').$$

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:

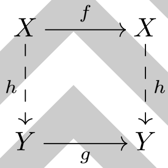

Sketch of proof of Corollary 6.3.17. By a theorem of Franks [61], if a compact Riemannian manifold M admits an expanding self-map, then the Riemannian universal covering $\widetilde{M}$ has polynomial volume growth; hence, by the Švarc-Milnor lemma (Proposition 6.2.14), the fundamental group $\pi_1(M)$ 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 [163, 61], this implies that every expanding self-map of M is topologically conjugate to an infra-nil-endomorphism. □

6.4 Groups of uniform exponential growth

In contrast to Chapter 6.3, we will now focus on groups of exponential growth. We will first introduce a stronger version of exponential growth (Chapter 6.4.1). We will then discuss the interesting relation between exponential growth and number theory (Chapter 6.4.2–6.4.4), as discovered by Breuillard [25, 27].

6.4.1 Uniform exponential growth

The rate of exponential growth can be measured as follows.

Proposition 6.4.1 (Exponential growth rate). *Let G be a finitely generated group and let $S \subset G$ be a finite generating set of G.*

1. *Then the sequence $\bigl((\beta_{G,S}(n))^{1/n}\bigr)_{n\in\mathbb{N}}$ is convergent. The limit*

$$\varrho_{G,S} := \lim_{n\to\infty} \bigl(\beta_{G,S}(n)\bigr)^{1/n} = \inf_{n\in\mathbb{N}_{>1}} \bigl(\beta_{G,S}(n)\bigr)^{1/n}$$

is the exponential growth rate of G with respect to S.

2. *Then $\varrho_{G,S} > 1$ if and only if G has exponential growth.*

Proof. The proof consists of a standard argument for sub-multiplicative sequences: We abbreviate $(a_n)_{n\in\mathbb{N}} := (\beta_{G,S}(n))_{n\in\mathbb{N}}$ and $a := \inf_{n\in\mathbb{N}_{>0}} a_n^{1/n}$.

Ad 1. The growth function $\beta_{G,S}$ is submultiplicative (Proposition 6.1.3). Therefore, inductively we obtain

$$a_{N\cdot m+r} \leq {a_N}^m \cdot {a_1}^r$$

for all $N, m, r \in \mathbb{N}_{>0}$. Let $N \in \mathbb{N}_{>0}$. Every $n \in \mathbb{N}_{>0}$ can be written in the form

$$n = N \cdot m_n + r_n$$

with $m_n \in \mathbb{N}$ and $r_n \in \{1, \dots, N\}$. Therefore, we have

$$\begin{aligned} a \leq {a_n}^{\frac{1}{n}} &\leq {a_N}^{\frac{m_n}{n}} \cdot {a_1}^{\frac{r_n}{n}} \\ &= ({a_N}^{\frac{1}{N}})^{\frac{N\cdot m_n}{n}} \cdot {a_1}^{\frac{r_n}{n}}. \end{aligned}$$

Because of $\lim_{n\to\infty} r_n/n = 0$ and $a_1 \neq 0$ the second factor converges to 1 for $n \to \infty$; moreover,

$$\frac{N \cdot m_n}{n} = \frac{n - r_n}{n}$$

converges to 1 for $n \to \infty$. Therefore, $\limsup_{n\to\infty} a_n \leq a_N^{1/N}$. Taking the infimum over all $N \in \mathbb{N}_{>0}$ proves

$$a \leq \liminf_{n\to\infty} {a_n}^{1/n} \leq \limsup_{n\to\infty} {a_n}^{1/n} \leq a.$$

In particular, $({a_n}^{1/n})_{n\in\mathbb{N}}$ is convergent with limit $a = \inf_{n\in\mathbb{N}_{>0}} {a_n}^{1/n}$.

Ad 2. Clearly, $\inf_{n\in\mathbb{N}_{>0}} {a_n}^{1/n} > 1$ if and only if $(a_n)_{n\in\mathbb{N}}$ has the growth type of $(x \mapsto e^x)$. In view of the first part we hence obtain $\varrho_{G,S} > 1$ if and only if G has exponential growth. □

Example 6.4.2 (Exponential growth rates of free groups)**.** Let $n \in \mathbb{N}_{\geq 1}$ and let $S \subset F_n$ be a free generating set. Then the calculation of the growth function $\beta_{F_n,S}$ in Example 6.1.2 shows that $\varrho_{F_n,S} = 2 \cdot n - 1$.

In general, the exact value of the exponential growth rate does depend on the finite generating set (Exercise 6.E.27).

Definition 6.4.3 (Uniform exponential growth)**.** A finitely generated group G has *uniform exponential growth* if

$$\inf\{\varrho_{G,S} \mid S \subset G \text{ is a finite generating set}\} > 1.$$

Example 6.4.4 (Uniform exponential growth of free groups)**.** The free group F_2 of rank 2 has uniform exponential growth: Let $S \subset F_2$ be a generating set. In particular, $|S| \geq 2$ and using the Nielsen-Schreier theorem (Corollary 4.2.8) it is not hard to see that there exists a subset $T \subset S$ with $|T| = 2$ that generates a free subgroup F of rank 2. Therefore, we obtain

$$\varrho_{F_2,S} \geq \varrho_{F,T} = 3,$$

where the last equality follows from Example 6.4.2.

It seems to be an open problem to decide whether having uniform exponential growth is a quasi-isometry invariant or not [78].

It is known that finitely generated solvable groups [137] and finitely generated linear groups [60] have uniform exponential growth whenever they have exponential growth. However, there also exist finitely generated groups of exponential growth that do *not* have uniform exponential growth [176].

6.4.2 Uniform uniform exponential growth

In the case of linear groups, one can ask for another level of uniformity, namely uniformity in the base field:

Conjecture 6.4.5 (Breuillard's growth conjecture [25]). For every $d \in \mathbb{N}$ there exists $\varepsilon(d) \in \mathbb{R}_{>0}$ with the following property: For every field K and every finite set $S \subset \mathrm{GL}(d, K)$

- either $\varrho_{\langle S \rangle_{\mathrm{GL}(d,K)}, S} = 1$ and $\langle S \rangle_{\mathrm{GL}(d,K)}$ is virtually nilpotent
- or $\varrho_{\langle S \rangle_{\mathrm{GL}(d,K)}, S} > 1 + \varepsilon(d)$.

Remark 6.4.6 (Uniformity in dimension?). In the growth conjecture of Breuillard, uniformity of the growth gap in the dimension is impossible: Grigorchuk and de la Harpe [70] constructed out of the Grigorchuk group of intermediate growth a sequence $(G_n)_{n \in \mathbb{N}}$ of groups with the following properties:

- For every $n \in \mathbb{N}$, there exists a $d_n \in \mathbb{N}$ such that G_n is isomorphic to a subgroup of $\mathrm{GL}(d_n, \mathbb{Z})$ that is generated by a set $S_n \subset G_n$ of four elements. In this example, $\lim_{n \to \infty} d_n = \infty$.
- For every $n \in \mathbb{N}$ the group G_n has exponential growth and

$$\lim_{n \to \infty} \varrho_{G_n, S_n} = 1.$$

While the growth conjecture is open in general, partial results are known:

Theorem 6.4.7 (Growth gap [25]). *For every $d \in \mathbb{N}$ there exists $\varepsilon(d) \in \mathbb{R}_{>0}$ with the following property: For every field K and every finite subset S of $\mathrm{GL}(d, K)$ that generates a subgroup of $\mathrm{GL}(d, K)$ that is* not *virtually solvable, we have*

$$\varrho_{\langle S \rangle_{\mathrm{GL}(d,K)}, S} > 1 + \varepsilon(d).$$

The growth gap theorem is a consequence of the uniform Tits alternative (Theorem 6.4.8 below).

6.4.3 The uniform Tits alternative

By the Tits alternative (Theorem 4.4.7), finitely generated linear groups are either virtually solvable or they contain a free subgroup of rank 2. The uniform Tits alternative by Breuillard adds control on "how quickly" one can find in the latter case a free subgroup of rank 2:

Theorem 6.4.8 (Uniform Tits alternative [25]). *For every $d \in \mathbb{N}$ there exists $N(d) \in \mathbb{N}$ with the following property: For every field K and every finite subset $S \subset \mathrm{GL}(d, K)$ with $S^{-1} = S$ and $e \in S$ we have*

- *either $\langle S \rangle_{\mathrm{GL}(d,K)}$ is virtually solvable*
- *or the set $S^{N(d)}$ contains two elements that generate a free subgroup of $\mathrm{GL}(d, K)$ of rank 2.*

Breuillard's proof of the uniform Tits alternative follows the blueprint of the original proof of the Tits alternative (Chapter 4.4.3). We will briefly indicate some of the main steps for $d = 2$ in the characteristic 0 case: Arguments from model theory show that in order to prove the uniform Tits alternative for all fields of characteristic 0 it is sufficient to prove the uniform Tits alternative for the algebraic closure $\overline{\mathbb{Q}}$ of $\mathbb{Q}$. Let $S \subset \mathrm{SL}(2, \overline{\mathbb{Q}})$ be a finite subset with $S^{-1} = S$ and $e \in S$ such that the subgroup $G := \langle S \rangle_{\mathrm{SL}(2,\overline{\mathbb{Q}})}$ is *not* virtually solvable. For the classical Tits alternative, one proceeds as follows:

- Find a diagonalisable matrix $a \in G$ with an eigenvalue of (some) norm greater than 1.
- Find a matrix $b \in G$ such that the eigenspaces of a and b are not related.
- Take $\ell \in \mathbb{N}$ large enough (this will ensure that the eigenvalues of a^ℓ and $b \cdot a^\ell \cdot b^{-1}$ are large).
- Apply the ping-pong lemma to a^ℓ and $b \cdot a^\ell \cdot b^{-1}$ acting on $\overline{\mathbb{Q}}^2$ (or the projective line over $\overline{\mathbb{Q}}$) to conclude that
$$\langle a^\ell, b \cdot a^\ell \cdot b^{-1} \rangle_{\mathrm{SL}(2,\overline{\mathbb{Q}})}$$
is free of rank 2.

In order to promote this to a proof of the *uniform* Tits alternative for $\mathrm{SL}(2, \overline{\mathbb{Q}})$, one needs to control the S-word length of a and b, the size of the eigenvalues of a, and the "distance" between eigenobjects of a and b; then a quantitative version of the ping-pong lemma allows to control ℓ. In this context, the control of matrices, eigenvalues, and eigenobjects of matrices is formulated in terms of *heights*, a measurement of complexity of algebraic numbers, i.e., of elements of $\overline{\mathbb{Q}}$; a crucial ingredient for control of eigenvalues and eigenobjects in the non-virtually solvable case is Breuillard's height gap theorem [25], which builds on results from diophantine geometry.

From the uniform Tits alternative, the growth gap theorem can be derived by elementary means:

Proof of Theorem 6.4.7. Let $d \in \mathbb{N}$ and let $N(d) \in \mathbb{N}$ be as provided by the uniform Tits alternative (Theorem 6.4.8). We set
$$\varepsilon(d) := \frac{1}{2} \cdot (3^{1/N(d)} - 1) > 0.$$

If K is a field and $S \subset \mathrm{GL}(d, K)$ is finite (with $S^{-1} = S$ and $e \in S$) and $G := \langle S \rangle_{\mathrm{GL}(d,K)}$ is *not* virtually solvable, then by the uniform Tits alternative the set $S^{N(d)}$ contains two elements a, b that generate a free subgroup $F \subset G$ of rank 2. Therefore, we obtain (in combination with Example 6.4.4)

$$\begin{aligned} \varrho_{G,S} &= \varrho_{G,S^{N(d)}}{}^{1/N(d)} \\ &\geq \varrho_{F,\{a,b\}}{}^{1/N(d)} \geq 3^{1/N(d)} \\ &> 1 + \varepsilon(d), \end{aligned}$$

as desired. □

In addition to the growth gap theorem, the uniform Tits alternative also has various other applications in group theory, such as uniform girth estimates for linear groups (i.e., improvements of Theorem 4.4.6), uniform escape from torsion in linear groups, uniform non-amenability of (non-amenable) linear groups, uniform diameter estimates for finite groups, and applications to the structure theory of approximate groups [25].

6.4.4 Application: The Lehmer conjecture

We will now digress briefly to number theory and a beautiful relation between growth of linear groups and heights in number theory.

Let $\overline{\mathbb{Q}}$ be the algebraic closure of $\mathbb{Q}$ (for concreteness, we may assume that $\overline{\mathbb{Q}} \subset \mathbb{C}$). Let $\alpha \in \overline{\mathbb{Q}}^{\times}$ and let $f \in \mathbb{Z}[X]$ be its (integral) minimal polynomial; i.e., $f(\alpha) = 0$, the polynomial f is not constant, and the coefficients of f are coprime. Over $\mathbb{C}$ we can factor f as

$$f = a_d \cdot (X - \alpha_1) \cdot \dots \cdot (X - \alpha_d),$$

where $d := \deg f$ and $a_d \in \mathbb{Z}$ is the leading coefficient of f, and $\alpha_1, \dots, \alpha_d \in \mathbb{C}$ are the roots of f; in particular, $\alpha \in \{\alpha_1, \dots, \alpha_d\}$. Then the *Mahler measure of* α is defined by

$$M(\alpha) := |a_d| \cdot \prod_{j=1}^{d} \max\bigl(1, |\alpha_j|\bigr),$$

where $|\cdot|$ denotes the ordinary absolute value on $\mathbb{C}$. The Mahler measure has several alternative descriptions, e.g., in terms of the height of algebraic numbers and as a certain integral over f [21].

Numerical experiments support the following gap phenomenon:

Conjecture 6.4.9 (Lehmer conjecture)**.** There exists an $\varepsilon \in \mathbb{R}_{>0}$ such that: If $\alpha \in \overline{\mathbb{Q}}^{\times}$, then

- either α is a root of unity
- or $M(\alpha) > 1 + \varepsilon$.

Breuillard discovered that this open problem in number theory is strongly linked to growth of linear groups [25, 27].

Theorem 6.4.10 (Growth conjecture $\Longleftrightarrow$ Lehmer conjecture)**.**

1. *If the growth conjecture holds for* $\mathrm{GL}(2, \overline{\mathbb{Q}})$*, then the Lehmer conjecture is true.*
2. *If the Lehmer conjecture holds, then for each* $d \in \mathbb{N}$ *the growth conjecture holds for* $\mathrm{GL}(d, \overline{\mathbb{Q}})$*.*

The link between growth and Mahler measure is provided by the following linear groups:

Example 6.4.11. Let $\alpha \in \overline{\mathbb{Q}}$. We then set

$$A(\alpha) := \begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix}, \qquad B := \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

and

$$S(\alpha) := \{e, A(\alpha), A(\alpha)^{-1}, B, B^{-1}\} \subset \mathrm{GL}(2, \overline{\mathbb{Q}}).$$

If α is not a root of unity, then the subgroup $\langle S(\alpha)\rangle_{\mathrm{GL}(2,\overline{\mathbb{Q}})}$ is not virtually nilpotent (Exercise 6.E.28). Moreover, one can show that

$$\varrho_{\langle S(\alpha)\rangle_{\mathrm{GL}(2,\overline{\mathbb{Q}})}, S(\alpha)} \leq M(\alpha).$$

In particular, these groups show that validity of the growth conjecture for $\mathrm{GL}(2, \overline{\mathbb{Q}})$ would imply the Lehmer conjecture. It should be noted that the groups $\langle S(\alpha)\rangle_{\mathrm{GL}(2,\overline{\mathbb{Q}})}$ are virtually solvable (Exercise 6.E.28). Hence, despite of the growth gap theorem (Theorem 6.4.7) the Lehmer conjecture remains open.

The converse implication that the Lehmer conjecture implies the growth conjecture over the field $\overline{\mathbb{Q}}$ requires careful entropy estimates [27], which are far beyond the scope of this book.

6.E Exercises

Growth functions

Quick check 6.E.1 (Balls in Abelian groups*).
1. Is there a finite generating set $S \subset \mathbb{Z}^2$ satisfying
$$\beta_{\mathbb{Z}^2,S}(42) = 2016\ ?$$
Hints. Parity!
2. Let $n \in \mathbb{N}$. Is there a finite generating set $S \subset \mathbb{Z}^n$ such that for all $r \in \mathbb{N}$ we have
$$\beta_{\mathbb{Z}^n,S}(r) = \{-r,\dots,r\}^n\ ?$$

Exercise 6.E.2 (Growth of $\mathbb{Z}^n$ *). Let $n \in \mathbb{N}$. Show that $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$.
Hints. Pick nice generating sets!

Exercise 6.E.3 (Exponential generalised growth functions*).
1. Show that $(x \mapsto a^x) \sim (x \mapsto a'^x)$ holds for all $a, a' \in \mathbb{R}_{>1}$.
2. Let $a \in \mathbb{R}_{>1}$ and $a' \in \mathbb{R}_{>0}$. Show that $(x \mapsto a^x) \succ (x \mapsto x^{a'})$.
3. Let $a \in \mathbb{R}_{>1}$ and $a' \in \mathbb{R}_{>0}$. Show that $(x \mapsto a^x) \not\prec (x \mapsto x^{a'})$.
4. Find a generalised growth function $f\colon \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ with $f \prec (x \mapsto e^x)$, $f \not\sim (x \mapsto e^x)$, and $(x \mapsto x^a) \prec f$ for all $a \in \mathbb{R}_{>0}$.

Exercise 6.E.4 (Quasi-dominance*). Show that the the quasi-dominance relation "$\prec$" is *not* a total order on the set of generalised growth functions.

Exercise 6.E.5 (Basic properties of growth functions*). Let G be a finitely generated group and let $S \subset G$ be a finite generating set.
1. Show that the growth function $\beta_{G,S}$ is sub-multiplicative:
$$\forall_{r,r' \in \mathbb{N}}\ \ \beta_{G,S}(r + r') \leq \beta_{G,S}(r) \cdot \beta_{G,S}(r').$$
2. Prove that $\beta_{G,S}$ is strictly increasing if G is infinite.

Hints. Look at paths that realise the distance ...

Growth types of groups

Exercise 6.E.6 (Growth type of the Heisenberg group**). Let H be the Heisenberg group (Exercise 2.E.32). We consider the presentation
$$H \cong \bigl\langle x, y, z \bigm| [x,z], [y,z], [x,y] = z \bigr\rangle.$$

We write $S := \{x, y, z\}$ and view S as a subset of H. Let m, n, $k \in \mathbb{Z}$.

1. Show that $d_S(x^m \cdot y^n \cdot z^k, e) \leq |m| + |n| + 6 \cdot \sqrt{|k|}$.
2. Show that $|m| + |n| \leq d_S(x^m \cdot y^n \cdot z^k, e)$ and $|k| \leq d_S(x^m \cdot y^n \cdot z^k, e)^2$.
3. Show that $1/2 \cdot (|m| + |n| + \sqrt{|k|}) \leq d_S(x^m \cdot y^n \cdot z^k, e)$.
4. Conclude that the growth function $\beta_{H,S}$ is quasi-equivalent to a polynomial of degree 4.

Exercise 6.E.7 (Growth of surface groups**)**.** Find as many proofs as possible that the surface group (Exercise 2.E.23)

$$\bigl\langle a_1, a_2, b_1, b_2 \bigm| [a_1, b_1] \cdot [a_2, b_2] \bigr\rangle$$

has exponential growth.

Exercise 6.E.8 (Semi-ping-pong** [77, Proposition VII.2])**.** Let G be a finitely generated group that acts on a set X. Moreover, let $a, b \in G$ and let $A, B \subset X$ be non-empty subsets with the following properties:

$$A \cap B = \emptyset, \qquad a \cdot (A \cup B) \subset B, \qquad b \cdot (A \cup B) \subset A.$$

Show that then G has exponential growth.

Hints. Show by induction that the canonical map $\{a, b\}^* \longrightarrow G$ is injective.

Exercise 6.E.9 (A matrix group** [77, Example VII.3])**.** We consider the matrices

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

in $\mathrm{GL}(2, \mathbb{Q})$ and the subgroup $G := \langle a, b \rangle_{\mathrm{GL}(2,\mathbb{Q})}$ of $\mathrm{GL}(2, \mathbb{Q})$. Show that G has exponential growth.

Hints. Apply the semi-ping-pong (Exercise 6.E.8) to (large powers of) a^{-1} and (large powers of) $a^{-1} \cdot b$ and the action of G on $\mathbb{R}$ given by

$$\begin{pmatrix} A & B \\ 0 & 1 \end{pmatrix} \cdot x := A \cdot x + B$$

for all $A, B \in \mathbb{Q}$, $x \in \mathbb{R}$. Good candidates for the ping-pong table are small neighbourhoods of 0 and 1, respectively.

Exercise 6.E.10 (Growth of $\mathrm{BS}(1, 2)$ **)**.**

1. Show that the Baumslag-Solitar group $\mathrm{BS}(1, 2)$ has exponential growth.
   Hints. Exercise 6.E.9 and the matrices from Exercise 2.E.21 will help.
2. Why doesn't this result contradict the normal form developed in Exercise 2.E.22 ?!

Groups of intermediate growth

Exercise 6.E.11 (Groups of intermediate growth**)**.** Let G be a finitely generated group with the following properties:

- The group G is *not* of exponential growth,
- and G is quasi-isometric to $G \times G$.

Show that this implies that G is of intermediate growth.

Exercise 6.E.12 (A criterion for subexponential growth**)**.** Let G be a finitely generated group with finite generating set S and suppose that there exists a finite index subgroup $L \subset G$ as well as $n \in \mathbb{N}_{\geq 2}$, $c \in (0,1)$, $b \in \mathbb{R}_{\geq 0}$ and an injective group homomorphism $\psi \colon L \longrightarrow G^n$ satisfying

$$\forall_{g \in L} \quad \sum_{j=1}^{n} d_S\bigl(\psi(g)_j, e\bigr) \leq c \cdot d_S(g) + b.$$

Show that G does *not* have exponential growth.

Exercise 6.E.13 (Intermediate growth of the Grigorchuk group***)**.** We use the notation from Definition 4.E.2 ff. Show that the group homomorphism

$$\begin{aligned} L_3 &\longrightarrow \prod_{w \in \{0,1\}^3} \mathrm{Gri} \\ g &\longmapsto \bigl(\varphi_{j_1} \circ \varphi_{j_2} \circ \varphi_{j_3}(g)\bigr)_{j_1 j_1 j_3 \in \{0,1\}^3} \end{aligned}$$

satisfies the hypotheses of Exercise 6.E.12. Use Exercise 4.E.36 and 6.E.11 to conclude that the Grigorchuk group Gri is a group of intermediate growth.

Growth and nilpotence/solvability

Exercise 6.E.14 (Quotients of the lower central series*)**.** Let G be a group and let $j \in \mathbb{N}$.

1. Show that $C_{(j+1)}(G) \subset C_{(j)}(G)$.
2. Show that $C_{(j+1)}(G)$ is a normal subgroup of $C_{(j)}(G)$ and of G.
3. Show that $C_{(j)}(G)/C_{(j+1)}(G)$ is a central subgroup of $G/C_{(j+1)}(G)$. *Hints.* A subgroup C of a group H is *central* if for all $g \in C$ and all $h \in H$ we have $g \cdot h = h \cdot g$.
4. Conclude that the quotient group $C_{(j)}(G)/C_{(j+1)}(G)$ is Abelian.

Exercise 6.E.15 (Quotients of the derived series*)**.** Let G be a group and let $j \in \mathbb{N}$.

1. Show that $G^{(j+1)} \subset G^{(j)}$.
2. Show that $G^{(j+1)}$ is a normal subgroup of $G^{(j)}$, and that the quotient group $G^{(j)}/G^{(j+1)}$ is Abelian.

Definition 6.E.1. Let $n \in \mathbb{N}_{>1}$. For a matrix $A \in \mathrm{GL}(n, \mathbb{Z})$ we consider the semi-direct product

$$G_A := \mathbb{Z}^n \rtimes_A \mathbb{Z}$$

with respect to the homomorphism $\mathbb{Z} \longrightarrow \operatorname{Aut} \mathbb{Z}^n$ given by the action of the powers of A on $\mathbb{Z}^n$ by matrix multiplication.

Exercise 6.E.16 (Solvable semi-direct products*)**.** Let $n \in \mathbb{N}_{>1}$, $A \in \mathrm{GL}(n, \mathbb{Z})$. Show that then the group G_A (Definition 6.E.1) is solvable.

Exercise 6.E.17 ((Non-)ilpotent semi-direct products**)**.** We consider the groups constructed in Definition 6.E.1.

1. Prove without using the polynomial growth theorem: For
$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$
the group G_A is solvable but *not* nilpotent. A little bit more challenging: Show that in this case G_A is *not* virtually nilpotent.
2. Show that G_A is virtually nilpotent for
$$A = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}.$$

Exercise 6.E.18 (Solvable groups of exponential growth**)**.** Let $n \in \mathbb{N}_{>1}$ and let $A \in \mathrm{GL}(n, \mathbb{Z})$ be a matrix that over $\mathbb{C}$ has an eigenvalue $\lambda \in \mathbb{C}$ with $|\lambda| \geq 2$. We consider the associated semi-direct product $G_A = \mathbb{Z}^n \rtimes_A \mathbb{Z}$ (Definition 6.E.1).

1. Show that there is an $x \in \mathbb{Z}^n$ with the following property: If $k \in \mathbb{N}$, then the 2^{k+1} elements $\sum_{j=0}^{k} \varepsilon_j \cdot A^j \cdot x$ of $\mathbb{Z}^n$ with $\varepsilon_0, \dots, \varepsilon_k \in \{0, 1\}$ are all different.
2. Conclude that G_A has exponential growth.
3. Show that G_A does *not* contain a free group of rank 2.
   Hints. Exercise 4.E.20 will help.

Exercise 6.E.19 (Linear groups and projections on $\mathbb{Z}$ ***)**.** Let G be a finitely generated group with subexponential growth.

1. Suppose that there exists an $n \in \mathbb{N}$ such that there is a homomorphism $G \longrightarrow \mathrm{GL}(n, \mathbb{C})$ with infinite image. Show that G contains a finite index subgroup that admits a surjective homomorphism to $\mathbb{Z}$.
   Hints. The Tits alternative (Theorem 4.4.7) implies that the image is virtually solvable.
2. Suppose that there exists an $n \in \mathbb{N}$ such that for every $N \in \mathbb{N}$ there is a homomorphism $\varphi_N \colon G \longrightarrow \mathrm{GL}(n, \mathbb{C})$ with $N \leq |\operatorname{im} \varphi_N| < \infty$. Prove that then G contains a finite index subgroup that admits a surjective homomorphism to $\mathbb{Z}$.
   Hints. Use Jordan's theorem on finite subgroups of $\mathrm{GL}(n, \mathbb{C})$ and then apply a counting/diagonalisation argument.

Exercise 6.E.20 (Finite generation and projections on $\mathbb{Z}$ **). Let G be a finitely generated group with subexponential growth that admits a surjective homomorphism $\pi\colon G \longrightarrow \mathbb{Z}$. Show that then also the kernel of π is finitely generated.

Hints. Choose an element $g \in G$ with $\pi(g) = 1$. Show that there is a finite subset $S \subset \ker\varphi$ such that $\{g\} \cup S$ generates G. For $s \in S$ and $n \in \mathbb{N}$ let

$$g_{n,s} := g^n \cdot s \cdot g^{-n} \in \ker\varphi.$$

Prove that there is an $N \in \mathbb{N}$ such that $\ker\varphi$ is generated by the (finite!) set $\{g_{n,s} \mid s \in S, n \in \{-N,\dots,N\}\}$.

Exercise 6.E.21 (Die Hilbertschen Probleme**). On August 8, 1900, David Hilbert gave his famous speech *Mathematische Probleme* (Mathematical Problems) at the International Congress of Mathematicians in Paris. These problems are now known as *Hilbert's problems*.

1. Take a random number n between 1 and 23. Describe Hilbert's n-th problem and the status of its solution.
2. Which of Hilbert's problems is your favourite? Why?

Exercise 6.E.22 (Growth of torsion groups***). Show that finitely generated torsion groups are either finite or do *not* have polynomial growth.

Growth series$^+$

A paradigm of combinatorics is to organise counting invariants by means of the corresponding power series. In our context, this leads to growth series and (the slightly more convenient) spherical growth series of finitely generated groups:

Definition 6.E.2 (Growth series). Let G be a finitely generated group with finite generating set S.

- The *growth series of G with respect to S* is the formal power series

$$B_{G,S} := \sum_{n=0}^{\infty} \beta_{G,S}(n) \cdot X^n \in \mathbb{Z}[\![X]\!]$$

- The *spherical growth series of G with respect to S* is the formal power series (where we set $\beta_{G,S}(-1) := 0$)

$$\Sigma_{G,S} := \sum_{n=0}^{\infty} \bigl(\beta_{G,S}(n) - \beta_{G,S}(n-1)\bigr) \cdot X^n \in \mathbb{Z}[\![X]\!].$$

Quick check 6.E.23 (Radius of convergence*). Let G be a finitely generated group and let $S \subset G$ be a finite generating set. How are the radius of convergence of the growth series and the exponential growth rate $\varrho_{G,S}$ related?

Exercise 6.E.24 (Growth series of $\mathbb{Z}$ *).

1. Determine the spherical growth series $\Sigma_{\mathbb{Z},\{1\}}$. Can you express this also as a rational function?
2. Determine the spherical growth series $\Sigma_{\mathbb{Z},\{2,3\}}$. Can you express this also as a rational function?

Exercise 6.E.25 (Growth series of (free) products*). Let G and H be finitely generated groups with free generating sets S and T, respectively.

1. Show that $\Sigma_{G\times H,S\times\{e\}\cup\{e\}\times T} = \Sigma_{G,S}\cdot\Sigma_{H,T}$ (Cauchy product of formal power series). How are the (non-spherical) growth series related?
2. Show that

$$\Sigma_{G*H,S\sqcup T}\cdot\bigl(1-(\Sigma_{G,S}-1)\cdot(\Sigma_{H,T}-1)\bigr) = \Sigma_{G,S}\cdot\Sigma_{H,T}.$$

How can one reformulate this relation more concisely (using inverses)? Use these properties to determine spherical growth series of $\mathbb{Z}^n$ and of the free group of rank n with respect to suitable generating sets. Can you express these series also as rational functions?

Uniform exponential growth

Quick check 6.E.26 (Uniform exponential growth*). Let G and H be finitely generated groups and let G be of uniform exponential growth.

1. Does then also $G\times H$ have uniform exponential growth?
2. Does then also $G * H$ have uniform exponential growth?

Exercise 6.E.27 (Exponential growth rates for different generating sets*). Let F be a free group of rank 2, freely generated by $\{a,b\}$, and let

$$S := \{a, b, a\cdot b\cdot a^{-1}, a^2\}\subset F.$$

Show that

$$\varrho_{F,\{a,b\}}\neq\varrho_{F,S}.$$

Exercise 6.E.28 (Linear groups for number theory**). Let $\alpha\in\overline{\mathbb{Q}}$. We then set

$$A(\alpha) := \begin{pmatrix}\alpha & 0\\ 0 & 1\end{pmatrix},\qquad B := \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}$$

and

$$S(\alpha) := \bigl\{A(\alpha), B\bigr\}\subset \mathrm{GL}(2,\overline{\mathbb{Q}}),$$

as well as $G(\alpha) := \langle S(\alpha)\rangle_{\mathrm{GL}(2,\overline{\mathbb{Q}})}$.

1. Show that $G(\alpha)$ is virtually solvable.
2. Show that $G(\alpha)$ is virtually nilpotent if and only if α is a root of unity.

Exercise 6.E.29 (Uniform exponential growth and quasi-isometries?!$^{\infty}$*)**.** Let G and H be finitely generated quasi-isometric groups, where G has uniformly exponential growth. Does then also H have uniformly exponential growth?!
Hints. This is an open problem!

Dehn functions and isoperimetric inequalities$^+$

In the language of Cayley graphs, generators correspond to edges. In contrast, relations have a distinct two-dimensional flair, as seen in presentation complexes (Outlook 3.2.5). Therefore, it is natural to associate a notion of area with relations. Algebraically, this can be formalised in the following way:

Definition 6.E.3 (Area of a relation)**.** Let $\langle S \,|\, R\rangle$ be a finite presentation of a group G and let $w \in F_{\text{red}}(S)$ be a reduced word that represents the trivial element in G. The *area of w (with respect to $\langle S \,|\, R\rangle$)* is defined as

$$\begin{aligned}\operatorname{Area}_{\langle S \,|\, R\rangle}(w) := \min\bigl\{n \in \mathbb{N} \bigm| \exists_{a_1,\dots,a_n \in F(S)}\ \exists_{r_1,\dots,r_n \in R \cup R^{-1}}\\ w = a_1 \cdot r_1 \cdot a_1^{-1} \cdot \dots \cdot a_n \cdot r_n \cdot a_n^{-1} \text{ in } F(S)\bigr\}.\end{aligned}$$

The Dehn function encodes the maximal area that can be "surrounded" by a given length:

Definition 6.E.4 (Dehn function)**.** Let $\langle S \,|\, R\rangle$ be a finite presentation of a group G and let $\pi\colon F_{\text{red}}(S) \longrightarrow G$ be the canonical projection. Then the *Dehn function of G with respect to the presentation $\langle S \,|\, R\rangle$* is given by

$$\begin{aligned}\operatorname{Dehn}_{\langle S \,|\, R\rangle}\colon \mathbb{N} &\longmapsto \mathbb{N}\\ n &\longmapsto \max\bigl\{\operatorname{Area}(s_1 \cdots s_k) \bigm| k \in \{0,\dots,n\}, s_1 \dots s_k \in F_{\text{red}}(S),\\ &\qquad \pi(s_1 \dots s_k) = e \text{ in } G\bigr\}.\end{aligned}$$

Exercise 6.E.30 (Simple Dehn functions*)**.**

1. Determine the Dehn function of $\langle\,|\,\rangle$.
2. Determine the Dehn function of $\langle x \,|\,\rangle$.
3. Determine the Dehn function of $\langle x, y \,|\,\rangle$.
4. Determine the Dehn function of $\langle x \,|\, x\rangle$.
5. Determine the Dehn function of $\langle x, y \,|\, y\rangle$ (approximately).

Definition 6.E.5 (Dehn equivalence)**.** Let $f, g\colon \mathbb{N} \longrightarrow \mathbb{R}_{\geq 0}$ be increasing functions.

- The function f is *Dehn dominated* by g if there exists a $c \in \mathbb{N}$ with
$$\forall_{n \in \mathbb{N}}\ \ f(n) \leq c \cdot g(c \cdot n + c) + c \cdot n + c.$$
We then write $f \prec_D g$.
- The function f is *Dehn equivalent* to g if $f \prec_D g$ and $g \prec_D f$. If this is the case, we write $f \sim_D g$.

Exercise 6.E.31 (Dehn equivalence*).
1. Prove that Dehn domination is a partial order on the set of increasing functions of type $\mathbb{N} \longrightarrow \mathbb{R}_{\geq 0}$.
2. Prove that Dehn equivalence is an equivalence relation on the set of increasing functions of type $\mathbb{N} \longrightarrow \mathbb{R}_{\geq 0}$.
3. Prove that $(n \mapsto n)$ is *not* Dehn equivalent to $(n \mapsto n^2)$.
4. Prove that $(n \mapsto n^2)$ is *not* Dehn equivalent to $(n \mapsto 2^n)$.
5. Are Dehn equivalent functions always quasi-equivalent?

Exercise 6.E.32 (Dehn functions of Abelian groups**).
1. Show that the Dehn function of $\langle x, y \mid [x, y] \rangle$ is Dehn equivalent to the function $(n \mapsto n^2)$.
2. Show that the Dehn function of $\langle x, y, z \mid [x, y], [y, z], [x, z] \rangle$ is Dehn equivalent to the function $(n \mapsto n^2)$ (!).

Exercise 6.E.33 (Dehn functions and change of presentation*). Let G be a finitely presented group and let $\langle S \mid R \rangle$ and $\langle S' \mid R' \rangle$ be finite presentations of G. Show that

$$\mathrm{Dehn}_{\langle S \mid R \rangle} \sim_D \mathrm{Dehn}_{\langle S' \mid R' \rangle}.$$

Hints. Rewrite S' and R' in terms of S and R, and vice versa.

Exercise 6.E.34 (Dehn functions and quasi-isometry***). Let G, G' be finitely presented groups with finite presentations $\langle S \mid R \rangle$ and $\langle S' \mid R' \rangle$, respectively. Show that $\mathrm{Dehn}_{\langle S \mid R \rangle}$ and $\mathrm{Dehn}_{\langle S' \mid R' \rangle}$ are Dehn equivalent if G and G' are quasi-isometric.

Hints. It is convenient, to think about words that represent trivial elements as "cycles" (they are not necessarily actual cycles, because they might revisit vertices and edges ...) in Cayley graphs.

Let $f \colon G \longrightarrow G'$ and $f' \colon G' \longrightarrow G'$ be mutually quasi-inverse quasi-isometries. Let $w \in F_{\mathrm{red}}(S)$ be a reduced word that represents the trivial element of G. Then $f^*(w)$ is a word over $S' \cup S'^{-1}$ (but not necessarily reduced; this needs some attention). Connecting subsequent vertices in the "cycle" $f^*(w)$ through quasi-geodesic paths in $\mathrm{Cay}(G', S')$ leads to a "cycle" w' in $\mathrm{Cay}(G', S')$. One then considers an R'-filling of w' with minimal area.

This can be pushed back to G via f'^*. Because f and f' are quasi-inverse, $f'^*(w')$ will almost be w. Let w'' be the "cycle" obtained from $f'^*(w')$ through quasi-geodesic path connection of subsequent vertices. The difference between w'' and w can be R-filled efficiently (in terms of the length of w).

One now needs to figure out how to interpret fillings geometrically and how to translate the f'^*-image of the chosen R'-filling of w' into an R-filling of w'' ...

Exercise 6.E.35 (Quasi-isometry invariance of finite presentability***). Prove that finite presentability is a geometric property of groups.

Hints. The same method as in Exercise 6.E.34 can be used.

Definition 6.E.6 (Isoperimetric inequalities of groups). Let G be a finitely presented group and let $f \colon \mathbb{N} \longrightarrow \mathbb{R}_{\geq 0}$ be an increasing function. We say that *G satisfies an f-isoperimetric inequality* if for one (hence every) finite presentation $\langle S \mid R \rangle$ of G we have $\operatorname{Dehn}_{\langle S \mid R \rangle} \prec_D f$. We say that G satisfies a *linear, quadratic, ... isoperimetric inequality* if f can be taken to be linear, quadratic, ..., respectively.

Exercise 6.E.36 (Isoperimetric inequalities for Abelian groups*).
1. Show that all virtually cyclic groups satisfy a linear isoperimetric inequality.
2. Show that $\mathbb{Z}^2$ and $\mathbb{Z}^3$ satisfy quadratic isoperimetric inequalities, but *not* a linear isoperimetric inequality.

Exercise 6.E.37 (Isoperimetric inequality for $\operatorname{BS}(1,2)$ ***). We consider the standard presentation $\langle S \mid R \rangle := \langle a, b \mid bab^{-1} = a^2 \rangle$ of $\operatorname{BS}(1,2)$.
1. Warm-up: For $n \in \mathbb{N}$ let $w_n := [a, b^n a b^{-n}]$. Show that w_n represents the trivial element in G and that $\operatorname{Area}_{\langle S \mid R \rangle}(w_n) \leq 2^n$ (Figure 6.4).
2. Show that $\operatorname{Dehn}_{\langle S \mid R \rangle} \prec_D (n \mapsto 2^n)$.
   Hints. This can be shown by a careful analysis of the proof for the normal form in Exercise 2.E.22.
3. Show that $\operatorname{Dehn}_{\langle S \mid R \rangle} \succ_D (n \mapsto 2^n)$.
   Hints. The basic idea is to look at w_n from the first part and, in w_n, to look at the last b^{-1} in b^{-n}. The proof then requires a careful case-by-case analysis. It helps to organise relations etc. in a graphical way. Happy puzzling!

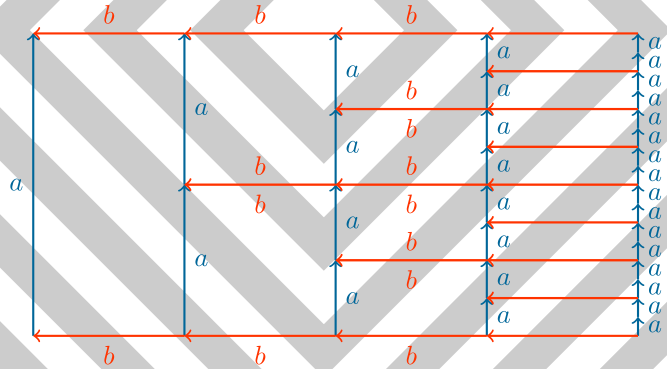

Figure 6.4.: A relation puzzle in $\operatorname{BS}(1,2)$

7

Hyperbolic groups

In the universe of groups (Figure 1.2), on the side opposite to Abelian, nilpotent, solvable, and amenable groups, we find free groups, and then further out negatively curved groups. This chapter is devoted to negatively curved groups.

The definition of negatively curved groups requires a notion of negative curvature that applies to Cayley graphs and that is invariant under change of finite generating sets, or more generally, under quasi-isometries. We will start with a quick reminder of classical curvature of plane curves and of surfaces (Chapter 7.1). We will then introduce Gromov's extension of the notion of negative curvature to large scale geometry via slim triangles (Chapter 7.2). In particular, this leads to a notion of negatively curved finitely generated groups: hyperbolic groups (Chapter 7.3).

The hyperbolicity condition for groups has far-reaching algebraic consequences: The word problem is solvable for hyperbolic groups (Chapter 7.4) and elements of infinite order in hyperbolic groups are well-behaved (Chapter 7.5). Chapter 7.6 contains a brief outlook on non-positively curved groups.

Overview of this chapter

7.1 Classical curvature, intuitively

Key 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. A particularly striking example of such a situation is the condition of 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. There are several ways of measuring such effects (using curves, triangles, angles, volumes, ...).

In the following, we will 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 will 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 [96].

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 γ at $\gamma(t)$. It can be shown that this set contains exactly one circle that at the point $\gamma(t)$ has the same acceleration vector as γ; this circle is the *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

$$\kappa_\gamma(t) = \bigl\|\ddot\gamma(t)\bigr\|_2.$$

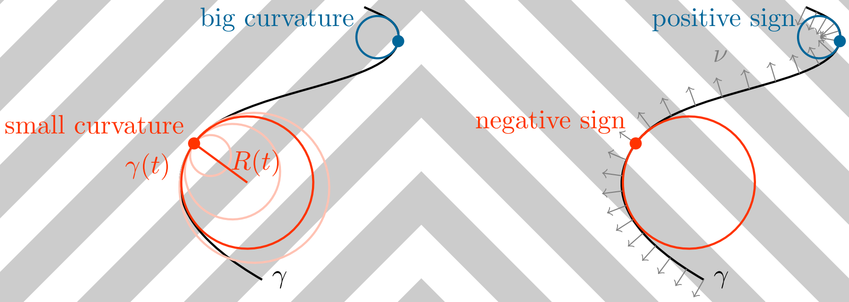

Figure 7.1.: Curvature of a plane curve

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.

7.1.2 Curvature of surfaces in $\mathbb{R}^3$

As second step, we consider 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 that is 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

$$\kappa^+_{S,\nu}(x) := \sup_V \widetilde\kappa_{\gamma_V,\nu|_{\gamma_V}}(x) \qquad \text{and} \qquad \kappa^-_{S,\nu}(x) := \inf_V \widetilde\kappa_{\gamma_V,\nu|_{\gamma_V}}(x).$$

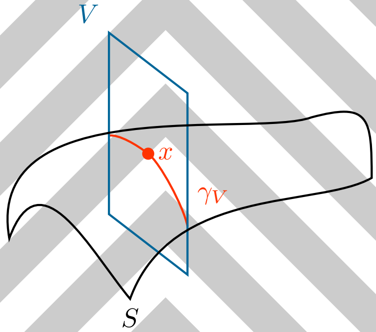

Figure 7.2.: Gaussian curvature of a surface via curves

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) [96, p. 6], the Gaussian curvatures of a surface are intrinsic (*Theorema Egregium* [96, Chapter 8]). More precisely, Gaussian curvature can be quantified as the dependence of parallel transport of tangent vectors along different curves.

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.
- An influential example for the history of geometry is the hyperbolic plane; an elementary introduction into the geometry of the hyperbolic plane is given in Appendix A.3.
  For example, one can calculate via the Theorema Egregium that the hyperbolic plane has everywhere negative Gaussian curvature (Theorem A.3.29).

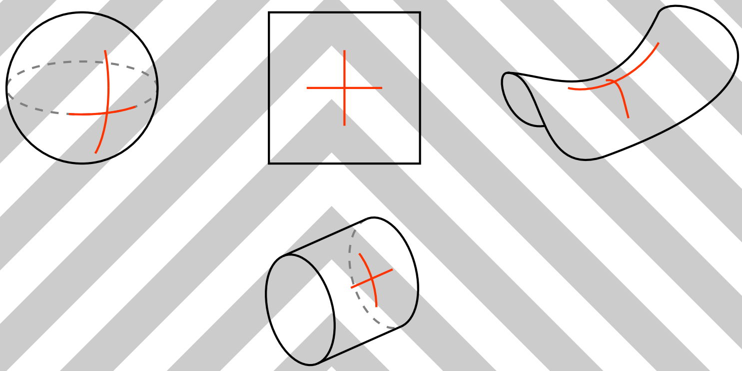

positive flat negative

Figure 7.3.: Examples of Gaussian curvatures of surfaces

Outlook 7.1.2 (Curvature of Riemannian manifolds)**.** How can one define curvature of higher dimensional Riemannian manifolds? Let M be a Riemannian manifold, and let $x \in M$. 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 invariant of Riemannian manifolds and can be described analytically in terms of 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 [96].

By construction, the Gaussian curvatures are defined in terms of the local structure of a surface, and so are not suited for a notion of curvature in large scale geometry. However, surprisingly, 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 geodesics in the surface in question.

In positively curved spaces, geodesic triangles are "fatter" than in Euclidean space, while in negatively curved spaces, geodesic triangles are "slimmer" than in Euclidean space. For example, all geodesic triangles in the hyperbolic plane are uniformly slim (Theorem A.3.27).

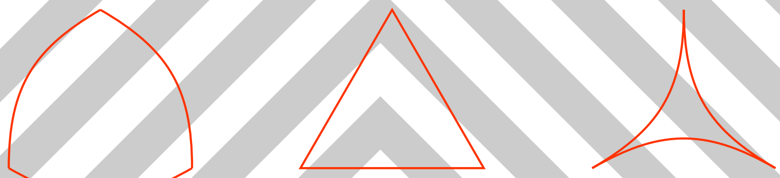

Figure 7.4.: Geodesic triangles in surfaces

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 [74]. Taking slim triangles as the *defining* property leads to the notion of (Gromov) hyperbolic spaces.

We first explain the notion of hyperbolicity for geodesic spaces (Chapter 7.2.1). As next step, we translate this notion to quasi-geometry (Chapter 7.2.2). Finally, in Chapter 7.2.3, we establish the quasi-isometry invariance of hyperbolicity.

7.2.1 Hyperbolic spaces

Taking slim geodesic triangles as defining property for negative curvature 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)

$$\begin{aligned} \operatorname{im}\gamma_0 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_1 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_2 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1\bigr). \end{aligned}$$

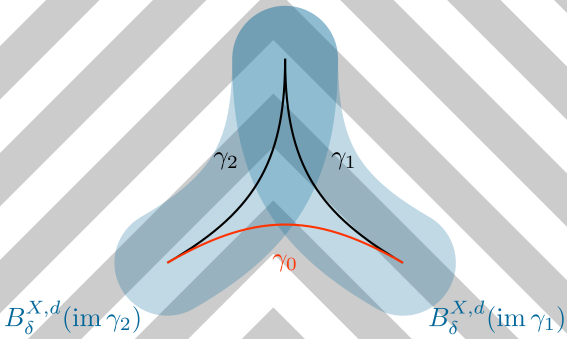

Figure 7.5.: A δ-slim triangle

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 $B_\delta^{X,d}(A) := \{x \in X \mid \exists_{a\in A}\ d(x,a) \leq \delta\}$.

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)**.**

- Every 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 plane $\mathbb{R}^2$ is *not* hyperbolic because for $\delta \in \mathbb{R}_{\geq 0}$, the Euclidean triangle with vertices $(0,0)$, $(0, 3\cdot\delta)$, and $(3\cdot\delta, 0)$ (with isometrically parametrised sides) 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 (Theorem A.3.27). More generally, if M is a closed connected Riemannian manifold of negative sectional curvature (e.g., a hyperbolic manifold), then the Riemannian universal covering of M is hyperbolic in the sense of Definition 7.2.2 [31, Chapter II.1.A, Proposition III.H.1.2].
- If T is a tree, then the geometric realisation $|T|$ of T (Chapter 5.3.2) is 0-hyperbolic because all geodesic triangles in $|T|$ are degenerate tripods (Exercise 7.E.3, see also Proposition 7.2.17 below); in a sense, hyperbolic spaces can be viewed as thickenings of metric trees (Exercise 7.E.9).

Caveat 7.2.4 (Quasi-isometry invariance of hyperbolicity)**.** From the definition it is not clear that hyperbolicity is a quasi-isometry invariant (among geodesic spaces), because the composition of a geodesic triangle with a quasi-isometry

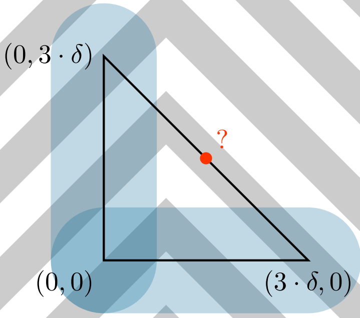

Figure 7.6.: The Euclidean plane $\mathbb{R}^2$ is *not* hyperbolic

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 a quasi-isometry invariant among geodesic spaces.

7.2.2 Quasi-hyperbolic spaces

We now translate the definition of hyperbolicity to quasi-geometry; so instead of geodesics we consider quasi-geodesics. This results in 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 $c, b \in \mathbb{R}_{>0}$, $\delta \in \mathbb{R}_{\geq 0}$.

- 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 be a metric space.

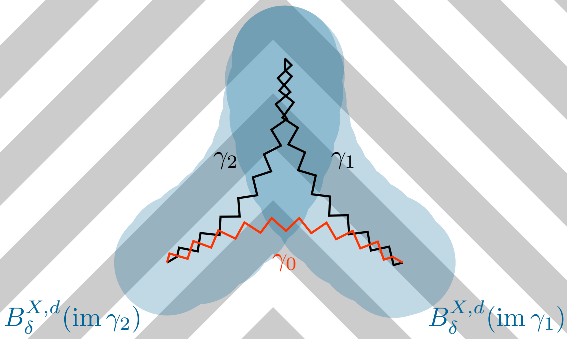

Figure 7.7.: A quasi-slim quasi-geodesic triangle

- 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-geodesic and all (c, b)-quasi-geodesic triangles in X are δ-slim.
- Let $c,\ b \in \mathbb{R}_{>0}$. The space X is called (c, b)-*quasi-hyperbolic* if for all $c',\ b' \in \mathbb{R}_{\geq 0}$ 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. All metric spaces of finite diameter are quasi-hyperbolic.

Caveat 7.2.8. In general, it is rather difficult to prove 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 write d_X and d_Y for the metrics on X and Y, respectively.

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

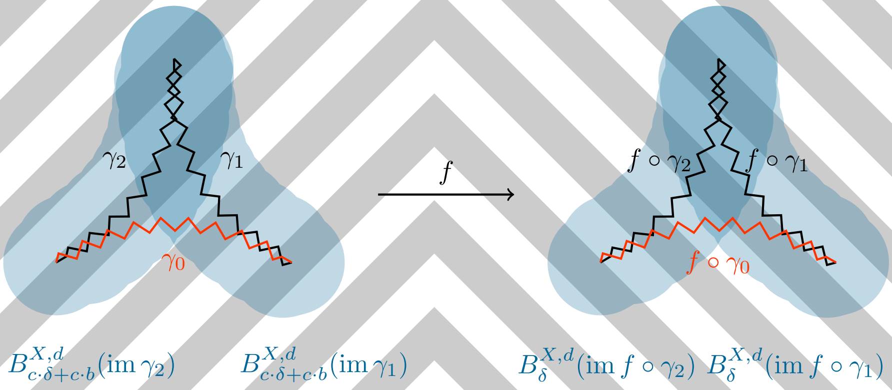

Figure 7.8.: Using a quasi-isometric embedding to translate quasi-geodesic triangles and neighbourhoods back and forth

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 (or the quantitative version of Exercise 5.E.4) 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 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. □

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); then we 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 δ, 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 7.E.2), but this quasi-geodesic ray does not have bounded

[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

distance from any geodesic ray. So, quasi-geodesics in $\mathbb{R}^2$ cannot be uniformly approximated by geodesics.

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

$$\begin{aligned} \operatorname{im} \gamma_0' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_1' \cup \operatorname{im} \gamma_2') \\ \operatorname{im} \gamma_1' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_0' \cup \operatorname{im} \gamma_2') \\ \operatorname{im} \gamma_2' &\subset B_\delta^{X,d}(\operatorname{im} \gamma_0' \cup \operatorname{im} \gamma_1'), \end{aligned}$$

and so

$$\begin{aligned} \operatorname{im} \gamma_0 &\subset B_{\Delta+\delta+\Delta}^{X,d}(\operatorname{im} \gamma_1 \cup \operatorname{im} \gamma_2) \\ \operatorname{im} \gamma_1 &\subset B_{\Delta+\delta+\Delta}^{X,d}(\operatorname{im} \gamma_0 \cup \operatorname{im} \gamma_2) \\ \operatorname{im} \gamma_2 &\subset B_{\Delta+\delta+\Delta}^{X,d}(\operatorname{im} \gamma_0 \cup \operatorname{im} \gamma_1). \end{aligned}$$

Therefore, X is $(c, b, 2 \cdot \Delta + \delta)$-quasi-hyperbolic. ☐

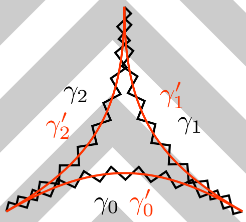

Figure 7.10.: Approximating quasi-geodesic triangles by geodesic triangles

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:

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, d) 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*

$$d\bigl(\gamma'(t), \operatorname{im}\gamma\bigr) \leq \delta \cdot \bigl|\log_2(L_X(\gamma))\bigr| + 1$$

for all $t \in [0, L']$. Here, $L_X(\gamma)$ denotes the length *of γ:*

$$L_X(\gamma) := \sup\biggl\{\sum_{j=0}^{k-1} d\bigl(\gamma(t_j), \gamma(t_{j+1})\bigr) \biggm| k \in \mathbb{N},\ t_0, \dots, t_k \in [0, L],$$
$$t_0 \leq t_1 \leq \dots \leq t_k\biggr\} \in \mathbb{R}_{\geq 0} \cup \{\infty\}.$$

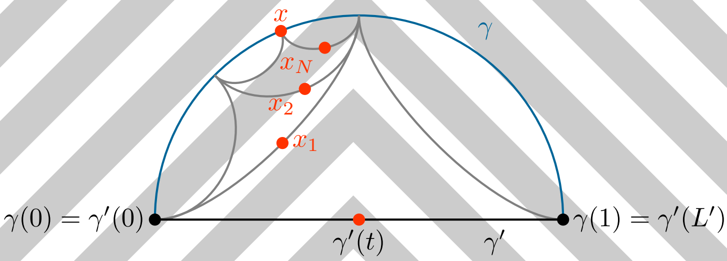

Figure 7.11.: Distance from geodesics to curves in hyperbolic spaces

Proof. Without loss of generality, we may assume that $L_X(\gamma) > 1$, that $L_X(\gamma) < \infty$, 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 (Exercise 7.E.5). Let $N \in \mathbb{N}$ with

$$\frac{L_X(\gamma)}{2^{N+1}} < 1 \leq \frac{L_X(\gamma)}{2^N}.$$

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 at most $L_X(\gamma)/2^N$ whose endpoints lie on $\operatorname{im}\gamma$ (see Figure 7.11). In particular, there is an $x \in \operatorname{im}\gamma$ with

$$\begin{aligned} d\bigl(\gamma'(t), x\bigr) &\leq d\bigl(\gamma'(t), x_N\bigr) + d(x_N, x) \leq \delta \cdot N + \frac{L_X(\gamma)}{2^{N+1}} \\ &\leq \delta \cdot \bigl|\log_2(L_X(\gamma))\bigr| + 1. \end{aligned}$$

□

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, d) 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*

$$L_X(\gamma'|_{[s,t]}) \leq c' \cdot d\bigl(\gamma'(s), \gamma'(t)\bigr) + b'.$$

2. *Moreover,*

$$\operatorname{im}\gamma' \subset B^{X,d}_{c+b}(\operatorname{im}\gamma) \qquad \textit{and} \qquad \operatorname{im}\gamma \subset B^{X,d}_{c+b}(\operatorname{im}\gamma').$$

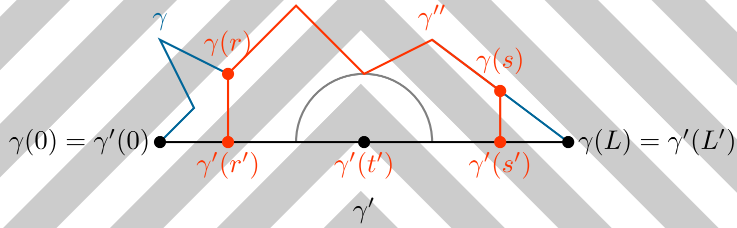

Figure 7.12.: In hyperbolic spaces, geodesics are close to quasi-geodesics

Proof. Let $I := \bigl([0,L) \cap \mathbb{Z}\bigr) \cup \{L\}$. As first step, we set $\gamma'|_I := \gamma|_I$. We then extend the definition of γ' to all of $[0,L]$ by inserting (appropriately reparametrised) 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 replacing 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.,

$$L_X(\gamma|_{[r,s]}) \leq c \cdot d\bigl(\gamma(r), \gamma(s)\bigr) + b$$

for all r, $s \in [0,L]$ with $r \leq s$.

As *first step*, we give an upper estimate for $\sup_{t \in [0,L]} d(\gamma'(t), \operatorname{im}\gamma)$ 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 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

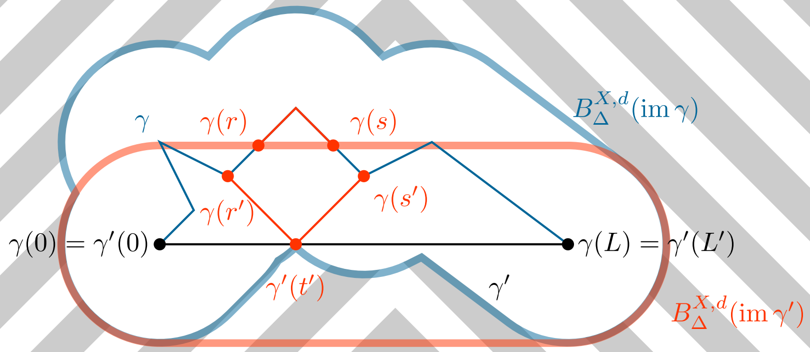

Figure 7.13.: In hyperbolic spaces, quasi-geodesics are close to geodesics

$$\Delta \leq d\bigl(\gamma'(t'), \operatorname{im}\gamma''\bigr) \leq \delta \cdot \bigl|\log_2 L_X(\gamma'')\bigr| + 1;$$

moreover, because γ satisfies the length estimate from Lemma 7.2.15 as described above, we have

$$\begin{aligned} L_X(\gamma'') &\leq L_X(\gamma|_{[r,s]}) + 2 \cdot \Delta \\ &\leq c \cdot d\bigl(\gamma(r), \gamma(s)\bigr) + b + 2 \cdot \Delta \\ &\leq c \cdot (\Delta + 2 \cdot \Delta + 2 \cdot \Delta + \Delta) + b + 2 \cdot \Delta. \end{aligned}$$

Hence,

$$\Delta \leq \delta \cdot \bigl|\log_2((6 \cdot c + 2) \cdot \Delta + b)\bigr| + 1.$$

Because the logarithm function $\log_2$ 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 $\sup_{t \in [0,L]} d(\gamma(t), \operatorname{im}\gamma')$ 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 $B_\Delta^{X,d}(\operatorname{im}\gamma')$. To this end, let $r,\ s \in [0, L]$ be such that $[r, s]$ is a (with respect to inclusion) maximal interval with

$$\gamma((r, s)) \subset X \setminus B_\Delta^{X,d}(\operatorname{im}\gamma').$$

If there is no such non-trivial interval, then there remains nothing to prove; hence, we assume $r \neq s$. By construction of Δ, we know $\operatorname{im}\gamma' \subset B_\Delta^{X,d}(\operatorname{im}\gamma)$, and so $\operatorname{im}\gamma' \subset B_\Delta^{X,d}(\operatorname{im}\gamma|_{[0,r]} \cup \operatorname{im}\gamma|_{[s,L]})$. Because γ is continuous and the interval $[0, L']$ is connected there is a $t' \in [0, L']$ such that there are $r' \in [0, r]$ and $s' \in [s, L]$ with

$$d\bigl(\gamma'(t'), \gamma(r')\bigr) \leq \Delta \qquad \text{and} \qquad d\bigl(\gamma'(t'), \gamma(s')\bigr) \leq \Delta$$

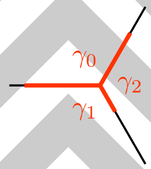

Figure 7.14.: Geodesic triangles (red) in trees are tripods

(see also Figure 7.13). Hence, we obtain

$$\begin{aligned} L_X(\gamma|_{[r,s]}) \leq L_X(\gamma|_{[r',s']}) &\leq c \cdot d\bigl(\gamma(r'), \gamma(s')\bigr) + b \\ &\leq 2 \cdot c \cdot \Delta + b, \end{aligned}$$

and thus

$$\gamma([r,s]) \subset B^{X,d}_{c\cdot\Delta+b/2+\Delta}(\operatorname{im}\gamma').$$

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

$$\operatorname{im}\gamma \subset B^{X,d}_{c\cdot\Delta+b/2+\Delta}(\operatorname{im}\gamma').$$

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

How can one check whether a graph is hyperbolic or not? Graphs, viewed as metric spaces, are not geodesic (unless they have at most one vertex). Therefore, one can either work with quasi-hyperbolicity or pass to the geometric realisation; the geometric realisation has the advantage that we can work with actual geodesics rather with potentially wild quasi-geodesics. Therefore, in the context of hyperbolicity, it is more common to use geometric realisations of graphs instead of graphs.

Corollary 7.2.16 (Hyperbolicity of graphs)**.** *Let X be a connected graph. Then X is quasi-hyperbolic 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 5.3.8). □

Proposition 7.2.17 (Hyperbolicity of trees)**.** *If T is a tree, then the geometric realisation $|T|$ of T is 0-hyperbolic. Hence, T is quasi-hyperbolic.*

Proof. Because the graph T does not contain any graph-theoretic cycles, one can show that every geodesic triangle in $|T|$ looks like a tripod, as depicted in Figure 7.14 (Exercise 7.E.3), which implies 0-hyperbolicity. $\square$

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 [74]:

Definition 7.3.1 (Hyperbolic group)**.** A finitely generated group G is *hyperbolic* if for some (and hence every) 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. $\square$

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, $\operatorname{SL}(2, \mathbb{Z})$ is hyperbolic, because $\operatorname{SL}(2, \mathbb{Z})$ is quasi-isometric to a free group of rank 2 (Example 5.4.8).
- Let M be a compact Riemannian manifold of negative sectional curvature (e.g., a hyperbolic manifold in the sense of Definition 5.4.11).

Then the fundamental group $\pi_1(M)$ is hyperbolic, because by the Švarc-Milnor lemma (Corollary 5.4.10) $\pi_1(M)$ 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 5.4.12).

- The group $\mathbb{Z}^2$ is *not* hyperbolic, because it is quasi-isometric to the Euclidean plane $\mathbb{R}^2$, which is a geodesic metric space that is not hyperbolic (Example 7.2.3).
- We will see that the Heisenberg group is *not* hyperbolic (Example 7.5.16) and that $\mathrm{BS}(1,2)$ is *not* hyperbolic (Exercise 7.E.24).

Caveat 7.3.4 (Non-compact hyperbolic manifolds)**.** If M is a connected complete hyperbolic Riemannian manifold of finite volume, then, in general, the fundamental group $\pi_1(M)$ is *not* hyperbolic. The geometric group theoretic notion capturing such fundamental groups (and their relation with the subgroups given by the fundamental groups of the cusps) are *relatively hyperbolic groups* [138].

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 7.3.5. Let $z \in \mathbb{H}^2$. The isometric action of $\mathrm{SL}(2,\mathbb{Z})$ on the hyperbolic plane $\mathbb{H}^2$ by Möbius transformations (Proposition A.3.11, Proposition A.3.14) 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 halfplane model (Appendix A.3): 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 (Remark A.3.16)

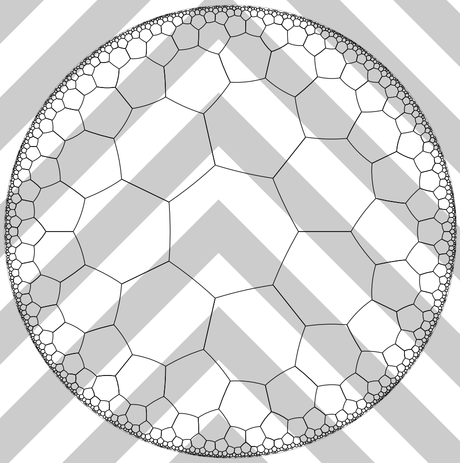

Figure 7.15.: A regular tiling of the hyperbolic plane (in the Poincaré disk model) by regular heptagons, where in each vertex exactly three heptagons meet

$$\begin{aligned} d_H(z, A^n \cdot z) &= d_{\mathbb{H}^2}(z, z+n) = \operatorname{arcosh}\Bigl(1 + \frac{n^2}{2 \cdot (\operatorname{Im} z)^2}\Bigr) \\ &\in O\bigl(\operatorname{arcosh}(n^2)\bigr) = O\bigl(\ln(n^2 + \sqrt{n^2 - 1})\bigr) = O(\ln n). \end{aligned}$$

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.

Example 7.3.6 (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 left-overs of the dough remain (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

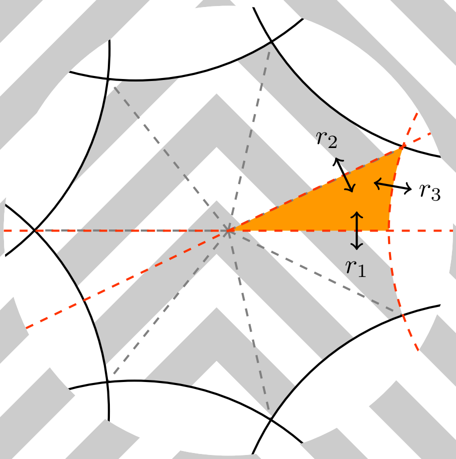

Figure 7.16.: Fundamental triangle of a regular hyperbolic heptagon and its reflections

$$\frac{1}{k} + \frac{1}{d} < \frac{1}{2}.$$

The *symmetry group* of the heptagonal 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

$$G := \langle r_1, r_2, r_3 \mid r_1^2, r_2^2, r_3^2, (r_1 r_2)^7, (r_1 r_3)^2, (r_2 r_3)^3 \rangle;$$

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 5.4.2) one can show that G is quasi-isometric to $\mathbb{H}^2$; hence, G is a hyperbolic group. Beautiful tilings of $\mathbb{H}^2$ based on quasi-regular tilings of $\mathbb{H}^2$ can be found in the work of Escher [59].

The group G is an example of a reflection group. More abstractly, reflection groups are special cases of Coxeter groups [35, 45] (Exercise 2.E.24). Coxeter groups play an important role in various areas of mathematics. For example, the *reflection group trick* by Davis and Januszkiewicz [45] in algebraic and geometric topology allows to construct aspherical manifolds with exotic properties.

By definition, the geometric property of being hyperbolic is modelled on the behaviour of (fundamental groups of) manifolds of negative sectional curvature in Riemannian geometry. On the other hand, hyperbolicity also has non-trivial algebraic consequences for groups (Chapter 7.4 and 7.5).

7.4 The word problem in hyperbolic groups

As first algebraic consequence, we show that the geometric condition of being hyperbolic implies solvability of the word problem.

Definition 7.4.1 (Word problem)**.** 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 $\langle S \mid R \rangle$ or not.

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 $(S \cup S^{-1})^*$. As usual, in such situations, we view S^{-1} as the set of formal inverses of S.

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 [28, 22, 12].

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 [150, 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 *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 producing a finite presentation with unsolvable word problem.

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).

7.4.1 Application: "Solving" the word problem

Gromov observed that hyperbolic groups have solvable word problem [74], 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 [75, 133] and Ol'shanskii [135] established the following:

Theorem 7.4.4 (Generic groups are hyperbolic). *In a 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 $\langle x, y \mid xx^{-1}\varepsilon, yy^{-1}\varepsilon, x^{-1}x\varepsilon, y^{-1}y\varepsilon \rangle$ 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 in 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).

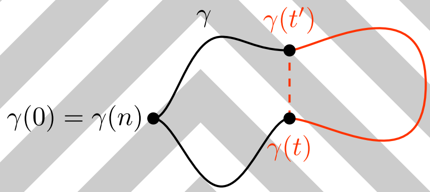

Figure 7.17.: Short-cuts in hyperbolic groups

– 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 $\langle S \,|\, R \rangle$; hence, w represents the trivial element of the group $\langle S \,|\, R \rangle$ if and only if the shorter word $w' v_j w''$ represents the trivial element of the group $\langle S \,|\, R \rangle$ (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 $\langle S \,|\, R \rangle$ or not. Hence, the word problem for $\langle S \,|\, R \rangle$ is solvable. □

In the setting of hyperbolic groups, short-cuts as required in the definition of Dehn presentations are enforced by 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 $\langle S \,|\, R \rangle$ is a Dehn presentation and $G \cong \langle S \,|\, R \rangle$.*

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 let $|\mathrm{Cay}(G,S)|$ be 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*

$$L_{|\mathrm{Cay}(G,S)|}(\gamma|_{[t,t']}) \leq 8 \cdot \delta$$

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 $D := \lceil 8 \cdot \delta \rceil + 2$, and let $\pi\colon F(S) \longrightarrow G$

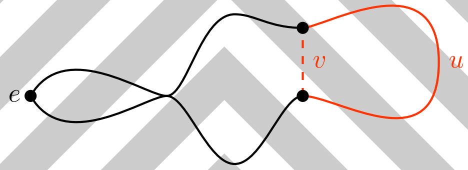

Figure 7.18.: Short-cuts in hyperbolic groups and Dehn presentations

be the canonical homomorphism. Modelling the short-cut lemma above in terms of group theory, leads to the (finite) set

$$R := \bigl\{uv^{-1} \bigm| u, v \in (S \cup S^{-1})^*, |u| \le D, |u| > d_S(e, \pi(u)),\\ \pi(v) = \pi(u), |v| = d_S(e, \pi(u))\bigr\};$$

here, $|\cdot|$ denotes the length of words in $(S \cup S^{-1})^*$. In particular,

$$R' := \bigl\{st\varepsilon \bigm| s, t \in S \cup S^{-1} \text{ with } \pi(st) = e\bigr\}$$

is contained in R. Clearly, the canonical homomorphism $\langle S \,|\, R\rangle \longrightarrow G$ is surjective (because S generates G).

This homomorphism is also injective and $\langle S \,|\, R\rangle$ is a Dehn presentation for G: Let $w \in (S \cup S^{-1})^*$ be a word such that $\pi(w) = e$. We prove that $w \in \langle R\rangle^{\triangleleft}_{F(S)}$ 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 assume 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 contains a subword of the form st with $s, t \in S \cup S^{-1}$ and $\pi(st) = e$, then $st\varepsilon$ is contained in $R' \subset R$ and st is the desired Dehn word. Removing st from w results in a word w' that is in $\langle R\rangle^{\triangleleft}_{F(S)}$ (by induction). Multiplying w' by a suitable conjugate of st shows that $w \in \langle R\rangle^{\triangleleft}_{F(S)}$.
- If w contains *no* subword st with $s, t \in S \cup S^{-1}$ and $\pi(st) = e$, then the word w (or a non-empty subword of w) translates into a graph-theoretic cycle in $\operatorname{Cay}(G, S)$ (see Definition 3.1.6). 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

  $$d_S\bigl(e, \pi(u)\bigr) < |u| \le D.$$

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). Hence, u is the desired Dehn subword of w. 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 $w'vw'' \in \langle R \rangle^{\triangleleft}_{F(S)}$. Therefore, we also have $w \in \langle R \rangle^{\triangleleft}_{F(S)}$ (by multiplying $w'vw''$ with the w'-conjugate of uv^{-1}). □

Corollary 7.4.10. *Every hyperbolic group admits a finite presentation.*

Proof. By Theorem 7.4.8, every hyperbolic group possesses a finite Dehn presentation, whence a finite presentation. □

In particular, we can now prove Theorem 7.4.3:

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. □

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. This is Exercise 7.E.7. □

Proof of Lemma 7.4.9. We first show that there is no $c \in \mathbb{R}_{>8\cdot\delta}$ such that γ is a c-local geodesic: *Assume* for a contradiction that there is a $c \in \mathbb{R}_{>8\cdot\delta}$ such that γ is a c-local geodesic in $|\operatorname{Cay}(G, S)|$. Because of $\gamma(0) = \gamma(n)$ and $c > 8 \cdot \delta$, it is clear that $n > 8 \cdot \delta$.

By Lemma 7.4.11, γ is $2 \cdot \delta$-close to every 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 $\operatorname{im} \gamma \subset B^{|\operatorname{Cay}(G,S)|, d_S}_{2\cdot\delta}\bigl(\gamma(0)\bigr)$. Therefore, we obtain

$$4 \cdot \delta \geq \operatorname{diam} B^{|\operatorname{Cay}(G,S)|, d_S}_{2\cdot\delta}\bigl(\gamma(0)\bigr) \geq d_S\bigl(\gamma(0), \gamma(5 \cdot \delta)\bigr) = 5 \cdot \delta,$$

which is a contradiction. So γ is not a c-local geodesic for $c \in \mathbb{R}_{>8\cdot\delta}$.

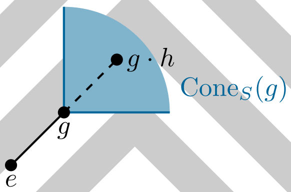

Figure 7.19.: Cone type, schematically; the drawing shows the intuitive visualisation $g \cdot \mathrm{Cone}_S(g)$ of the cone type instead of $\mathrm{Cone}_S(g)$.

Hence, there exist $t, t' \in [0, n]$ with $|t-t'| \leq 8 \cdot \delta$ and $d\bigl(\gamma(t), \gamma(t')\bigr) \neq |t-t'|$; in particular $\gamma|_{[t,t']}$ is not a geodesic. Moreover, because γ is the geometric realisation of a cycle in $\mathrm{Cay}(G, S)$, it follows that

$$L_{|\mathrm{Cay}(G,S)|}(\gamma|_{[t,t']}) = |t - t'| \leq 8 \cdot \delta. \qquad \square$$

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 every infinite hyperbolic group contains an element of infinite order (Chapter 7.5.1) and that centralisers of elements of infinite order in hyperbolic groups are "small" (Chapter 7.5.2). Consequently, hyperbolic groups cannot contain $\mathbb{Z}^2$ as a subgroup. We will mostly follow the arguments by Bridson and Haefliger [31].

7.5.1 Existence

In general, an infinite finitely generated group does not necessarily contain an element of infinite order; for example, the Grigorchuk group is of this type (Exercise 4.E.35, Exercise 4.E.37). However, in the case of hyperbolic groups, the geometry forces the existence of elements of infinite order:

Theorem 7.5.1. *Every infinite hyperbolic group contains an element of infinite order.*

For the proof, 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)

$$\mathrm{Cone}_S(g) := \bigl\{h \in G \bigm| d_S(e, g \cdot h) \geq d_S(e, g) + d_S(e, h)\bigr\}.$$

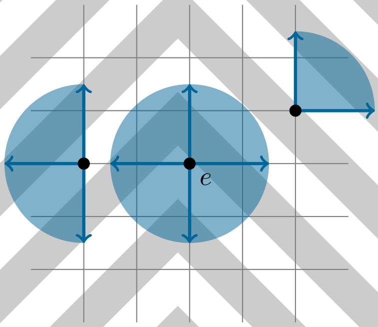

Figure 7.20.: Cone types of $\mathbb{Z}^2$

Example 7.5.3 (Cone types).

- Let F be a finitely generated free group of rank $n \in \mathbb{N}$, 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 $\operatorname{Cone}_S(e) = F$, and

$$\begin{aligned}\operatorname{Cone}_S(s) = \{w \mid\ & w \text{ is a reduced word over } S \cup S^{-1} \\ & \text{that does not start with } s^{-1}\}\end{aligned}$$

  for all $s \in S \cup S^{-1}$.
- The group $\mathbb{Z}^2$ has only finitely many (nine) different cone types with respect to the generating 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).

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

$$P_r^S(g) := \bigl\{h \in B_r^{G,S}(e) \bigm| d_S(e, g \cdot h) \leq d_S(e,g)\bigr\}$$

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 $|\operatorname{Cay}(G,S)|$ is δ-hyperbolic. Let

$$r := 2 \cdot \delta + 2.$$

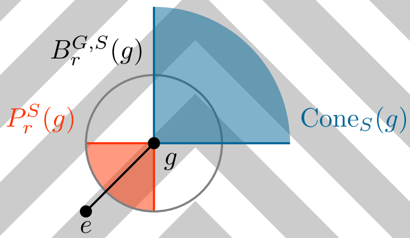

Figure 7.21.: The past of an element, schematically

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 $h \in \mathrm{Cone}_S(g')$.

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 $h \notin \mathrm{Cone}_S(g')$. 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': On the one hand, we have (by choice of k_1)

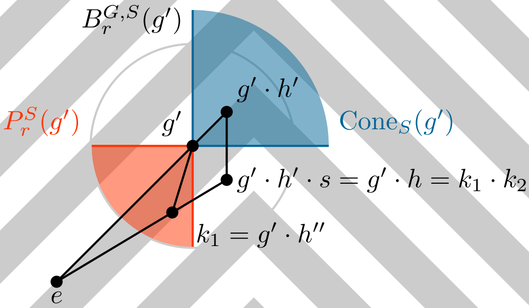

Figure 7.22.: The past of an element determines its cone type

$$d_S(e, g' \cdot h'') = d_S(e, k_1) \leq d_S(e, g').$$

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 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 by the choice of k_1 and k_2 we have $d_S(e, g' \cdot h) = d_S(e, k_1) + d_S(e, k_2)$. 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.

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.

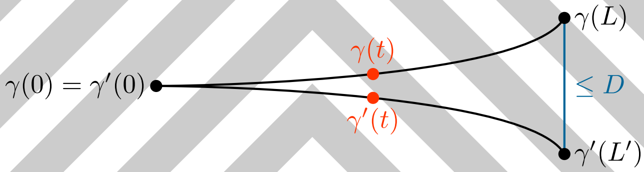

Figure 7.23.: Geodesics in hyperbolic spaces starting at the same point with close endpoints are uniformly close

Therefore, there are only finitely many cone types with respect to S in G, as claimed. □

Lemma 7.5.5 (Geodesics in hyperbolic spaces starting at the same point). *Let $\delta, D \in \mathbb{R}_{\geq 0}$ 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*

$$\gamma(0) = \gamma'(0) \qquad \text{and} \qquad d\bigl(\gamma(L), \gamma'(L')\bigr) \leq D.$$

Then γ and γ' are uniformly $(2 \cdot \delta + D)$-close, i.e.,

$$\forall_{t \in [0, \min(L, L')]} \ d\bigl(\gamma(t), \gamma'(t)\bigr) \leq 2 \cdot \delta + D \qquad \text{and} \qquad |L - L'| \leq D.$$

Proof. This follows from a suitable application of the slim triangles condition (Exercise 7.E.6), see also Figure 7.23. □

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

$$k := \bigl|\bigl\{\mathrm{Cone}_S(g) \bigm| g \in G\bigr\}\bigr| \in \mathbb{N}.$$

Because G is infinite and $\mathrm{Cay}(G, S)$ is 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 the following conditions hold:

- $h \neq e$,
- $d_S(e, g) = d_S(e, g') + d_S(e, h) + d_S(e, g'')$, and
- $\mathrm{Cone}_S(g') = \mathrm{Cone}_S(g' \cdot h)$.

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: We will show by induction over the exponent $n \in \mathbb{N}_{>0}$ 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, for all $n \in \mathbb{N}_{>0}$ the elements $g' \cdot h^n$ and g' must be different. Thus, h has infinite order. □

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). □

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. *Every 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 of infinite order (by Theorem 7.5.1). □

Corollary 7.5.8 (Quasi-isometry rigidity of $\mathbb{Z}$)**.** *Every 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, the group G contains an element g of infinite order. One then shows (Exercise 7.E.28):

- The subgroup $\langle g \rangle_G$ generated by g is quasi-dense in G because G and $\mathbb{Z}$ are quasi-isometric.
- Hence, $\langle g \rangle_G$ has finite index in G. □

7.5.2 Centralisers

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 7.5.15).

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. We will show that this is impossible in hyperbolic groups.

The key insight for this and many other results on hyperbolic groups is that elements of infinite order give rise to quasi-geodesic lines; one also says that these elements are *undistorted* or *loxodromic*:

Theorem 7.5.9 (Homogeneous quasi-geodesic lines in hyperbolic groups)**.** *Let G be a hyperbolic group and let $g \in G$ be an element of infinite order. Then the map*

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longmapsto g^n \end{aligned}$$

is a quasi-isometric embedding.

The proof of this theorem will be given in Chapter 7.5.3 below. Using these (quasi-)geodesic lines, we can prove that the hyperbolic geometry indeed forces centralisers of elements of infinite order to be small:

Theorem 7.5.10 (Centralisers in hyperbolic groups)**.** *Let G be a hyperbolic group and let $g \in G$ be an element of infinite order. Then 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 the sake of completeness, we recall the notion of centraliser:

Definition 7.5.11 (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.,

$$C_G(g) := \{h \in G \mid h \cdot g = g \cdot h\}.$$

The centraliser of a group element always is a subgroup of the group in question.

Example 7.5.12 (Centralisers).

- 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 *centre of* G.

Caveat 7.5.13. 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

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longrightarrow g^n \end{aligned}$$

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 happens in the Heisenberg group (Caveat 6.2.13).

We now prove Theorem 7.5.10, using Theorem 7.5.9. This is a geometric argument involving invariant geodesics, inspired by Preissmann's theorem (and its proof) in Riemannian geometry [57, Theorem 10.2.2]:

In view of Theorem 7.5.9, 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 hypothesis 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.

In fact, we will first prove the following, slightly more general, statement about elements that quasi-commute with the given element g (which will also be valuable in exhibiting free subgroups of hyperbolic groups in Chapter 8.3.4):

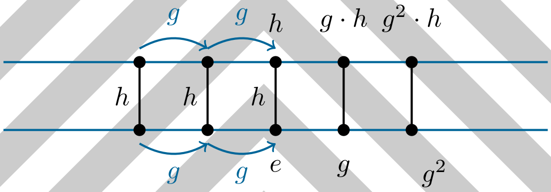

Figure 7.24.: Elements in the centraliser $C_G(g)$ lead to flat strips

Lemma 7.5.14 (Close conjugates). *Let G be a hyperbolic group, let $g \in G$ be of infinite order, and let $S \subset G$ be a finite generating set of G. Then there is a constant $\Delta \in \mathbb{R}_{>0}$ with the following property: If $k \in G$ and $\varepsilon \in \{-1, 1\}$ satisfy*

$$\sup_{n \in \mathbb{Z}} d_S(k \cdot g^n \cdot k^{-1}, g^{\varepsilon \cdot n}) < \infty,$$

then

$$d_S\bigl(k, \langle g \rangle_G\bigr) \leq \Delta.$$

Proof. We first need to make some of the constants explicit: By Theorem 7.5.9, there exists a constant $c \in \mathbb{R}_{\geq 1}$ such that the map

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longmapsto g^n \end{aligned}$$

is a (c, c)-quasi-isometric embedding. Because G is hyperbolic, there exists a $\delta \in \mathbb{R}_{>0}$ such that G is (c, c, δ)-hyperbolic with respect to d_S (Exercise 7.E.13). We set

$$\Delta := 2 \cdot \delta$$

(however, one should note that c, whence δ, depends on g).

We now start with the actual proof: Let $k \in G$ and let $\varepsilon \in \{-1, 1\}$ with

$$C := \sup_{n \in \mathbb{Z}} d_S(k \cdot g^n \cdot k^{-1}, g^{\varepsilon \cdot n}) < \infty.$$

By Theorem 7.5.9, we can choose $n \in \mathbb{N}$ so big that

$$d_S(e, g^n) > C + 2 + 2 \cdot \delta + d_S(e, k).$$

We consider a quasi-geodesic quadrilateral with the vertices $k \cdot g^{-n}$, $k \cdot g^n$, $g^{\varepsilon \cdot n}$, $g^{-\varepsilon \cdot n}$. To this end we pick $(1,1)$-quasi-geodesics γ from $g^{-\varepsilon \cdot n}$ to $k \cdot g^n$, as well as γ_+ from $k \cdot g^n$ to $g^{\varepsilon \cdot n}$ and γ_- from $k \cdot g^{-n}$ to $g^{-\varepsilon \cdot n}$. As "bottom" and "top" quasi-geodesics, we use the segments of $m \mapsto g^m$ and $m \mapsto k \cdot g^m$ (which by left-invariance of d_S is a (c, c)-quasi-geodesic embedding). This situation

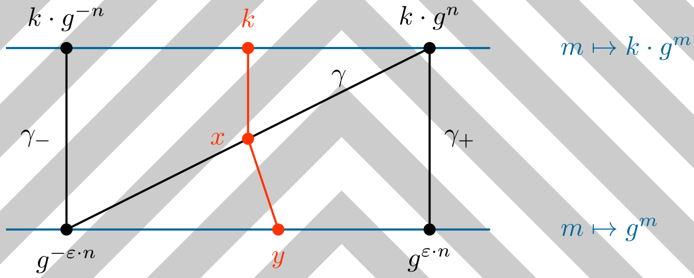

Figure 7.25.: Comparing the quasi-geodesic lines $n \mapsto k \cdot g^n$ and $n \mapsto g^n$

is illustrated in Figure 7.25. We now argue similarly as in Lemma 7.5.5 and Lemma 7.4.11:

The conjugating element k lies on the "top" quasi-geodesic. Hence, by hyperbolicity, there is an x in $\operatorname{im} \gamma$ or $\operatorname{im} \gamma_-$ that is δ-close to k. We can rule out the case of $\operatorname{im} \gamma_-$ as follows: For all $x \in \operatorname{im} \gamma_-$ we have by the triangle inequality and the fact that γ_- is $(1,1)$-quasi-geodesic:

$$\begin{aligned} d_S(x,k) &\geq d_S(k \cdot g^{-n}, k) - d_S(x, k \cdot g^{-n}) \\ &\geq d_S(g^{-n}, e) - d_S(g^{-\varepsilon \cdot n}, k \cdot g^{-n}) - 2. \end{aligned}$$

By the choice of n, we know $d_S(g^{-n}, e) = d_S(e, g^n) > C + 2 + 2 \cdot \delta + d_S(e,k)$. Moreover,

$$\begin{aligned} d_S(g^{-\varepsilon \cdot n}, k \cdot g^{-n}) &\leq d_S(g^{-\varepsilon \cdot n}, k \cdot g^{-n} \cdot k^{-1}) + d_S(k \cdot g^{-n} \cdot k^{-1}, k \cdot g^{-n}) \\ &\leq C + d_S(e,k). \end{aligned}$$

Putting these estimates together, we obtain $d_S(x,k) > 2 \cdot \delta \geq \delta$. Hence, there is an $x \in \operatorname{im} \gamma$ with $d_S(k,x) \leq \delta$.

Analogously, we can use hyperbolicity in the "lower" quasi-geodesic triangle to show that there is a point $y \in \langle g \rangle_G$ with $d_S(x,y) \leq \delta$ (by ruling out $\operatorname{im} \gamma_+$). Therefore,

$$d_S\bigl(k, \langle g \rangle_G\bigr) \leq d_S(k,y) \leq d_S(k,x) + d_S(x,y) \leq 2 \cdot \delta = \Delta,$$

as claimed. □

Proof of Theorem 7.5.10. Let S be a finite generating set and let $\Delta \in \mathbb{R}_{>0}$ be a constant for g as provided by Lemma 7.5.14. For every $h \in C_G(g)$ we have

$$\sup_{n \in \mathbb{Z}} d_S(h \cdot g^n \cdot h^{-1}, g^n) = \sup_{n \in \mathbb{Z}} d_S(g^n, g^n) = 0 < \infty$$

and so

$$d_S(h, \langle g \rangle_G) \leq \Delta.$$

In other words, $\langle g \rangle_G$ is Δ-dense in $C_G(g)$ with respect to d_S, which implies that $\langle g \rangle_G$ has finite index in $C_G(g)$ (Exercise 5.E.23). □

Corollary 7.5.15. *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

$$\mathbb{Z}^2 \cong H = C_H(h) \subset C_G(h),$$

which contradicts that the centraliser $C_G(h)$ of h is virtually $\langle h \rangle_G$ (Theorem 7.5.10). □

Example 7.5.16 (Heisenberg group, $\mathrm{SL}(n, \mathbb{Z})$ and hyperbolicity)**.** By Corollary 7.5.15, the Heisenberg group is *not* hyperbolic: the subgroup of the Heisenberg group generated by the matrices

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

is isomorphic to $\mathbb{Z}^2$. Thus, for all $n \in \mathbb{N}_{\geq 3}$, the matrix groups $\mathrm{SL}(n, \mathbb{Z})$ also are not hyperbolic.

The proofs of Theorem 7.5.10 and Theorem 7.5.9 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 (Chapter 8.3). For example, these techniques then also show that "generic" elements in hyperbolic groups fail to commute in the strongest possible way (Theorem 8.3.13).

Moreover, the question whether group elements/isometries act by translation on (quasi-)geodesic lines is already present in the classical classification of isometries of the hyperbolic plane:

Remark 7.5.17 (The conic trichotomy)**.** Orientation preserving isometries of the hyperbolic plane $\mathbb{H}^2$ are Möbius transformations (Theorem A.3.23). Therefore, non-trivial orientation preserving isometries of $\mathbb{H}^2$ can be classified into the following three types [18, 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. Such an axis is then unique.
- *Parabolic.* A non-trivial orientation preserving isometry of $\mathbb{H}^2$ is called *parabolic* if it has no fixed point in $\mathbb{H}^2$ and if it admits no axis.

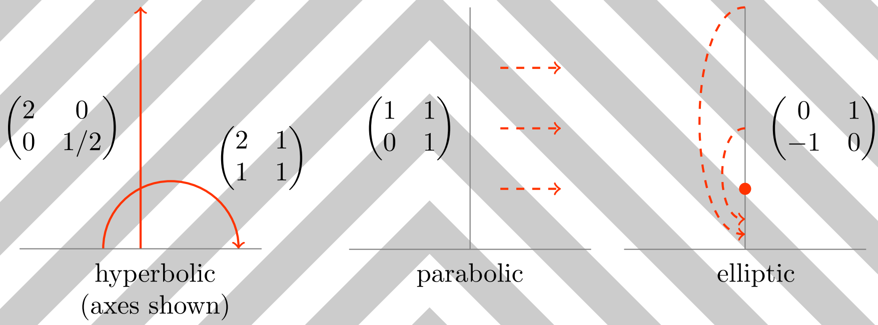

Figure 7.26.: The three types of orientation preserving isometries of $\mathbb{H}^2$, in the halfplane model

- *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 halfplane model for $\mathbb{H}^2$, prototypical examples of these three types of isometries are the following (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 (in the halfplane model). Both of these isometries are hyperbolic. The first one has the imaginary axis as axis; the latter one has the geodesic line in $\mathbb{H}^2$ as axis whose image is the semi-circle that meets the real line in the points $1/2 \cdot (1 + \sqrt{5})$ and $1/2 \cdot (1 - \sqrt{5})$.
  More generally, the Möbius transformation associated with a matrix $A \in \mathrm{SL}(2,\mathbb{R}) \setminus \{E_2, -E_2\}$ is hyperbolic if and only if A is diagonalisable over $\mathbb{R}$; this is equivalent to $|\mathrm{tr}\, A| > 2$ (Exercise 7.E.25).
- *Parabolic.* The matrix

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

yields the parabolic Möbius transformation

$$z \longmapsto z + 1.$$

More generally, the Möbius transformation associated with a matrix $A \in \mathrm{SL}(2,\mathbb{R}) \setminus \{E_2, -E_2\}$ is parabolic if and only if A has real eigenvalues but is not diagonalisable over $\mathbb{R}$; this is equivalent to $|\mathrm{tr}\, A| = 2$ (Exercise 7.E.25).

- *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.

More generally, the Möbius transformation associated with a matrix $A \in \mathrm{SL}(2,\mathbb{R}) \setminus \{E_2, -E_2\}$ is elliptic if and only if A has no real eigenvalues; this is equivalent to $|\mathrm{tr}\, A| < 2$ (Exercise 7.E.25).

Using the language of boundaries of hyperbolic spaces/groups (Chapter 8.3), one can express this trichotomy also purely in terms of fixed points (Outlook 8.3.7).

7.5.3 Quasi-convexity

In order to complete the proof of Theorem 7.5.10 we still need to provide a proof of Theorem 7.5.9, i.e., that elements of infinite order in hyperbolic groups give rise to quasi-geodesic lines. Also the proof of Theorem 7.5.9 in turn heavily uses centralisers. We first 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 every geodesic whose endpoints lie in this subset must already be contained completely in this subset (Figure 7.27).

Definition 7.5.18 (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

$$\operatorname{im} \gamma \subset B_c^{X,d}(C).$$

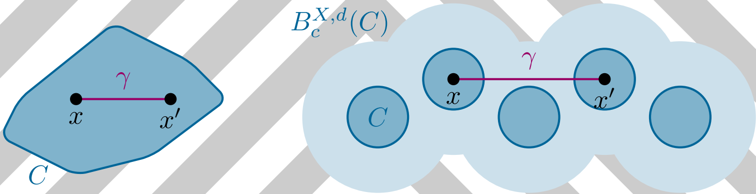

Figure 7.27.: Convexity and quasi-convexity, schematically

One could also formulate a notion of quasi-convexity in quasi-geodesic spaces, but for simplicity we will only consider the geodesic case.

Definition 7.5.19 (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 7.5.20 (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. This is the contents of Exercise 7.E.26 and Exercise 7.E.27; it should be noted that the second part is non-trivial in the sense that the analogous statement for general quasi-convex subspaces is *not* true (Caveat 7.5.21). □

Caveat 7.5.21. The intersection of 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

$$C := 2 \cdot \mathbb{Z} \qquad \text{and} \qquad C' := (2 \cdot \mathbb{Z} + 1) \cup \{n^2 \mid n \in \mathbb{Z}\}$$

are quasi-convex subsets of $\mathbb{R}$. But the intersection $C \cap C' = \{n^2 \mid n \in 2 \cdot \mathbb{Z}\}$ clearly is *not* quasi-convex in $\mathbb{R}$.

Proposition 7.5.22 (Quasi-convexity of centralisers)**.** *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 every 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 $\gamma(0) = e$. Furthermore, we set $h := \gamma(L) \in C_G(g)$. 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 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 two 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 in $\overline{h}$. Then $g \cdot \overline{\gamma}$ 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 and so (see above)

$$\overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) \in B_G^{S,c}(e)$$

for all $s \in [0,t]$ with $\overline{\gamma}(s) \in G$.

If $d_S(e, \overline{h}) > |B_c^{G,S}(e)|$, then by the pigeon-hole principle and the fact that geodesics in $|\mathrm{Cay}(G,S)|$ basically follow a shortest path in $\mathrm{Cay}(G,S)$ 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

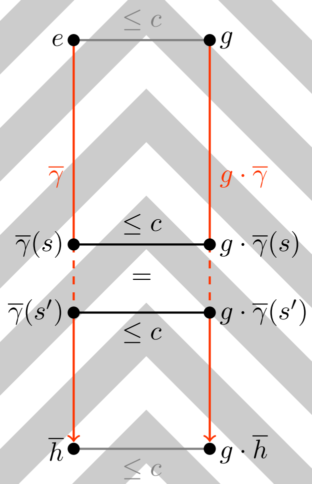

Figure 7.28.: Finding a shorter conjugating element

$$\begin{aligned}\overline{\overline{h}}^{-1} \cdot g \cdot \overline{\overline{h}} &= \overline{h}^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) \cdot \overline{\gamma}(s')^{-1} \cdot \overline{h} \\ &= \overline{h}^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot g \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot \overline{h} \\ &= \overline{h}^{-1} \cdot g \cdot \overline{h}.\end{aligned}$$

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

$$\begin{aligned}d_S\bigl(k, \gamma(t)\bigr) = d_S(\overline{h} \cdot \overline{\overline{h}}^{-1}, \overline{h}) = d_S(\overline{\overline{h}}^{-1}, e) = d_S(e, \overline{\overline{h}}) \\ \leq \bigl|B_c^{G,S}(e)\bigr|.\end{aligned}$$

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. □

As promised, these quasi-convexity considerations allow us to complete the proof of Theorem 7.5.9:

Proof of Theorem 7.5.9. In view of Proposition 7.5.22, the centraliser $C_G(g)$ is a quasi-convex subgroup of G. In particular, $C_G(g)$ is finitely generated by Proposition 7.5.20, 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 7.5.20). In particular, also $C(C_G(g))$ is a hyperbolic group (Proposition 7.3.2).

On the other hand, $C(C_G(g))$ is Abelian and contains $\langle g \rangle_G \cong \mathbb{Z}$; because $C(C_G(g))$ 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. □

7.5.4 Application: Products and negative curvature

In view of Corollary 7.5.15, most non-trivial products of finitely generated groups are *not* hyperbolic.

Corollary 7.5.23. *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 7.5.15 and rule out $\mathbb{Z}^2$ as a subgroup. □

Example 7.5.24 (Heisenberg manifold)**.** In particular, the closed connected smooth manifold given as the quotient of the 3-dimensional Heisenberg group with $\mathbb{R}$-coefficients by the Heisenberg group (i.e., the *Heisenberg manifold*) does *not* admit a Riemannian metric of negative sectional curvature (the Heisenberg group contains $\mathbb{Z}^2$, Example 7.5.16).

Outlook 7.5.25 (Splittings and presentability by products)**.** A geometric counterpart of the non-existence of $\mathbb{Z}^2$-subgroups is the following classical theorem by Gromoll and Wolf [71] (for simplicity, we only state the version where the fundamental group has trivial centre):

Theorem 7.5.26 (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. *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 [92, 93]:

Theorem 7.5.27 (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.*

7.6 Non-positively curved groups

We conclude this chapter with a very brief discussion non-positively curved groups, so-called CAT(0)-groups. Hyperbolic metric spaces are geodesic spaces whose geodesic triangles are slim. This can be reformulated as geodesic triangles being not much fatter than geodesic triangles in trees (Exercise 7.E.9).

In order to define a notion of non-positive curvature for metric spaces, we replace the comparison space: We compare geodesic triangles with triangles in the Euclidean plane (Figure 7.29) instead of trees.

Definition 7.6.1 (CAT(0)-inequality, CAT(0)-space)**.**

- A geodesic triangle $(\gamma_0 \colon [0, L_1] \to X, \gamma_1 \colon [0, L_1] \to X, \gamma_2 \colon [0, L_2] \to X)$ in a metric space (X, d) satisfies the CAT(0)-*inequality* if the following holds: Let $(\gamma'_0 \colon [0, L_1] \to X, \gamma'_1 \colon [0, L_1] \to X, \gamma'_2 \colon [0, L_2] \to X)$ be a geodesic triangle in the Euclidean plane $(\mathbb{R}^2, d_2)$ with the same side lengths (such triangles are unique up to a Euclidean isometry). Then

$$\forall_{j,k \in \{0,1,2\}} \ \forall_{s \in [0,L_j]} \ \forall_{t \in [0,L_k]} \ d\bigl(\gamma_j(s), \gamma_k(t)\bigr) \leq d_2\bigl(\gamma'_j(s), \gamma'_k(t)\bigr).$$

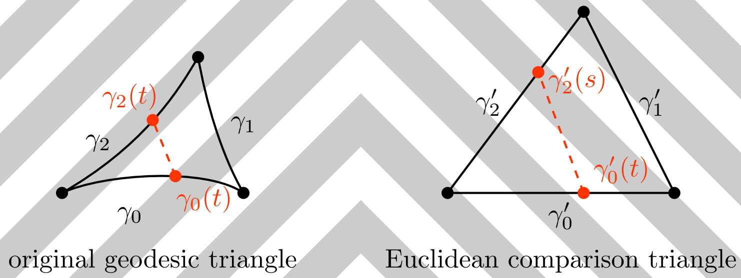

Figure 7.29.: The CAT(0)-inequality

- A CAT(0)-*space* is a geodesic metric space (X, d) such that all geodesic triangles in X satisfy the CAT(0)-inequality.

Remark 7.6.2 (CAT)**.** The name CAT(0) refers to the pioneers of comparison geometry: Cartan, Alexandrov, and Toponogov. The number 0 denotes the upper curvature bound. Using other simply connected two-dimensional model spaces of constant Gaussian/sectional curvature than the Euclidean plane, one obtains the general notion of CAT(κ)-spaces. For instance, CAT(-1)-spaces are those geodesic metric spaces whose geodesic triangles are at most as fat as geodesic triangles in the hyperbolic plane $\mathbb{H}^2$.

Example 7.6.3 (CAT(0)-Spaces)**.**

- The Euclidean plane $(\mathbb{R}^2, d_2)$ is a CAT(0)-space (by definition). More generally, for every $n \in \mathbb{N}$, the n-dimensional Euclidean space $(\mathbb{R}^n, d_2)$ is a CAT(0)-space (because every geodesic triangle in $\mathbb{R}^n$ lies in a Euclidean subspace of dimension 2).
- The hyperbolic plane $\mathbb{H}^2$ is a CAT(0)-space: Geodesic triangles in the hyperbolic plane are slimmer than their Euclidean comparison triangles (Exercise 7.E.33). More generally, for every $n \in \mathbb{N}_{\geq 2}$, the n-dimensional hyperbolic space $\mathbb{H}^n$ is a CAT(0)-space.
- The round sphere S^2 is *not* a CAT(0)-space: All non-degenerate geodesic triangles in S^2 are fatter than their Euclidean comparison triangles (Exercise 7.E.30).
- Geometric realisations of trees are 0-hyperbolic and therefore also CAT(0)-spaces.

Caveat 7.6.4. The property of being a CAT(0)-space is *not* only a global property, but also a local property (also small triangles have to satisfy the CAT(0)-inequality). This has several consequences:

- In general, hyperbolic metric spaces are *not* CAT(0)-spaces! For example, the sphere S^2 is a hyperbolic metric space, but it is *not* a CAT(0)-space.

- Being a CAT(0)-space is *not* a quasi-isometry invariant property for geodesic metric spaces. For example, the sphere S^2 is quasi-isometric to the one-point space, which clearly is a CAT(0)-space.

In view of Caveat 7.6.4, we have to be careful when trying to define the notion of CAT(0)-groups (see also Exercise 7.E.32). Instead of Cayley graphs, one uses group actions:

Definition 7.6.5 (CAT(0)-group)**.** A group is a CAT(0)-*group* if it admits a proper cocompact isometric action on a non-empty CAT(0)-space.

Example 7.6.6 (CAT(0)-groups)**.**

- For every $n \in \mathbb{N}$ the group $\mathbb{Z}^n$ is a CAT(0)-group (as witnessed by the translation action of $\mathbb{Z}^n$ on the Euclidean n-space). In particular, not every CAT(0)-group is hyperbolic.
- All finite groups are CAT(0)-groups (as witnessed by the trivial action on the one-point space).
- All finitely generated free groups are CAT(0)-groups (as witnessed by the translation action on the geometric realisations of Cayley graphs with respect to free generating sets).
- The fundamental groups of oriented closed connected surfaces are CAT(0)-groups. In the case of sphere, the fundamental group is trivial and thus CAT(0). In the case of the torus, the fundamental group is $\mathbb{Z}^2$ and thus CAT(0). In the case of higher genus, we choose a hyperbolic Riemannian metric and consider the corresponding deck transformation action on $\mathbb{H}^2$.
- All Coxeter groups are CAT(0)-groups [45, Chapter 12].

However, it remains an open problem to determine whether all hyperbolic groups are CAT(0)-groups or not.

Caveat 7.6.7. Being CAT(0) is *not* a quasi-isometry invariant among finitely generated groups: Let M be a closed connected Seifert 3-manifold with hyperbolic base-surface S and suppose that M is *not* finitely covered by a product of a surface and S^1; such manifolds indeed exist [77, IV.B.48]. We consider $G := \pi_1(M)$.

- Then G is *not* a CAT(0)-group [31, Theorem II.7.27]. The manifold M has $\widetilde{\mathrm{PSL}}$-geometry [7, Table.1 on p. 19][77, IV.B.48] and so G is quasi-isometric to $\widetilde{\mathrm{PSL}}$ (by the Švarc-Milnor lemma).
- On the other hand, $\widetilde{\mathrm{PSL}}$ is quasi-isometric to $\mathbb{H}^2 \times \mathbb{R}$ [77, IV.B.48], which is a CAT(0)-space, and thus (again by the Švarc-Milnor lemma) quasi-isometric to the CAT(0)-group $\pi_1(S) \times \mathbb{Z}$.

It turns out that CAT(0)-spaces share many of the properties of simply connected manifolds of non-positive sectional curvature and that CAT(0)-groups share many of the properties of fundamental groups of closed connected Riemannian manifolds of non-positive sectional curvature [31, Part II, Chapter III.Γ]. Also, many of the results for hyperbolic groups have suitable counterparts in the world of CAT(0)-groups:

- All CAT(0)-groups are finitely presented [31, Theorem III.Γ.1.1].
- The word problem is solvable for every CAT(0)-group [31, Theorem III.Γ.1.4].
- There is also a version of the splitting theorem (Theorem 7.5.26) for CAT(0)-spaces and CAT(0)-groups [31, Theorem II.6.21].
- We have seen that hyperbolic groups cannot contain $\mathbb{Z}^2$ as a subgroup (Corollary 7.5.15) because such a subgroup would lead to a flat plane in a hyperbolic space. Similarly, by the *flat torus theorem* [31, Chapter II.7], the maximal dimension of flat subspaces of CAT(0)-spaces controls the maximal rank of free Abelian groups in CAT(0)-groups. Moreover, this also leads to the *solvable subgroup theorem* [31, Theorem II.7.8] stating that all virtually solvable subgroups of CAT(0)-groups are in fact finitely generated virtually Abelian.

Outlook 7.6.8 (Application to 3-manifolds)**.** One of the most stunning recent applications of geometric group theory and CAT(0)-techniques (in the form of CAT(0)-cube complexes) is Agol's proof of Waldhausen's conjecture that all compact aspherical 3-manifolds are virtually Haken and of Thurston's conjecture that all hyperbolic 3-manifolds are virtually fibred [1]. This result revolutionised the theory of 3-manifolds and their fundamental groups.

Only very few groups are fundamental groups of closed surfaces [115]; in contrast, every finitely presented group is the fundamental group of some closed 4-manifold (and also in higher dimensions). Dimension 3 is a fascinating intermediate stage: The class of fundamental groups of compact 3-manifolds is rich enough for interesting examples, but still small and geometric enough to allow for good control [7]. One classical application of the study of fundamental groups of compact 3-manifolds with torus boundary is knot theory.

The definition of CAT(0)-spaces and CAT(0)-groups is based on comparison geometry. Alternatively, one can also capture non-positive curvature in a more combinatorial way. This leads to so-called *systolic complexes* and *systolic groups* [85].

7.E Exercises

(Quasi-)Hyperbolic spaces

Quick check 7.E.1 (Hyperbolicity?*).

1. Let (X, d) be a geodesic metric space with the following property: There is a constant $\delta \in \mathbb{R}_{\geq 0}$ such that for all geodesics γ and γ' in X that have the same start and end points we have the inclusions $\operatorname{im}\gamma \subset B_\delta^{X,d}(\operatorname{im}\gamma')$ and $\operatorname{im}\gamma' \subset B_\delta^{X,d}(\operatorname{im}\gamma)$. Is then X necessarily hyperbolic?
2. Is every hyperbolic space also a 2017-hyperbolic space?

Exercise 7.E.2 (A weird geodesic ray in the Euclidean plane*). Show that

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}^2 \\ t &\longmapsto t \cdot \big(\sin(\ln(1+t)), \cos(\ln(1+t))\big) \end{aligned}$$

is a quasi-isometric embedding with respect to the standard metrics on $\mathbb{R}$ and $\mathbb{R}^2$ respectively. (It follows that the stability theorem for quasi-geodesics does not hold in the Euclidean space $\mathbb{R}^2$).

Exercise 7.E.3 (Metric trees are hyperbolic**). Let X be a tree. Show that the geometric realisation $|X|$ of X is 0-hyperbolic.
Hints. A systematic way of organising this is via $\mathbb{R}$-trees (Exercise 7.E.4).

Exercise 7.E.4 ($\mathbb{R}$-trees**). A metric space (X, d) is an $\mathbb{R}$-*tree* if the following conditions are satisfied:

- For all $x, y \in X$ there exists a unique geodesic from x to y. We denote this geodesic by $[x, y]$ and its image by $|[x, y]|$.
- For all $x, y, z \in X$ with $|[y, x]| \cap |[x, z]| = \{x\}$ we have

$$\big|[y, x]\big| \cup \big|[x, z]\big| = \big|[y, z]\big|.$$

The second condition can also be reformulated in terms of the geodesics and not only of their images (how?!).

1. Let T be a tree. Prove that the geometric realisation $|T|$ is an $\mathbb{R}$-tree.
2. Is every $\mathbb{R}$-tree the geometric realisation of a tree?
3. Prove that every $\mathbb{R}$-tree is 0-hyperbolic.
4. Prove that every 0-hyperbolic metric space is an $\mathbb{R}$-tree.

Exercise 7.E.5 (Distance between geodesics and curves in hyperbolic spaces**). Complete the sketch proof of the Christmas tree lemma (Lemma 7.2.14): Let $\delta \in \mathbb{R}_{\geq 0}$ and let (X, d) be a δ-hyperbolic space. Let $\gamma\colon [0, L] \longrightarrow X$ be a continuous curve and let $\gamma'\colon [0, L'] \longrightarrow X$ be a geodesic from $\gamma(0)$ to $\gamma(L)$. Show for all $t \in [0, L']$ that

$$d\bigl(\gamma'(t), \operatorname{im}\gamma\bigr) \leq \delta \cdot \bigl|\log_2 L_X(\gamma)\bigr| + 1.$$

Exercise 7.E.6 (Geodesics in hyperbolic spaces starting at the same point**)**.** Let $\delta, D \in \mathbb{R}_{\geq 0}$, let (X, d) be a δ-hyperbolic space and let $\gamma\colon [0, L] \longrightarrow X$, $\gamma'\colon [0, L'] \longrightarrow X$ be geodesics in X with

$$\gamma(0) = \gamma'(0) \qquad \text{and} \qquad d\bigl(\gamma(L), \gamma'(L')\bigr) \leq D.$$

Show that γ and γ' are uniformly $(2 \cdot \delta + D)$-close, i.e.,

$$\forall_{t \in [0, \min(L, L')]} \ d\bigl(\gamma(t), \gamma'(t)\bigr) \leq 2 \cdot \delta + D \qquad \text{and} \qquad |L - L'| \leq D.$$

Exercise 7.E.7 (Local geodesics in hyperbolic spaces**)**.** Let (X, d) be a δ-hyperbolic space and let $c \in \mathbb{R}_{>8\delta}$. Let $\gamma\colon [0, L] \longrightarrow X$ be a *c-local geodesic*. Prove that if $\gamma'\colon [0, L'] \longrightarrow X$ is a geodesic that satisfies $\gamma(0) = \gamma'(0)$ and $\gamma(L) = \gamma(L')$, then

$$\operatorname{im}\gamma \subset B^{X,d}_{2 \cdot \delta}(\operatorname{im}\gamma').$$

Hints. Consider a point in $\operatorname{im}\gamma$ that has maximal distance from $\operatorname{im}\gamma'$ and then look at a suitable geodesic quadrilateral that connects $\operatorname{im}\gamma$ and $\operatorname{im}\gamma'$ and that contains this point on one of its sides.

Characterisations of hyperbolicity

Exercise 7.E.8 (Tripod triples*)**.** Let $a, b, c \in \mathbb{R}_{\geq 0}$ with $a + b \geq c$, $b + c \geq a$, $c + a \geq b$. Show that there are unique $\overline{a} \in [0, a]$, $\overline{b} \in [0, b]$, $\overline{c} \in [0, c]$ with

$$\overline{a} + \overline{b} = a, \qquad \overline{b} + \overline{c} = b, \qquad \overline{c} + \overline{a} = c$$

(Figure 7.30). We call $(\overline{a}, \overline{b}, \overline{c})$ the *tripod triple associated with* (a, b, c).

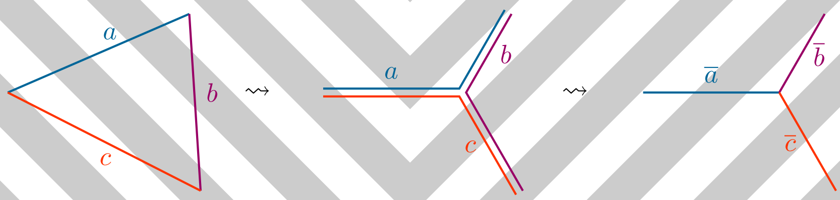

Figure 7.30.: Tripod triple

Exercise 7.E.9 (Slim triangles, thin triangles, insize**). Let (X, d) be a geodesic metric space, let $\Delta = (\gamma_0 \colon [0, L_0] \to X, \gamma_1 \colon [0, L_1] \to X, \gamma_2 \colon [0, L_2] \to X)$ be a geodesic triangle in X, and let $(\overline{L}_0, \overline{L}_1, \overline{L}_2)$ be the tripod triple associated with the triple (L_0, L_1, L_2) of side lengths of Δ (Exercise 7.E.8).

- Let $\delta \in \mathbb{R}_{\geq 0}$. The geodesic triangle Δ is *δ-thin* if

$$\begin{aligned} \forall_{t \in [0, \overline{L}_0]} \ & d\bigl(\gamma_0(t), \gamma_2(L_2 - t)\bigr) \leq \delta, \\ \forall_{t \in [0, \overline{L}_1]} \ & d\bigl(\gamma_1(t), \gamma_0(L_0 - t)\bigr) \leq \delta, \\ \forall_{t \in [0, \overline{L}_2]} \ & d\bigl(\gamma_2(t), \gamma_1(L_1 - t)\bigr) \leq \delta. \end{aligned}$$

- The *insize* of the geodesic triangle Δ is the diameter of the set

$$\bigl\{\gamma_0(\overline{L}_0), \gamma_1(\overline{L}_1), \gamma_2(\overline{L}_2)\bigr\}.$$

Prove that the following conditions are equivalent:

1. The space X is hyperbolic.
2. There exists a $\delta \in \mathbb{R}_{\geq 0}$ such that every geodesic triangle in X is δ-thin.
3. There exists a $\delta \in \mathbb{R}_{\geq 0}$ such that the insize of every geodesic triangle in X is at most δ.

Hints. A convenient way to proceed is to prove the implications "$2 \Longrightarrow 1$", "$1 \Longrightarrow 3$", "$3 \Longrightarrow 1$".

Exercise 7.E.10 (The Gromov product**). Let (X, d) be a metric space. For $x, y, z \in X$ the *Gromov product of x and y with respect to z* is defined by

$$(x \cdot y)_z := \frac{1}{2} \cdot \bigl(d(x, z) + d(y, z) - d(x, y)\bigr).$$

1. How is the Gromov product related to tripod triples (Exercise 7.E.8) associated with the side lengths of geodesic triangles?
2. Show that a geodesic metric space (X, d) is hyperbolic if and only if there exists a $\delta \in \mathbb{R}_{>0}$ with

$$\forall_{x,y,z,w \in X} \ (x \cdot y)_w \geq \min\bigl((x \cdot z)_w, (y \cdot z)_w\bigr) - \delta.$$

Hints. For instance, one can use the characterisation of hyperbolicity via the thin triangles condition (Exercise 7.E.9).

Exercise 7.E.11 (Hyperbolicity of graphs***). Let $X = (V, E)$ be a connected graph and let d be the induced metric on V. Let $L := (L_{x,y})_{x,y \in V}$ be a family of connected subgraphs of X. We say that X satisfies the *slim triangles condition with respect to L* if there exists a constant $C \in \mathbb{R}_{>0}$ satisfying the following conditions:

- For all $x, y \in V$, the vertices x and y belong to $L_{x,y}$.
- For all $x, y \in V$ with $d(x, y) \leq 1$, the d-diameter of $L_{x,y}$ is at most C.
- For all $x, y, z \in V$ the subgraph $L_{x,y}$ is contained in the C-neighbourhood (with respect to d) of the union of $L_{x,z}$ and $L_{z,y}$.

The goal of this exercise is to establish the following hyperbolicity criterion by Masur and Schleimer [116]:

1. Let X be a connected graph such that $|X|$ is hyperbolic. Show that X admits a family of subgraphs with respect to which it satisfies the slim triangles condition.
2. Suppose that X satisfies the slim triangles condition with respect to L. Prove that then $|X|$ is a hyperbolic metric space.
   Hints. It suffices to show that all geodesics in the graph between vertices $x, y \in V$ are uniformly close to the selected graph $L_{x,y}$ (why?). In order to prove this approximation of geodesics, one can first proceed as in the proof of the Christmas tree lemma (Lemma 7.2.14) and then as in the proof of the stability theorem (Theorem 7.2.11).
3. Look up in the literature how the *curve graph* is defined. Sketch the proof of Bowditch [24] of hyperbolicity of the curve graph via the criterion from the second part.

Hyperbolic groups

Quick check 7.E.12 (Growth of hyperbolic groups*).

1. Does every infinite finitely generated hyperbolic group have exponential growth?
2. Is every finitely generated group of exponential growth hyperbolic?

Exercise 7.E.13 (Quasi-hyperbolic groups*). Let G be a finitely generated group with finite generating set $S \subset G$. Show that G is a hyperbolic group if and only if (G, d_S) is $(1,1)$-quasi-hyperbolic.

Exercise 7.E.14 (Products and hyperbolic groups*). Characterise (i.e., give necessary and sufficient conditions) when the product $G \times H$ of two finitely generated groups G and H is hyperbolic.

Exercise 7.E.15 (0-Hyperbolic groups*). Let G be a finitely generated group that does not contain elements of order 2 and that has a finite generating set $S \subset G$ for which $|\mathrm{Cay}(G,S)|$ is 0-hyperbolic. Prove that G is free.

Exercise 7.E.16 (Free products and hyperbolicity**). Let G and H be finitely generated groups.

1. Let $G * H$ be hyperbolic. Are then also G and H hyperbolic?
2. Let G and H be hyperbolic. Is then also $G * H$ hyperbolic?

Exercise 7.E.17 (Geometric structures on manifolds*).

1. Does there exist a closed connected hyperbolic manifold whose fundamental group is isomorphic to $\mathrm{Out}(F_{2017})$?
2. Does there exists a closed connected flat manifold whose fundamental group is isomorphic to $F_{2017} \times F_{2017}$?

Exercise 7.E.18 (Residual finiteness$^{\infty}$*). Is every finitely generated hyperbolic group residually finite (Definition 4.E.1)?
Hints. This is an open problem!

The word problem

Exercise 7.E.19 (Elements of finite order in hyperbolic groups**). Let G be a finitely generated hyperbolic group. Show that G contains only finitely many conjugacy classes of elements of finite order.
Hints. Let $\langle S \mid R \rangle$ be a Dehn presentation of G and let $g \in G$ be an element of order $n \in \mathbb{N}_{>1}$. Let $w \in (S \cup S^{-1})^*$ be a word of minimal length that represents g in G. One then applies the Dehn property of $\langle S \mid R \rangle$ to the word w^n. How does this help to bound the length of w ?

Exercise 7.E.20 (Hyperbolic groups satisfy a linear isoperimetric inequality*). Show that finitely generated hyperbolic groups satisfy a linear isoperimetric inequality (Definition 6.E.6).
Hints. Use a Dehn presentation and Dehn's algorithm.

Exercise 7.E.21 (...and vice versa***). Let G be a finitely presented group that satisfies a linear isoperimetric inequality. Show that G is hyperbolic.
Hints. Proceed by contradiction. The proof then requires some serious puzzling and area/counting estimates [31, Chapter III.H.2].

Exercise 7.E.22 (The word problem and Dehn functions**). Let $\langle S \mid R \rangle$ be a finite presentation. Show that the word problem for $\langle S \mid R \rangle$ is solvable if and only if the associated Dehn function $\mathrm{Dehn}_{\langle S \mid R \rangle} \colon \mathbb{N} \longrightarrow \mathbb{N}$ (Definition 6.E.4) is a computable function.
Hints. If you are not comfortable with terms from computability theory, you should just solve this exercise with an intuitive notion of algorithmic computability.

Exercise 7.E.23 (The word problem in $\mathrm{BS}(1,2)$ **). Show that the word problem in $\mathrm{BS}(1,2)$ is solvable for the standard presentation $\langle a, b \mid bab^{-1} = a^2 \rangle$.
Hints. Look at the proof for the normal form in Exercise 2.E.22.

Elements of infinite order

Exercise 7.E.24 (Distorted elements in $\mathrm{BS}(1,2)$ *).

1. Give an example of an element of $\mathrm{BS}(1,2)$ that has infinite order and is distorted.
   Hints. The normal form of Exercise 2.E.22 might be helpful.
2. Conclude that $\mathrm{BS}(1,2)$ is *not* hyperbolic.

Exercise 7.E.25 (The conic trichotomy**). Let $A \in \mathrm{SL}_2(2, \mathbb{R}) \setminus \{E_2, -E_2\}$.

1. Show that the Möbius transformation on $\mathbb{H}^2$ associated with A is hyperbolic if and only if A is diagonalisable over $\mathbb{R}$. Furthermore, prove that this is equivalent to $|\operatorname{tr} A| > 2$.
2. Show that the Möbius transformation on $\mathbb{H}^2$ associated with A is parabolic if and only if A has real eigenvalues but is not diagonalisable over $\mathbb{R}$. Furthermore, prove that this is equivalent to $|\operatorname{tr} A| = 2$.
3. Show that the Möbius transformation on $\mathbb{H}^2$ associated with A is elliptic if and only if A has no real eigenvalues. Furthermore, prove that this is equivalent to $|\operatorname{tr} A| < 2$.

Exercise 7.E.26 (Quasi-convex subgroups**). Let G be a finitely generated group and let $S \subset G$ be a finite generating set.

1. Show that if $H \subset G$ is quasi-convex with respect to S, then H is finitely generated and the inclusion homomorphism $H \longrightarrow G$ is a quasi-isometric embedding.
2. Does the converse also hold? I.e., is every finitely generated subgroup H of G such that the inclusion $H \longrightarrow G$ is a quasi-isometric embedding quasi-convex with respect to S?

Exercise 7.E.27 (Intersections of quasi-convex subgroups***). Let G be a finitely generated group and let $S \subset G$ be a finite generating set of G. Show that if H, $H' \subset G$ are quasi-convex subgroups with respect to S, then also the intersection $H \cap H'$ is quasi-convex in G with respect to S.
Hints. Let $c \in \mathbb{R}_{\geq 0}$ be chosen in such a way that H and H' are c-quasi-convex in $|\mathrm{Cay}(G, S)|$. Let

$$C := \left|B_c^{G,S}(e)\right|^2.$$

Show that $H \cap H'$ is C-quasi-convex in $|\mathrm{Cay}(G, S)|$ using the following method: Let $h \in H \cap H'$ and let $\gamma\colon [0, L] \longrightarrow |\mathrm{Cay}(G, S)|$ be a geodesic joining e and h. For $t \in [0, L]$ with $\gamma(t) \in G$ then consider the set

$$\begin{aligned} M := \bigl\{g \in G \bigm| \gamma(t) \cdot g \in H \cap H',\ & \text{and for all } k \in G \\ & \text{that lie on a geodesic in } |\mathrm{Cay}(G,S)| \text{ from } e \text{ to } g \\ & \text{we have } d_S(\gamma(t) \cdot k, H) \leq c \text{ and } d_S(\gamma(t) \cdot k, H') \leq c\bigr\}. \end{aligned}$$

Show that M is non-empty and that every $d_S(e, \cdot\,)$-minimal element g in M satisfies $d_S(e, g) \leq C$.

Exercise 7.E.28 (Quasi-isometry rigidity of $\mathbb{Z}$ **). Let G be a finitely generated group with $G \sim_{\mathrm{QI}} \mathbb{Z}$. Let $g \in G$ be an element of infinite order (such an element exists by Corollary 7.5.7), and let $S \subset G$ be a finite generating set of G.

1. Show that there is a $c \in \mathbb{R}_{\geq 0}$ with the following property: For all $h \in G$ there is an $n \in \mathbb{Z}$ with $d_S(h, g^n) \leq c$.
2. Conclude that the infinite cyclic subgroup $\langle g \rangle_G$ has finite index in G.
Hints. Exercise 5.E.23 can be applied.

Exercise 7.E.29 (Lamplighter group**). Let

$$G := \bigoplus_{\mathbb{Z}} \mathbb{Z}/2 \rtimes_\alpha \mathbb{Z},$$

where $\alpha\colon \mathbb{Z} \longrightarrow \operatorname{Aut}\bigl(\bigoplus_{\mathbb{Z}} \mathbb{Z}/2\bigr)$ is given by translation.

1. Show that G is finitely generated.
2. Show that G is *not* hyperbolic.
3. Does G contain a subgroup that is isomorphic to $\mathbb{Z}^2$?

Non-positively curved groups

Exercise 7.E.30 (The round sphere*). Show that the round sphere S^2 is *not* a CAT(0)-space.

Exercise 7.E.31 (Geodesics and contractibility of CAT(0)-spaces*). Let (X, d) be a CAT(0)-space.

1. Let $x, y \in X$. Prove that there is a *unique* geodesic from x to y.
2. Prove that X is contractible (with respect to the metric topology).

Exercise 7.E.32 (CAT(0)-Cayley graphs*). Let $X = (V, E)$ be a connected graph.

1. When is V with the metric induced by X a CAT(0)-space?
2. When is the geometric realisation $|X|$ of X a CAT(0)-space?
3. Would it be a good idea to define that a finitely generated group G is a CAT(0)-group if there exists a finite generating set $S \subset G$ of G such that $\operatorname{Cay}(G, S)$ or $|\operatorname{Cay}(G, S)|$ is a CAT(0)-space?

Exercise 7.E.33 (The hyperbolic plane is non-positively curved**). Prove that the hyperbolic plane $\mathbb{H}^2$ is a CAT(0)-space.

Exercise 7.E.34 (Products of CAT(0)-spaces**).

1. Let (X, d_X) and (Y, d_Y) be CAT(0)-spaces. Show that also the product $X \times Y$ with the metric

$$\begin{aligned} (X \times Y) \times (X \times Y) &\longrightarrow X \times Y \\ \bigl((x, y), (x', y')\bigr) &\longmapsto \sqrt{d_X(x, x')^2 + d_Y(y, y')^2} \end{aligned}$$

   is a CAT(0)-space.
2. Conclude that the direct product of finitely many CAT(0)-groups is a CAT(0)-group.

8

Ends and boundaries

This chapter is a brief introduction into geometry "at infinity" and its applications to 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; more concisely, a boundary mechanism should be a functor promoting maps and properties from the wild world of quasi-isometries to the potentially tamer world of topology. In particular, boundaries are quasi-isometry invariants; surprisingly, in many cases, boundaries know enough about the underlying metric spaces to allow for interesting rigidity results.

In Chapter 8.1, we will formulate a wish-list for boundary constructions and we will outline a basic construction principle. In Chapters 8.2 and 8.3, we will discuss two such constructions in more detail, namely ends of groups and the Gromov boundary of hyperbolic groups. As sample applications, we will mention Stallings's decomposition theorem for groups with infinitely many ends, the ubiquity of free subgroups in hyperbolic groups, and Mostow rigidity.

Overview of this chapter

8.1 Geometry at infinity

What should a boundary 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 (Remark 5.1.12) to the category of (preferably compact) topological spaces. More explicitly, a notion of boundary should associate

- to every metric space X a (preferably compact) topological space X_∞,
- and to every 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 have finite distance from each other, 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 a trivial example:

Example 8.1.1 (Trivial boundary). We can assign to every metric space the topological space consisting of exactly one point, 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

$$X_\infty := \frac{\text{a suitable set of rays } [0,\infty) \longrightarrow X}{\text{a suitable equivalence relation on these rays}},$$

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 (Chapter 8.2) and the Gromov boundary of hyperbolic groups (Chapter 8.3).

8.2 Ends

The ends of a space can be viewed as the set of principal “regions” leading “to infinity.” Formally, the concept is described via rays and connected components when bounded pieces are removed. We will first give the definition in the most straightforward, geodesic, case (Chapter 8.2.1). We will then explain how this definition can be extended in a meaningful way to the quasi-geodesic case; in particular, this will also prove quasi-isometry invariance of ends (Chapter 8.2.2). After that we will focus on the case of finitely generated groups (Chapter 8.2.3).

8.2.1 Ends of geodesic spaces

As first step, we define ends of geodesic spaces as equivalence classes of proper rays (Figure 8.1).

Definition 8.2.1 (Ends of a geodesic 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 preimage $\gamma^{-1}(B) \subset [0,\infty)$ is bounded.
- Two proper rays γ, $\gamma'\colon [0,\infty) \longrightarrow X$ *represent the same end of* X if for every bounded subset $B \subset X$ 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 $\operatorname{Ends}(X)$ through convergence of sequences of ends in X 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 the ends $x_0, x_1, \dots$ and x respectively such that the following condition is satisfied:
  For every bounded set $B \subset X$ there is 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 $X \setminus B$.
  A subset $A \subset \operatorname{Ends}(X)$ is *closed* if the following holds: If $(x_n)_{n\in\mathbb{N}}$ is a sequence in A that converges to an end $x \in \operatorname{Ends}(X)$, then $x \in A$.

It should be noted that we used the term “proper” in the previous definition in the metric sense, not in the topological sense.

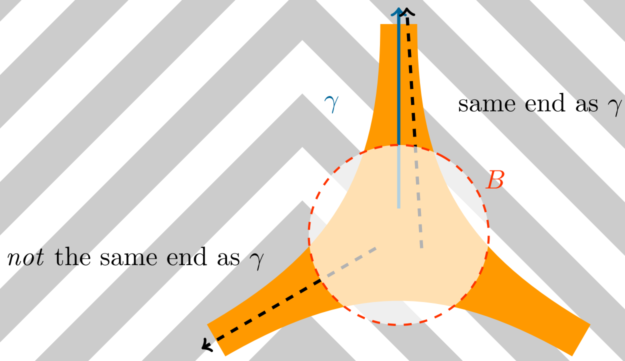

Figure 8.1.: The ends of a space, schematically

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. We work in a setup similar to the one by Bridson and Haefliger [31, Chapter I.8].

Example 8.2.3 (Ends of spaces).

- If X is a geodesic 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 (Figure 8.2). Moreover, a straightforward topological argument shows that every end of $\mathbb{R}$ can be represented by one of these two rays (Exercise 8.E.3); hence,

$$\operatorname{Ends}(\mathbb{R}) = \bigl\{\operatorname{end}(t \mapsto t), \operatorname{end}(t \mapsto -t)\bigr\}.$$

- If $n \in \mathbb{N}_{\geq 2}$, then the 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 $B \subset B_r^{\mathbb{R}^n, d_2}(0)$, and $\mathbb{R}^n \setminus B_r^{\mathbb{R}^n, d_2}(0)$ has exactly one path-component.
- Similarly, $|\operatorname{Ends}(\mathbb{H}^n)| = 1$ for all $n \in \mathbb{N}_{\geq 2}$.
- The subspace $\mathbb{R} \times \{0\} \cup \{0\} \times \mathbb{R}$ of the Euclidean plane $\mathbb{R}^2$ (with respect to the ℓ^1-metric) 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 the geometric realisation of T has infinitely many ends; more precisely, as a topological space, $\operatorname{Ends}(T)$ is a Cantor set [66, Chapter 13.4, 13.5][31, Exercise 8.31].

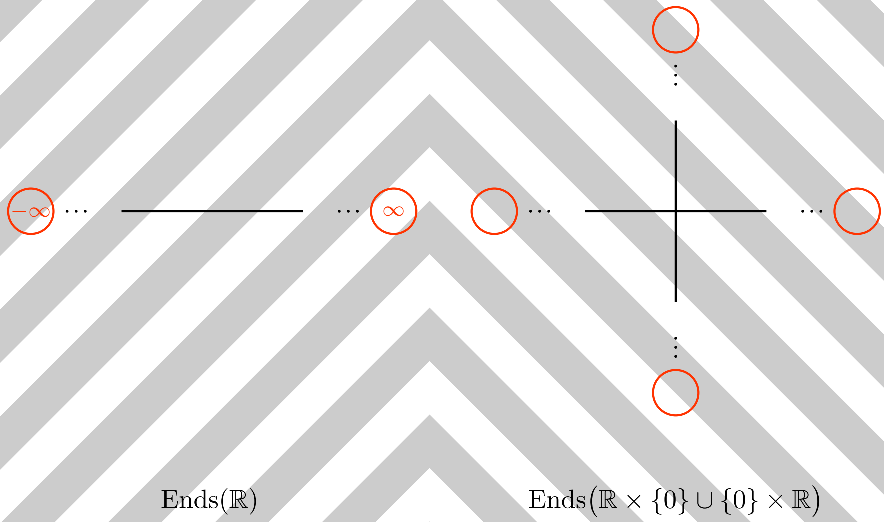

Figure 8.2.: Examples of ends (indicated by circles)

A first step towards understanding the space of ends of regular trees is the following simple case:

Example 8.2.4 (Ends of combs)**.** We consider the two combs X and Y in Figure 8.3. In both cases, we equip these subspaces of $\mathbb{R}^2$ with the path-metric associated with the Euclidean metric of $\mathbb{R}^2$; i.e., the distance of two points is the minimal Euclidean length of paths in X and Y respectively that connect the two given points. In this way, X and Y are geodesic metric spaces; hence, $\operatorname{Ends}(X)$ and $\operatorname{Ends}(Y)$ are defined.

- In X, we consider for $n \in \mathbb{N}$ the proper ray

$$\begin{aligned} \gamma_n \colon [0,\infty) &\longrightarrow X \\ t &\longmapsto (n,t). \end{aligned}$$

  Then the ends represented by the rays $(\gamma_n)_{n\in\mathbb{N}}$ converge in $\operatorname{Ends}(X)$ to the end represented by the proper ray

$$\begin{aligned} \gamma \colon [0,\infty) &\longrightarrow X \\ t &\longmapsto (t,0). \end{aligned}$$

- In Y, we consider for $n \in \mathbb{N}$ the proper ray

$$\begin{aligned} \gamma_n \colon [0,\infty) &\longrightarrow Y \\ t &\longmapsto \Bigl(\frac{1}{n}, t\Bigr). \end{aligned}$$

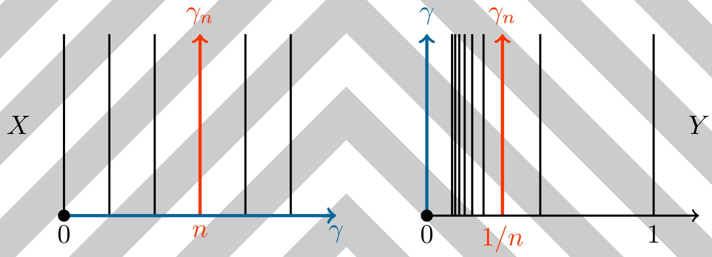

Figure 8.3.: Two combs

Then the ends represented by $(\gamma_n)_{n\in\mathbb{N}}$ do *not* converge in $\operatorname{Ends}(Y)$ to the end represented by the proper ray

$$\begin{aligned}\gamma\colon [0,\infty) &\longrightarrow Y\\ t &\longmapsto (0,t).\end{aligned}$$

8.2.2 Ends of quasi-geodesic spaces

We now extend the definition of ends to quasi-geodesic spaces; as always, we have to be careful with the handling of constants.

Definition 8.2.5 (Ends of quasi-geodesic spaces)**.** Let $c \in \mathbb{R}_{>0}$, $b \in \mathbb{R}_{\geq 0}$ and let X be a (c,b)-quasi-geodesic metric space.

- A *proper* (c,b)*-quasi-ray* in X is a map $\gamma\colon [0,\infty) \longrightarrow X$ that is proper in the metric sense (for every bounded set $B \subset X$, the preimage $\gamma^{-1}(B)$ is bounded in $[0,\infty)$) and that satisfies the estimate
$$\forall_{t,t'\in[0,\infty)}\ d\bigl(\gamma(t),\gamma(t')\bigr) \leq c\cdot|t-t'|+b.$$
- Two proper quasi-rays *represent the same quasi-end of* X, if, far out, they lie in the same quasi-path component. I.e., proper (c,b)-quasi-rays rays $\gamma,\gamma'\colon [0,\infty) \longrightarrow X$ represent the same quasi-end of X if for every bounded subset $B \subset X$ there exists a $t \in [0,\infty)$ such that $\gamma([t,\infty))$ and $\gamma'([t,\infty))$ lie in the same (c,b)-quasi-path-component (that is these points can be connected by (c,b)-quasi-paths).
Here, by a (c,b)*-quasi-path* we mean a (not necessarily continuous!) map $\gamma\colon [0,T] \longrightarrow X$ that satisfies
$$\forall_{t,t'\in[0,T]}\ d\bigl(\gamma(t),\gamma'(t)\bigr) \leq c\cdot|t-t'|+b.$$

- If $\gamma\colon [0,\infty) \longrightarrow X$ is a proper (c,b)-quasi-ray, then we write $\operatorname{end}_Q(\gamma)$ for the set of all proper (c,b)-quasi-rays that represent the same quasi-end as γ.
- We call

$$\operatorname{Ends}_Q(X) := \bigl\{\operatorname{end}_Q(\gamma) \bigm| \gamma\colon [0,\infty) \longrightarrow X \text{ is a proper } (c,b)\text{-quasi-ray}\bigr\}$$

  the *space of quasi-ends of* X (more precisely, the space of (c,b)-quasi-ends of X).
- We define a topology on $\operatorname{Ends}_Q(X)$ through convergence of sequences of quasi-ends in X to a point in $\operatorname{Ends}_Q(X)$: Let $(x_n)_{n\in\mathbb{N}} \subset \operatorname{Ends}_Q(X)$, and let $x \in \operatorname{Ends}_Q(X)$. We say that $(x_n)_{n\in\mathbb{N}}$ *converges to* x if there exist proper (c,b)-quasi-rays $(\gamma_n)_{n\in\mathbb{N}}$ and γ in X representing the quasi-ends $x_0, x_1, \dots$ and x respectively such that the following condition is satisfied:
  For every bounded set $B \subset X$ there is 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 (c,b)-quasi-path-component of $X \setminus B$.

Remark 8.2.6 (Quasi-ends and constants)**.** The initial choice of constants in the definition of quasi-ends does not affect the resulting space of quasi-ends: Let $c \in \mathbb{R}_{>0}$, $b \in \mathbb{R}_{\geq 0}$ and let (X,d) be a (c,b)-quasi-geodesic space. If $c' \in \mathbb{R}_{\geq c}$ and $b' \in \mathbb{R}_{\geq b}$, then every (c',b')-quasi-end can be represented by a proper (c,b)-quasi-ray and two proper (c,b)-quasi-rays represent the same (c,b)-quasi-end if and only if they represent the same (c',b')-quasi-ends. (Exercise 8.E.4).

Particularly nice examples of proper quasi-rays are quasi-geodesic rays. For proper geodesic metric spaces, every (quasi-)end can be represented by geodesic rays and the space of quasi-ends coincides with the space of ends:

Proposition 8.2.7 (Ends of geodesic spaces)**.** *Let X be a geodesic metric space and let $x \in X$.*

1. *If X is proper, then every end can be represented by a geodesic ray that starts at x.*
2. *There is a canonical homeomorphism* $\operatorname{Ends}(X) \cong \operatorname{Ends}_Q(X)$.

Proof. Ad 1. The basic idea is as follows: Let $\gamma\colon [0,\infty) \longrightarrow X$ be a proper ray. For every $n \in \mathbb{N}$ we pick a geodesic γ_n from x to $\gamma(n)$. Using the Arzelá-Ascoli theorem, one can then find a subsequence of these geodesics (extended constantly to all of $[0,\infty)$) that converges to a geodesic ray that starts at x, which will represent the same end as γ [31, Lemma I.8.28, Proposition I.8.29].

Ad 2. This can be shown by connecting the dots in proper quasi-rays and quasi-paths by geodesic segments (Exercise 8.E.5). □

In view of Proposition 8.2.7, we will in the following also use the symbol Ends to denote the space of quasi-ends of a quasi-geodesic metric space.

Proposition 8.2.8 (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 have finite distance from each other, 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 (Exercise 8.E.6). In order to prove well-definedness in the first part, we need the freedom to represent ends by more general 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 describe ends by some sort of quasi-rays. □

8.2.3 Ends of groups

In particular, we obtain a notion of ends for finitely generated groups:

Definition 8.2.9 (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. (Up to canonical homeomorphism, this does *not* depend on the choice of the finite generating set.)

In view of Proposition 8.2.8, the space of ends of a finitely generated group is a quasi-isometry invariant. From Example 8.2.3 we obtain:

Example 8.2.10 (Ends of groups).
- 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 $\mathbb{R}$).
- Finitely generated 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 8.2.11 (Possible numbers of ends of groups). *Let G be a finitely generated group. Then G has 0, 1, 2 or infinitely many ends.*

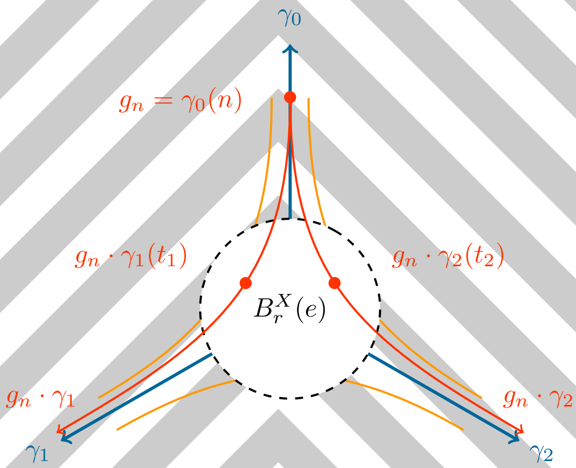

Figure 8.4.: There is no group with three ends

Proof. We proceed as Bridson and Haefliger [31, Theorem 8.32]; the notation is illustrated in Figure 8.4. *Assume* for a contradiction that G has a finite number of ends that is bigger than 2. By the functoriality of the ends construction, the space $\operatorname{Ends}(G)$ inherits a G-action from the left translation action of G on itself. By passing to a suitable finite index subgroup, we may assume without loss of generality that this G-action on $\operatorname{Ends}(G)$ is trivial (Exercise 8.E.10).

Let $S \subset G$ be a finite generating set, let $X := |\operatorname{Cay}(G,S)|$, and let $\gamma_0, \gamma_1, \gamma_2 \colon [0,\infty) \longrightarrow X$ be proper rays that represent three different ends of G. In view of Proposition 8.2.7, we may assume that $\gamma_0, \gamma_1, \gamma_2$ are geodesic rays that start at e. Because $\gamma_0, \gamma_1, \gamma_2$ represent three different ends, we will find $r \in \mathbb{R}_{>0}$ such that

$$\gamma_0\bigl((r,\infty)\bigr), \quad \gamma_1\bigl((r,\infty)\bigr), \quad \gamma_2\bigl((r,\infty)\bigr)$$

lie in three different path-components of $X \setminus B_r^X(e)$. Because X is a geodesic space, we obtain in particular: if $t, t' \in \mathbb{R}_{>2\cdot r}$, then

$$d\bigl(\gamma_1(t), \gamma_2(t)\bigr) > 2 \cdot r$$

(every path, whence every geodesic, between $\gamma_1(t)$ and $\gamma_2(t)$ must take a detour through the ball $B_r^X(e)$, which disconnects $\gamma_1((r,\infty))$ from $\gamma_2((r,\infty))$).

We now let the elements $g_n := \gamma_0(n) \in G$ with $n \in \mathbb{N}$ act on these rays. Because the G-action on $\operatorname{Ends}(G)$ is trivial, the rays $g_n \cdot \gamma_1$ and $g_n \cdot \gamma_2$ represent the same ends as γ_1 and γ_2 respectively. Let $n > 3 \cdot r$ and $j \in \{1,2\}$. Then the element $g_n = \gamma_0(n) \in \gamma_0((r,\infty))$ lies in a different path-component of $X \setminus B_r^X(e)$ than $\gamma_j((r,\infty))$. Because $g_n \cdot \gamma_j$ represents the same end as γ_j, the ray $g_n \cdot \gamma_j$ has to pass through $B_r^X(e)$ (in order to eventually reach the same component of $X \setminus B_r^X(e)$). In view of $n > 3 \cdot r$ we hence find a $t_j \in \mathbb{R}_{>2\cdot r}$ with $g_n \cdot \gamma_j(t_j) \in B_r^X(e)$. In particular, we obtain

$$d\bigl(\gamma_1(t_1), \gamma_2(t_2)\bigr) = d\bigl(g_n \cdot \gamma_1(t_1), g_n \cdot \gamma_2(t_2)\bigr) \leq r + r = 2 \cdot r.$$

However, this contradicts the estimate obtained in the first part of the proof.

Hence, there is no finitely generated group that has finitely many ends but more than two ends. □

Example 8.2.12. There is *no* finitely generated group that is quasi-isometric to the cross $\mathbb{R} \times \{0\} \cup \{0\} \times \mathbb{R}$ (with the ℓ^1-metric), because such a group would have exactly four ends (which is impossible by Theorem 8.2.11).

Moreover, using the Arzelá-Ascoli theorem, one can show that the space of ends of a finitely generated group is compact [31, 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 [31, Theorem I.8.32].

As we have seen above, every 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:

Definition 8.2.13 (Splitting over a finite group)**.** A finitely generated group G *splits over a finite group* if G is isomorphic to a group of the following type:

- an amalgamated free product $G_1 *_A G_2$, where A is a finite group, G_1 and G_2 are finitely generated groups, and

$$[G_1 : A] \geq 2, \quad [G_2 : A] \geq 2, \quad [G_1 : A] + [G_2 : A] \geq 5,$$

- or an HNN-extension $H*_\vartheta$, where ϑ is an isomorphism between finite subgroups of H that have index at least 2 in H.

Theorem 8.2.14 (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 splits over a finite group.*

Sketch of proof. Let G be a finitely generated group.

Ad 1. If G is finite, then $\operatorname{Ends}(G) = \emptyset$ (Example 8.2.10). Conversely, if G is infinite, then G contains at least one infinite proper quasi-ray (Exercise 3.E.11), which implies that $\operatorname{Ends}(G)$ is non-empty.

Ad 2. If G is virtually $\mathbb{Z}$, then $\operatorname{Ends}(G) \cong \operatorname{Ends}(\mathbb{Z})$ (by quasi-isometry invariance), which consists of exactly two elements (Example 8.2.10). Conversely, if G has exactly two ends, then one can show that G contains an element of infinite order, which generates a quasi-dense (whence finite index) subgroup of G (Exercise 8.E.11).

Ad 3. If G splits over a finite group (and hence is a non-trivial amalgamated free product or a non-trivial HNN-extension of the above type), then arguments similar to the case of free groups show that G has infinitely many ends.

Conversely, if G has infinitely many ends, one can concoct a tree on which G acts with finite stabilisers (the ends of G being the shadow of the branching of this tree). Then Bass-Serre theory shows that G is of the claimed shape [66, Chapter 13.6][53]. □

The second part, 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. Furthermore, one can also derive quasi-isometry rigidity of virtually free groups from decomposition results of this type [53]. From a more pessimistic point of view, Stallings's decomposition theorem tells us that most interesting groups will have exactly one end.

8.3 Gromov boundary

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. One example of such a construction is the Gromov boundary.

8.3.1 Gromov boundary of quasi-geodesic spaces

We will refine the construction of ends by looking at the points and directions of the rays directly instead of at looking only at the location with respect to path-components at infinity.

Definition 8.3.1 (Gromov boundary)**.** Let X be a quasi-geodesic 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 through convergence of sequences in ∂X 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.

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 $\operatorname{Mor}_{\mathsf{QMet}}([0,\infty),X)$ via pre-composition. Moreover, the canonical map

$$\begin{aligned} \operatorname{Mor}_{\mathsf{QMet}}([0,\infty),X)\big/\operatorname{QI}([0,\infty)) &\longrightarrow \partial X \\ [\gamma] &\longmapsto [\gamma] \end{aligned}$$

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 $\sup_{t\in[0,\infty)} d(\gamma(t),\gamma'\circ f(t)) \leq c$.

Proposition 8.3.3 (Quasi-isometry invariance of the Gromov boundary)**.** *Let X and Y be quasi-geodesic 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, continuous, and injective.
2. *If $f,\ g \colon X \longrightarrow Y$ are quasi-isometric embeddings that have finite distance from each other, then $\partial f = \partial g$.*

Hence, the Gromov boundary ∂ 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 $\partial f \colon \partial X \longrightarrow \partial Y$ is a homeomorphism.*

Proof. The first part is a straightforward computation (Exercise 8.E.14); the second part follows directly from the definitions. □

8.3.2 Gromov boundary of hyperbolic spaces

While the definitions above technically make sense for every (quasi-geodesic) metric space, in general, such a boundary is far too big (as there are too many quasi-geodesics) and rather difficult to keep under control (Exercise 8.E.22). However, for proper hyperbolic metric spaces, the Gromov boundary can be expressed in terms of geodesic rays instead of quasi-geodesic rays:

Theorem 8.3.4 (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*
$$\gamma'(0) = x \qquad \textit{and} \qquad \sup_{t\in[0,\infty)} d\bigl(\gamma(t), \gamma'(t)\bigr) < \infty.$$
4. *In particular, for all $x \in X$ the canonical maps*
$$\text{geodesic rays in } X \big/ \text{finite distance} \longrightarrow \partial X$$
$$\text{geodesic rays in } X \text{ starting in } x \big/ \text{finite distance} \longrightarrow \partial X$$
are bijective.

Proof. The second part follows by applying 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 claims (Exercise 8.E.15). The last part just subsumes the other parts. □

Example 8.3.5 (Gromov boundary of spaces).

- If X is a metric space of finite diameter, then clearly $\partial X = \emptyset$.
- 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 [87].
- 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} [18, Proposition A.5.10] (Exercise 8.E.18).

As a first example application, we use quasi-isometry invariance of the Gromov boundary (Proposition 8.3.3) to prove quasi-isometry invariance of hyperbolic dimension:

Corollary 8.3.6 (Quasi-isometry invariance of hyperbolic dimension). *Let n, $m \in \mathbb{N}_{\geq 2}$. 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 Example 8.3.5 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 [50, Corollary IV.2.3]; hence, we get $n - 1 = m - 1$, and so $n = m$. □

Outlook 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 [31, Definition III.H.3.5ff]; moreover, $\overline{X}$ in this topology is compact.

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 every isometry f of $\mathbb{H}^2$ yields a homeomorphism $\overline{f}\colon D^2 \longrightarrow D^2$. By the Brouwer fixed point theorem [50, Corollary IV.2.6], the latter map always has a fixed point. It is then possible to reformulate the conic trichotomy (Remark 7.5.17) for non-trivial orientation preserving isometries of $\mathbb{H}^2$ as follows [18]:

- 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.

8.3.3 Gromov boundary of groups

The quasi-isometry invariance of the Gromov boundary allows to define the Gromov boundary for hyperbolic groups:

Definition 8.3.8 (Gromov boundary of a group). Let G be a finitely generated group. The *Gromov boundary* of G is defined as $\partial G := \partial \mathrm{Cay}(G, S)$, 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.

Proposition 8.3.3 shows that the Gromov boundary of finitely generated groups is a quasi-isometry invariant and that the Gromov boundary of a group coincides with the Gromov boundary of the geometric realisations of its Cayley graphs. In the case of hyperbolic groups, we can hence describe the Gromov boundary in terms of geodesic rays on the geometric realisations (Theorem 8.3.4). Moreover, Example 8.3.5 yields:

Example 8.3.9 (Gromov boundary of groups)**.**

- 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 $\mathbb{R}$. We will see a converse of this fact in Proposition 8.3.12.
- 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 [13].
- 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 $\mathbb{H}^2$.
  We did know this already from the study of ends. However, using Gromov boundaries, we can do even better:
  There is no quasi-isometric embedding $\mathbb{H}^2 \longrightarrow F$: *Assume* for a contradiction that there is a quasi-isometric embedding $f \colon \mathbb{H}^2 \longrightarrow F$. Then the induced map $\partial f \colon \partial \mathbb{H}^2 \longrightarrow \partial F$ is continuous and injective. However, because $\partial \mathbb{H}^2 \cong S^1$ is connected and the Cantor set ∂F is totally disconnected, it follows that ∂f is constant. This contradicts injectivity of ∂f; hence, there is no such map f.
- 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 [168].

The geometry of Gromov boundaries of hyperbolic groups and spaces is quite rich [87]. In the following, we will focus merely on two aspects: how to find free groups in hyperbolic groups (Chapter 8.3.4) and how to prove rigidity results by boundary methods (Chapter 8.4).

8.3.4 Application: Free subgroups of hyperbolic groups

Using the language of the Gromov boundary of hyperbolic groups, we derive standard results on the ubiquity of free subgroups in hyperbolic groups.

Definition 8.3.10 (Boundary points of group elements, independence)**.** Let G be a hyperbolic group and let $g \in G$ be an element of infinite order. Then

$$\begin{aligned} \gamma(g)^+ \colon [0,\infty) &\longrightarrow G \\ t &\longmapsto g^{\lfloor t \rfloor}, \\ \gamma(g)^- \colon [0,\infty) &\longrightarrow G \\ t &\longmapsto g^{-\lfloor t \rfloor} \end{aligned}$$

are quasi-geodesic rays (Theorem 7.5.9). We write

$$g^\infty := \bigl[\gamma(g)^+\bigr], \quad g^{-\infty} := \bigl[\gamma(g)^-\bigr]$$

for the corresponding points in the Gromov boundary ∂G of G.

Elements $g, h \in G$ of infinite order are *independent* if in the Gromov boundary ∂G we have $\{g^\infty, g^{-\infty}\} \cap \{h^\infty, h^{-\infty}\} = \emptyset$.

Remark 8.3.11. Let G be a hyperbolic group and let $g \in G$ be an element of infinite order. Because

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longmapsto g^n \end{aligned}$$

is a quasi-isometric embedding (Theorem 7.5.9), it follows that $g^\infty \neq g^{-\infty}$.

This notation of boundary points of group elements is already useful in proving the following characterisation of elementary hyperbolic groups via the Gromov boundary.

Proposition 8.3.12 (Elementary hyperbolic groups and Gromov boundary)**.** *Let G be a hyperbolic group. Then $|\partial G| = 2$ if and only if G is virtually $\mathbb{Z}$.*

Proof. If G is virtually $\mathbb{Z}$, then G is quasi-isometric to $\mathbb{Z}$ and hence $\partial G \cong \partial \mathbb{Z}$. In particular, $|\partial G| = |\partial \mathbb{Z}| = 2$ (Example 8.3.9).

Conversely, we suppose that $|\partial G| = 2$. There are several ways to prove that then G is virtually $\mathbb{Z}$. One possibility is via ends: Because $|\partial G| = 2$, it is not hard to see from the definition that ∂G has exactly two connected components. Therefore, G also has exactly two ends (Exercise 8.E.20) and so G is virtually $\mathbb{Z}$ (Theorem 8.2.14, Exercise 8.E.11).

Alternatively, we can proceed as follows: Because of $\partial G \neq \emptyset$, the group G is infinite; hence, the hyperbolic group G contains an element g of infinite order (Theorem 7.5.1). Therefore, $\partial G = \{g^\infty, g^{-\infty}\}$. We will now show that $\langle g \rangle_G \cong \mathbb{Z}$ has finite index in G; equivalently, it suffices to show that $\langle g \rangle_G$ is quasi-dense in G (Exercise 5.E.23).

Assume for a contradiction that $\langle g \rangle_G$ is *not* quasi-dense in G, i.e., for a finite generating set $S \subset G$ and every $n \in \mathbb{N}$ there exists $x_n \in G$ with

$$d_S\bigl(x_n, \langle g \rangle_G\bigr) \geq n.$$

We can then connect e and x_n through a $(1,1)$-quasi-geodesic. Because the balls in $\operatorname{Cay}(G,S)$ are finite, a counting/inductive selection argument shows that there exists a $(1,1)$-quasi-geodesic ray $\gamma\colon [0,\infty) \longrightarrow G$ that passes through infinitely many elements of the set $\{x_n \mid n \in \mathbb{N}\}$ (this is a discrete version of the Arzelá-Ascoli theorem). By construction,

$$[\gamma] \notin \{g^\infty, g^{-\infty}\},$$

which contradicts $\partial G = \{g^\infty, g^{-\infty}\}$. Hence, $\langle g\rangle_G$ is quasi-dense in G, as desired. □

More generally, the boundary points of group elements determine the algebraic relations between the given group elements to a large extent:

Theorem 8.3.13 (Mini Tits alternative for hyperbolic groups)**.** *Let G be a hyperbolic group and let $g, h \in G$ be elements of infinite order. Then:*

1. *If g and h are* not *independent, then the subgroup $\langle g,h\rangle_G$ is virtually $\mathbb{Z}$ and $\{g^\infty, g^{-\infty}\} = \{h^\infty, h^{-\infty}\}$ in ∂G.*
2. *If g and h are independent, then there are $m, n \in \mathbb{N}$ such that $\langle g^m, h^n\rangle_G$ is free of rank 2.*

Before we start with the actual proof, we will first look at these statements from an intuitive, geometric, point of view. If g and h are independent, then the geodesic lines given by the powers of g and h respectively, grow further and further apart. Therefore, far out on these lines, one can set up a ping-pong situation and therefore find powers of g and h that generate a free subgroup of rank 2.

Conversely, if g and h are *not* independent, then g and h (or inverses of these elements) act by translation on the same geodesic ray (up to finite distance). However, in negative curvature, there is (up to finite error) only a one-dimensional family of isometries that induces translations on any given geodesic line. Therefore, it is plausible that $\langle g,h\rangle_G$ is virtually $\mathbb{Z}$.

Proof. Ad 1. Let $\{g^\infty, g^{-\infty}\} \cap \{h^\infty, h^{-\infty}\} \neq \emptyset$; without loss of generality, we may assume $g^\infty = h^\infty$. Then Lemma 8.3.14 below shows that there exists an $n \in \mathbb{Z} \setminus \{0\}$ such that $h^n \cdot g = g \cdot h^n$. Thus, $\langle g,h\rangle_G$ is contained in the centraliser $C_G(h^n)$. Because h^n has infinite order, the centraliser $C_G(h^n)$ is virtually $\mathbb{Z}$ (Theorem 7.5.10). In particular, also the infinite group $\langle g,h\rangle_G$ is virtually $\mathbb{Z}$.

Moreover, because $\langle g,h\rangle_G$ is virtually $\mathbb{Z}$ and because the powers of g and h give rise to quasi-geodesic lines in G, which have to fit inside of $\langle g,h\rangle_G$, we also have $\{g^\infty, g^{-\infty}\} = \{h^\infty, h^{-\infty}\}$ (Exercise 8.E.26).

Ad 2. Let g and h be independent and let $S \subset G$ be a finite generating set of G. Then, by Lemma 8.3.15 below, there exists $R \in \mathbb{N}_{>0}$ such that the sets

$$A := \bigl\{x \in G \bigm| d_S(x, \langle g\rangle_G) < d_S(x, \{g^{-R}, \dots, g^R\})\bigr\},$$
$$B := \bigl\{x \in G \bigm| d_S(x, \langle h\rangle_G) < d_S(x, \{h^{-R}, \dots, h^R\})\bigr\}$$

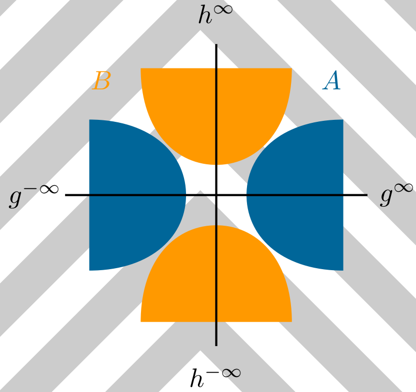

Figure 8.5.: Setting up the ping-pong table, using two diverging geodesic lines

are disjoint (Figure 8.5). We will now apply the ping-pong lemma to this situation: We already know that A and B are non-empty and $B \not\subset A$. Moreover, for all $n \in \mathbb{Z} \setminus \{0\}$ we have

$$h^{n \cdot 3 \cdot R} \cdot A \subset B,$$

which can be seen as follows: Let $x \in A$. Then $A \cap B = \emptyset$ implies that $x \notin B$. By definition, therefore all points in $\langle h \rangle_G$ that are closest to x lie in $\{h^{-R}, \dots, h^R\}$. As the word metric d_S is left-invariant, the points in $\langle h \rangle_G$ closest to $h^{n \cdot 3 \cdot R} \cdot x$ will lie in $\{h^{n \cdot 3 \cdot R - R}, \dots, h^{n \cdot 3 \cdot R + R}\}$, which is disjoint from $\{h^{-R}, \dots, h^R\}$. Therefore, $h^{n \cdot 3 \cdot R} \cdot x \in B$.

The same argument shows that for all $n \in \mathbb{Z} \setminus \{0\}$ we have

$$g^{n \cdot 3 \cdot R} \cdot B \subset A.$$

Hence, we can apply the ping-pong lemma (Theorem 4.3.1) and obtain that the group $\langle g^{3 \cdot R}, h^{3 \cdot R} \rangle_G$ is free of rank 2. □

It remains to establish the two lemmas used in the proof:

Lemma 8.3.14. *Let G be a hyperbolic group and let $g, h \in G$ be elements of infinite order with $g^{\infty} = h^{\infty}$ in ∂G. Then there exists an $n \in \mathbb{Z} \setminus \{0\}$ such that h^n commutes with g.*

Proof. Let $S \subset G$ be a finite generating set of G. In view of Theorem 8.3.4, there exists a constant $c \in \mathbb{R}_{>0}$ such that

$$\forall_{n \in \mathbb{N}} \ d_S(g^n, h^n) \le c.$$

We then obtain for all $m \in \mathbb{N}$ that

$$\begin{aligned} d_S(h^{-m}\cdot g\cdot h^m, g) &= d_S(g\cdot h^m, h^m\cdot g)\\ &\le d_S(g\cdot h^m, h^{m+1}) + d_S(h^m\cdot h, h^m\cdot g)\\ &\le d_S(g\cdot h^m, h^{m+1}) + c\\ &\le d_S(g\cdot h^m, g\cdot g^m) + d_S(g^{m+1}, h^{m+1}) + c\\ &\le c + c + c, \end{aligned}$$

and so $\{h^{-m}\cdot g\cdot h^m \mid m\in\mathbb{N}\}\subset B^{G,S}_{3\cdot c}(g)$. Because this ball is finite, there exist $m,k\in\mathbb{N}$ with $m\neq k$ such that $h^{-m}\cdot g\cdot h^m = h^{-k}\cdot g\cdot h^k$. In other words, $h^{m-k}\cdot g = g\cdot h^{m-k}$. ☐

Lemma 8.3.15. *Let G be a hyperbolic group with finite generating set $S\subset G$ and let $g,h\in G$ be independent elements of infinite order. Then there exists an $R\in\mathbb{N}_{>0}$ such that the sets*

$$\begin{aligned} A &:= \bigl\{x\in G \bigm| d_S(x,\langle g\rangle_G) < d_S(x,\{g^{-R},\dots,g^R\})\bigr\},\\ B &:= \bigl\{x\in G \bigm| d_S(x,\langle h\rangle_G) < d_S(x,\{h^{-R},\dots,h^R\})\bigr\} \end{aligned}$$

are disjoint.

Proof. Because $\{g^\infty,g^{-\infty}\}\cap\{h^\infty,h^{-\infty}\}=\emptyset$, there exists $r\in\mathbb{R}_{>0}$ such that every geodesic in $|\mathrm{Cay}(G,S)|$ joining points in $\langle g\rangle_G$ and $\langle h\rangle_G$ has to pass through the ball $B^{G,S}_r(e)$ (Exercise 8.E.17). Let $\delta\in\mathbb{R}_{\ge 0}$ be a hyperbolicity constant for $|\mathrm{Cay}(G,S)|$ and let $R\in\mathbb{N}_{>0}$ with

$$\forall_{n\in\mathbb{N}_{>R}}\ \ d_S(e,g^n) > 2\cdot(r+\delta).$$

Then the sets A and B as defined in the lemma are disjoint:

Assume for a contradiction that there exists $x\in A\cap B$. Let $g_x\in\langle g\rangle_G$ and $h_x\in\langle h\rangle_G$ be closest points to x in $\langle g\rangle_G$ and $\langle h\rangle_G$, respectively. Moreover, let γ be a geodesic in $|\mathrm{Cay}(G,S)|$ joining g_x and h_x. In particular, γ passes through a point $z\in B^{G,S}_r(e)$. Because $|\mathrm{Cay}(G,S)|$ is δ-hyperbolic, we may assume without loss of generality that there is a point z' on a geodesic from x to g_x with $d_S(z,z')\le\delta$. Using the definition of A and g_x, we obtain

$$\begin{aligned} 2\cdot(r+\delta) &< d_S(e,g_x)\\ &\le d_S(e,z') + d_S(z',g_x) \le d_S(e,z') + d_S(z',e)\\ &\le 2\cdot\bigl(d_S(e,z) + d_S(z,z')\bigr)\\ &\le 2\cdot(r+\delta), \end{aligned}$$

which is a contradiction. Hence, $A\cap B=\emptyset$. ☐

Example 8.3.16 (Independent elements in $\mathrm{SL}(2,\mathbb{Z})$)**.** The group $G:=\mathrm{SL}(2,\mathbb{Z})$ is hyperbolic (Example 7.3.3). We consider the elements

$$g := \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad h := \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \in \mathrm{SL}(2, \mathbb{Z}),$$

which both are of infinite order. Arguments similar to those in the proof of Proposition 4.4.2 show that $\{g, h\}$ is a generating set of $\mathrm{SL}(2, \mathbb{Z})$. In particular, the subgroup $\langle g, h \rangle_G$ is *not* free (alternatively, one can easily find concrete non-trivial relations between g and h). On the other hand, $\{g^\infty, g^{-\infty}\} \cap \{h^\infty, h^{-\infty}\} = \emptyset$ and $\langle g^2, h^2 \rangle_G$ is free (Example 4.4.1). This example also shows that in general passage to higher powers is needed in the second part of Theorem 8.3.13.

We can now deduce that most hyperbolic groups contain a free group of rank 2:

Corollary 8.3.17 (Ubiquity of free groups in hyperbolic groups)**.** *Let G be a hyperbolic group. Then either G is virtually cyclic or G contains a free group of rank 2 (and hence has exponential growth).*

Proof. Clearly, the two alternatives exclude each other. We now consider the case that G is *not* virtually cyclic and we prove that then G has to contain a free group of rank 2. Because G is not virtually cyclic, G is infinite; in particular, G contains an element g of infinite order (Theorem 7.5.1). In view of Theorem 8.3.13, it suffices to find an element $h \in G$ of infinite order that is independent of g. Because G is not virtually cyclic, there exist elements $k \in G$ of arbitrarily large distance to $\langle g \rangle_G$. Therefore, Lemma 7.5.14 implies that there is a $k \in G$ such that the conjugate $h := k \cdot g \cdot k^{-1}$ satisfies for all $\varepsilon \in \{-1, 1\}$:

$$\sup_{n \in \mathbb{Z}} d_S(h^n, g^{\varepsilon \cdot n}) = \sup_{n \in \mathbb{Z}} d_S(k \cdot g^n k^{-1}, g^{\varepsilon \cdot n}) = \infty.$$

With g also h has infinite order and so using Theorem 8.3.4 we can reformulate the previous expression as

$$\{h^\infty, h^{-\infty}\} \neq \{g^\infty, g^{-\infty}\}.$$

By the first part of Theorem 8.3.13, this already implies that g and h are independent; therefore, the second part of Theorem 8.3.13 can be applied. $\square$

Outlook 8.3.18 (Acylindrically hyperbolic groups)**.** A wide-ranging generalisation of hyperbolic groups are acylindrically hyperbolic groups [139]: The notion of acylindrically hyperbolic groups is based on the observation that many features of hyperbolic groups do not require a hyperbolic Cayley graph, but only a suitable action on a hyperbolic metric space. Hyperbolic groups can be characterised as the finitely generated groups that admit proper, cocompact actions on proper hyperbolic metric spaces (Exercise 8.E.30). One now replaces proper actions by acylindrical actions:

- An isometric action of a group G on a metric space (X, d) is *acylindrical* if: For every $\varepsilon \in \mathbb{R}_{>0}$, there exist $r, n \in \mathbb{N}$ such that for all $x, y \in X$ with $d(x, y) \geq r$ we have

$$\bigl|\{g \in G \mid d(x, g \cdot x) \leq \varepsilon \text{ and } d(y, g \cdot y) \leq \varepsilon\}\bigr| \leq n.$$

- A group G is *acylindrically hyperbolic* if there exists a (not necessarily finite) generating set $S \subset G$ such that $|\mathrm{Cay}(G, S)|$ is hyperbolic, the left translation action of G on $|\mathrm{Cay}(G, S)|$ is acylindrical, and $\partial|\mathrm{Cay}(G, S)|$ contains more than two points.

Because $|\mathrm{Cay}(G, S)|$ is *not* a proper metric space if S is infinite, one needs a version of the Gromov boundary that is robust enough for this situation, e.g., the description via the Gromov product (Exercise 8.E.21).

For example, finitely generated hyperbolic groups are acylindrically hyperbolic if and only if they are *not* virtually cyclic (Exercise 8.E.32, Exercise 8.E.33). Moreover, the class of acylindrically hyperbolic groups subsumes the following, geometrically relevant, classes of groups [139]:

- Most mapping class groups of surfaces are acylindrically hyperbolic.
- Outer automorphism groups of free groups of rank at least 2 are acylindrically hyperbolic.
- ...

For example, analogously to non-elementary hyperbolic groups, acylindrically hyperbolic groups contain free groups of rank 2 (this is a far-reaching generalisation of Corollary 8.3.17).

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 $\mathrm{Isom}(\mathbb{H}^n)$. If Γ and Λ are isomorphic groups, then they are conjugate in $\mathrm{Isom}(\mathbb{H}^n)$, i.e., there exists a $g \in \mathrm{Isom}(\mathbb{H}^n)$ with*

$$g \cdot \Gamma \cdot g^{-1} = \Lambda.$$

The corresponding statement for flat manifolds does *not* hold: Scaling the flat metric on a torus 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. 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 object [18, Chapter B.4].

Sketch proof of Mostow rigidity. Why are the geometric version and the algebraic version of Mostow rigidity equivalent? The universal covering of hyperbolic n-manifolds is hyperbolic n-space $\mathbb{H}^n$. In particular, hyperbolic manifolds have a contractible universal covering and so are classifying spaces for the fundamental group. Standard arguments in algebraic topology concerning the homotopy theory of classifying spaces then show that hyperbolic manifolds are homotopy equivalent if and only if they have isomorphic fundamental groups [48, Chapter 8.8]. On the other hand, covering theory shows that a connected hyperbolic n-manifold M with fundamental group $\Gamma \subset \operatorname{Isom}(\mathbb{H}^n)$ is isometric to the quotient $\Gamma \setminus \mathbb{H}^n$ and that isometries between hyperbolic n-manifolds lift to isometries of $\mathbb{H}^n$. Now the equivalence between the geometric and the algebraic version follows from a straightforward calculation.

We will now sketch a proof of the geometric version of Mostow rigidity: Let $f\colon M \longrightarrow N$ be a homotopy equivalence between closed connected hyperbolic n-manifolds. 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 $\partial\mathbb{H}^n$ is conformal (i.e., locally angle-preserving) with respect to the canonical homeomorphism $\partial\mathbb{H}^n \cong S^{n-1}$; here, the condition that $n \geq 3$ enters. One way to show that $\partial\widetilde{f}$ is conformal is Gromov's proof via simplicial volume and regular ideal simplices [126, 18, 146].

By a classical result from hyperbolic geometry, every conformal map on $\partial\mathbb{H}^n$ can be obtained as the boundary map of an isometry of $\mathbb{H}^n$ [18, Proposition A.5.13]; 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 [63, 53].

Outlook 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 [106].

8.E Exercises

Ends of spaces

Quick check 8.E.1 (Ends of subspaces*). Let X be a geodesic metric space and let $Y \subset X$ be a subspace that is geodesic with respect to the subspace metric. We write $i\colon Y \longrightarrow X$ for the inclusion map.

1. Does the map i always induce an injection $\operatorname{Ends}(Y) \longrightarrow \operatorname{Ends}(X)$?
2. Does the map i always induce a surjection $\operatorname{Ends}(Y) \longrightarrow \operatorname{Ends}(X)$?

Exercise 8.E.2 (Topology via convergence*). Let X be a set and let C be a subset of $X^{\mathbb{N}} \times X$. We say that a set $A \subset X$ is *C-closed*, if the following holds: For all $\bigl((x_n)_{n\in\mathbb{N}}, x\bigr) \in C$ we have

$$\bigl(\forall_{n\in\mathbb{N}}\ \ x_n \in A\bigr) \Longrightarrow x \in A.$$

Here, we view C as a specification of when a sequence "converges" to a point in X.

1. Show that the set
$$T_C := \bigl\{U \subset X \bigm| \text{the set } X \setminus U \text{ is } C\text{-closed}\bigr\}$$
is a topology on X.
2. Let $((x_n)_{n\in\mathbb{N}}, x) \in C$. Show that $(x_n)_{n\in\mathbb{N}}$ converges to x with respect to the topology T_C.
3. Does the converse also hold?!

Exercise 8.E.3 (Ends of the real line*).

1. Show that the proper rays
$$\begin{aligned} [0,\infty) &\longrightarrow \mathbb{R} \\ t &\longmapsto t \\ t &\longmapsto -t \end{aligned}$$
do *not* represent the same ends in $\operatorname{Ends}(\mathbb{R})$.
2. Show that every end of $\mathbb{R}$ can be represented by one of these two proper rays.
3. What is the topology on $\operatorname{Ends}(\mathbb{R})$?

Exercise 8.E.4 (Quasi-ends and constants**). Let $c \in \mathbb{R}_{>0}$, $b \in \mathbb{R}_{\geq 0}$ and let (X,d) be a (c,b)-quasi geodesic space. Moreover, let $c' \in \mathbb{R}_{\geq c}$ and $b' \in \mathbb{R}_{\geq b}$.

1. Show that every (c',b')-quasi-end is represented by a proper (c,b)-quasi-ray.

2. Show that two proper (c, b)-quasi-rays represent the same (c, b)-quasi-end if and only if they represent the same (c', b')-quasi-end.

Exercise 8.E.5 (Quasi-ends**)**.** Let X be a geodesic metric space. Fill in the details of the proof of Proposition 8.2.7.

1. Let X be proper and $x \in X$. Fill in the details for the Arzelá-Ascoli argument that shows that every end in $\operatorname{Ends}(X)$ is represented by a geodesic ray that starts at x. Do the same for $\operatorname{Ends}_Q(X)$.
2. Construct a canonical homeomorphism $\operatorname{Ends}(X) \cong \operatorname{Ends}_Q(X)$ by "connecting the dots" of proper quasi-rays and quasi-paths through geodesics.

Exercise 8.E.6 (QI-invariance of ends**)**.** Let X and Y be quasi-geodesic metric spaces.

1. Let $f \colon X \longrightarrow Y$ be a quasi-isometric embedding. Show that the map

$$\begin{aligned} \operatorname{Ends}_Q(f) \colon \operatorname{Ends}_Q(X) &\longrightarrow \operatorname{Ends}_Q(Y) \\ \operatorname{end}_Q(\gamma) &\longmapsto \operatorname{end}_Q(f \circ \gamma) \end{aligned}$$

is well-defined and continuous.
2. Let $f, g \colon X \longrightarrow Y$ be quasi-isometric embeddings that have finite distance from each other. Show that then $\operatorname{Ends}_Q(f) = \operatorname{Ends}_Q(g)$.

Exercise 8.E.7 (Ends via π_0 ***)**.**

1. Look up the definition of the path-components functor π_0 in algebraic topology.
2. Formulate the definition of Ends for proper metric spaces in terms of π_0. *Hints.* For a streamlined formulation, it might be convenient to think about the inverse system $(X \setminus K)_{K \in K(X)}$, where X is a topological space and $K(X)$ is the set of all compact subsets of X.
3. Conclude: Proper homotopy equivalences $X \longrightarrow Y$ between proper geodesic metric spaces induce homeomorphisms $\operatorname{Ends}(X) \longrightarrow \operatorname{Ends}(Y)$.

Ends of groups

Quick check 8.E.8 (Ends of groups*)**.**

1. Does the Heisenberg group have infinitely many ends?
2. Does every group of exponential growth have infinitely many ends?

Exercise 8.E.9 (Free group vs. hyperbolic plane*)**.** Use ends to prove that the free group of rank 2 is *not* quasi-isometric to $\mathbb{H}^2$.

Exercise 8.E.10 (Groups that act trivially on their ends*)**.** Let G be a finitely generated group with finitely many ends. By functoriality, the space $\operatorname{Ends}(G)$ inherits a G-action from the left translation of G on itself.

1. Prove that there is a finite index subgroup H of G such that the restriction of the G-action to H is the trivial H-action on $\operatorname{Ends}(G)$.
   Hints. Look at the kernel of the action.
2. Conclude that H acts trivially on $\operatorname{Ends}(H)$ and that H has the same number of ends as G.

Exercise 8.E.11 (Groups with two ends**)**.** Let G be a finitely generated group that has exactly two ends.

1. Show that G contains an element of infinite order.
   Hints. Look at an element of G whose translation action on the ends is trivial and that is far away from the neutral element.
2. Show that the subgroup of G generated by an element of infinite order has finite index in G. In particular, G is quasi-isometric to $\mathbb{Z}$ and virtually $\mathbb{Z}$.

Gromov boundary of spaces

Quick check 8.E.12 (Finite Hausdorff distance*)**.** We consider the subsets

$$X := \mathbb{R} \times \{0\}, \qquad Y := \mathbb{R} \times \{1\}, \qquad Z := \{0\} \times \mathbb{R}$$

of the Euclidean plane $\mathbb{R}^2$.

1. Do X and Y have finite Hausdorff distance?
2. Do X and Z have finite Hausdorff distance?

Exercise 8.E.13 (Small Gromov boundary**)**.**

1. Give an example of a hyperbolic metric space X that is unbounded but has empty boundary ∂X.
   Hints. Look at a sufficiently spiky tree.
2. Let G be a hyperbolic finitely generated group. Show that G is finite if and only if $\partial G = \emptyset$.

Exercise 8.E.14 (Induced maps on Gromov boundaries**)**.** Let X and Y be quasi-geodesic metric spaces and let $f \colon X \longrightarrow Y$ be a quasi-isometric embedding.

1. Show that the map

$$\begin{aligned} \partial f \colon \partial X &\longrightarrow \partial Y \\ [\gamma] &\longmapsto [f \circ \gamma] \end{aligned}$$

   is well-defined.
2. Show that the map $\partial f \colon \partial X \longrightarrow \partial Y$ is continuous and injective.
   Hints. Which statement about convergence of points in ∂X will you need to prove? How can you modify the involved quasi-geodesic rays in such a way that convergence holds after applying the (not necessarily continuous!) quasi-isometric embedding f ?

Exercise 8.E.15 (Quasi-geodesic rays in hyperbolic spaces**)**.** Let X be a proper hyperbolic metric space.

1. Let $\gamma\colon \mathbb{N} \longrightarrow X$ be a quasi-geodesic ray. Show that there exists a geodesic ray $\gamma'\colon \mathbb{R}_{\geq 0} \longrightarrow X$ such that $\operatorname{im}\gamma$ and $\operatorname{im}\gamma'$ have finite Hausdorff distance.
   Hints. When in need, look at Figure 8.6 and call the Arzelá-Ascoli theorem for help!
2. Let $x \in X$ and let $\gamma \in \mathbb{R}_{\geq 0} \longrightarrow X$ be a geodesic ray. Show that there exists a geodesic ray $\gamma'\colon \mathbb{R}_{\geq 0} \longrightarrow X$ with
$$\gamma'(0) = x \qquad \text{and} \qquad \sup_{t\in\mathbb{R}_{\geq 0}} d\big(\gamma(t), \gamma'(t)\big) < \infty.$$

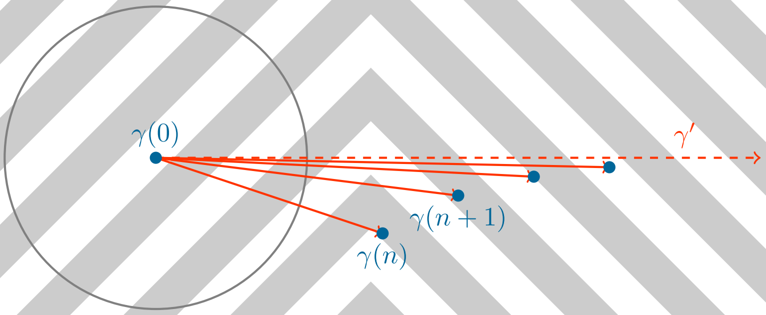

Figure 8.6.: Using geodesics to approximate quasi-geodesic rays

Exercise 8.E.16 (Topology of the Gromov boundary in the geodesic description**)**.** Let X be a proper hyperbolic metric space. How can the topology on ∂X be described in terms of geodesic rays?

Exercise 8.E.17 (Diverging rays**)**.** Let X be a proper geodesic hyperbolic metric space and let $\gamma, \eta\colon [0,\infty) \longrightarrow X$ be geodesic rays with $\gamma(0) = \eta(0)$ and $\sup_{t\in[0,\infty)} d(\gamma(t), \eta(t)) = \infty$.

1. Let $c \in \mathbb{R}_{>0}$. Show that then there is a $T \in [0,\infty)$ with
$$\forall_{s,t\in[T,\infty)}\ d\big(\gamma(t), \eta(s)\big) \geq c.$$
   Hints. For instance, one can use Lemma 7.5.5.
2. Show that there exists $r \in \mathbb{R}_{>0}$ such that every geodesic joining a point on $\operatorname{im}\gamma$ with a point on $\operatorname{im}\eta$ passes through the ball $B_r^{X,d}(\gamma(0))$.
3. Formulate and prove a version of the second part for quasi-geodesic rays instead of geodesic rays.

Exercise 8.E.18 (Boundary of the hyperbolic plane**). Prove that the Gromov boundary $\partial \mathbb{H}^2$ is homeomorphic to S^1. Try to give a proof in the halfplane model as well as a proof in the Poincaré disk model (Figure 8.7). Illustrate your arguments with suitable pictures!
Hints. Use the classification of geodesic rays starting at a given point in the hyperbolic plane! (Appendix A.3)

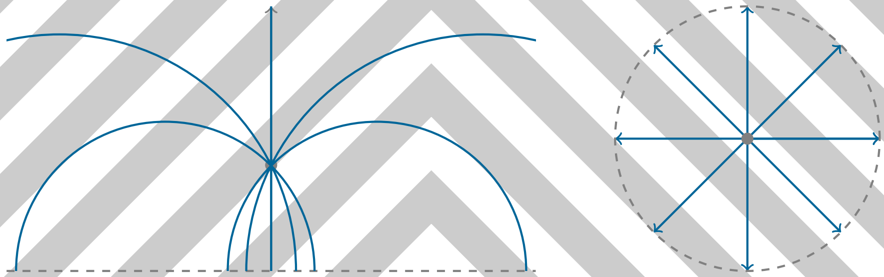

Figure 8.7.: The boundary of the hyperbolic plane, schematically

Exercise 8.E.19 (Visibility of the Gromov boundary**). Let X be a proper hyperbolic metric space and let $x_+, x_- \in \partial X$ with $x_+ \neq x_-$. Prove that there exists a geodesic line $\gamma \colon \mathbb{R} \longrightarrow X$ that satisfies

$$[\gamma_+] = x_+ \in \partial X \qquad \text{and} \qquad [\gamma_-] = x_- \in \partial X,$$

where $\gamma_+ := \gamma|_{[0,\infty)}$ and

$$\begin{aligned} \gamma_- \colon [0,\infty) &\longrightarrow X \\ t &\longmapsto \gamma(-t). \end{aligned}$$

Hints. Use the Arzelá-Ascoli theorem.

Exercise 8.E.20 (Gromov boundary vs. ends***). Let X be a proper hyperbolic metric space. Show that the canonical map

$$\partial X \longrightarrow \operatorname{Ends}(X)$$

induces a bijection from the set of connected components of ∂X to $\operatorname{Ends}(X)$.

Exercise 8.E.21 (Gromov product and Gromov boundary***). Let X be a proper hyperbolic metric space and let $x \in X$.

1. Let $\gamma \colon \mathbb{R}_{\geq 0} \longrightarrow X$ be a geodesic ray. Show that then the Gromov products (Exercise 7.E.10) satisfy

$$\lim_{m,n\to\infty} \bigl(\gamma(n) \cdot \gamma(m)\bigr)_x \to \infty.$$

2. Let $\gamma, \gamma' \colon \mathbb{R}_{\geq 0} \longrightarrow X$ be geodesic rays. Show that γ and γ' represent the same point in the Gromov boundary ∂X if and only if
$$\lim_{n,m\to\infty} \bigl(\gamma(n) \cdot \gamma'(m)\bigr)_x = \infty.$$
3. Let $(x_n)_{n\in\mathbb{N}}$ be a sequence in X with $\lim_{n,m\to\infty}(x_n \cdot x_m)_x = \infty$. Show that there exists a (quasi-)geodesic ray $\gamma \colon \mathbb{R}_{\geq 0} \longrightarrow X$ with
$$\lim_{n,m\to\infty} \bigl(x_n \cdot \gamma(m)\bigr)_x = \infty.$$
4. How can the topology on the Gromov boundary ∂X be described in terms of the Gromov product?

Exercise 8.E.22 (Boundary of the Euclidean plane$^{\infty}$*)**.** Give a reasonable description of $\partial\mathbb{R}^2$ (where we endow $\mathbb{R}^2$ with the Euclidean metric). Beware! The Euclidean plane has lots of quasi-geodesic rays ...
Hints. This seems to be an open problem!

Exercise 8.E.23 (Asymptotic cones***)**.**
1. Look up the definition of *asymptotic cones* in the literature.
2. Why/how can one view asymptotic cones as "geometry at infinity"?
3. How can hyperbolic groups be characterised via asymptotic cones?
4. How can asymptotic cones be used for rigidity results?

Gromov boundary of groups

Quick check 8.E.24 (Translation action on the Gromov boundary*)**.** Let G be a finitely generated group. Then the left translation action of G on itself induces a continuous G-action on the Gromov boundary ∂G.
1. Is the left translation action of F_2 on ∂F_2 free?
2. Let $x \in \partial F_2$. Is then $F_2 \cdot x$ the whole boundary ∂F_2 ?

Exercise 8.E.25 (Groups with small Gromov boundary?!*)**.** Show that there is *no* hyperbolic group G with $|\partial G| = 1$.

Exercise 8.E.26 (Gromov boundary and virtually $\mathbb{Z}$ subgroups*)**.** Let G be a hyperbolic group, let $g, h \in G$ be elements of infinite order with the property that $\langle g, h\rangle_G$ is virtually $\mathbb{Z}$. Show that $\{g^{\infty}, g^{-\infty}\} = \{h^{\infty}, h^{-\infty}\}$ holds in ∂G.

Exercise 8.E.27 (Centre of hyperbolic groups**)**.** Let G be a finitely generated hyperbolic group that is *not* virtually $\mathbb{Z}$. Show that the centre of G is finite.
Hints. There are several approaches. One is to look at a free subgroup of rank 2 in G and at centralisers of free generators.

Exercise 8.E.28 (Geometric structures on manifolds**)**.**
1. Does there exist a closed connected hyperbolic manifold whose fundamental group is isomorphic to F_{2017} ?

2. Does there exist a closed connected hyperbolic manifold whose fundamental group has $S^1 \times S^1$ as boundary?
   Hints. This requires some algebraic topology.

Exercise 8.E.29 (Combinatorial horoballs***). Let $X = (V, E)$ be a connected, locally finite graph. The graph $H(X)$ is defined as follows: The set of vertices of $H(X)$ is $V \times \mathbb{N}$. A vertex of the form (x, n) with $x \in V$ and $n \in \mathbb{N}$ is a vertex of *level* n. The set of edges of $H(X)$ is given by

$$\bigl\{\{(x,n),(x,n+1)\} \bigm| x \in V,\ n \in \mathbb{N}\bigr\} \qquad \text{(vertical edges)}$$
$$\cup \bigl\{\{(x,n),(y,n)\} \bigm| \{x,y\} \in E,\ n \in \mathbb{N}\bigr\} \qquad \text{(horizontal edges)}.$$

The *combinatorial horoball of* X is the metric space $\|H(X)\|$ whose underlying set is the geometric realisation of $H(X)$ and whose metric is the path-metric associated with the following lengths of edges (Figure 8.8):
- All vertical edges have length 1.
- All horizontal edges of level n have length $1/2^n$.

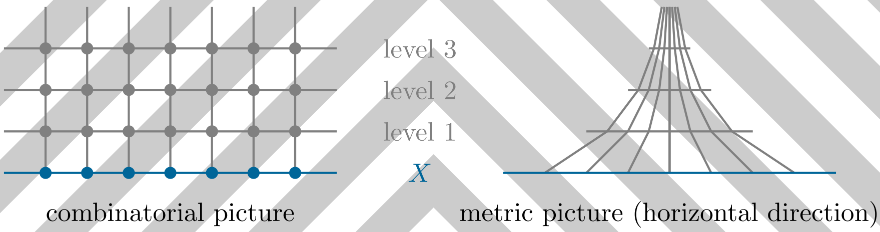

Figure 8.8.: Combinatorial horoballs, schematically

1. Describe the geodesics in $\|H(X)\|$.
   Hints. Similarly to $\mathbb{H}^2$, in $\|H(X)\|$ the path ⎾ is shorter than ⎿⏌ . . .
2. Prove that $\|H(X)\|$ is a proper hyperbolic metric space.
3. Compute the Gromov boundary $\partial\|H(X)\|$.

Exercise 8.E.30 (Proper actions on hyperbolic spaces**). Let G be a finitely generated group.
1. Show that G is hyperbolic if and only if G admits a proper cocompact isometric action on a (non-empty) proper hyperbolic metric space.
2. Does G always admit a cocompact isometric action on a (non-empty) proper hyperbolic metric space?
3. Show that G admits a proper isometric action on a (non-empty) proper hyperbolic metric space.
   Hints. Let G act on the combinatorial horoball (Exercise 8.E.29) associated with a Cayley graph of G with respect to a finite generating set of G.

Acylindrically hyperbolic groups

Quick check 8.E.31 (Acylindrical actions*).

1. Is every action of a group on a metric space of finite diameter acylindrical?
2. Does every group admit an acylindrical action on some metric space?

Exercise 8.E.32 (Hyperbolic acylindrically hyperbolic groups*). Show that every hyperbolic group that is not virtually cyclic is acylindrically hyperbolic.

Exercise 8.E.33 (Elementary hyperbolic groups**). Let G be virtually $\mathbb{Z}$.

1. Let $S \subset G$ be a generating set. Show that $\operatorname{diam}\bigl(\operatorname{Cay}(G,S), d_S\bigr)$ is finite if and only if S is infinite.
2. Conclude: The group G is *not* acylindrically hyperbolic.

Exercise 8.E.34 (Acylindrical hyperbolicity and quasi-isometries?!$^\infty$*). Let G and H be finitely generated quasi-isometric groups, where G is acylindrically hyperbolic. Does this imply that also H is acylindrically hyperbolic?
Hints. This seems to be an open problem!

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 (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) and bilipschitz equivalence rigidity of non-amenable groups (Chapter 9.4). Moreover, we will briefly sketch (co)homological characterisations (Chapter 9.2.4).

Overview of this chapter

9.1 Amenability via means

We will introduce amenable groups via invariant means. A mean can be viewed 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 turn $\ell^\infty(X,\mathbb{R})$ into a real vector space. If G is a group, then every left G-action on X induces a left G-action on $\ell^\infty(X,\mathbb{R})$ 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* if there exists 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.

9.1.1 First examples of amenable groups

Example 9.1.2 (Amenability of finite groups)**.** Finite groups are amenable: If G is a finite group, then the averaging operator

$$\begin{aligned} \ell^\infty(G,\mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto \frac{1}{|G|}\cdot \sum_{g\in G} f(g) \end{aligned}$$

is a G-invariant mean on $\ell^\infty(G,\mathbb{R})$.

Proposition 9.1.3 (Amenability of Abelian groups)**.** *Every Abelian group is amenable.*

The proof relies on the Markov-Kakutani fixed point theorem from functional analysis [144, Proposition 0.14]:

Theorem 9.1.4 (Markov-Kakutani fixed point theorem). *Let V be a locally convex $\mathbb{R}$-vector space (e.g., a normed $\mathbb{R}$-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. Then 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 set

$$M := \bigl\{ m \in L \bigm| m(1) = 1 \text{ and } m \text{ is positive} \bigr\}.$$

Clearly, M is a closed, convex subspace of L, and $G \cdot M \subset M$. Moreover, M is non-empty (it contains evaluation at e) and compact by the Banach-Alaoglu theorem [95, 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. □

Classically, the most prominent example of a non-amenable group is the free group of rank 2. The idea behind the following proof will be systematically exploited in the context of Banach-Tarski type decomposition paradoxa (Chapter 9.2.2).

Proposition 9.1.5 (Non-amenability of free groups). *The free group F_2 of rank 2 is* 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

$$A \cup a^{-1} \cdot A = F_2.$$

Denoting characteristic functions of subsets with $\chi_{\dots}$, we obtain

$$\begin{aligned} 1 = m(1) = m(\chi_{F_2}) &\leq m(\chi_A + \chi_{a^{-1} \cdot A}) \\ &= m(\chi_A) + m(a^{-1} \cdot \chi_A) \\ &= 2 \cdot m(\chi_A), \end{aligned}$$

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

$$\begin{aligned}1 = m(1) \geq m(\chi_{A \cup b \cdot A \cup b^2 \cdot A}) &= m(\chi_A) + m(\chi_{b \cdot A}) + m(\chi_{b^2 \cdot A}) \\ &= m(\chi_A) + m(b \cdot \chi_A) + m(b^2 \cdot \chi_A) \\ &= 3 \cdot m(\chi_A) \\ &\geq \frac{3}{2},\end{aligned}$$

which is impossible. 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 $G = \bigcup_{i \in I} G_i$. 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 the 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 the groups N and Q are amenable with left-invariant means m_N and m_Q, respectively. Without loss of generality, we may assume $N \subset G$ and $Q = G/N$. 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 well-defined 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. □

Corollary 9.1.7 (Amenability of locally amenable groups)**.** *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 subgroups of G are amenable.

Conversely, let all finitely generated subgroups of G be 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 (Amenability of solvable groups)**.** *If a group is solvable, then it is also 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.

Outlook 9.1.9 (Elementary amenable groups)**.** The class of so-called *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. However, not every amenable group is elementary amenable; this can for example be seen via the Grigorchuk groups [69].

Corollary 9.1.10. *Groups that contain a free subgroup of rank* 2 *are* not *amenable.*

Proof. This follows from the fact that free groups of rank 2 are non-amenable (Proposition 9.1.5) and that amenability is inherited by subgroups (Proposition 9.1.6). □

Corollary 9.1.11 (Amenability vs. Hyperbolicity). *Let G be a hyperbolic group. Then either G is virtually cyclic or G is* not *amenable.*

Proof. If G is not virtually cyclic, then G contains a free subgroup of rank 2 (Corollary 8.3.17). So, Corollary 9.1.10 implies that G is not amenable. □

Remark 9.1.12 (The von Neumann problem). The notion of amenability was originally introduced by John von Neumann [130]. In view of Corollary 9.1.10, he asked whether every non-amenable group contained a free subgroup of rank 2.

A first candidate in this direction seemed to be Thompson's group F (Example 2.2.21); while it is known that F does *not* contain a free subgroup of rank 2 [33], even now (2017) it remains an open problem to decide whether this group is amenable or not.

Von Neumann's question was answered negatively by Ol'shanskii [134] who constructed a non-amenable torsion group; in particular, such a group cannot contain a free subgroup of rank 2. In contrast, the von Neumann problem has a positive answer for many well-behaved classes of groups such as linear groups (Exercise 9.E.7).

Outlook 9.1.13 (Geometric von Neumann problems). We conclude with a brief overview of geometric versions of the von Neumann problem. While the situation is rather involved in the case of groups (Remark 9.1.12), it does simplify in more geometric contexts and leads to positive answers:

Theorem 9.1.14 (The von Neumann problem for Cayley graphs [160]). *Let $k \in \mathbb{N}_{>2}$. Then a finitely generated group is non-amenable if and only if it admits a Cayley graph with respect to a finite generating set that has a regular spanning tree of degree k.*

Theorem 9.1.15 (The von Neumann problem for actions [175]). *A UDBG space X is non-amenable (in the sense of Definition 9.2.9) if and only if X admits a free action by a free group of rank 2 by bilipschitz maps at bounded distance from the identity.*

Theorem 9.1.16 (The von Neumann problem in measurable group theory [65]). *Let G be a countable discrete non-amenable group. Then there exists a measurable ergodic essentially free action of F_2 on $([0,1]^G, \lambda^{\otimes G})$ such that almost every G-orbit of the Bernoulli shift action of G on $[0,1]^G$ decomposes into F_2-orbits.*

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:

- almost invariant subsets (Følner sequences),
- paradoxical decompositions and the Banach-Tarski paradox,
- (co)homological characterisations.

A more thorough treatment of amenable groups can be found in the books by Paterson [144] and Runde [152]; representation theoretic aspects are explained also in the book by Bekka, de la Harpe, and Valette [17].

9.2.1 Følner sequences

We begin with a geometric characterisation of amenability via Følner sets. We first formulate the notion of Følner sets for UDBG spaces (Definition 5.6.11).

Definition 9.2.1 (Følner sequence)**.** Let X be a UDBG space.

- If $F \subset X$ and $r \in \mathbb{N}$, then the *r-boundary of F in X* is given 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

$$\lim_{n\to\infty} \frac{|\partial_r^X F_n|}{|F_n|} = 0.$$

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 subsets $F \subset G$ we have

$$\partial_1^{G,d_S}(F) = \bigl\{g \in G \setminus F \bigm| \exists_{s \in S \cup S^{-1}} \;\; g \cdot s \in F\bigr\} = F \cdot (S \cup S^{-1}) \setminus F.$$

Hence, $\partial_1^{G,d_S}(F)$ being "small" in comparison to F means that we have $g \cdot (S \cup S^{-1}) \subset F$ for "many" $g \in F$.

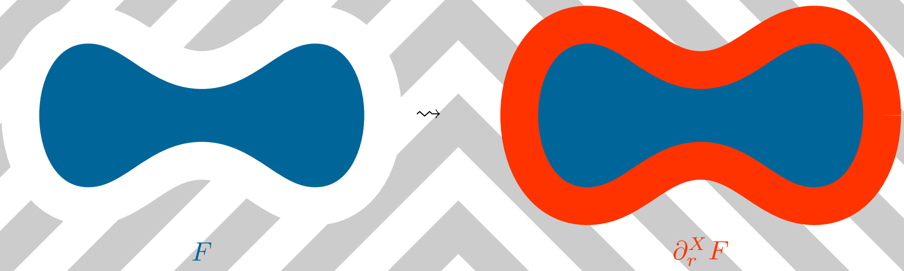

Figure 9.1.: The r-boundary of a set, schematically

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 $(B_n^{F,S}(e))_{n \in \mathbb{N}}$ 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 (Theorem 9.2.6 or Exercise 9.E.15).

More generally, if concentric balls in a UDBG space do *not* contain a Følner subsequence, then they have to grow fast enough (in order to allow for enough space for a "thick" boundary):

Proposition 9.2.3 (Subexponential growth yields Følner sequences)**.** *Let X be a UDBG space and let $x_0 \in X$. For $n \in \mathbb{N}$ we consider the ball $F_n := B_n^X(x_0)$. If the growth function*

$$\begin{aligned} \beta \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ n &\longmapsto |F_n| \end{aligned}$$

of X (based at the point x_0) has subexponential growth (i.e., $\beta \prec (x \mapsto 2^x)$ but $\beta \not\sim (x \mapsto 2^x)$), then $(F_n)_{n \in \mathbb{N}}$ contains a Følner subsequence for X.

Proof. Because β has subexponential growth, we have

$$\forall_{r \in \mathbb{N}} \ \forall_{N \in \mathbb{N}} \ \forall_{\varepsilon \in \mathbb{R}_{>0}} \ \exists_{n \in \mathbb{N}_{\geq N}} \ \frac{\beta(n+r)}{\beta(n)} < 1 + \varepsilon.$$

Looking at "$r = j, N = n_{j-1} + 1, \varepsilon = 1/j$" we can inductively find a strictly increasing sequence $(n_j)_{j \in \mathbb{N}}$ with

$$\forall_{j \in \mathbb{N}} \ \frac{\beta(n_j + j)}{\beta(n_j)} < 1 + \frac{1}{j}.$$

We now prove that the subsequence $(F_{n_j})_{j \in \mathbb{N}}$ is a Følner sequence for X: Let $r \in \mathbb{N}$. By definition of the r-boundary, we have $\partial_r^X F_n \subset B_{n+r}^X(x_0) \setminus B_n^X(x_0)$

for all $n \in \mathbb{N}$ and thus

$$\left|\partial_r^X F_n\right| \leq \left|B_{n+r}^X(x_0) \setminus B_n^X(x_0)\right| = \beta(n+r) - \beta(n).$$

Therefore, for all $j \in \mathbb{N}_{\geq r}$ we obtain the estimate

$$\begin{aligned}\frac{\left|\partial_r^X F_{n_j}\right|}{|F_{n_j}|} &\leq \frac{\beta(n_j+r) - \beta(n_j)}{\beta(n_j)} = \frac{\beta(n_j+r)}{\beta(n_j)} - 1 \leq \frac{\beta(n_j+j)}{\beta(n_j)} - 1 \\ &\leq 1 + \frac{1}{j} - 1 = \frac{1}{j},\end{aligned}$$

which tends to 0 for $j \to \infty$. Thus, $(F_{n_j})_{j\in\mathbb{N}}$ is a Følner sequence for X. ☐

Corollary 9.2.4 (Subexponential growth yields Følner sequences, group case)**.** *Let G be a finitely generated group of subexponential growth and let $S \subset G$ be a finite generating set of G. Then (G, d_S) admits a Følner sequence.*

Proof. This is an immediate consequence of Proposition 9.2.3. ☐

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 9.2.5. *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. Every 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 9.2.6 (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 the UDBG space (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 [39, Chapter J]. The almost invariance of the Følner sets translates into invariance in the limit. Hence, G is amenable. Alternatively, such a mean can also be obtained using weak*-limits in $\ell^\infty(G,\mathbb{R})'$ [39, Theorem 4.9.2].

Conversely, let G be amenable. Recall that $\ell^1(G,\mathbb{R})$ is weak*-dense in the double dual $\ell^1(G,\mathbb{R})'' = \ell^\infty(G,\mathbb{R})'$ [151, Exercise I.3.5, Section I.4.5]. Then every invariant mean on $\ell^\infty(G,\mathbb{R})$ is an element of $\ell^\infty(G,\mathbb{R})'$ 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 Følner sets in the sense of Proposition 9.2.5 [144, Proposition (0.8), Lemma (4.7)][17, Appendix G]. □

Because we know an explicit Følner sequence for $\mathbb{Z}$ we could attempt to use the "recipe" in the proof of Theorem 9.2.6 to produce an explicit 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 9.2.7. *Every finitely generated group of subexponential growth is amenable.*

Proof. Finitely generated groups of subexponential growth admit a Følner sequence (Corollary 9.2.4) and hence are amenable (Theorem 9.2.6). □

The converse of this corollary does not hold in general:

Caveat 9.2.8 (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 also amenable (Corollary 9.1.8).

In view of the characterisation of amenable groups in terms of Følner sequences (Theorem 9.2.6), we are hence led to the following definition of amenability for spaces:

Definition 9.2.9 (Amenable spaces)**.** A UDBG space is *amenable* if it admits a Følner sequence.

9.2.2 Paradoxical decompositions

Another geometric characterisation of amenability is based on decomposition paradoxa. These decomposition properties are the main ingredient in

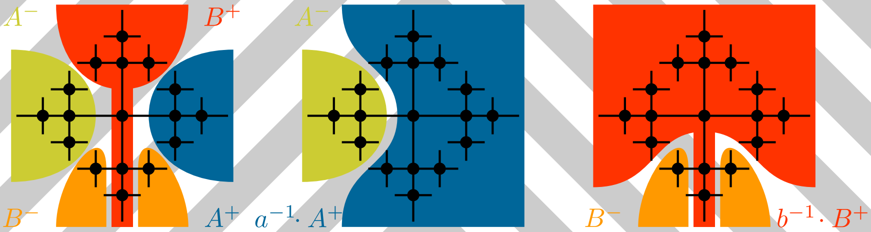

Figure 9.2.: A paradoxical decomposition of a free group of rank 2

the Banach-Tarski paradox (Chapter 9.2.3). In short, a paradoxical decomposition of a group is a decomposition into finitely many disjoint subsets such that these subsets can be rearranged by translations into two copies of the group:

Definition 9.2.10 (Paradoxical group)**.** A group G is *paradoxical* if it admits a paradoxical decomposition. A *paradoxical decomposition of* G is a pair $\bigl((A_g)_{g\in K}, (B_h)_{h\in L}\bigr)$ where $K, L \subset G$ are finite and $(A_g)_{g\in K}$, $(B_h)_{h\in L}$ are families of subsets of G with the property that

$$G = \Bigl(\bigcup_{g\in K} A_g\Bigr) \cup \Bigl(\bigcup_{h\in L} B_h\Bigr), \qquad G = \bigcup_{g\in K} g \cdot A_g, \qquad G = \bigcup_{h\in L} h \cdot B_h$$

are *disjoint* unions.

Proposition 9.2.11 (Non-Abelian free groups are paradoxical)**.** *Free groups of rank at least* 2 *are paradoxical.*

Proof. We use the description of free groups in terms of reduced words. In order to keep notation simple, we consider the case of rank 2 (higher ranks basically work in the same way). Let F be a free group of rank 2, freely generated by $\{a, b\}$. We then define the following subsets of F (Figure 9.2):

1. Let A^+ be the set of all reduced words starting with a positive power of a.
2. Let A^- be the set of all reduced words starting with a negative power of a.
3. Let B^+ be the set containing the neutral element, all powers of b as well as all reduced words starting with a positive power of b.
4. Let B^- be the set of all reduced words starting with a negative power of b, excluding the powers of b.

Then

$$F = A^+ \cup A^- \cup B^+ \cup B^-, \qquad F = A^- \cup a^{-1} \cdot A^+ \qquad F = B^- \cup b^{-1} \cdot B^+$$

are disjoint unions (Figure 9.2). So $((A_e, A_{a^{-1}})_{\{e,a^{-1}\}}, (B_e, B_{b^{-1}})_{\{e,b^{-1}\}})$ is a paradoxical decomposition of F, where $A_e := A^-$, $A_{a^{-1}} := A^+$, $B_e := B^-$, and $B_{b^{-1}} := B^+$. □

The proof of the previous proposition is similar to the proof of Proposition 9.1.5; more generally, paradoxical groups are not amenable (and vice versa):

Theorem 9.2.12 (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 $((A_g)_{g\in K}, (B_h)_{h\in L})$ 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

$$1 = m(\chi_G) = \sum_{g\in K} m(\chi_{g\cdot A_g}) = \sum_{g\in K} m(g\cdot \chi_{A_g}) = \sum_{g\in K} m(\chi_{A_g}),$$
$$1 = \sum_{h\in L} m(\chi_{B_h}),$$

and hence

$$1 = m(\chi_G) = \sum_{g\in K} m(\chi_{A_g}) + \sum_{h\in L} m(\chi_{B_h}) = 1 + 1 = 2,$$

which is a contradiction. Therefore, G is *not* amenable

Conversely, let G be non-amenable. Then there is "enough space" in G to perform a combinatorial construction – based on a version of Hall's marriage theorem (Theorem 9.4.3) – leading to a paradoxical decomposition [39, Theorem 4.9.1] (Exercise 9.E.18). □

9.2.3 Application: The Banach-Tarski paradox

Actions of paradoxical groups lead to paradoxical decompositions of sets, the most famous example being the Banach-Tarski paradox (Theorem 9.2.17).

Definition 9.2.13 (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 pair $\bigl((A_g)_{g\in K}, (B_h)_{h\in L}\bigr)$ where $K, L \subset G$ are finite and $(A_g)_{g\in K}$, $(B_h)_{h\in L}$ are families of subsets of Y with the property that

$$Y = \Bigl(\bigcup_{g\in K} A_g\Bigr) \cup \Bigl(\bigcup_{h\in L} B_h\Bigr), \qquad Y = \bigcup_{g\in K} g\cdot A_g, \qquad Y = \bigcup_{h\in L} h\cdot B_h$$

are *disjoint* unions.

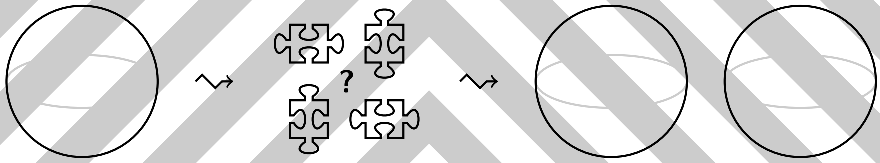

Figure 9.3.: The Hausdorff paradox, schematically

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 the group action into two copies of the set. 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 decompositions of groups into paradoxical decomposition of sets:

Proposition 9.2.14 (Paradoxical groups induce paradoxical decomposition). *Let G be a paradoxical group acting freely on a non-empty set X. Then X is G-paradoxical with respect to this action.*

Proof. Let $((A_g)_{g\in K}, (B_h)_{h\in L})$ be a paradoxical decomposition of G. Using the axiom of choice, we find a subset $R \subset X$ that contains exactly one point of every G-orbit. Then a straightforward calculation shows that

$$\big((A_g \cdot R)_{g\in K}, (B_h \cdot R)_{h\in L}\big)$$

is a G-paradoxical decomposition of X. □

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 9.2.15), which leads to the Hausdorff paradox (Figure 9.3).

Proposition 9.2.15. *The subgroup of the special orthogonal group* $\mathrm{SO}(3)$ *generated by*

$$\begin{pmatrix} \frac{3}{5} & \frac{4}{5} & 0 \\ -\frac{4}{5} & \frac{3}{5} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \textit{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{3}{5} & -\frac{4}{5} \\ 0 & \frac{4}{5} & \frac{3}{5} \end{pmatrix}$$

is free of rank 2.

Proof. This can be shown by a careful application of the ping-pong lemma (Exercise 4.E.17). □

Corollary 9.2.16 (Hausdorff paradox). *There is a countable set $D \subset S^2$ such that $S^2 \setminus D$ is* $\mathrm{SO}(3)$*-paradoxical with respect to the canonical* $\mathrm{SO}(3)$*-action on S^2.*

Proof. By Proposition 9.2.15, 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,

$$D := \{g \cdot x_{g,j} \mid g \in G \setminus \{e\}, j \in \{1,2\}\} \subset S^2$$

is a countable set and the SO(3)-action on S^2 restricts to a free G-action on $S^2 \setminus D$. Because the group G is paradoxical (Proposition 9.2.11) the complement $S^2 \setminus D$ is G-paradoxical (Proposition 9.2.14). 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 [152, Chapter 0.1]:

Theorem 9.2.17 (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 $g_1, \dots, g_n, h_1, \dots, h_m \in$* 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. The Hausdorff paradox and the Banach-Tarski paradox rely on (some basic version of) the axiom of choice (for the set S^2) [82, p. 134].

Further generalisations include the Banach-Tarski paradoxes for balls in $\mathbb{R}^3$ and for bounded subsets in $\mathbb{R}^3$ [152, 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 discuss characterisations in terms of uniformly finite homology and bounded cohomology, respectively.

Outlook 9.2.18 (Amenability and uniformly finite homology). We begin with the characterisation via uniformly finite homology [19, 175]. Uniformly finite homology is introduced in Chapter 5.E. In short, for every normed ring R with unit (e.g., $\mathbb{Z}$ or $\mathbb{R}$) uniformly finite homology provides a sequence of quasi-isometry invariant functors $H_n^{\mathrm{uf}}(\,\cdot\,; R)$ from the category of UDBG spaces to the category of R-modules. There are three popular ways to describe uniformly finite homology:

- via explicit geometric chains (Exercise 5.E.31, 5.E.32, Definition 5.E.1),
- via coarsening of locally finite homology (Exercise 5.E.35),
- via group homology (Exercise 5.E.36, in the case of finitely generated groups).

Using the description by explicit geometric chains for uniformly homology, we can in particular define the fundamental class of UDBG spaces:

Definition 9.2.19 (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 $C_0^{\text{uf}}(X;R)$. The corresponding class

$$[X]_R := \Bigl[\sum_{x \in X} 1 \cdot x\Bigr] \in H_0^{\text{uf}}(X;R)$$

is the *fundamental class of X in $H_0^{\text{uf}}(\,\cdot\,,R)$*.

Theorem 9.2.20 (Amenability and the fundamental class [19, 175])**.** *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 $H_0^{\text{uf}}(X;\mathbb{Z})$.*
3. *We have $[X]_{\mathbb{R}} = 0$ in $H_0^{\text{uf}}(X;\mathbb{R})$.*

The proof of the implications "$1 \Longrightarrow 2$" and "$1 \Longrightarrow 3$" admits a nice interpretation in terms of Ponzi schemes [19] (Exercise 9.E.29). Conversely, one can prove the implications "$2 \Longrightarrow 1$" and "$3 \Longrightarrow 1$" via Følner sequences and suitable averaging maps (Exercise 9.E.27).

Outlook 9.2.21 (Amenability and bounded cohomology)**.** There is a complementary characterisation of amenability in terms of bounded cohomology and ℓ^1-homology: these theories are functional analytic versions of ordinary group (co)homology, obtained by taking the dual and the ℓ^1-completion respectively of the simplicial chain complex of the group [84, 123, 99] (Appendix A.2). Bounded cohomology and ℓ^1-homology have a wide range of applications in geometric and measurable group theory [124] as well as in geometric topology [73, 100].

Theorem 9.2.22 (Amenability and bounded cohomology/ℓ^1-homology [86, 131, 99])**.** *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 $k \in \mathbb{N}_{>0}$ we have $H_b^k(G;V') \cong 0$.*
3. *For all Banach G-modules V and all $k \in \mathbb{N}_{>0}$ we have $H_k^{\ell^1}(G;V) \cong 0$.*

The easiest part of this theorem is to prove that bounded cohomology with trivial $\mathbb{R}$-coefficients of amenable groups is trivial – by using an invariant mean to define a transfer map to bounded cohomology of the trivial group (Exercise 9.E.30).

Outlook 9.2.23 (Amenability and L^2-invariants)**.** Also L^2-(co)homology and L^2-invariants exhibit an interesting behaviour in the presence of amenability [104].

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 quasi-isometric UDBG spaces. Then X is amenable if and only if Y is amenable.*

Proof. We only need to prove inheritance of Følner sequences under quasi-isometries. 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 (Exercise 5.E.7):

- The map $f\colon X \longrightarrow Y$ is a (c,c)-quasi-isometric embedding with c-dense image.
- For all finite sets $F \subset Y$ we have

$$\bigl|f^{-1}(B^Y_c(F))\bigr| \geq \frac{1}{C} \cdot |F| \qquad \text{and} \qquad \bigl|f^{-1}(F)\bigr| \leq C \cdot |F|.$$

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, as we will show now: Let $r \in \mathbb{N}$. 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

$$\bigl|\partial^X_r(\widetilde{F}_n)\bigr| \leq \bigl|f^{-1}(f(\partial^X_r(\widetilde{F}_n)))\bigr| \leq \bigl|f^{-1}(\partial^Y_{c\cdot(r+2)}(F_n))\bigr| \leq C \cdot \bigl|\partial^Y_{c\cdot(r+2)}(F_n)\bigr|$$

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. □

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 9.2.6) and the quasi-isometry invariance of amenability for UDBG spaces (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 9.2.20).

9.4 Quasi-isometry vs. 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 infinite groups that are quasi-isometric but *not* bilipschitz equivalent [55, 56]. 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 the following, we will give an elementary proof of this bilipschitz equivalence rigidity result for non-amenable UDBG spaces [77, 39]; a refined approach is sketched briefly in Outlook 9.4.10.

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:

- Bijective quasi-isometries between UDBG spaces are bilipschitz equivalences; thus, it suffices to deform f into a bijective quasi-isometry.
- Because Y (and hence also X) are non-amenable, Hall's marriage theorem guarantees that there is enough space to deform f and a quasi-isometry inverse $Y \longrightarrow X$ of f into injective quasi-isometries.
- These injective maps can be glued via the Schröder-Bernstein theorem to form a bijective quasi-isometry.

We will 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, similar to the case of finitely generated groups (Exercise 5.E.5). ☐

A standard tool in combinatorics for finding injections satisfying additional constraints are marriage theorems.

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. Then there exists a (W, M, F)-marriage*, i.e., an injective map $\mu\colon W \longrightarrow M$ with*

$$\forall_{w \in W} \quad \mu(w) \in F(w).$$

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 beautiful proofs by induction [49, Theorem 2.1.2] (Exercise 3.E.8); the general marriage theorem can then be obtained by applying Zorn's lemma to suitably extendable partial marriages [39, Theorem H.3.2] (Exercise 3.E.8).

The last ingredient is the Schröder-Bernstein theorem [164, Theorem 9.2.1] 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, let $g\colon Y \longrightarrow X$ be a quasi-isometry quasi-inverse to f.

We first 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

$$\forall_{V \in P^{\text{fin}}(X)} \quad |V| \leq E \cdot \bigl|f(V)\bigr|.$$

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 9.E.13). Therefore, the map

$$X \longrightarrow P^{\text{fin}}(Y)$$
$$x \longmapsto B_r^Y\bigl(f(x)\bigr)$$

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, the same arguments as above yield 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 (Exercise 5.E.3).

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 bijective quasi-isometry. Proposition 9.4.2 then implies that $\overline{f} \colon X \longrightarrow Y$ is a bilipschitz equivalence. ☐

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 the map f is at finite distance of a bilipschitz equivalence $G \longrightarrow H$.*

Proof. Finitely generated groups are UDBG spaces with respect to word metrics of finite generating sets. Hence, we can apply Theorem 9.4.1. ☐

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.

A direct application is 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 $G \sim_{\text{QI}} G'$ and let G be non-amenable. Then*

$$G * H \sim_{\text{QI}} G' * H.$$

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 (the corresponding statement for quasi-isometries does *not* hold in general, because one loses control over the additive error term when taking free products (Exercise 9.E.20)). In particular, we obtain $G * H \sim_{\mathrm{QI}} G' * H$. □

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 5.4.9).

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 the free product $1 * \mathbb{Z}/2 \cong \mathbb{Z}/2$ 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 basically the only cases where quasi-isometry is not inherited through free products [143].

Outlook 9.4.10 (Bilipschitz equivalence rigidity via uniformly finite homology)**.** The bilipschitz rigidity theorem (Theorem 9.4.1) can be refined as follows [175]:

Theorem 9.4.11 (Bilipschitz equivalence rigidity for UDBG spaces, via uniformly finite homology)**.** *Let X and Y be UDBG spaces and let $f\colon X \longrightarrow Y$ be a quasi-isometry. Then the following are equivalent:*

1. *The map f is at finite distance of a bilipschitz equivalence.*
2. *The map f is compatible with the fundamental classes:*

$$H_0^{\mathrm{uf}}(f;\mathbb{Z})[X]_{\mathbb{Z}} = [Y]_{\mathbb{Z}}.$$

From this one can easily recover the result of Theorem 9.4.1: Let X and Y be non-amenable UDBG spaces. Then every quasi-isometry $X \longrightarrow Y$ is at finite distance of a bilipschitz equivalence; this can be seen as follows:

As X and Y are non-amenable, we have $[X]_{\mathbb{Z}} = 0 = [Y]_{\mathbb{Z}}$ (Theorem 9.2.20). In particular, every quasi-isometry $f\colon X \longrightarrow Y$ maps $[X]_{\mathbb{Z}}$ to $[Y]_{\mathbb{Z}}$. By Theorem 9.4.11, f is hence at finite distance of a bilipschitz equivalence.

The refined formulation of Theorem 9.4.11 allows in concrete examples to separate quasi-isometry from bilipschitz equivalence of finitely generated (amenable) groups [55, 56].

9.E Exercises

Basic properties of amenable groups

Quick check 9.E.1 (Amenable groups?!*).
1. Is the infinite dihedral group D_∞ amenable?
2. Is the modular group $\mathrm{SL}(2,\mathbb{Z})$ amenable?
3. Is the group $\mathrm{Hom}\bigl(\mathrm{SL}(2,\mathbb{Z}),\mathbb{Z}\bigr)$ amenable?
4. Is the lamplighter group over $\mathbb{Z}$ amenable?

Quick check 9.E.2 (Amenable group presentations?!*).
1. Is the group $\langle x,y,z \,|\, xy\rangle$ amenable?
2. Is the group $\langle x,y,z \,|\, xy, yz\rangle$ amenable?
3. Is the group $\langle x,y,z \,|\, xyz\rangle$ amenable?
4. Is the group $\langle x,y,z \,|\, x^2y^2, y^2, z^2\rangle$ amenable?

Quick check 9.E.3 (Means?*).
1. Let G be a group and let $g \in G$. Is the map
$$\begin{aligned} \ell^\infty(G,\mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto f(g) \end{aligned}$$
a G-invariant mean?
2. Is the map
$$\begin{aligned} \ell^\infty(\mathbb{Z},\mathbb{R}) &\longrightarrow \mathbb{R} \\ f &\longmapsto \sup_{x\in\mathbb{Z}} f(x) \end{aligned}$$
a $\mathbb{Z}$-invariant mean?

Exercise 9.E.4 (Means of group constructions*). Carry out the calculations omitted in the proof of Proposition 9.1.6.

Exercise 9.E.5 (Big symmetric groups*).
1. Let X be an infinite set. Show that then the symmetric group S_X is *not* amenable.
2. Show that the group of permutations of $\mathbb{N}$ with finite support is amenable.

Exercise 9.E.6 (Amenable free products?**). Characterise by necessary and sufficient conditions for which finitely generated groups G and H the free product $G * H$ is amenable.

Hints. For which $m,n \in \mathbb{Z}$ does $\mathbb{Z}/m * \mathbb{Z}/n$ contain a free subgroup of rank 2 ?

Exercise 9.E.7 (Characterisation of amenability for linear groups*). Let K be a field, let $n \in \mathbb{N}$ and let G be a finitely generated subgroup of $\mathrm{GL}(n, K)$. Prove that G is amenable if and only if G has no free subgroup of rank 2.
Hints. Use the Tits alternative (Theorem 4.4.7).

Exercise 9.E.8 (Non-amenability of large groups*). A group is *large* if it has a finite index subgroup that admits a surjective homomorphism onto the free group of rank 2. Prove that large groups are *not* amenable.

Exercise 9.E.9 (Amenable radical**).

1. Prove that every group contains a maximal (with respect to inclusion) normal amenable subgroup, the *amenable radical*.
2. What is the amenable radical of F_{2017} ?

Følner sequences

Quick check 9.E.10 (Simplified Følner conditions?!*).

1. Let X be a UDBG space and let $(F_n)_{n\in\mathbb{N}}$ be a sequence of finite non-empty subsets of X with
$$\lim_{n\to\infty} \frac{|\partial_1^X(F_n)|}{|F_n|} = 0.$$
Is X then amenable?
2. Let G be a finitely generated group, let $S \subset G$ be a finite generating set of G, and let $(F_n)_{n\in\mathbb{N}}$ be a sequence of finite non-empty subsets of G with
$$\lim_{n\to\infty} \frac{|\partial_1^{G,d_S}(F_n)|}{|F_n|} = 0.$$
Is G then amenable?

Quick check 9.E.11 (Amenable spaces?!*).

1. We consider the set $\{(x, y) \in \mathbb{Z}^2 \mid x \cdot y = 0\} \subset \mathbb{R}^2$ with the ℓ^1-metric. Is this UDBG space amenable?
2. Do there exist amenable UDBG spaces with infinitely many ends?
Hints. How about a big comb?

Quick check 9.E.12 (More amenable groups?*).

1. Is the first Grigorchuk group amenable?
2. Are all subgroups of automorphism groups of trees amenable?

Exercise 9.E.13 (Thickened subsets in non-amenable spaces**). Let X be a non-amenable UDBG space and let $E \in \mathbb{R}_{\geq 1}$. Show that there exists an $r \in \mathbb{N}$ with
$$\forall_{F\in P^{\mathrm{fin}}(X)} \ \bigl|B_r^X(F)\bigr| \geq E \cdot |F|.$$

Exercise 9.E.14 (Exhausting Følner sequences**). Let X be an amenable UDBG space and let $(F_n)_{n\in\mathbb{N}}$ be a Følner sequence of X.

1. If $|X| = \infty$, show that $\lim_{n\to\infty} |F_n| = \infty$.
2. Let $x_0 \in X$. Show that $(\{x_0\} \cup F_n)_{n\in\mathbb{N}}$ also is a Følner sequence.
3. Show that X has a Følner sequence $(\widetilde{F}_n)_{n\in\mathbb{N}}$ with the following properties:
   - For all $n \in \mathbb{N}$ we have $\widetilde{F}_n \subset \widetilde{F}_{n+1}$,
   - and $\bigcup_{n\in\mathbb{N}} \widetilde{F}_n = X$.

   Hints. Construct the new Følner sequence inductively by combining $(\{x_0\} \cup F_n)_{n\in\mathbb{N}}$ with larger and larger balls. Be careful to choose radii and indices in the right order!

Exercise 9.E.15 (Følner sequences in trees?**). Let $r \in \mathbb{N}_{\geq 3}$, let $T = (V, E)$ be a regular tree of degree r, and consider the metric d on V induced by T. Show that the UDBG space (V, d) is *not* amenable by analysing the size of boundaries of finite subsets.

Exercise 9.E.16 (Kazhdan's property (T) and amenability**). Let G be a group (for simplicity, we only consider discrete groups). A *unitary representation of G* on a Hilbert space H is an action of G on H by unitary operators.

- Let $Q \subset G$ and $\varepsilon \in \mathbb{R}_{>0}$. A (Q, ε)*-invariant vector* of a unitary representation of G on H is a vector $x \in H$ with

$$\forall_{g\in Q} \;\; \|g \cdot x - x\| < \varepsilon \cdot \|x\|.$$

- A unitary representation of G *has almost invariant vectors* if for every finite set $Q \subset G$ and every $\varepsilon \in \mathbb{R}_{>0}$ there is a (Q, ε)-invariant vector.
- A subset $Q \subset G$ is a *Kazhdan set* if there exists an $\varepsilon \in \mathbb{R}_{>0}$ with: Every unitary representation of G with a (Q, ε)-invariant vector has a non-zero invariant vector.
- The group G has *property (T)* if G contains a finite Kazhdan set.

The goal of this exercise is to compare property (T) with amenability:

1. Show that the left translation action of a finitely generated amenable group G on $\ell^2(G, \mathbb{C})$ has almost invariant vectors.
2. Show that every finitely generated amenable group with property (T) is finite.
3. Show that all finite groups have property (T) (and are amenable).
4. Look up in the literature how the interplay of amenability and property (T) is used in the normal subgroup theorem by Margulis.

Bilipschitz equivalence rigidity

Quick check 9.E.17 (Bilipschitz equivalence of groups*).

1. Are the groups F_{2017} and $\mathbb{Z} \times F_{2017}$ bilipschitz equivalent?
2. Are the groups F_{2017} and $F_{2017} \times \mathbb{Z}/2017$ bilipschitz equivalent?

Exercise 9.E.18 (Non-amenable groups are paradoxical**). Let G be a finitely generated non-amenable group (via Følner sequences).

1. Show that there exists a surjection $f: G \longrightarrow G$ that is at finite distance from the identity with $|f^{-1}(g)| = 2$ for all $g \in G$; in particular, there is a finite set $K \subset G$ with the following property: For all $g \in G$ we have
$$g \cdot f(g)^{-1} \in K.$$
Hints. Use the fact that the projection $G \times \mathbb{Z}/2 \longrightarrow G$ is a quasi-isometry and apply Theorem 9.4.1.
2. Conclude that G is paradoxical.
Hints. Let f and K be as in the previous step. Use the axiom of choice to pick two complementing sections $f_1, f_2: G \longrightarrow G$ of f and consider the sets $(f_1^{-1}(k))_{k \in K}$ as well as $(f_2^{-1}(k))_{k \in K}$.

How can one now generalise this result to all (not necessarily finitely generated) groups?

Exercise 9.E.19 (Quasi-Isometry vs. bilipschitz equivalence for $\mathbb{Z}^n$ **). Let G be a finitely generated group that is quasi-isometric to $\mathbb{Z}^n$ for some $n \in \mathbb{N}$.

1. Show that there exists an $m \in \mathbb{N}$ such that there is an injective quasi-isometric embedding $G \longrightarrow \mathbb{Z}^n \times \mathbb{Z}/m$.
2. Show that there is an injective quasi-isometric embedding $G \longrightarrow \mathbb{Z}^n$.
3. Show that there is an injective quasi-isometric embedding $\mathbb{Z}^n \longrightarrow G$ by passing to a suitable finite index subgroup of $\mathbb{Z}^n$.
4. Why do the previous steps not immediately imply via a Schröder-Bernstein argument that G and $\mathbb{Z}^n$ are bilipschitz equivalent?
Hints. There is a serious issue! [37]
5. Prove that G and $\mathbb{Z}^n$ are bilipschitz equivalent, using quasi-isometry rigidity of $\mathbb{Z}^n$.

Exercise 9.E.20 (Free products of quasi-isometries*). Let G, G', H be finitely generated groups and let $f: G \longrightarrow G'$ be a map.

1. Show in a concrete example that $f * \mathrm{id}_H : G * H \longrightarrow G' * H$ is *not* necessarily a quasi-isometric embedding if f is a quasi-isometric embedding.
2. Show in a concrete example that $f * \mathrm{id}_H : G * H \longrightarrow G' * H$ does *not* necessarily have quasi-dense image if $f: G \longrightarrow H$ has quasi-dense image.

Amenable actions$^+$

Following the general principle of extending geometric notions from groups to group actions, one can define and investigate amenability for group actions [68]:

Definition 9.E.1 (Amenable action). Let G be a group and let X be a set. An action of G on X is *amenable*, if there exists a G-invariant mean on $\ell^\infty(X, \mathbb{R})$.

Caveat 9.E.2. There are also other notions of amenable actions [180, 5], in the context of ergodic theory. These are different from the one above!

Quick check 9.E.21 (Amenable action basics*).
1. For which groups G is the left translation action on G amenable?
2. Which groups admit free amenable actions?
3. Are all actions of amenable groups amenable?
4. For which groups is the trivial action on the one-point space amenable?

Definition 9.E.3 (Faithful action). An action of a group G on a set X is *faithful* if the associated homomorphism $G \longrightarrow S_X$ is injective.

Quick check 9.E.22 (Faithful action basics*).
1. Is every free action faithful?
2. Is every faithful action free?
3. How can one define faithful actions in general categories?

Quick check 9.E.23 (Faithful amenable actions*). Show that every group admits a faithful amenable action.

Exercise 9.E.24 (Amenable actions via Følner sets** [149]).
1. Give a definition for Følner sequences for group actions (of countable groups on countable sets).
2. State and prove a characterisation of amenable actions (of countable groups on countable sets) in terms of Følner sequences.

Exercise 9.E.25 (Amenable actions without finite orbits** [68, Lemma 2.16]). Let G be a countable group that is *not* finitely generated. Show that there exists an amenable action of G that has no finite orbits.
Hints. Use the characterisation of amenable actions in terms of Følner sets (Exercise 9.E.24) and let G act on the disjoint union of all coset spaces of G by finitely generated subgroups.

Exercise 9.E.26 (Amenable actions of non-amenable groups***). Prove that the free group $F(\{a,b\})$ of rank 2 admits a faithful, transitive amenable action, following Glasner and Monod [68]: Let $X := \mathbb{Z} \times \{0,1\}$ and

$$b^n \cdot (z,i) := (z+n,i)$$

for all $n \in \mathbb{Z}$, $(z,i) \in \mathbb{Z} \times \{0,1\}$. For an injection $J\colon \mathbb{Z} \longrightarrow \mathbb{Z} \times \{0,1\}$ the action by powers of a on X is defined by

$$a^n \cdot (z,i) := \begin{cases} (z,i) & \text{if } (z,i) \notin \operatorname{im} J \\ (J(m+n),i) & \text{if } (z,i) = J(m) \text{ with } m \in \mathbb{Z} \end{cases}$$

for all $n \in \mathbb{Z}$, $(z,i) \in \mathbb{Z} \times \{0,1\}$. We only consider injections J that satisfy $J(0) := (0,0)$ and $J(n) := (n,1)$ for all $n \in \mathbb{Z}_{<0}$ as well as $J(\mathbb{N}) \subset \mathbb{Z} \times \{0\}$.
1. Show that every such action is amenable.

2. Show that every such action is transitive.
3. Show that there exists such an injection J that induces a faithful action.
Hints. Construct J on $\mathbb{N}_{>0}$ step by step, by induction over (some enumeration of) $F(\{a,b\})$.

However, not every finitely generated group admits a faithful, transitive amenable action [68].

(Co)Homology and amenability$^+$

Exercise 9.E.27 (Fundamental class in uniformly finite homology**)**.** Let X be an amenable UDBG space. Show that then the fundamental classes

$$[X]_{\mathbb{Z}} := \Bigl[\sum_{x\in X} 1\cdot x\Bigr] \in H_0^{\mathrm{uf}}(X;\mathbb{Z}),$$
$$[X]_{\mathbb{R}} := \Bigl[\sum_{x\in X} 1\cdot x\Bigr] \in H_0^{\mathrm{uf}}(X;\mathbb{R}).$$

are non-trivial.
Hints. Use a Følner sequence of the space X to define appropriate averaging maps $H_0^{\mathrm{uf}}(X;\mathbb{Z}) \longrightarrow \mathbb{R}$ and $H_0^{\mathrm{uf}}(X;\mathbb{R}) \longrightarrow \mathbb{R}$.

Exercise 9.E.28 (The doubling trick in uniformly finite homology**)**.** Let X be a UDBG space. For a subset $A \subset X$ we write

$$[A]_{\mathbb{R}} := \Bigl[\sum_{x\in A} 1\cdot x\Bigr] \in H_0^{\mathrm{uf}}(X;\mathbb{R})$$

for the corresponding class in uniformly finite homology (with $\mathbb{R}$-coefficients).
1. Let $A, B \subset X$ be disjoint sets. Show that

$$[A\cup B]_{\mathbb{R}} = [A]_{\mathbb{R}} + [B]_{\mathbb{R}} \in H_0^{\mathrm{uf}}(X;\mathbb{R}).$$

2. Let X be non-amenable. Conclude that $[X]_{\mathbb{R}} = 0 \in H_0^{\mathrm{uf}}(X;\mathbb{R})$ (using Theorem 9.4.1).
Hints. Use that the projection $X \times \{0,1\} \longrightarrow X$ is a quasi-isometry (provided that $X \times \{0,1\}$ is equipped with a suitable metric).

Exercise 9.E.29 (Ponzi schemes in uniformly finite homology***)**.** Let X be a non-amenable UDBG space.
1. For each $x \in X$ construct a chain (“tail”) $t_x \in C_1^{\mathrm{uf}}(X;\mathbb{R})$ with $\partial_1 t_x = x$.
2. Interpret this situation as a flow of money; hence, $x \in X$ will gain money, through the flow of money specified by the tail t_x. All the other vertices that are involved in t_x will have a balanced account!

3. Prove that $[X]_\mathbb{R} = 0$ in $H_0^{\mathrm{uf}}(X;\mathbb{R})$ without using Exercise 9.E.28 or Theorem 9.4.1.
   Hints. Use non-amenability to show that there is enough space in X to construct tails as in the first part with the additional property that also "$\sum_{x\in X} t_x$" gives a well-defined chain in $C_1^{\mathrm{uf}}(X;\mathbb{R})$.
4. Interpret the previous step as a Ponzi scheme!

Exercise 9.E.30 (Bounded cohomology of amenable groups**)**.** Let G be a group. Then the *bounded cochain complex of* G is the cochain complex $C_b^*(G;\mathbb{R})$ obtained from the complex $C_*(G;\mathbb{R})$ (Appendix A.2) by taking the topological dual with respect to the ℓ^1-norm on $C_*(G;\mathbb{R})$ given by the obvious basis.

1. Show that the boundary operator on $C_*(G;\mathbb{R})$ indeed induces a well-defined coboundary operator on the cochain complex $C_b^*(G;\mathbb{R})$.

The cohomology of $C_b^*(G;\mathbb{R})$ is called *bounded cohomology of* G (with trivial $\mathbb{R}$-coefficients), which is denoted by $H_b^*(G;\mathbb{R})$. One can easily extend this construction to a functor with respect to group homomorphisms.

2. Show that $H_b^k(\{e\};\mathbb{R}) \cong 0$ for all $k \in \mathbb{N}_{>0}$.
3. Let G be an amenable group and let $k \in \mathbb{N}_{>0}$. Prove that $H_b^0(G;\mathbb{R}) \cong 0$.
   Hints. Imitate the proof of the corresponding result for ordinary cohomology of finite groups: Use an invariant mean on $\ell^\infty(G,\mathbb{R})$ to define a suitable transfer homomorphism $H_b^k(\{e\};\mathbb{R}) \longrightarrow H_b^k(G;\mathbb{R})$.

Exercise 9.E.31 (Non-amenability and G-theory$^\infty$*)**.** Let R be a (potentially non-commutative) ring with unit. Then we define the Abelian group $G_0(R)$ as follows: Let F_R be the set(!) of isomorphism classes of finitely generated R-modules; in contrast to K-theory, only finite generation but no projectivity is assumed. Then $G_0(R)$ is defined as the quotient of the free Abelian group $\bigoplus_{[M]\in F_R} \mathbb{Z}\cdot[M]$ by the submodule generated by the relations

$$[C] = [A] + [B]$$

for every short exact sequence

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0$$

of finitely generated R-modules.

If G is a group, then we denote the complex group ring of G by $\mathbb{C}G$ (Appendix A.2).

1. Show that the element $[\mathbb{C}F_2]$ is trivial in $G_0(\mathbb{C}F_2)$.
2. Let G be a group that contains a free subgroup of rank 2. Show that then the element $[\mathbb{C}G]$ is trivial in $G_0(\mathbb{C}G)$.
3. Let G be a non-amenable group. Show that then the element $[\mathbb{C}G]$ is trivial in $G_0(\mathbb{C}G)$.
   Hints. This is an open problem! [105, Section 7] It is known that the converse is true: If G is amenable, then the element $[\mathbb{C}G]$ is *not* trivial in $G_0(\mathbb{C}G)$ [105, Remark 7.11].

Part IV

Reference material

A

Appendix

Overview

A.1 The fundamental group

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 homotopy invariant functor

$$\begin{aligned} \pi_1 \colon \text{path-connected pointed topological spaces} &\longrightarrow \text{groups} \\ \text{basepoint preserving continuous maps} &\longrightarrow \text{group homomorphisms.} \end{aligned}$$

A.1.1 Construction and examples

Geometrically, the fundamental group measures how many “holes” a space has that can be detected by loops in the space in question. The group structure on such loops is given by concatenation of loops. This idea can be turned into a precise definition – with two small modifications:

- In order to be able to concatenate loops, we need to fix a basepoint in the space. A *pointed space* is a pair (X, x_0) consisting of a topological space X and a point $x_0 \in X$, the *basepoint*.
- In order to obtain an associative composition by concatenation and inverses, we have to identify loops that can be continuously deformed into each other through pointed homotopies.
  Basepoint preserving maps $f, g \colon (Y, y_0) \longrightarrow (Z, z_0)$ are *homotopic* (in the pointed sense), if f can be continuously deformed into g while fixing the basepoint, i.e., if there is a continuous map $h \colon Y \times [0,1] \longrightarrow Z$ satisfying
  $$h(\,\cdot\,, 0) = f \qquad \text{and} \qquad h(\,\cdot\,, 1) = g$$
  and
  $$\forall_{t \in [0,1]} \;\; h(y_0, t) = z_0.$$
  In this case, we write $f \simeq_* g$. Such basepoint preserving maps are *homotopy equivalences* if they admit an inverse up to pointed homotopy. Pointed spaces are *homotopy equivalent* if there exists a homotopy equivalence between them. A space is *contractible* if it is homotopy equivalent to a single point.

Definition A.1.1 (Fundamental group). The *fundamental group*

$$\pi_1(X, x_0) := \operatorname{map}_*\bigl((S^1, 1), (X, x_0)\bigr)/ \simeq_*$$

of a pointed space (X, x_0) is defined as the set of all loops in X based at x_0, modulo basepoint preserving homotopies; the composition of two

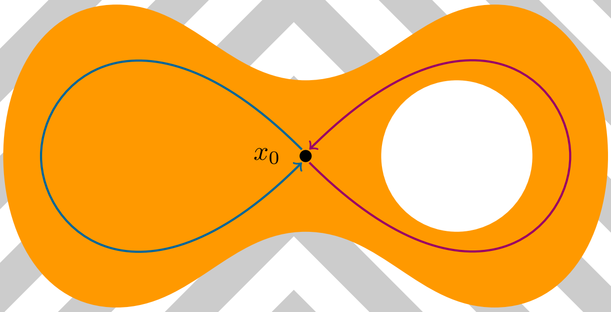

Figure A.1.: The definition of the fundamental group, schematically; the right hand loop will discover the hole, while the left hand loop represents the trivial element in π_1 (it is homotopic to the constant loop).

classes of loops is given by taking the class represented by concatenation (and reparametrisation) of these loops.

If $f\colon (X, x_0) \longrightarrow (Y, y_0)$ is a continuous map, then the induced homomorphism $\pi_1(f)\colon \pi_1(X, x_0) \longrightarrow \pi_1(Y, y_0)$ is defined by composing f with loops in X.

This definition is illustrated in Figure A.1. By definition, the fundamental group functor is homotopy invariant: If $f,\ g\colon (X, x_0) \longrightarrow (X', x_0')$ are basepoint preserving continuous maps that are homotopic (via a basepoint preserving homotopy), then the induced group homomorphisms

$$\pi_1(f),\ \pi_1(g)\colon \pi_1(X, x_0) \longrightarrow \pi_1(X', x_0')$$

coincide.

Homotopy invariance and functoriality of π_1 show that if (X, x_0) and (X', x_0') are homotopy equivalent pointed topological spaces, then $\pi_1(X, x_0)$ and $\pi_1(X', x_0')$ are isomorphic groups. In particular, π_1 can detect that certain spaces are *not* homotopy equivalent: If $\pi_1(X, x_0) \not\cong \pi_1(X', x_0')$, then (X, x_0) and (X, x_0) cannot be homotopy equivalent (and so also cannot be homeomorphic).

Caveat A.1.2. In general, if spaces have isomorphic fundamental groups, then they need not be homotopy equivalent.

It turns out that the fundamental group of a path-connected space does *not* depend (up to non-canonical isomorphism) on the chosen basepoint. Therefore, one often does not mention the basepoints in the notation explicitly and writes $\pi_1(X)$ instead of $\pi_1(X, x_0)$.

The prototypical calculation of a fundamental group is $\pi_1(S^1) \cong \mathbb{Z}$; the fundamental group measures how often a loop wraps around the "hole" in S^1. In contrast, all loops in the sphere S^2 can be deformed into the constant loop, whence $\pi_1(S^2)$ is the trivial group. Some basic examples of fundamental groups are listed in Figure A.2. Moreover, the following dictionary between topological and group theoretical constructions is helpful in the calculation of fundamental groups:

- **Products.** The fundamental group functor is compatible with products: The projection maps induce an isomorphism
$$\pi_1\bigl(X \times Y, (x_0, y_0)\bigr) \cong \pi_1(X, x_0) \times \pi_1(Y, y_0).$$
(This also holds for products over arbitrary index sets.)
- **Glueings.** The fundamental group functor is compatible with (many) pushouts: By the Seifert and van Kampen theorem, there is a canonical isomorphism between $\pi_1\bigl((X, a_0) \cup_{(A,a_0)} (Y, a_0)\bigr)$ and the pushout of $\pi_1(X, x_0)$ and $\pi_1(Y, y_0)$ over $\pi_1(A, a_0)$, provided that $X \cap Y = A$ and both X and Y are path-connected open subsets of $X \cup_A Y$.
- **Self-glueings.** The fundamental group of mapping tori along π_1-injective maps leads to ascending HNN-extensions. More general π_1-injective self-glueings result in more general HNN-extensions on the level of fundamental groups.
- **Fibrations.** Fundamental groups of fibration sequences of path-connected topological spaces correspond roughly to extensions of groups.

Details on the fundamental group and proofs of these claims can be found in most books on algebraic topology [115, 81, 48].

The fundamental group can also be used to introduce a higher version of path-connectedness: A topological space is path-connected, if every pair of points can be joined by a continuous path. A topological space is simply connected if it is path-connected and if every pair of points can be joined by a continuous path in an essentially unique way:

Definition A.1.3 (Simply connected)**.** A topological space is *simply connected* if it is path-connected and if its fundamental group is trivial.

More generally, the paradigm of algebraic topology is to find good homotopy invariants of topological spaces with the goal of classifying (large classes of) topological spaces up to homotopy equivalence; the fundamental group is just one example of this type. Further examples include, for instance, singular homology and cohomology, the Euler characteristic, higher homotopy groups, ...

A.1.2 Covering theory

In geometric group theory, the most important aspect of the fundamental group is that it is a source of convenient group actions and that it serves as

Topology	**Group theory** (via π_1)
$\bullet$, $\mathbb{R}$, $\mathbb{R}^2$, $\mathbb{H}^2$, ...	trivial group
$S^1 =$ 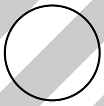	$\mathbb{Z}$
$S^1 \longrightarrow S^1,\ z \mapsto z^2$	$\mathbb{Z} \longrightarrow \mathbb{Z}, z \mapsto 2 \cdot z$
$S^1 \vee S^1 =$ 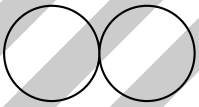	$\mathbb{Z} * \mathbb{Z}$
torus: $S^1 \times S^1 =$ 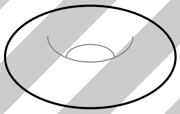	$\mathbb{Z} \times \mathbb{Z}$
sphere: $S^2 =$ 	trivial group
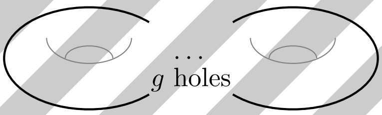	$\Bigl\langle a_1, \dots, a_g, b_1, \dots, b_g \Bigm\vert \prod_{j=1}^{g} [a_j, b_j] \Bigr\rangle$
projective plane: $\mathbb{R}\mathrm{P}^2$	$\mathbb{Z}/2$

Figure A.2.: Basic dictionary for the fundamental group

a sort of Galois group in covering theory. Covering maps are maps such that the preimages over small neighbourhoods in the target space look like several sheets of the given neighbourhood; in technical terms:

Definition A.1.4 (Covering map). A *covering map* is a continuous map between topological spaces that is a locally trivial fibre bundle with discrete fibre.

For every "nice" (i.e., path-connected, locally path-connected, and semi-locally simply connected) path-connected pointed topological space (X, x_0) there is a path-connected pointed space $(\widetilde{X}, \widetilde{x}_0)$, the *universal covering space of* (X, x_0) with the following properties:

- The space $(\widetilde{X}, \widetilde{x}_0)$ is simply connected.

- The space $\widetilde{X}$ admits a free, properly discontinuous group action of the fundamental group $\pi_1(X, x_0)$ 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.
- All other path-connected pointed coverings of (X, x_0) correspond to intermediate quotients of $(\widetilde{X}, \widetilde{x}_0)$ by subgroups of $\pi_1(X, x_0)$.

In particular, the fundamental group $\pi_1(X, x_0)$ coincides with the deck transformation group (i.e., the automorphism group) of the universal covering $\widetilde{X} \longrightarrow X$. A thorough treatment of covering theory is given in Massey's book [115].

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). 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; simple examples of this type are given in the sketch proof of Corollary 6.2.15 and Remark 4.2.11.

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. Therefore, fundamental groups can be viewed as (subgroups of) isometry groups. This is a helpful point of view in Riemannian geometry and geometric group theory, in particular, in the context of the Švarc-Milnor lemma (Chapter 5.4).

Furthermore, for every group there is a path-connected topological space that has the given group as fundamental group (one construction is sketched in Outlook 3.2.5); moreover, this can be arranged in such a way that the universal covering of this space is contractible (equivalently, all higher homotopy groups are trivial). Spaces with this property are called *classifying spaces* or *Eilenberg-MacLane 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 [34] (Appendix A.2, Outlook 3.2.5).

A.2 Group (co)homology

Group (co)homology is an algebraic tool that encodes linear information of groups. Generally speaking, (co)homology theories measure exactness of algebraic constructions; group (co)homology measures the exactness of taking tensor products and homomorphism spaces over the integral group ring. Group (co)homology has a wide range of applications, including, for example, structure theory of groups, algebraic number theory, topology of group actions, and geometry of groups.

We will briefly sketch one of the constructions of group (co)homology and summarise basic properties and applications of group (co)homology. For more details we refer to the excellent book by Brown [34].

A.2.1 Construction

The first basic linearisation of a group is obtained through the group ring:

Definition A.2.1 (Group ring)**.** Let R be a commutative ring with unit. Then the associated *group ring* RG (sometimes also denoted by $R[G]$) is defined as follows:

- The underlying Abelian group is $\bigoplus_G R$. Moreover, one identifies elements of G with the corresponding standard R-basis vectors in $\bigoplus_G R$.
- The multiplicative structure is the R-bilinear extension of the group multiplication

$$\begin{aligned} G \times G &\longrightarrow G \\ (g,h) &\longmapsto g \cdot h. \end{aligned}$$

More explicitly: If $\sum_{g \in G} a_g \cdot g, \sum_{g \in G} b_g \cdot g \in RG$ (where the families $(a_g)_{g \in G}$ and $(b_g)_{g \in G}$ in R have only finitely many non-zero entries), then

$$\Bigl(\sum_{g \in G} a_g \cdot g\Bigr) \cdot \Bigl(\sum_{g \in G} b_g \cdot g\Bigr) = \sum_{g \in G} \Bigl(\sum_{h \in G} a_h \cdot b_{h^{-1} \cdot g}\Bigr) \cdot g.$$

For example, the group ring $\mathbb{C}[\mathbb{Z}]$ is nothing but the ring of Laurent polynomials $\mathbb{C}[T, T^{-1}]$ with complex coefficients.

Caveat A.2.2. If G is a non-Abelian group and the ring R is non-trivial, then the group ring RG is non-commutative! Therefore, it is necessary to pay attention to the difference between left and right modules over RG.

A quick, explicit, way of defining group (co)homology is through the simplicial chain complex associated with the group: We take the full (infinite) simplex whose vertices are the group elements, then we take the corresponding simplicial chain complex (including degenerate simplices). This chain complex has trivial homology (because the underlying space is a simplex, which is contractible). But if we twist the simplicial chain complex by modules over the group ring, we obtain interesting objects:

Definition A.2.3 (Group (co)homology)**.** Let G be a group.

- The *simplicial resolution* $C_*(G)$ *of* G is the following ($\mathbb{N}$-indexed) chain complex of $\mathbb{Z}$-modules: For each $n \in \mathbb{N}$ we set
$$C_n(G) := \bigoplus_{G^{n+1}} \mathbb{Z};$$
thus, elements of $C_n(G)$ can be written as (finite) linear combinations of the form $\sum_{g \in G^{n+1}} a_g \cdot g$.
The boundary operator $\partial_* \colon C_*(G) \longrightarrow C_{*-1}(G)$ is defined by setting $\partial_0 := 0 \colon C_0(G) \longrightarrow 0$ and for all $n \in \mathbb{N}_{>0}$
$$\begin{aligned} \partial_n \colon C_n(G) &\longrightarrow C_{n-1}(G) \\ \sum_{g \in G^n} a_g \cdot g &\longmapsto \sum_{g \in G^n} a_g \cdot \sum_{j=0}^{n} (-1)^j \cdot (g_0, \dots, g_{j-1}, g_{j+1}, \dots, g_n). \end{aligned}$$
- Let A be a right $\mathbb{Z}G$-module. Then we define the *simplicial chain complex of* G *with* A*-coefficients* by
$$C_*(G; A) := A \otimes_{\mathbb{Z}G} C_*(G)$$
(with the boundary operator $\partial^A_* := \mathrm{id}_A \otimes_{\mathbb{Z}G} \partial_*$). Elements in $\ker \partial^A_n$ are called *n-cycles*. For $n \in \mathbb{N}$ the *n-th homology of* G *with* A*-coefficients* is the $\mathbb{Z}$-module
$$H_n(G; A) := \frac{\ker\bigl(\partial^A_n \colon C_n(G; A) \to C_{n-1}(G; A)\bigr)}{\mathrm{im}\bigl(\partial^A_{n+1} \colon C_{n+1}(G; A) \to C_n(G; A)\bigr)}.$$
- Let A be a left $\mathbb{Z}G$-module. Then we define the *simplicial cochain complex of* G *with* A*-coefficients* by
$$C^*(G; A) := \mathrm{Hom}_{\mathbb{Z}G}\bigl(C_*(G), A\bigr)$$
(with the coboundary operator δ^*_A induced by the A-dual of ∂_*). Elements in $\ker \delta^n_A$ are called *n-cocycles*. For $n \in \mathbb{N}$ the *n-th cohomology of* G *with* A*-coefficients* is the $\mathbb{Z}$-module
$$H^n(G; A) := \frac{\ker\bigl(\delta^n_A \colon C^n(G; A) \to C^{n+1}(G; A)\bigr)}{\mathrm{im}\bigl(\delta^{n-1}_A \colon C^{n-1}(G; A) \to C^n(G; A)\bigr)}.$$

Remark A.2.4 (Functoriality of group (co)homology)**.** Group (co)homology is functorial with respect to compatible transformations of groups and coefficient modules [34, Chapter II.6, Chapter III.8].

The definition above is only one of many descriptions of group (co)homology. The most common, generally applicable, descriptions are the following:

- **Simplicial picture.** This is the description that we used as a definition. It is easy to write down, but it is usually not helpful for concrete computations because the chain modules tend to be huge.
- **Algebraic picture.** Alternatively, group (co)homology can be described as the derived functors associated with tensor products and homomorphism modules over group rings [34, Chapter III]. Hence, group (co)homology can flexibly be computed via resolutions over the group ring.
- **Geometric picture.** From a more topological/geometric perspective, group cohomology can be described in terms of classifying spaces of groups (Appendix A.1). More precisely, if G is a group and BG is a classifying space of G (Appendix A.1), then group (co)homology of G is nothing but ordinary (simplicial, singular, cellular) (co)homology of BG with twisted coefficients [34, Chapter II.4, Chapter III.1]. Therefore, "good" models of classifying spaces provide a means to compute group (co)homology geometrically.

A.2.2 Applications

This diversity of descriptions is the reason for the versatility of group (co)homology in applications. Classical applications of group (co)homology include, for example:

- If G is a group, then $H_1(G;\mathbb{Z})$ is nothing but the abelianisation of G [34, Chapter II.3].
- The second group cohomology allows to classify group extensions with Abelian kernel [34, Chapter IV].
- In algebraic number theory, group cohomology is used to study field extensions and their Galois groups; for example, the Hilbert 90 Theorem admits a description and proof in terms of the first cohomology of the Galois group [128, Chapter VI.2].
- Group (co)homology and its interplay with classifying spaces allows to introduce several notions of dimension and finiteness conditions for groups [34, Chapter VIII].
- The study of groups with so-called periodic cohomology is intimately related to the question of which finite groups admit free continuous actions on spheres [34, Chapter VI].

- If G is a finitely generated infinite group and R is a principal ideal domain, then $1 + \operatorname{rk}_R H^1(G; RG)$ equals the number of ends of G [66, Theorem 13.5.5] (ends of groups are defined in Chapter 8.2.3).

For geometric group theory the following observation is essential: If we pick the right type of coefficient modules, then the corresponding group (co)homology will be functorial with respect to interesting types of morphisms between groups. Therefore, group (co)homology leads to algebraic invariants for various geometric notions of equivalences of groups. Basic examples of this technique are:

- If we equip the ring $\mathbb{Z}$ with the trivial group action, then we obtain group (co)homology with coefficients in the trivial module $\mathbb{Z}$. This theory is functorial with respect to all group homomorphisms. Hence, $H_*(\,\cdot\,;\mathbb{Z})$ and $H^*(\,\cdot\,;\mathbb{Z})$ are invariant under group isomorphisms.
- Taking $\ell^\infty(\,\cdot\,;\mathbb{Z})$ or $\ell^\infty(\,\cdot\,;\mathbb{R})$ as coefficient modules leads to group homology that is, for finitely generated groups, functorial with respect to quasi-isometric embeddings (!). In this way, group (co)homology provides us with algebraic quasi-isometry invariants. In fact, these functors are nothing but uniformly finite homology (Chapter 5) [19][46, Appendix A]. A systematic study of coefficient systems that lead to quasi-isometry invariants was carried out by Xin Li [98].
- Taking group von Neumann algebras as coefficients leads to L^2-(co)homology of groups. This measure-theoretic setting is related to measure equivalence of groups. Thus, the vanishing of L^2-Betti numbers is a measure equivalence invariant of groups [64].

A.3 The hyperbolic plane

The hyperbolic plane is one of the origins of modern geometry and a source of instructive examples. The geometry of the hyperbolic plane is "dual" to the intuitively more accessible spherical geometry and can be used to show the independence of the parallel postulate from the other axioms of Euclid.

We will recall the construction of the hyperbolic plane via the halfplane model and we will then sketch how one can develop its basic metric properties from scratch, using the language of elementary Riemannian geometry. In contrast with many of the examples from Riemannian geometry mentioned in this book, this appendix will not require any previous experience in Riemannian geometry.

A.3.1 Construction of the hyperbolic plane

We construct the hyperbolic plane as a Riemannian manifold. To this end, we consider the halfplane model. In short the construction reads as follows: The hyperbolic plane $\mathbb{H}^2$ is the open upper halfplane in $\mathbb{R}^2$, equipped with the Riemannian metric

$$\frac{dx^2 + dy^2}{y^2}.$$

We will now give more details: A *Riemannian manifold* is a smooth manifold together with a smooth family of scalar products on the tangent spaces. Using the local notion of lengths of vectors in tangent spaces, one can introduce the length of smooth curves (by integration) and a notion of angles between smooth curves that start at the same point. Minimising the length of smooth curves between two points gives a metric on the underlying manifold.

Definition A.3.1 (Upper halfplane)**.** We write

$$H := \bigl\{(x,y) \in \mathbb{R}^2 \bigm| y > 0\bigr\} \subset \mathbb{R}^2$$

for the *upper halfplane*. Depending on the context, we will also view H as a subset of $\mathbb{C}$, using the following identifications:

$$\begin{aligned} H &\longrightarrow \{z \in \mathbb{C} \mid \operatorname{Im} z > 0\} \\ (x,y) &\longmapsto x + i \cdot y \\ (\operatorname{Re} z, \operatorname{Im} z) &\longleftarrow z \end{aligned}$$

Because $H \subset \mathbb{R}^2$ is open, the set H inherits the structure of a smooth manifold from $\mathbb{R}^2$. Because the tangent bundle of $\mathbb{R}^2$ is trivial also the tangent bundle of H is trivial and we can hence canonically identify the tangent space $T_z H$ for $z \in H$ with $\mathbb{R}^2$.

Definition A.3.2 (Hyperbolic plane)**.** The *hyperbolic plane* $\mathbb{H}^2$ is the Riemannian manifold (H, g_H), where:

- We endow the open subset $H \subset \mathbb{R}^2$ with the smooth structure of $\mathbb{R}^2$,
- and we consider the Riemannian metric g_H on H given by the scalar products

$$\begin{aligned} g_{H,(x,y)} \colon \mathbb{R}^2 \times \mathbb{R}^2 &\longrightarrow \mathbb{R} \\ (v, v') &\longmapsto \frac{1}{y^2} \cdot \langle v, v' \rangle \end{aligned}$$

  for all $(x, y) \in H$ on $T_{(x,y)} H = \mathbb{R}^2$, where $\langle \cdot, \cdot \rangle$ denotes the standard scalar product on $\mathbb{R}^2$. Moreover, we will also write $\langle \cdot, \cdot \rangle_{H,z}$ for $g_{H,z}$, and we will denote the norm on $T_z H = \mathbb{R}^2$ induced by $g_{H,z}$ by $\| \cdot \|_{H,z}$.

A.3.2 Length of curves

The Riemannian metric allows us to define the length of smooth curves by integration of the length of the speed vectors:

Definition A.3.3 (Hyperbolic length of a curve)**.** Let $\gamma \colon [T_0, T_1] \longrightarrow H$ be a smooth curve. Then the *hyperbolic length of* γ is defined by

$$L_{\mathbb{H}^2}(\gamma) := \int_{T_0}^{T_1} \big\| \dot\gamma(t) \big\|_{H,\gamma(t)} \, dt \in \mathbb{R}_{\geq 0}.$$

The definition of the Riemannian metric g_H implies that curves that are "further up" in the upper halfplane will seem to have a shorter length than curves that are "further down."

We will now give two basic estimates for the hyperbolic length of curves that are the foundation for most results on the metric geometry of the hyperbolic plane:

Proposition A.3.4 (Trivial estimate)**.** *Let* $\gamma \colon [T_0, T_1] \longrightarrow H$ *be a smooth curve and let*

$$\begin{aligned} m &:= \min\{\operatorname{Im} \gamma(t) \mid t \in [T_0, T_1]\} \in \mathbb{R}_{>0}, \\ M &:= \max\{\operatorname{Im} \gamma(t) \mid t \in [T_0, T_1]\} \in \mathbb{R}_{>0}. \end{aligned}$$

Then we have

$$\frac{1}{m} \cdot L_{(\mathbb{R}^2,d_2)}(\gamma) \geq L_{\mathbb{H}^2}(\gamma) \geq \frac{1}{M} \cdot L_{(\mathbb{R}^2,d_2)}(\gamma) \geq \frac{1}{M} \cdot d_2\big(\gamma(T_0), \gamma(T_1)\big),$$

where $L_{(\mathbb{R}^2,d_2)}$ *denotes the Euclidean length of* γ.

Proof. This is a straightforward integration estimate; the last inequality follows from the classical fact that the Euclidean length $L_{(\mathbb{R}^2,d_2)}$ can be calculated both analytically (by integration of the Euclidean length of the speed vectors) and metrically (as in Lemma 7.2.14) [142, Chapter 1]. □

Proposition A.3.5 (Vertical estimate). *Let $\gamma\colon [T_0,T_1] \longrightarrow H$ be a smooth curve and let $p := i \cdot \operatorname{Im}\colon H \longrightarrow H$ be the projection onto the imaginary part.*

1. *We then have $L_{\mathbb{H}^2}(\gamma) \geq L_{\mathbb{H}^2}(p\circ\gamma)$ and equality holds if and only if $\operatorname{Re}\gamma$ is constant.*
2. *Furthermore,*

$$L_{\mathbb{H}^2}(p \circ \gamma) \geq \bigl|\ln \operatorname{Im}\gamma(T_1) - \ln \operatorname{Im}\gamma(T_0)\bigr|$$

and equality holds if and only if the differential of $\operatorname{Im}\gamma$ does not change its sign.

Proof. The map p is smooth and for all $z \in H$ we have

$$\begin{aligned} T_z p\colon \mathbb{R}^2 &\longrightarrow \mathbb{R}^2 \\ (x,y) &\longmapsto y. \end{aligned}$$

The two claims now follow from straightforward integration estimates. □

Proposition A.3.6 (Metric on the hyperbolic plane). *We define*

$$\begin{aligned} d_H\colon H \times H &\longrightarrow \mathbb{R}_{\geq 0} \\ (z,z') &\longmapsto \inf\bigl\{L_{\mathbb{H}^2}(\gamma) \bigm| \gamma \text{ is a smooth curve in } H \text{ from } z \text{ to } z'\bigr\}. \end{aligned}$$

Then d_H is a metric on H.

Proof. Symmetry of d_H is immediate from the definition. For the triangle inequality, one uses smooth approximation of the concatenation of smooth curves. It remains to show that $d_H(z,z') = 0$ happens only for $z = z'$; in order to check this, one can apply the trivial and vertical estimates for hyperbolic lengths and then use that the Euclidean metric d_2 on $\mathbb{R}^2$ has this property. □

A similar argument shows that the topology on H induced by the metric d_H coincides with the subspace topology of $H \subset \mathbb{R}^2$. Moreover, a careful analysis of this type also shows that hyperbolic length of smooth curves coincides with metric length with respect to d_H [147, Chapter 13].

Theorem A.3.7 (Hyperbolic vs. metric length). *Let $\gamma\colon [T_0,T_1] \longrightarrow H$ be a smooth curve. Then*

$$L_{\mathbb{H}^2}(\gamma) = L_{(H,d_H)}(\gamma),$$

where the metric length $L_{(H,d_H)}(\gamma)$ is given by (cf. Lemma 7.2.14)

$$L_{(H,d_H)}(\gamma) := \sup \biggl\{ \sum_{j=0}^{k-1} d_H\bigl(\gamma(t_j), \gamma(t_{j+1})\bigr) \biggm| k \in \mathbb{N},\ t_0, \dots, t_k \in [T_0, T_1],\ t_0 \leq t_1 \leq \dots \leq t_k \biggr\}.$$

In order to get a better understanding of the metric geometry of the hyperbolic plane it is essential to classify the geodesics and isometries of the hyperbolic plane. We start with a special case; the general case will be developed in the next section.

Proposition A.3.8 (Vertical geodesics). *Let $y \in \mathbb{R}_{>1}$. Then there is exactly one smooth geodesic in (H, d_H) from i to $i \cdot y$, namely*

$$\begin{aligned} \gamma\colon [0, \ln y] &\longrightarrow H \\ t &\longmapsto i \cdot e^t. \end{aligned}$$

We have $L_{\mathbb{H}^2}(\gamma) = \ln y$ and hence $d_H(i, i \cdot y) = \ln y$.

Proof. These facts can be derived from the vertical estimate (Proposition A.3.5) and Theorem A.3.7. □

Analogously, we can classify all smooth geodesics between points on the imaginary axis and smooth geodesic lines passing through two points on the imaginary axis. We will see a posteriori (Remark A.3.19) that the same statements also hold if we drop the smoothness hypothesis.

A.3.3 Symmetry and geodesics

We will now determine the isometry group and all geodesics of the hyperbolic plane. Taking advantage of the point of view of geometric group theory, that groups, their actions, and the corresponding geometry are intertwined, we will use a bootstrap mechanism that allows us to simultaneously classify isometries and geodesics.

Riemannian isometries

We start with Riemannian isometries; these are an analytic source of metric isometries.

Definition A.3.9 (Riemannian isometry group). A *Riemannian isometry of* $\mathbb{H}^2$ is a smooth diffeomorphism $f\colon H \longrightarrow H$ that satisfies

$$\forall_{z \in H}\ \forall_{v,v' \in T_z H}\ \bigl\langle T_z f(v), T_z f(v')\bigr\rangle_{H,f(z)} = \langle v, v'\rangle_{H,z}.$$

The *Riemannian isometry group of* $\mathbb{H}^2$ is the set $\operatorname{Isom}(\mathbb{H}^2)$ of all Riemannian isometries of $\mathbb{H}^2$ with the group multiplication given by composition of maps.

Proposition A.3.10 (Riemannian isometries are isometries)**.** *Every Riemannian isometry of* $\mathbb{H}^2$ *is a metric isometry of* (H, d_H)*. In particular, we obtain an injective group homomorphism*

$$\operatorname{Isom}(\mathbb{H}^2) \longrightarrow \operatorname{Isom}(H, d_H).$$

Proof. This follows from the definition of d_H in terms of Riemannian lengths of smooth curves and the chain rule (applied to Riemannian isometries and their inverses). ☐

In fact, we will see in Theorem A.3.23 that this homomorphism is not only injective, but even an isomorphism; i.e., every metric isometry of the hyperbolic plane is a Riemannian isometry. (More generally, the analogous statement holds for all complete Riemannian manifolds.)

Möbius transformations

As next step, we consider an explicit class of isometries of the hyperbolic plane, the Möbius transformations:

Proposition A.3.11 (Möbius transformations)**.** *For*

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}(2, \mathbb{R})$$

we define the associated Möbius transformation *by*

$$\begin{aligned} f_A \colon H &\longrightarrow H \\ z &\longmapsto \frac{a \cdot z + b}{c \cdot z + d}. \end{aligned}$$

Then:

1. *For all* $z \in H$ *we have* $\operatorname{Im} f_A(z) = \frac{1}{|c \cdot z + d|^2} \cdot \operatorname{Im} z$.
2. *The map* f_A *is a well-defined smooth diffeomorphism.*
3. *We have* $f_{E_2} = \mathrm{id}_H = f_{-E_2}$.
4. *For all* $A, B \in \mathrm{SL}(2, \mathbb{R})$ *we have* $f_{A \cdot B} = f_A \circ f_B$.

Proof. These are straightforward calculations. ☐

Example A.3.12 (Simple Möbius transformations)**.**

- Let $b \in \mathbb{R}$. Then the Möbius transformation associated with the matrix

  $$\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}(2, \mathbb{R})$$

  is the horizontal translation $z \longmapsto z + b$ on H by b.

- The Möbius transformation associated with the matrix
$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \in \mathrm{SL}(2,\mathbb{R})$$
is the map $z \longmapsto -1/z$ (Figure A.3), which is its own inverse. This map is related to circle inversion in elementary geometry.

Elementary row and column operations show that the set
$$\left\{\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}\right\} \cup \left\{\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} \;\middle|\; b \in \mathbb{R}\right\}$$
generates $\mathrm{SL}(2,\mathbb{R})$.

The groups $\mathrm{SL}(2,\mathbb{R})$ and $\mathrm{PSL}(2,\mathbb{R}) = \mathrm{SL}(2,\mathbb{R})/\{E_2, -E_2\}$ act through isometries on the hyperbolic plane:

Proposition A.3.13 (Möbius transformations are isometries)**.** *If $A \in \mathrm{SL}(2,\mathbb{R})$, then the associated Möbius transformation $f_A \colon H \longrightarrow H$ is a Riemannian isometry of $\mathbb{H}^2$. In particular, we obtain an injective group homomorphism*
$$\begin{aligned} \mathrm{PSL}(2,\mathbb{R}) = \mathrm{SL}(2,\mathbb{R})/\{E_2, -E_2\} &\longrightarrow \mathrm{Isom}(H, d_H) \\ [A] &\longmapsto f_A \end{aligned}$$
and Möbius transformations map geodesics in (H, d_H) to geodesics.

Proof. This is a lengthy, but elementary, calculation. One way to simplify the proof that Möbius transformations are Riemannian isometries is to use the generating set of Example A.3.12. □

Transitivity of Möbius transformations

We will now prove that the action of $\mathrm{SL}(2,\mathbb{R})$ by Möbius transformations on the hyperbolic plane is transitive in a strong sense; hence, the hyperbolic plane is a symmetric space.

Proposition A.3.14 (Transitivity of Möbius transformations)**.**
1. *The action of $\mathrm{SL}(2,\mathbb{R})$ on H by Möbius transformations is transitive.*
2. *The stabiliser group of i with respect to this action is $\mathrm{SO}(2)$ (i.e., the group of all 2×2-rotation matrices).*
3. *For all $z, z' \in H$ there exists $A \in \mathrm{SL}(2,\mathbb{R})$ with*
$$f_A(z) = i \qquad \textit{and} \qquad \mathrm{Re}\big(f_A(z')\big) = 0, \qquad \mathrm{Im}\big(f_A(z')\big) > 1.$$

Proof. Ad 1. It suffices to show that every point in H can be moved by a Möbius transformation to the point i. Let $z \in H$ with $x := \mathrm{Re}\, z$, $y := \mathrm{Re}\, z$. Then the Möbius transformation associated with the matrix

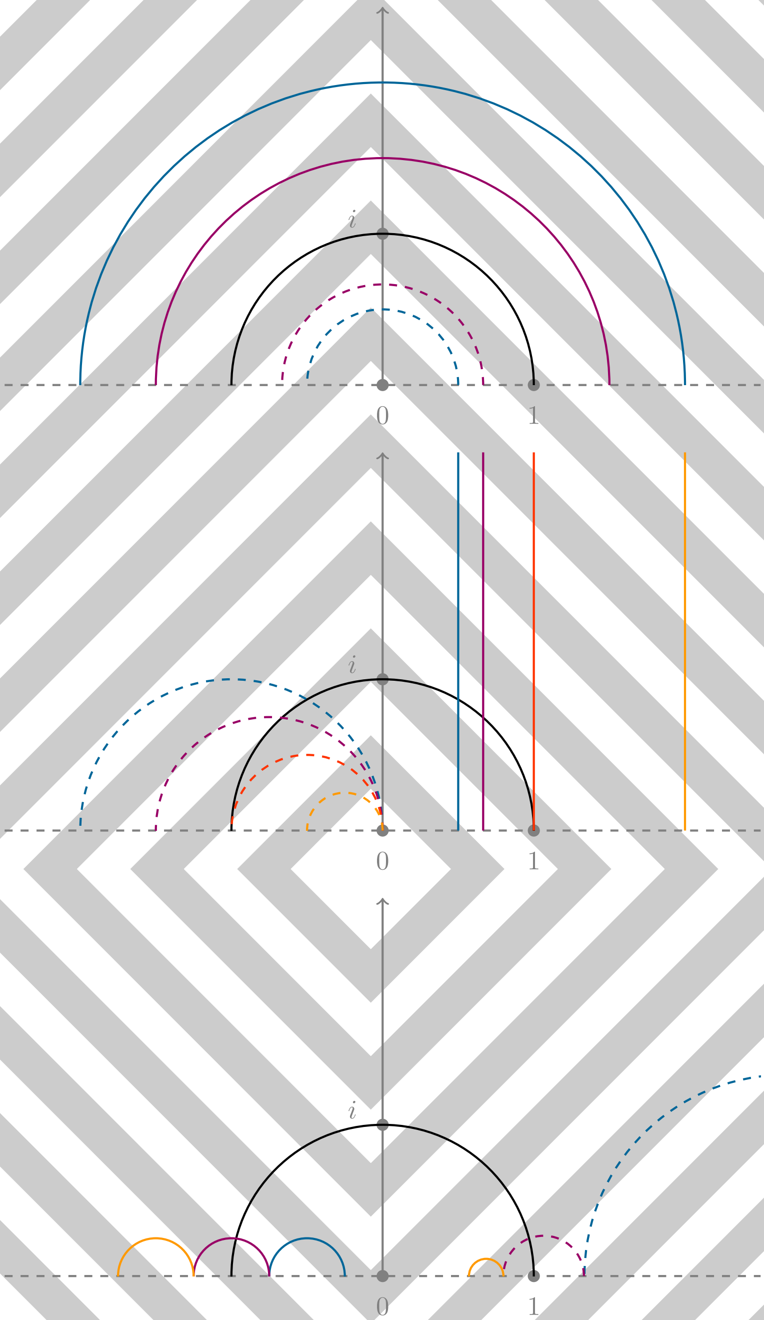

Figure A.3.: The Möbius transformation $z \longmapsto -1/z$ on H; in every picture, objects of the same colour are mapped to each other

$$\begin{pmatrix} 1/\sqrt{y} & 0 \\ 0 & \sqrt{y} \end{pmatrix} \cdot \begin{pmatrix} 1 & -x \\ 0 & 1 \end{pmatrix}$$

(which first translates z to the imaginary axis and then scales the imaginary axis by $1/\sqrt{y}$) maps z to i.

Ad 2. This is a simple calculation.

Ad 3. This can be derived from the previous parts (and either a nasty calculation or a continuity argument on the angles of rotations). □

Geodesics

Combining our knowledge on vertical geodesics and the transitivity of Möbius transformations allows us to classify all geodesics of the hyperbolic plane:

Theorem A.3.15 (Characterisation of hyperbolic geodesics)**.** *Let $z, z' \in H$ with $z \neq z'$.*

1. *Then there exists precisely one [smooth] geodesic in (H, d_H) from z to z'. In particular, the metric space (H, d_H) is geodesic.*
2. *Up to reparametrisation on $\mathbb{R}$ there exists precisely one [smooth] geodesic line in (H, d_H) that contains z and z'.*

More precisely: If $A \in \mathrm{SL}(2, \mathbb{R})$ with $\mathrm{Re}(f_A(z)) = 0 = \mathrm{Re}(f_A(z'))$, then $f_{A^{-1}} \circ (t \mapsto i \cdot e^t)$ is a geodesic line through z and z' and the [smooth] geodesic from z to z' is a segment of this line.

Proof. We use the Möbius transformation action on (H, d_H), which is an isometric action (Proposition A.3.13). Then the theorem is just a combination of Proposition A.3.14 and Proposition A.3.8. □

Remark A.3.16 (Hyperbolic metric, explicitly)**.** Using the description of hyperbolic geodesics from Theorem A.3.15 and the computation of vertical distances (Proposition A.3.5) allows to cook up an explicit formula for the metric d_H on H: For all $z, z' \in H$ we have

$$d_H(z, z') = \operatorname{arcosh}\Bigl(1 + \frac{|z - z'|^2}{2 \cdot \operatorname{Im} z \cdot \operatorname{Im} z'}\Bigr).$$

Here, the *area hyperbolic cosine* arcosh is the inverse

$$\begin{aligned} \operatorname{arcosh} \colon \mathbb{R}_{\geq 1} &\longrightarrow \mathbb{R} \\ x &\longmapsto \ln\bigl(x + \sqrt{x^2 - 1}\bigr) \end{aligned}$$

of the *hyperbolic cosine* function

$$\begin{aligned} \cosh \colon \mathbb{R} &\longrightarrow \mathbb{R}_{\geq 1} \\ x &\longmapsto \frac{e^x + e^{-x}}{2} = \cos(i \cdot x). \end{aligned}$$

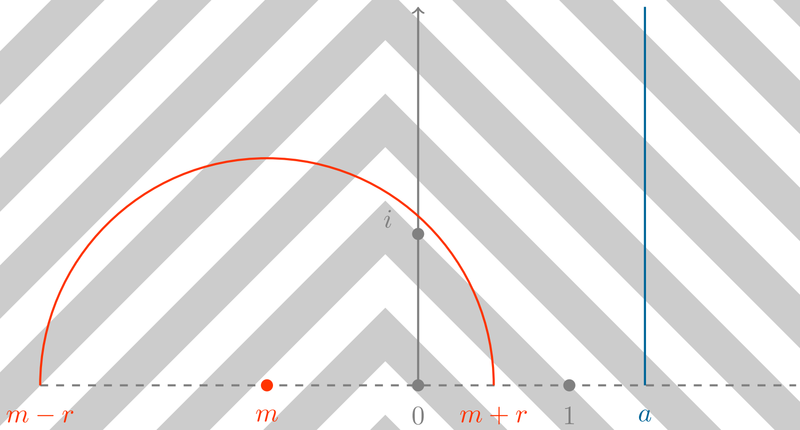

Figure A.4.: Generalised semi-circles in the upper halfplane

To complete our understanding of hyperbolic geodesics, we introduce the notion of generalised semi-circles (Figure A.4):

Definition A.3.17 (Generalised semi-circle). A *generalised semi-circle* is a subset $K \subset H$ of the following form:

- There exist $m \in \mathbb{R}$ and $r \in \mathbb{R}_{>0}$ with $K = \{z \in H \mid |z - m| = r\}$ or
- there is an $a \in \mathbb{R}$ with $K = \{a + i \cdot t \mid t \in \mathbb{R}_{>0}\}$.

Looking at the generators for $\mathrm{SL}(2, \mathbb{R})$ from Example A.3.12 shows that Möbius transformations map generalised semi-circles to generalised semi-circles.

Corollary A.3.18 (Hyperbolic geodesics are generalised semi-circles). *[Smooth] geodesics and [smooth] geodesic lines in (H, d_H) are exactly the (correctly parametrised) segments of generalised semi-circles in H.*

Proof. This follows from the characterisation of hyperbolic geodesics in Theorem A.3.15 and the compatibility of Möbius transformations with generalised semi-circles. □

In particular, one can easily find configurations of geodesic lines in the hyperbolic plane that do *not* satisfy the parallel postulate by Euclid.

Remark A.3.19 (Smoothness of hyperbolic geodesics). We can now also prove that the characterisation of vertical geodesics (whence for all geodesics in (H, d_H)) also holds for general metric geodesics (and not only for the smooth ones); in particular, all geodesics in (H, d_H) are smooth.

Let $y \in \mathbb{R}_{>1}$ and let $\eta\colon [0, L] \longrightarrow H$ be a metric geodesic in (H, d_H) from i to $i \cdot y$; in particular, $L = \ln y$.

- As first step, one shows that $\operatorname{Re}\eta = 0$. *Assume* for a contradiction that there is a $t \in (0, L)$ with $\operatorname{Re}\eta(t) \neq 0$. By Theorem A.3.15, we can connect $\eta(t)$ with i and with $i \cdot y$ through smooth geodesics. One applies the projection from Proposition A.3.5 and then uses the calculations from Theorem A.3.7, Proposition A.3.5, and Proposition A.3.8 to arrive at a contradiction.
- By the first step, η lives on the imaginary line. As geodesic, the map η is injective; hence, $\operatorname{Im}\eta$ is injective (and monotonous). Together with the knowledge on vertical distances we hence obtain

$$\forall_{t \in [0,L]} \ \ \eta(t) = i \cdot e^t.$$

Angles

In order to complete the description of the isometry group of the hyperbolic plane it is convenient to work with angles. As in the case of the length of curves and of isometries, there are two notions of angles: a Riemannian one and a metric one.

Definition A.3.20 (Hyperbolic angle)**.** Let $\gamma_1 \colon [0, L_1] \longrightarrow H$, $\gamma_2 \colon [0, L_2] \longrightarrow H$ be smooth curves in $\mathbb{H}^2$ with $\gamma_1(0) = \gamma_2(0)$ and $\dot\gamma_1(0) \neq 0 \neq \dot\gamma_2(0)$. Then the *hyperbolic angle between* γ_1 *and* γ_2 is defined by

$$\begin{aligned} \sphericalangle_H(\gamma_1, \gamma_2) &:= \sphericalangle_H\bigl(\dot\gamma_1(0), \dot\gamma_2(0)\bigr) \\ &:= \arccos \frac{\langle \dot\gamma_1(0), \dot\gamma_2(0)\rangle_{H,\gamma_1(0)}}{\|\dot\gamma_1(0)\|_{H,\gamma_1(0)} \cdot \|\dot\gamma_2(0)\|_{H,\gamma_2(0)}} \in [0, \pi]. \end{aligned}$$

The halfplane model is *conformal* in the following sense: If γ_1 and γ_2 are smooth curves in H with $\gamma_1(0) = \gamma_2(0)$ and $\dot\gamma_1(0) \neq 0 \neq \dot\gamma_2(0)$, then

$$\sphericalangle_H(\gamma_1, \gamma_2) = \sphericalangle(\gamma_1, \gamma_2)$$

because $\langle \cdot\,, \cdot \rangle_{H,z}$ is obtained from the Euclidean scalar product by scaling (and hence defines the same angles). Hyperbolic angles drawn in the halfplane model therefore can be read off as the corresponding Euclidean angles.

Proposition A.3.21 (Hyperbolic angle, metric version)**.** *Let* $\gamma_1 \colon [0, L_1] \longrightarrow H$*,* $\gamma_2 \colon [0, L_2] \longrightarrow H$ *be geodesics in* (H, d_H) *with* $\gamma_1(0) = \gamma_2(0)$*. Then*

$$\sphericalangle_H(\gamma_1, \gamma_2) = \lim_{t \to 0} \arccos\Bigl(1 - \frac{d_H\bigl(\gamma_1(t), \gamma_2(t)\bigr)^2}{2 \cdot t^2}\Bigr).$$

Proof. By the characterisation of geodesics in (H, d_H) (Theorem A.3.15, Remark A.3.19) we know that γ_1 and γ_2 are smooth and have non-zero derivative at 0; hence, $\sphericalangle_H(\gamma_1, \gamma_2)$ is defined. Using transitivity of the Möbius transformation action and the fact that the right hand side in the theorem is invariant under isometries, we can restrict to the case that $\gamma_1(0) = i = \gamma_2(0)$. Because $\langle\,\cdot\,,\cdot\,\rangle_{H,i} = \langle\,\cdot\,,\cdot\,\rangle$, we then obtain

$$\sphericalangle_H(\gamma_1,\gamma_2) = \sphericalangle_H\bigl(\dot\gamma_1(0),\dot\gamma_2(0)\bigr) = \sphericalangle\bigl(\dot\gamma_1(0),\dot\gamma_2(0)\bigr).$$

Moreover, by the characterisation of hyperbolic geodesics, we have that $\|\dot\gamma_j(0)\|_{H,i} = \|\dot\gamma_j(0)\|_2 = 1$ for $j \in \{1,2\}$. Therefore, polarisation shows that

$$\sphericalangle_H(\gamma_1,\gamma_2) = \arccos\Bigl(1 - \frac{1}{2}\cdot \bigl\|\dot\gamma_1(0) - \dot\gamma_2(0)\bigr\|_2^2\Bigr).$$

Using the trivial and the vertical estimate (Proposition A.3.4 and A.3.5) as well as the description of the derivatives as differential quotient, one can now derive the claimed formula for the hyperbolic angle. □

Corollary A.3.22 (Hyperbolic isometries are conformal)**.** *Let* $f \in \operatorname{Isom}(H,d_H)$ *and let* γ_1,γ_2 *be geodesics in in* (H,d_H) *with* $\gamma_1(0) = \gamma_2(0)$*. Then*

$$\sphericalangle_H(f\circ\gamma_1, f\circ\gamma_2) = \sphericalangle_H(\gamma_1,\gamma_2).$$

Proof. Because f is an isometry, the curves $f\circ\gamma_1$ and $f\circ\gamma_2$ are geodesics in (H,d_H) with $f\circ\gamma_1(0) = f\circ\gamma_2(0)$. As hyperbolic angles can be expressed in terms of the metric (Proposition A.3.21), we obtain

$$\sphericalangle_H(f\circ\gamma_1, f\circ\gamma_2) = \sphericalangle_H(\gamma_1,\gamma_2). \qquad \square$$

The isometry group

We now complete the classification of isometries of the hyperbolic plane:

Theorem A.3.23 (The hyperbolic isometry group)**.** *The group* $\operatorname{Isom}(H,d_H)$ *is generated by*

$$\{f_A \mid A \in \mathrm{SL}(2,\mathbb{R})\} \cup \{z \mapsto -\overline{z}\}.$$

In particular, every isometry of (H,d_H) *is a (smooth) Riemannian isometry and so* $\operatorname{Isom}(H,d_H) = \operatorname{Isom}(\mathbb{H}^2)$*. The map*

$$\begin{aligned} \mathrm{PSL}(2,\mathbb{R}) &\longrightarrow \operatorname{Isom}^+(H,d_H) \\ A &\longmapsto f_A \end{aligned}$$

is an isomorphism. Here, $\operatorname{Isom}^+(H,d_H)$ *denotes the group of all orientation-preserving isometries of* (H,d_H)*.*

Proof. Let $f \in \operatorname{Isom}(H,d_H)$.

- In view of the transitivity of the Möbius transformation action and the description of the hyperbolic metric on the imaginary line, we may assume that

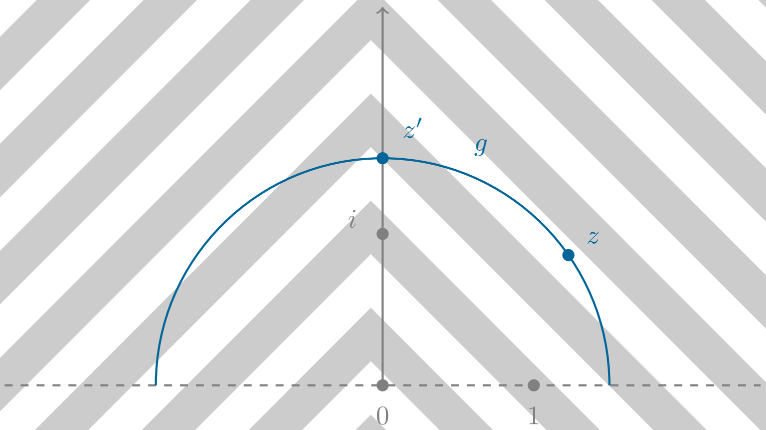

Figure A.5.: Orthogonal geodesic lines and isometries

$$f(i) = i \qquad \text{and} \qquad f(2 \cdot i) = 2 \cdot i.$$

- Because f is an isometry, it maps geodesic lines to geodesic lines. Hence, the characterisation of vertical geodesic lines (Proposition A.3.8) shows that $f(i \cdot y) = i \cdot y$ holds for all $y \in \mathbb{R}_{>0}$.
- The map f is a homeomorphism and so f permutes the path-connected components $P := \{z \in H \mid \operatorname{Re} z > 0\}$ and $N := \{z \in H \mid \operatorname{Re} z < 0\}$ of $H \setminus \mathbb{R}_{>0} \cdot i$. Applying the hyperbolic isometry $(z \mapsto -\overline{z})$ (reflection at the imaginary axis) allows us to assume that

$$f(P) = P \qquad \text{and} \qquad f(N) = N.$$

- We will now prove that $f = \operatorname{id}_H$: Let $z \in H$. If $\operatorname{Re} z = 0$, then we already know $f(z) = z$. We will now treat the case $\operatorname{Re} z > 0$ (the other case can be treated in the same way). By the characterisation of hyperbolic geodesics in terms of generalised semi-circles (Corollary A.3.18), there exists (up to reparametrisation) exactly one geodesic line g in (H, d_H) through z that intersects the imaginary axis (which is a geodesic line!) orthogonally; let z' be this intersection (Figure A.5).
  We now apply f to this situation. Because f maps geodesic lines to geodesic lines, because f maps the imaginary axis to itself, and is conformal (Corollary A.3.22), the geodesic line $f \circ g$ is orthogonal to the imaginary axis and goes through z'. Then the classification of hyperbolic geodesics shows that $f \circ g(\mathbb{R}) = g(\mathbb{R})$. Because of $f(P) = P$ and

$$d_H\bigl(f(z), f(z')\bigr) = d_H(z, z')$$

  we obtain $f(z) = z$ (by travelling from z' along g).

This shows that every isometry of (H, d_H) is a composition of a Möbius transformation (and possibly the reflection at the imaginary axis). Clearly, all Möbius transformations are orientation preserving and the reflection at the imaginary axis is orientation reversing; hence, also the claim on $\operatorname{Isom}^+(H, d_H)$ follows. □

A.3.4 Hyperbolic triangles

Geodesic triangles in the hyperbolic plane are uniformly slim; this is one the key properties of global negative curvature as studied in Chapter 7.

The Gauß-Bonnet theorem for hyperbolic triangles

In contrast with Euclidean triangles, hyperbolic triangles have bounded area and angle sum less than π. Using the Riemannian volume form on $\mathbb{H}^2$ one obtains a notion of area on the hyperbolic plane. More explicitly:

Definition A.3.24 (Hyperbolic area)**.** Let $f\colon H \longrightarrow \mathbb{R}_{\geq 0}$ be a Borel measurable map. Then we define the *integral of f over $\mathbb{H}^2$* by

$$\begin{aligned}\int_H f \ \mathrm{dvol}_H &:= \int_H f(x,y) \cdot \sqrt{\det G_{H,(x,y)}} \ d(x,y) \\ &= \int_H \frac{1}{y^2} \cdot f(x,y) \ d(x,y) \in \mathbb{R}_{\geq 0} \cup \{\infty\},\end{aligned}$$

where

$$G_{H,(x,y)} := \begin{pmatrix} g_{H,(x,y)}(e_1,e_1) & g_{H,(x,y)}(e_1,e_2) \\ g_{H,(x,y)}(e_2,e_1) & g_{H,(x,y)}(e_2,e_2) \end{pmatrix} = \begin{pmatrix} 1/y^2 & 0 \\ 0 & 1/y^2 \end{pmatrix}$$

and $e_1, e_2 \in T_{H,(x,y)} H = \mathbb{R}^2$ are the standard coordinate vectors.

If $A \subset H$ is a measurable set, we define the *hyperbolic area of A* by

$$\mu_{\mathbb{H}^2}(A) := \int_H \chi_A \ \mathrm{dvol}_H \in \mathbb{R}_{\geq 0} \cup \{\infty\},$$

where

$$\begin{aligned}\chi_A\colon H &\longrightarrow \mathbb{R}_{\geq 0} \\ z &\longmapsto \begin{cases} 1 & \text{if } z \in A \\ 0 & \text{if } z \notin A \end{cases}\end{aligned}$$

is the characteristic function of $A \subset H$.

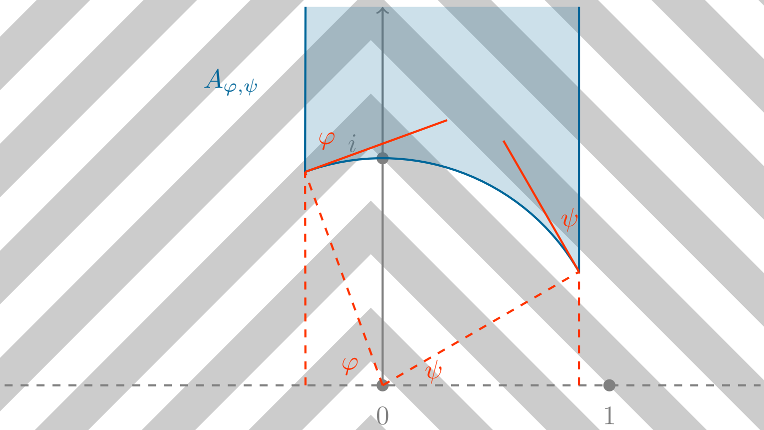

Figure A.6.: Proof of the Gauß-Bonnet theorem for hyperbolic triangles in the case that one vertex lies on the boundary of the hyperbolic plane.

Proposition A.3.25 (Isometries are area-preserving). *Let $A \subset H$ be a measurable set and let $f \in \operatorname{Isom}(H, d_H)$. Then $f(A)$ is measurable and*

$$\mu_{\mathbb{H}^2}\bigl(f(A)\bigr) = \mu_{\mathbb{H}^2}(A).$$

Proof. The isometry f is a Riemannian isometry of $\mathbb{H}^2$ (Theorem A.3.23). The transformation formula then shows that f is area-preserving. □

We can now define the area of geodesic triangles in the hyperbolic plane as follows: Let $\Delta := (\gamma_0\colon [0, L_0] \to H, \gamma_1\colon [0, L_1] \to H, \gamma_2\colon [0, L_2] \to H)$ be a geodesic triangle in (H, d_H). Using the characterisation of geodesics in the hyperbolic plane, we deduce that γ_0, γ_1 and γ_2 only meet where they should and that they bound a compact, measurable set A_Δ in H. We then define the *hyperbolic area of* Δ by

$$\mu_{\mathbb{H}^2}(\Delta) := \mu_{\mathbb{H}^2}(A_\Delta).$$

Theorem A.3.26 (Gauß-Bonnet theorem for hyperbolic triangles). *Let Δ be a geodesic triangle in (H, d_H) with angles α, β, γ and suppose that the image of Δ is not contained in a single geodesic line. Then*

$$\mu_{\mathbb{H}^2}(\Delta) = \pi - (\alpha + \beta + \gamma).$$

In particular: The angle sum in hyperbolic geodesic triangles is less than π and the hyperbolic area is bounded from above by π.

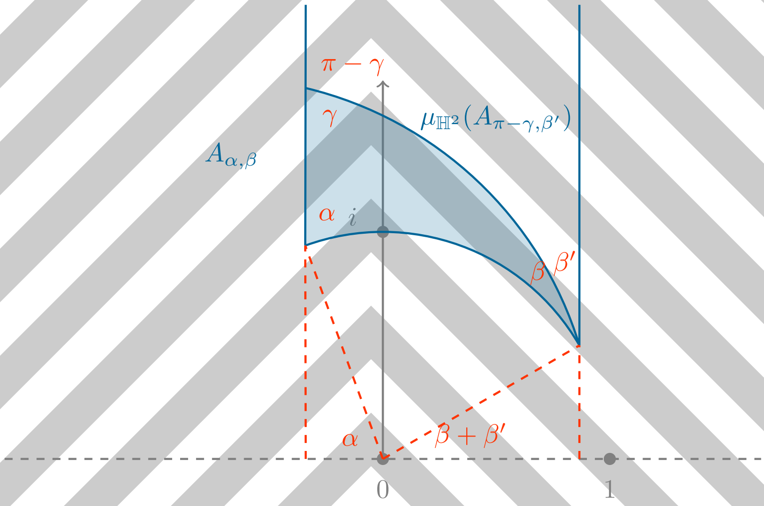

Figure A.7.: Proof of the Gauß-Bonnet theorem for hyperbolic triangles

Proof. We first determine the area of generalised geodesic triangles that have one vertex on the boundary of the hyperbolic plane; in a second step, we demonstrate how the case of ordinary geodesic triangles can be derived from the first case by a simple decomposition argument.

① Let $\varphi, \psi \in [0, \pi]$ with $\varphi + \psi < \pi$ and let

$$A_{\varphi,\psi} := \bigl\{(x, y) \in H \bigm| x \in [\cos(\pi - \varphi), \cos\psi],\ y \geq \sqrt{1 - x^2}\bigr\}$$

(Figure A.6). Then $A_{\varphi,\psi}$ is a closed (whence measurable) subset of H and a straightforward calculation shows that

$$\begin{aligned}\mu_{\mathbb{H}^2}(A_{\varphi,\psi}) &= \int_H \chi_{A_{\varphi,\psi}}\ \mathrm{dvol}_H = \int_{\cos(\pi-\varphi)}^{\cos\psi} \int_{\sqrt{1-x^2}}^{\infty} \frac{1}{y^2}\, dy\, dx \\ &= \int_{\cos(\pi-\varphi)}^{\cos\psi} \frac{1}{\sqrt{1-x^2}}\, dx = \int_{\pi-\varphi}^{\psi} -1\, dt \\ &= \pi - (\varphi + \psi).\end{aligned}$$

② We now return to the case of our geodesic triangle: Using the transitivity of the Möbius transformation action we may assume without loss of generality that the image of Δ lies above the geodesic opposite of γ and that the semi-circle of the geodesic opposite of γ has radius 1 and centre 0. Because the halfplane model is conformal, we can determine the necessary angles by Euclidean considerations. In particular, we obtain $\alpha + \beta < \pi$. Therefore, A_Δ is contained in $A_{\alpha,\beta+\beta'}$ (where

β' is defined as in Figure A.7). Moreover, the closure of the complement $A_{\alpha,\beta+\beta'} \setminus (A_\Delta)$ is isometric to $A_{\pi-\gamma,\beta'}$.
Therefore, Step ① and Proposition A.3.25 imply

$$\begin{aligned}\mu_{\mathbb{H}^2}(\Delta) &= \mu_{\mathbb{H}^2}(A_{\alpha,\beta+\beta'}) - \mu_{\mathbb{H}^2}(A_{\pi-\gamma,\beta'}) \\ &= \pi - (\alpha+\beta+\beta') - \pi + (\pi - \gamma + \beta') \\ &= \pi - (\alpha+\beta+\gamma),\end{aligned}$$

as claimed. □

For example, one can use the Gauß-Bonnet theorem for hyperbolic triangles to prove that the hyperbolic plane and the Euclidean plane are *not* locally isometric.

Hyperbolic triangles are slim

Hyperbolic geodesic triangles are slim in the following sense:

Theorem A.3.27 (Hyperbolic triangles are slim)**.** *There is a constant $C \in \mathbb{R}_{\geq 0}$ such that every geodesic triangle in (H, d_H) is C-slim, i.e.: For every geodesic triangle $(\gamma_0 \colon [0, L_0] \to H, \gamma_1 \colon [0, L_1] \to H, \gamma_2 \colon [0, L_2] \to H)$ in (H, d_H) and every $t \in [0, L_0]$, there exists*
- *an $s \in [0, L_1]$ with $d_H\bigl(\gamma_0(t), \gamma_1(s)\bigr) \leq C$*
- *or an $s \in [0, L_2]$ with $d_H\bigl(\gamma_0(t), \gamma_2(s)\bigr) \leq C$.*

In other words, the geodesic metric space (H, d_H) is hyperbolic in the sense of Definition 7.2.2. In fact, the hyperbolic plane is the name-giving example of hyperbolic metric spaces.

For the proof of Theorem A.3.27, we will combine the Gauß-Bonnet theorem with the following observation on area growth:

Proposition A.3.28 (Exponential growth of hyperbolic area)**.** *For all $r \in \mathbb{R}_{>10}$ we have*

$$\mu_{\mathbb{H}^2}\bigl(B_r^{H,d_H}(i)\bigr) \geq e^{\frac{r}{10}} \cdot (1 - e^{-\frac{r}{2}})$$

In particular, the area of hyperbolic disks growths exponentially in the radius.

Proof. Let $r \in \mathbb{R}_{\geq 10}$. Elementary estimates show that the set

$$Q_r := \bigl\{x + i \cdot y \bigm| x \in [0, e^{r/10}],\ y \in [1, e^{r/2}]\bigr\}$$

lies in $B_r^{H,d_H}(i)$. In particular, we obtain

$$\mu_{\mathbb{H}^2}\bigl(B_r^{H,d_H}(i)\bigr) \geq \mu_{\mathbb{H}^2}(Q_r) = e^{r/10} \cdot (1 - e^{-r/2}).$$ □

Proof of Theorem A.3.27. In view of the exponential growth of hyperbolic area (Proposition A.3.28) there exists a $C \in \mathbb{R}_{>0}$ satisfying

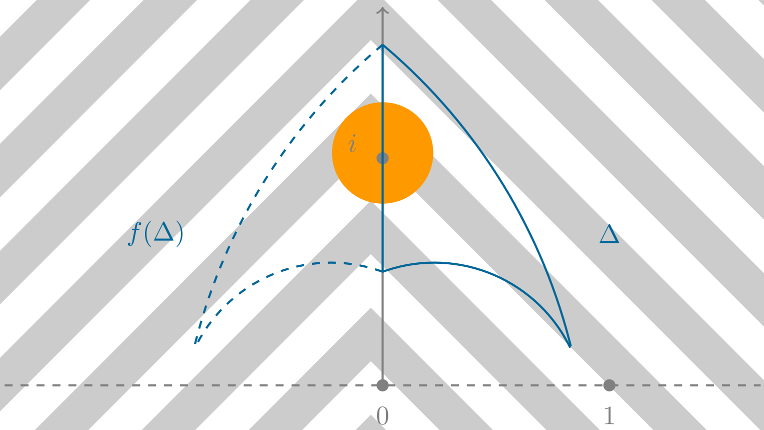

Figure A.8.: Proving that geodesic triangles in the hyperbolic plane are slim

$$\mu_{\mathbb{H}^2}\bigl(B_C^{H,d_H}(i)\bigr) \geq 4 \cdot \pi.$$

Let $\Delta := (\gamma_0, \gamma_1, \gamma_2)$ be a geodesic triangle in (H, d_H) and let $x \in \operatorname{im} \gamma_0$. Without loss of generality we may assume that the image of Δ is not contained in a single geodesic line. By transitivity of the Möbius transformation action we may assume in addition that $\operatorname{im} \gamma_0$ lies on the imaginary axis and $x = i$.

Assume for a contradiction that there is no point $y \in \operatorname{im} \gamma_1 \cup \operatorname{im} \gamma_2$ with $d_H(x, y) \leq C$. Then we have

$$B_C^{H,d_H}(i) \subset A_\Delta \cup \operatorname{im} \gamma_0 \cup f(A_\Delta),$$

where $f \colon z \mapsto -\overline{z}$ (Figure A.8). Because f is an isometry and hence area-preserving (Proposition A.3.25), we obtain from the Theorem of Gauß-Bonnet (Theorem A.3.26) that

$$\begin{aligned} 4 \cdot \pi &\leq \mu_{\mathbb{H}^2}\bigl(B_C^{H,d_H}(i)\bigr) \\ &\leq \mu_{\mathbb{H}^2}(\Delta) + \mu_{\mathbb{H}^2}(\operatorname{im} \gamma_0) + \mu_{\mathbb{H}^2}\bigl(f(A_\Delta)\bigr) \\ &= \mu_{\mathbb{H}^2}(\Delta) + 0 + \mu_{\mathbb{H}^2}(\Delta) \\ &< 2 \cdot \pi, \end{aligned}$$

which is absurd.

Hence, there is a $y \in \operatorname{im} \gamma_1 \cup \operatorname{im} \gamma_2$ with $d_H(x, y) \leq C$. □

A.3.5 Curvature

We will now indicate how the Gaussian curvature of the hyperbolic plane can be computed in the halfplane model.

Theorem A.3.29. *The Gaussian curvature of the hyperbolic plane is the constant function* -1.

Sketch of proof of Theorem A.3.29. By the *Theorema Egregium*, the Gaussian curvature indeed is an intrinsic invariant of Riemannian surfaces and can be expressed in terms of local data [96]. Roughly speaking these formulae quantify curvature as the dependence of parallel transport of tangent vectors along different curves (Chapter 7.1). The Gaussian curvature of $\mathbb{H}^2$ is the smooth function $K\colon H \longrightarrow \mathbb{R}$ given in local coordinates by the general formula

$$K = \frac{R_{1221}}{g_{12} \cdot g_{21} - g_{11} \cdot g_{22}}$$

for Riemannian surfaces, where we use the following notation (and calculate in the standard coordinate system of the upper halfplane H):

- As usual in Riemannian geometry, for $j, k \in \{1, 2\}$ we write g_{jk} for the jk-component of the Riemannian metric. In our case, g_{jk} is the function $z \mapsto g_{H,z}(e_j, e_k) = (G_{H,z})_{j,k}$ (Definition A.3.24). The coefficients of the pointwise inverse of the matrix valued function $z \mapsto G_{H,z}$ are denoted by g^{jk}. We have for all $z = (x, y) \in H$ that

$$G_{H,z} = \begin{pmatrix} 1/y^2 & 0 \\ 0 & 1/y^2 \end{pmatrix} \qquad \text{and} \qquad G_{H,z}^{-1} = \begin{pmatrix} y^2 & 0 \\ 0 & y^2 \end{pmatrix}.$$

  The denominator of K is hence given by $(x, y) \longmapsto -1/y^4$.
- The geometric relation between tangent spaces at different points and parallel transport is encoded in the so-called Levi-Civita connection. Locally, we can express the Levi-Civita connection in terms of the Riemannian metric by the Christoffel symbols: The *Christoffel symbols* are functions $H \longrightarrow \mathbb{R}$ indexed by $i, k, \ell \in \{1, 2\}$ and given by

$$\begin{aligned} {\Gamma^i}_{k\ell} &= \frac{1}{2} \cdot \sum_{m=1}^{2} g^{im} \cdot (\partial_\ell g_{mk} + \partial_k g_{m\ell} - \partial_m g_{k\ell}) \\ &= \frac{1}{2} \cdot g^{ii} \cdot (\partial_\ell g_{ik} + \partial_k g_{i\ell} - \partial_i g_{k\ell}). \end{aligned}$$

  We hence have for all $(x, y) \in H$ that

$$\begin{aligned} {\Gamma^2}_{11}(x, y) &= \frac{1}{y} & \qquad {\Gamma^2}_{22}(x, y) &= -\frac{1}{y} \\ {\Gamma^1}_{12}(x, y) &= -\frac{1}{y} & \qquad {\Gamma^1}_{21}(x, y) &= -\frac{1}{y} \end{aligned}$$

and all other Christoffel symbols are zero.

- The function R_{1221} is the coefficient of the Riemann curvature tensor, given by

$$R_{1221} = \sum_{h=1}^{2} g_{1h} \cdot R^h{}_{221} = g_{11} \cdot R^1{}_{221}.$$

The version $R^1{}_{221}$ of the Riemann curvature tensor is given by

$$R^1{}_{221} = \partial_2 \Gamma^1{}_{12} - \partial_1 \Gamma^1{}_{22} + \sum_{h=1}^{2} (\Gamma^1{}_{2h} \cdot \Gamma^h{}_{12} - \Gamma^1{}_{1h} \cdot \Gamma^h{}_{22}).$$

Therefore, for all $(x, y) \in H$ we obtain

$$\begin{aligned} R_{1221}(x,y) &= \frac{1}{y^2} \cdot R^1{}_{221}(x,y) = \frac{1}{y^2} \cdot \Bigl(\frac{1}{y^2} - 0 + \frac{1}{y^2} - 0 + 0 - \frac{1}{y^2} \Bigr) \\ &= \frac{1}{y^4}. \end{aligned}$$

In total, this leads to

$$K(x,y) = \frac{R_{1221}(x,y)}{g_{12}(x,y) \cdot g_{12}(x,y) - g_{11}(x,y) \cdot g_{22}(x,y)} = \frac{\frac{1}{y^4}}{-\frac{1}{y^4}} = -1$$

for all $(x, y) \in H$, as claimed. □

A.3.6 Other models

Models of the hyperbolic plane are metric spaces that are isometric to the hyperbolic plane. Popular models of the hyperbolic plane are:

- the halfplane model (which we used as defining model),
- the Poincaré disk model,
- the Klein disk model, and
- the hyperboloid model.

The definition of these models and their comparison is, for instance, explained in the comprehensive book by Ratcliffe [146]. A cunning way of illustrating these comparisons is realised by Segerman's 3D-models and suitable lighting of these models [158, Chapter 4]. For example, the Poincaré disk model can be obtained from the halfplane model as follows:

Example A.3.30 (Poincaré disk model)**.** The *Cayley transform*

$$\begin{aligned} C \colon H &\longrightarrow E := \bigl\{ z \in \mathbb{C} \bigm| |z| < 1 \bigr\} \\ z &\longmapsto \frac{z - i}{z + i} \end{aligned}$$

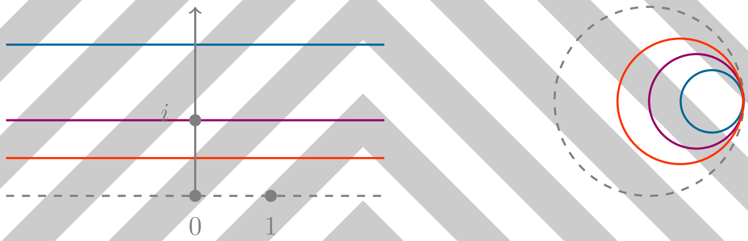

Figure A.9.: The Cayley transform

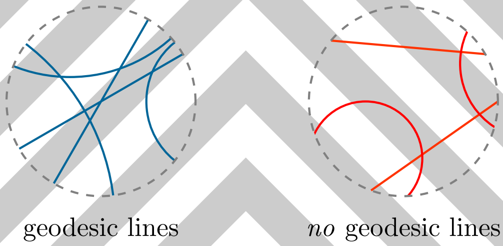

Figure A.10.: Geodesic lines in the Poincaré disk model

is a smooth diffeomorphism between H and $E \subset \mathbb{C}$ (Figure A.9). Translating the Riemannian metric g_H on H via C to E leads to the Riemannian metric g_E on E described by the term

$$\frac{4 \cdot (dx^2 + dy^2)}{(1 - (x^2 + y^2))^2}$$

Let d_E be the metric on E obtained from lengths of curves with respect to g_E. By construction, $C \colon H \longrightarrow E$ is an isometry with respect to the metrics d_H and d_E. In particular, (E, d_E) is a model of the hyperbolic plane, the *Poincaré disk model*.

The Poincaré disk model has the following properties:

- The Cayley transform is conformal. Therefore, also the Poincaré disk model is a conformal model of the hyperbolic plane; i.e., hyperbolic angles drawn in the disk model coincide with the Euclidean angles.
- Images of geodesic lines in the Poincaré disk model are the diameters of E and those circle arcs that intersect the boundary $\{z \in \mathbb{C} \mid |z| = 1\}$ orthogonally (Figure A.10).

Impressive illustrations of the hyperbolic plane in the Poincaré disk model are the tilings by congruent polygons in Escher's woodcuts *Cirkellimiet I* and *Cirkellimiet IV* [59].

A.4 An invitation to programming

An instructive way to experience and comprehend a mathematical field is to try to feed that subject into a computer – e.g., for illustrative or computational purposes. The following is a random selection of programming tasks encouraging the reader to get more involved with this aspect of geometric group theory. Happy hacking!

Programming task A.4.1 (Verified group theory).

1. Choose a proof assistant (such as Coq [42]).
2. Model the basic terminology of group theory in this proof assistant (or find and understand a library doing this).
3. Prove basic facts of group theory in the proof assistant, e.g., uniqueness of the neutral element (or find and understand a library doing this).

Programming task A.4.2 (Abelianisation).

1. Choose a programming language (such as Haskell [80]) or a computer algebra system (such as SageMath [153]).
2. Model finite presentations of groups in this language.
3. Model the standard isomorphism types of finitely generated Abelian groups in this language.
4. Implement abelianisation (Exercise 2.E.18) in this setup.
   Hints. This will also require the computation of Smith normal forms.

Programming task A.4.3 (Girth in free groups).

1. Choose a programming language or a computer algebra system.
2. Model (finitely generated) free groups in this language.
3. Implement a function that generates the (potentially infinite) list of free generating sets of a given finitely generated free group.
   Hints. The automorphism groups of free groups are generated by the so-called Nielsen transformations.
4. Model graphs in this language.
5. Implement a function that, given a finitely generated free group F and a finite generating set S of this free group, computes the girth of $\mathrm{Cay}(F, S)$.
   Hints. The girth is infinite if and only if the generating set is free. Simultaneously search for the generating set in the list of free generating sets and (try to) list short loops.

Programming task A.4.4 (Regular graphs of large girth).

1. Choose a programming language or a computer algebra system.
2. Model graphs in this language.
3. Construct regular graphs of large girth as in the proof of Theorem 4.4.6.

Programming task A.4.5 (Regular polyhedra).

1. Choose a 3D-modelling language (such as OpenSCAD [136] for 3D-printing or Povray [145] for nice rendering).
2. Create 3D-models of regular polyhedra (both convex regular polyhedra and non-convex regular polyhedra) by first modelling a single face and then letting the isometry group act on this base face.
   Hints. Figure A.11 displays an octahedron, rendered in Povray.

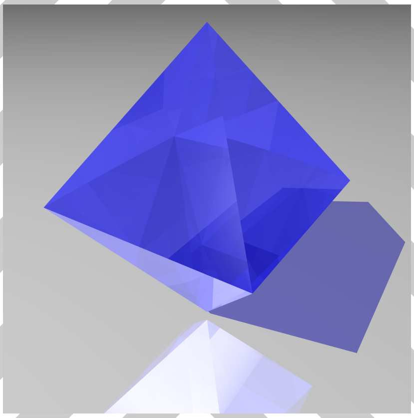

Figure A.11.: Octahedron

Programming task A.4.6 (Hyperbolic tilings).

1. Choose a programming language or a computer algebra system.
2. Model points and geodesics in the Poincaré disk model in this language.
3. Model/implement hyperbolic reflections in this setup. How can such reflections be evaluated on points and geodesics?
4. Implement a function that, given $n \in \mathbb{N}_{\geq 3}$ and an angle $\alpha \in (0, \pi - \pi/n)$, constructs a regular hyperbolic n-gon with vertex angle α in this setup.
5. Implement a function that, given numbers $n, k \in \mathbb{N}_{\geq 3}$ with $1/n + 1/k \leq 1/2$, generates the (infinite) list of all geodesics in a tiling of the hyperbolic plane by regular n-gons such that at every vertex exactly k of these n-gons meet.
   Hints. Let the symmetry group of the tiling act on a base polygon.
6. Implement a function that, given numbers $n, k \in \mathbb{N}_{\geq 3}$ with $1/n + 1/k \leq 1/2$, generates a graphical representation (e.g., via the tikz-package [169]) of (a sufficiently large portion of) a tiling of the hyperbolic plane by regular n-gons such that at every vertex exactly k of these n-gons meet.
   Hints. Figure A.12 displays examples of such tilings.

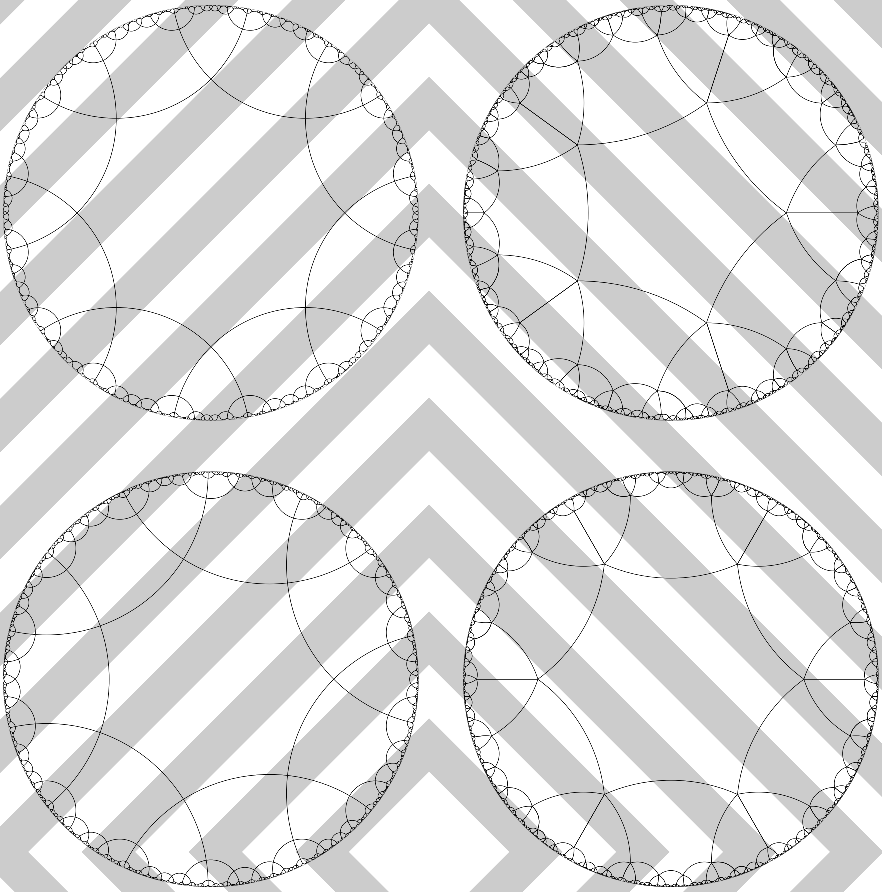

Figure A.12.: Examples of regular tilings of the hyperbolic plane

Programming task A.4.7 (The first Grigorchuk group)**.**

1. Choose a programming language or a computer algebra system.
2. Model the rooted binary tree in this language.
3. Model the generators a, b, c, d of the Grigorchuk group Gri (Definition 4.E.2) in this setup; i.e., it should be possible to evaluate these elements on words in $\{0,1\}^*$.
4. Implement a function that, given an element (as a word in the standard generators and their inverses) of the Grigorchuk group Gri and a word in $\{0,1\}^*$, computes the result of applying this group element to the given word.
5. Implement the child homomorphism of the Grigorchuk group.

Programming task A.4.8 (Verified quasi-geometry).

1. Choose a proof assistant.
2. Model the basic terminology of quasi-geometry in this proof assistant (such as metric spaces, isometric embeddings, bilipschitz embeddings, quasi-isometric embeddings, ...).
3. Prove basic facts of quasi-geometry in the proof assistant, e.g., inheritance properties of quasi-isometric embeddings or independence of the bilipschitz equivalence type of the choice of finite generating sets of finitely generated groups.
4. Model hyperbolic groups in this setup and prove basic facts of hyperbolic groups in this setup.

Programming task A.4.9 (Growth functions).

1. Choose a programming language or a computer algebra system.
2. Model finite presentations of groups in this language.
3. Use a suitable concept in this programming language to model what it means that the word problem for a given finite presentation is solvable.
4. Write a function that, given a finite presentation $\langle S \mid R \rangle$ with solvable word problem and a radius $n \in \mathbb{N}$, computes the value $\beta_{\langle S \mid R \rangle, S}(n)$ of the corresponding growth function.

Programming task A.4.10 (Dehn's algorithm).

1. Choose a programming language or a computer algebra system.
2. Model finite presentations of groups in this language.
3. Model Dehn presentations in this setup.
4. Implement Dehn's algorithm (Proposition 7.4.7) in this setup.

Programming task A.4.11 (The word problem in residually finite groups).

1. Choose a programming language or a computer algebra system.
2. Model finite presentations of groups in this language.
3. Implement a function that, given a finite presentation $\langle S \mid R \rangle$, generates the (potentially infinite) list of all words/elements of $\langle R \rangle^{\triangleleft}_{\langle S \mid R \rangle}$.
4. Model finite symmetric groups in this language.
5. Model group homomorphisms between finitely generated groups in this language. How can such a group homomorphism be evaluated on group elements?
6. Implement a function that, given a finite presentation $\langle S \mid R \rangle$ and $n \in \mathbb{N}$, generates the list of all group homomorphisms $\langle S \mid R \rangle \longrightarrow S_n$.
   Hints. Generators and relations satisfy a universal property!
7. Implement a function that, given a finite presentation $\langle S \mid R \rangle$ of a residually finite group (Definition 4.E.1) and a word $w \in (S \cup S^{-1})^*$, determines whether the group element of $\langle S \mid R \rangle$ represented by w is trivial or not.
   Hints. Simultaneously search for w in $\langle R \rangle^{\triangleleft}_{\langle S \mid R \rangle}$ and for non-triviality of w through a homomorphism to a finite symmetric group.

Bibliography

[1] Ian Agol. The virtual Haken conjecture, with an appendix by Ian Agol, Daniel Groves, and Jason Manning, *Doc. Math.*, 18, pp. 1045–1087, 2013. Cited on page: 249

[2] Alfred V. Aho, Ravi Sethi, Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*, Addison Wesley, 1986. Cited on page: 46

[3] Martin Aigner. *Diskrete Mathematik*, fifth edition, Vieweg Verlag, 2004. Cited on page: 83

[4] Noga Alon. The chromatic number of random Cayley graphs, *European J. Combin.*, 34, pp. 1232–1243, 2013. Cited on page: 74

[5] Claire Anantharaman-Delaroche, Jean N. Renault. *Amenable Groupoids*, with a foreword by Georges Skandalis and Appendix B by E. Germain, volume 36 of *Monographies de L'Enseignement Mathématique*, 2000. Cited on page: 313

[6] Mark A. Armstrong. *Groups and Symmetry*, Undergraduate Texts in Mathematics, Springer, 1988. Cited on page: 77, 83

[7] Matthias Aschenbrenner, Stefan Friedl, Henry Wilton. *3-Manifold Groups*, *EMS Series of Lectures in Mathematics*, EMS, 2015. Cited on page: 248, 249

[8] Tim Austin. Rational group ring elements with kernels having irrational dimension, *Proc. Lond. Math. Soc*, 107(6), pp. 1424–1448, 2013. Cited on page: 31

[9] Australian Academy of Science. *Bernhard Hermann Neumann 1909–2002*, http://sciencearchive.org.au/fellows/memoirs/neumann-b.html Cited on page: 38

[10] Australian Academy of Science. *Professor Bernhard Neumann (1909–2002), Mathematician*, http://www.sciencearchive.org.au/scientists/interviews/n/bn.html Cited on page: 38

[11] László Babai. Chromatic number and subgraphs of Cayley graphs. *Theory and applications of graphs (Proc. Internat. Conf., Western Mich. Univ., Kalamazoo, Mich., 1976)*, volume 642 of *Springer Lecture Notes in Mathematics*, pp. 10–22, Springer, 1978. Cited on page: 73

[12] 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: 224

[13] Arthur Bartels, Wolfgang Lück, Shmuel Weinberger. On hyperbolic groups with spheres as boundary, *J. Differential Geom.*, 86(1), pp. 1–16, 2010. Cited on page: 271

[14] Hyman Bass. The degree of polynomial growth of finitely generated nilpotent groups, *Proc. London Math. Soc.*, 25, pp. 603–614, 1972. Cited on page: 182

[15] Gilbert Baumslag. Wreath products and finitely presented groups, *Math. Z.*, 75, pp. 22–28, 1961. Cited on page: 30

[16] Gilbert Baumslag, Donald Solitar. Some two-generator one-relator non-Hopfian groups, *Bull. Amer. Math. Soc.*, 68, pp. 199–201, 1962. Cited on page: 28

[17] Bachir Bekka, Pierre de la Harpe, Alain Valette. *Kazhdan's Property (T)*, volume 11 of *New Mathematical Monographs*, Cambridge University Press, 2009. Cited on page: 295, 298

[18] Riccardo Benedetti, Carlo Petronio. *Lectures on Hyperbolic Geometry*, *Universitext*, Springer, 1992. Cited on page: 141, 239, 269, 270, 278

[19] 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: 155, 302, 303, 328

[20] Béla Bollobás. *Random graphs*, second edition, volume 73 of *Cambridge Studies in Advanced Mathematics*, Cambridge University Press, 2001. Cited on page: 100

[21] Enrico Bombieri, Walter Gubler. *Heights in Diophantine Geometry*, volume 4 of *New Mathematical Monographs*, Cambridge University Press, 2007. Cited on page: 192

[22] George S. Boolos, John P. Burgess, Richard C. Jeffrey. *Computability and logic*, fifth edition, Cambridge University Press, 2007. Cited on page: 224

[23] Brian H. Bowditch. Continuously many quasiisometry classes of 2-generator groups, *Comment. Math. Helv.*, 73, pp. 232–236, 1998. Cited on page: 149

[24] Brian H. Bowditch. Uniform hyperbolicity of the curve graphs, *Pacific J. Math.*, 269(2), pp. 269–280, 2014. Cited on page: 253

[25] Emmanuel Breuillard. Diophantine Geometry and uniform growth of finite and infinite groups, volume III of *Proceedings of the International Congress of Mathematicians, Seoul 2014*, pp. 27–50, 2014. Cited on page: 188, 190, 191, 192, 193

[26] Emmanuel Breuillard, Ben Green, Terence Tao. The structure of approximate groups, *Publ. Math. Inst. Hautes Études Sci.*, 116, pp. 115–221. 2012. Cited on page: 179

[27] Emmanuel Breuillard, Péter P. Varjú. Entropy of Bernoulli convolutions and uniform exponential growth for linear groups, preprint, arXiv:1510.04043 [math.CA], 2015. Cited on page: 188, 193

[28] Douglas S. Bridges. *Computability. A mathematical sketchbook*, volume 146 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 224

[29] Martin R. Bridson. Non-positive curvature and complexity for finitely presented groups, volume 2 of *International Congress of Mathematicians, Madrid 2006*, pp. 961–987, EMS, 2006. Cited on page: 2, 3, 5

[30] 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

[31] Martin R. Bridson, André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*, Springer, 1999. Cited on page: 5, 60, 150, 186, 209, 229, 248, 249, 254, 260, 263, 265, 266, 270

[32] Martin R. Bridson, Stephen M. Gersten. The optimal isoperimetric inequality for torus bundles over the circle. *Quart. J. Math. Oxford Ser. (2)*, 47/185, pp. 1–23, 1996. Cited on page: 186

[33] Matthew G. Brin, Craig G. Squier. Groups of piecewise linear homeomorphisms of the real line, *Invent. Math.*, 79(3), pp. 485–498, 1985. Cited on page: 294

[34] Kenneth S. Brown. *Cohomology of Groups*, volume 87 of *Graduate Texts in Mathematics*, Springer, 1982. Cited on page: 33, 60, 139, 324, 325, 327

[35] Kenneth S. Brown. *Buildings*, reprint of the 1989 original, *Springer Monographs in Mathematics*, Springer, 1998. Cited on page: 223

[36] Ulrich Bunke, Alexander Engel. Homotopy theory with bornological coarse spaces, preprint, arXiv:1607.03657 [math.AT], 2016. Cited on page: 152

[37] Dmitri Burago, Bruce Kleiner. Separated nets in Euclidean space and Jacobians of bi-Lipschitz maps, *Geom. Funct. Anal.*, 8(2), pp. 273–282, 1998. Cited on page: 312

[38] 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: 28

[39] Tullio Ceccherini-Silberstein, Michel Coornaert. *Cellular Automata and Groups*, *Springer Monographs in Mathematics*, Springer, 2010. Cited on page: 111, 298, 300, 305, 306

[40] Ian Chiswell. *A Course in Formal Languages, Automata and Groups*, *Universitext*, Springer, 2009. Cited on page: 46

[41] Matt Clay, Dan Margalit (eds.). *Office Hours with a Geometric Group Theorist*, Princeton University Press, 2017. Cited on page: 5

[42] The Coq Proof Assistant. https://coq.inria.fr/ Cited on page: 349

[43] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein. *Introduction to Algorithms*, third edition, MIT Press, 2009. Cited on page: 54

[44] Yves de Cornulier, Romain Tessera, Alain Valette. Isometric group actions on Hilbert spaces: growth of cocycles. *Geom. Funct. Anal.*, 17(3), pp. 770–792, 2007. Cited on page: 186

[45] 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: 44, 60, 86, 223, 248

[46] Francesca Diana. *Aspects of Uniformly Finite Homology*, PhD thesis, Universität Regensburg, 2015.
http://nbn-resolving.de/urn:nbn:de:bvb:355-epub-312248
Cited on page: 328

[47] Warren Dicks. Simplified Mineyev, preprint, 2011.
Available online at http://mat.uab.es/∼dicks/SimplifiedMineyev.pdf
Cited on page: 94

[48] Tammo tom Dieck. *Algebraic Topology*, *EMS Textbooks in Mathematics*, EMS, 2000. Cited on page: 153, 278, 322

[49] Reinhard Diestel. *Graph theory*, third edition, volume 173 of *Graduate Texts in Mathematics*, Springer, 2005. Cited on page: 54, 306

[50] Albrecht Dold. *Lectures on Algebraic Topology*, reprint of the 1972 edition, *Classics in Mathematics*, Springer, 1995. Cited on page: 178, 179, 270

[51] Alexander Dranishnikov. On macroscopic dimension of rationally essential manifolds, *Geom. Topol.*, 15(2), pp. 1107–1124, 2011. Cited on page: 155

[52] 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: 179, 182, 183, 184

[53] Cornelia Druţu, Michael Kapovich. *Geometric Group Theory*, with an appendix by Bogdan Nica, volume 63 of *Colloquium Publications*, AMS, 2017 (to appear). Cited on page: 5, 102, 150, 179, 182, 183, 267, 279

[54] 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: 183

[55] Tullia Dymarz. Bijective quasi-isometries of amenable groups, *Geometric Methods in Group Theory*, volume 372 of *Contemp. Math.*, pp. 181–188, AMS, 2005. Cited on page: 305, 308

[56] 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: 155, 305, 308

[57] Patrick B. Eberlein. *Geometry of Nonpositively Curved Manifolds*, University of Chicago Press, 1996. Cited on page: 236

[58] Alexander Engel. Wrong way maps in uniformly finite homology and homology of groups, *J. Homotopy Relat. Struct.*, online first, 2017.
DOI https://doi.org/10.1007/s40062-017-0187-x Cited on page: 155

[59] Maurits C. Escher. *Cirkellimiet IV*, 1960. Available online at
http://www.mcescher.nl/galerij/erkenning-succes/cirkellimiet-iv/ Cited on page: 223, 348

[60] Alex Eskin, Shahar Mozes, Hee Oh. On uniform exponential growth for linear groups, *Invent. Math.*, 160(1), pp. 1–30, 2005. Cited on page: 190

[61] John Franks. Anosov diffeomorphisms, *1970 Global Analysis*, volume XIV of *Proc. Sympos. Pure Math. (Berkeley, 1968)*, pp. 61–93, AMS. Cited on page: 188

[62] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture: with an appendix by Warren Dicks, volume 1100 of *Memoirs of the AMS*, AMS, 2015. Cited on page: 94

[63] Alex Furman. A survey of measured group theory, in *Geometry, Rigidity, and Group Actions* (edited by B. Farb and D. Fisher), Chicago University Press, 2011. Cited on page: 148, 279

[64] Damien Gaboriau. Invariants l^2 de relations d'équivalence et de groupes, *Publ. Math. Inst. Hautes Études Sci.*, 95, pp. 93–150, 2002. Cited on page: 328

[65] Damien Gaboriau. A measurable-group-theoretic solution to von Neumann's problem, *Invent. Math.*, 177(3), pp. 533–540, 2009. Cited on page: 294

[66] Ross Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*, Springer, 2008. Cited on page: 92, 260, 267, 328

[67] Steve M. Gersten. Quasi-isometry invariance of cohomological dimension, *C. R. Acad. Sci. Paris Sér. I Math.*, 316(5), pp. 411–416, 1993. Cited on page: 144

[68] Yair Glasner, Nicolas Monod. Amenable actions, free products and a fixed point property, *Bull. Lond. Math. Soc.*, 39(1), pp. 138–150, 2007. Cited on page: 312, 313, 314

[69] 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: 111, 114, 149, 175, 293

[70] Rostislav I. Grigorchuk, Pierre de la Harpe. Limit behaviour of exponential growth rates for finitely generated groups, *Essays on Geometry and Related Topics*, *Enseignement Math.*, 38, pp. 351–370, 2001. Cited on page: 190

[71] 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: 246

[72] Mikhael Gromov. Groups of polynomial growth and expanding maps, with an Appendix by Jacques Tits, *Publ. Math. Inst. Hautes Études Sci.*, 53, pp. 53–78, 1981. Cited on page: 179, 182, 183, 187

[73] Michael Gromov. Volume and bounded cohomology. *Publ. Math. Inst. Hautes Études Sci.*, 56, pp. 5–99, 1983. Cited on page: 303

[74] 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, 208, 220, 225

[75] Misha Gromov. Random walk in random groups, *Geom. Funct. Anal.*, 13(1), pp. 73–146, 2003. Cited on page: 225

[76] Philip Hall. Finiteness conditions for soluble groups, *Proc. London Math. Soc.*, 4, pp. 419–436, 1954. Cited on page: 30, 31

[77] Pierre de la Harpe. *Topics in Geometric Group Theory*, Chicago University Press, 2000. Cited on page: 5, 30, 42, 97, 107, 114, 139, 175, 177, 195, 248, 305

[78] Pierre de la Harpe. Uniform growth in groups of exponential growth. *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part II (Haifa, 2000)*, *Geom. Dedicata*, 95, pp. 1–17, 2002. Cited on page: 190

[79] John M. Harris, Jeffrey L. Hirst, Michael J. Mossinghoff. *Combinatorics and Graph Theory*, second edition, *Undergraduate Texts in Mathematics, Springer*, 2008. Cited on page: 54

[80] Haskell, an advanced, purely functional programming language. https://www.haskell.org/ Cited on page: 349

[81] Allen Hatcher. *Algebraic Topology*, Cambridge University Press, 2002. http://www.math.cornell.edu/~hatcher/AT/ATpage.html Cited on page: 60, 322

[82] Horst Herrlich. *Axiom of Choice*, volume 1876 of *Springer Lecture Notes in Mathematics*, Springer, 2006. Cited on page: 302

[83] 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 *London Math. Soc. Lecture Note Ser.*, Cambridge University Press, 1995. Cited on page: 152

[84] Nikolai V. Ivanov. Foundations of the theory of bounded cohomology. *J. Soviet Math.*, 37, pp. 1090–1114, 1987. Cited on page: 303

[85] Tadeusz Januszkiewicz, Jacek Świątkowski. Simplicial Nonpositive Curvature, *Publ. Math. Inst. Hautes Études Sci.*, 104(1), pp. 1–85, 2006. Cited on page: 249

[86] Barry E. Johnson. *Cohomology in Banach Algebras*, volume 127 of *Memoirs of the AMS*, AMS, 1972. Cited on page: 303

[87] Ilya Kapovich, Nadia Benakli. Boundaries of hyperbolic groups, *Combinatorial and Geometric Group Theory (New York, 2000)*, pp. 39–93, volume 296 of *Contemp. Math.*, AMS, 2002. Cited on page: 269, 271

[88] Christian Kassel, Vladimir Turaev. *Braid Groups*, volume 247 of *Graduate Texts in Mathematics*, Springer, 2008. Cited on page: 45

[89] John L. Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*, Springer, 1975. Cited on page: 146

[90] Bruce Kleiner. A new proof of Gromov's theorem on groups of polynomial growth, *J. Amer. Math. Soc.*, 23(3), pp. 815–829, 2010. Cited on page: 179

[91] Johannes Köbler, Uwe Schöning, Jacobo Torán. *Graph Isomorphism Problem: The Structural Complexity*, Birkhäuser, 1993. Cited on page: 55

[92] Dieter Kotschick, Clara Löh. Fundamental classes not representable by products, *J. Lond. Math. Soc.*, 79(3), 545–561, 2009. Cited on page: 246

[93] Dieter Kotschick, Clara Löh. Groups not presentable by products, *Groups Geom. Dyn.*, 7(1), pp. 181–204, 2013. Cited on page: 246

[94] Serge Lang. *Algebra*, revised third edition, volume 211 of *Graduate Texts in Mathematics*, Springer, 2002. Cited on page: 10, 180

[95] Serge Lang. *Real and Functional Analysis*, volume 142 of *Graduate Texts in Mathematics*, Springer, 1993. Cited on page: 291

[96] John M. Lee. *Riemannian Manifolds. An Introduction to Curvature*, volume 176 of *Graduate Texts in Mathematics*, Springer, 1997. Cited on page: 5, 139, 140, 141, 204, 206, 207, 346

[97] Cai Heng Li. On isomorphisms of finite Cayley graphs—a survey, *Discrete Math.*, 256(1–2), pp. 301–334, 2002. Cited on page: 59

[98] Xin Li. Quasi-isometry, Kakutani equivalence, and applications to cohomology, preprint, arXiv:1604.07375 [math.GR], 2016 Cited on page: 144, 146, 328

[99] Clara Löh. Isomorphisms in ℓ^1-homology, *Münster J. Math.*, 1, pp. 237–266, 2008. Cited on page: 303

[100] Clara Löh. Simplicial volume, *Bull. Man. Atl.*, pp. 7–18, 2011. Cited on page: 179, 303

[101] 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: 33, 60, 179

[102] Clara Löh. Which finitely generated Abelian groups admit isomorphic Cayley graphs? *Geom. Dedicata*, 164(1), pp. 97–111, 2013. Cited on page: 59, 126

[103] Clara Löh, Matthias Mann. Which finitely generated Abelian groups admit equal growth functions? preprint, arXiv: 1309.3381 [math.GR], 2013. Cited on page: 175

[104] Wolfgang Lück. *L^2-Invariants: Theory and Applications to Geometry and K-Theory*, volume 44 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*, 3. Folge, Springer, 2002. Cited on page: 303

[105] Wolfgang Lück. *L^2-Invariants from the algebraic point of view*, in *Geometric and Cohomological Methods in Group Theory*, volume 358 of *London Math. Soc. Lecture Note Ser.*, pp. 63–161, Cambridge University Press, 2009. Cited on page: 315

[106] Wolfgang Lück. Aspherical Manifolds, *Bull. Man. Atl.*, pp. 1-17, 2012. Cited on page: 279

[107] Roger C. Lyndon, Paul E. Schupp. *Combinatorial Group Theory*, reprint of the 1977 edition, *Classics in Mathematics*, Springer, 2001. Cited on page: 37, 38

[108] Russell Lyons, Yuval Peres. *Probability on Trees and Networks*, Cambridge University Press, 2016. Cited on page: 60

[109] Saunders MacLane. *Categories for the Working Mathematician*, second edition, volume 5 of *Graduate Texts in Mathematics*, Springer, 1998. Cited on page: 153

[110] Wilhelm Magnus, Abraham Karrass, Donald Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*, Interscience Publishers, 1966. Cited on page: 31, 181

[111] Avinoam Mann. *How groups grow*, volume 395 of *London Math. Soc. Lecture Note Ser.*, Cambridge University Press, 2012. Cited on page: 179

[112] Michał Marcinkowski, Piotr W. Nowak. Aperiodic tilings of manifolds of intermediate growth, *Groups Geom. Dyn.*, 8(2), pp. 479–483, 2014. Cited on page: 155

[113] Grigorii A. Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1), pp. 71–78, 1982. Cited on page: 100

[114] Andrei A. Markov. Insolubility of the problem of homeomorphy, (Russian) *Proc. Internat. Congress Math. 1958*, pp. 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: 29

[115] William S. Massey. *Algebraic Topology: An Introduction*, volume 56 of *Graduate Texts in Mathematics*, Springer, 1977. Cited on page: 5, 24, 35, 78, 80, 91, 93, 136, 140, 178, 179, 249, 322, 324

[116] Howard Masur, Saul Schleimer. The geometry of the disk complex, *J. Amer. Math. Soc.*, 26(1), pp. 1–62, 2013. Cited on page: 253

[117] John Meier. *Groups, Graphs and Trees, An Introduction to the Geometry of Infinite Groups*, volume 73 of *London Mathematical Society Student Texts*, Cambridge University Press, 2008. Cited on page: 46

[118] Kostya Medynets, Roman Sauer, Andreas Thom. Cantor systems and quasi-isometry of groups, preprint, arXiv:1508.07578 [math.DS], 2015. Cited on page: 146

[119] Donald W. Miller. On a theorem of Hölder, *The American Mathematical Monthly*, 65(4), pp. 252–254, 1958. Cited on page: 18

[120] John Milnor. A note on curvature and the fundamental group, *Journal of Differential Geometry*, 2, pp. 1–7, 1968. Cited on page: 176, 182

[121] Igor Mineyev. Submultiplicativity and the Hanna Neumann conjecture, *Ann. of Math.*, 175(1), pp. 393–414, 2012. Cited on page: 94

[122] Igor Mineyev. Groups, graphs, and the Hanna Neumann conjecture, *J. Topol. Anal.*, 4(1), pp. 1–12, 2012. Cited on page: 94

[123] Nicolas Monod. *Continuous Bounded Cohomology of Locally Compact Groups*, volume 1758 of *Lecture Notes in Mathematics*, Springer, 2001. Cited on page: 303

[124] Nicolas Monod. An invitation to bounded cohomology, volume 2 of *International Congress of Mathematicians*, pp. 1183–1211, EMS, 2006. Cited on page: 303

[125] Deane Montgomery, Leo Zippin. Small subgroups of finite-dimensional groups, *Ann. of Math.*, 56, pp. 213–241, 1952. Cited on page: 183

[126] Hans J. Munkholm. Simplices of maximal volume in hyperbolic space, Gromov's norm, and Gromov's proof of Mostow's rigidity theorem (following Thurston). In *Topology Symposium, Siegen 1979*, volume 788 of *Lecture Notes in Mathematics*, pp. 109–124. Springer, 1980. Cited on page: 278

[127] Jan Mycielski. Sur le coloriage des graphs. *Colloq. Math.*, 3, pp. 161–162, 1955. Cited on page: 100

[128] Jürgen Neukirch, Alexander Schmidt, Kay Wingberg. *Cohomology of Number Fields*, second edition, volume 323 of *Grundlehren der Mathematischen Wissenschaften*, Springer, 2008. Cited on page: 327

[129] Bernhard H. Neumann, Hanna Neumann, Peter Neumann. Wreath products and varieties of groups, *Math. Z.*, 80(1), pp. 44–62, 1962. Cited on page: 38

[130] John von Neumann. Zur allgemeinen Theorie des Maßes, *Fund. Math.*, 13(1), pp. 73–111, 1929. Cited on page: 294

[131] 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: 303

[132] Piotr W. Nowak, Guoliang Yu. *Large Scale Geometry, EMS Textbooks in Mathematics*, EMS, 2012. Cited on page: 152, 155

[133] Yann Ollivier. Sharp phase transition theorems for hyperbolicity of random groups, *Geom. Funct. Anal.*, 14(3), pp. 595–679, 2004. Cited on page: 225

[134] 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: 294

[135] Alexander Yu. Olshanskiĭ. Almost every group is hyperbolic, *Internat. J. Algebra Comput.*, 2(1), pp. 1–17, 1992. Cited on page: 225

[136] OpenSCAD, the Programmers Solid 3D Modeller. http://www.openscad.org/ Cited on page: 350

[137] Denis V. Osin. Algebraic entropy of elementary amenable groups *Geom. Dedicata*, 107, pp. 133–151, 2004. Cited on page: 190

[138] Denis V. Osin. *Relatively Hyperbolic Groups: Intrinsic Geometry, Algebraic Properties, and Algorithmic Problems*, volume 843 of *Memoirs of the AMS*, 2006. Cited on page: 221

[139] Denis V. Osin. Acylindrically hyperbolic groups, *Trans. Amer. Math. Soc.*, 368(2), pp. 851–888, 2016. Cited on page: 276, 277

[140] 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: 38

[141] Narutaka Ozawa. A functional analysis proof of Gromov's polynomial growth theorem, preprint, arXiv:1510.04223 [math.GR], 2015. Cited on page: 179

[142] Athanase Papadopoulos. *Metric Spaces, Convexity and Nonpositive Curvature*, volume 6 of *IRMA Lectures in Mathematics and Theoretical Physics*, EMS, 2005. Cited on page: 331

[143] Panos Papasoglu, Kevin Whyte. Quasi-isometries between groups with infinitely many ends, *Comment. Math. Helv.*, 77(1), pp. 133–144, 2002. Cited on page: 308

[144] Alan L.T. Paterson. *Amenability*, volume 29 of *Mathematical Surveys and Monographs*, AMS, 1988. Cited on page: 290, 295, 298

[145] Povray, the Persistence of Vision Raytracer. http://www.povray.org/ Cited on page: 350

[146] John G. Ratcliffe. *Foundations of Hyperbolic Manifolds*, volume 149 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 179, 278, 347

[147] Willi Rinow. *Die Innere Geometrie der Metrischen Räume*, volume 105 of *Die Grundlehren der mathematischen Wissenschaften*, Springer, 1961. Cited on page: 331

[148] John Roe. *Index Theory, Coarse Geometry, and Topology of Manifolds*, volume 90 of *CBMS Regional Conf. Ser. in Math.*, AMS, 1996. Cited on page: 152

[149] Joseph M. Rosenblatt. A generalization of Følner's condition, *Math. Scand.*, 33, pp. 153–170, 1973. Cited on page: 313

[150] Joseph J. Rotman. *An Introduction to the Theory of Groups*, fourth edition, volume 148 of *Graduate Texts in Mathematics*, Springer, 1999. Cited on page: 10, 29, 37, 64, 83, 224

[151] Walter Rudin. *Functional Analysis*, second edition, *International Series in Pure and Applied Mathematics*, McGraw-Hill, 1991. Cited on page: 298

[152] Volker Runde. *Amenability*, volume 1774 of *Springer Lecture Notes in Mathematics*, Springer, 2002. Cited on page: 295, 302

[153] SageMath, open-source mathematical software system. http://www.sagemath.org/ Cited on page: 349

[154] Parameswaran Sankaran. On homeomorphisms and quasi-isometries of the real line, *Proc. Amer. Math. Soc.*, 134(7), pp. 1875–1880, 2006. Cited on page: 122

[155] 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: 28

[156] 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: 146, 148

[157] Roman Sauer. Homological invariants and quasi-isometry, *Geom. Funct. Anal.*, 16(2), pp. 476–515, 2006. Cited on page: 144

[158] Henry Segerman. *Visualizing Mathematics with 3D Printing*, Johns Hopkins University Press, 2016. Cited on page: 347

[159] Jean-Pierre Serre. *Trees*, corrected second printing of the 1980 English translation, *Springer Monographs in Mathematics*, Springer, 2003. Cited on page: 5, 35, 38, 64, 92, 95, 100

[160] Brandon Seward. Burnside's Problem, spanning trees and tilings, *Geom. Topol.*, 18(1), pp. 179–210, 2014. Cited on page: 294

[161] Yehuda Shalom. Harmonic analysis, cohomology, and the large-scale geometry of amenable groups, *Acta Math.*, 192(2), pp. 119-185, 2004. Cited on page: 144, 186

[162] 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: 179

[163] Michael Shub. Expanding maps. *1970 Global Analysis*, volume XIV of *Proc. Sympos. Pure Math. (Berkeley, 1968)*, pp. 273–276, AMS. Cited on page: 188

[164] Raymond M. Smullyan, Melvin Fitting. *Set Theory and the Continuum Problem*, Dover, 2010. Cited on page: 13, 306

[165] Joel Spencer. What's not inside a Cayley graph, *Combinatorica*, 3, pp. 239–241, 1983. Cited on page: 73

[166] John R. Stallings. On torsion-free groups with infinitely many ends, *Ann. of Math.*, 88, pp. 312–334, 1968. Cited on page: 38

[167] John Stillwell. *Geometry of Surfaces*, *Universitext*, Springer, 1992. Cited on page: 141

[168] Jacek Świątkowski. The dense amalgam of metric compacta and topological characterization of boundaries of free products of groups, *Groups Geom. Dyn.*, 10(1), pp. 407–471, 2016. Cited on page: 271

[169] Till Tantau. pgf – A Portable Graphic Format for TEX, https://ctan.org/tex-archive/graphics/pgf/base Cited on page: 350

[170] Terence Tao. *An Epsilon of Room, I: Real Analysis*, volume 117 of *Graduate Studies in Mathematics*, American Mathematical Society, 2010. Cited on page: 108

[171] Terence Tao. *Hilbert's Fifth Problem and Related Topics*, volume 153 of *Graduate Studies in Mathematics*, American Mathematical Society, 2014. Cited on page: 183

[172] Jacques Tits. Free subgroups in linear groups, *J. Algebra*, 20, pp. 250–270, 1972. Cited on page: 102

[173] 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: 18

[174] 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: 29

[175] Kevin Whyte. Amenability, bi-Lipschitz equivalence, and the von Neumann conjecture, *Duke Math. J.*, 99(1), pp. 93–112, 1999. Cited on page: 155, 294, 302, 303, 308

[176] John S. Wilson. On exponential growth and uniformly exponential growth for groups, *Invent. Math.*, 155(2), pp. 287–303, 2004. Cited on page: 190

[177] Dave Witte Morris, Joy Morris, Gabriel Verret. Isomorphisms of Cayley graphs on nilpotent groups, *New York J. Math.*, 22, pp. 453–467, 2016. Cited on page: 59, 126

[178] Wolfgang Woess. Graphs and groups with tree-like properties, *J. Combin. Theory*, B 47(3), pp. 361–371, 1989. Cited on page: 150

[179] Joseph A. Wolf. Growth of finitely generated solvable groups and curvature of Riemannian manifolds, *J. Differential Geometry*, 2, pp. 421–446, 1968. Cited on page: 181, 182

[180] Robert J. Zimmer. *Ergodic theory and semisimple groups*, volume 81 of *Monographs in Mathematics*, Birkhäuser, 1984. Cited on page: 313

Index of notation

Symbols

A

Index

A

D

E

F

H

N

O

P

S

T

U

V

W