Geometric Group Theory

Sommersemester 2022
Universität Regensburg

Clara Löh

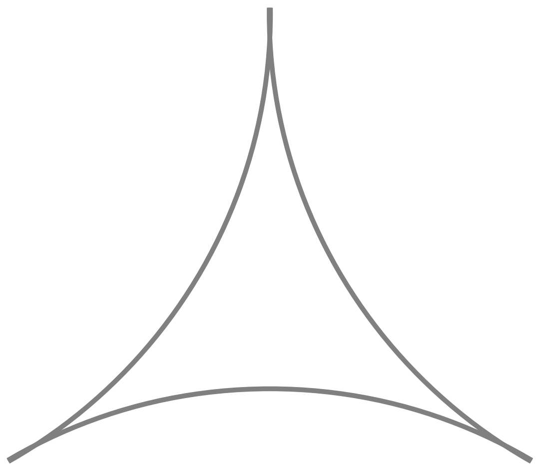

These lecture notes are based on:
C. Löh. Geometric Group Theory. An Introduction. Universitext, Springer, 2017.

Version of July 29, 2022
clara.loeh@mathematik.uni-regensburg.de
https://loeh.app.ur.de
Fakultät für Mathematik, Universität Regensburg, 93040 Regensburg

Contents

Guide to the Literature

This course and these lecture notes are mostly based on the following book (which in turn is based on previous incarnations of this course):

- Clara Löh. *Geometric Group Theory. An Introduction.* Universitext, Springer, 2017.

There are many sources that cover the standard topics (all with their own advantages and disadvantages). Therefore, you should individually compose your own favourite selection of books.

Geometric Group Theory

- Martin R. Bridson, André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*, Springer, 1999.
- Matt Clay, Dan Margalit (eds.). *Office Hours with a Geometric Group Theorist*, Princeton University Press, 2017.
- Cornelia Druţu, Michael Kapovich. *Geometric Group Theory*, with an appendix by Bogdan Nica, volume 63 of *Colloquium Publications*, AMS, 2017.
- Misha Gromov. Hyperbolic groups. in *Essays in Group Theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pp. 75–263, Springer, 1987.
- Pierre de la Harpe. *Topics in Geometric Group Theory*, Chicago University Press, 2000.
- Jean-Pierre Serre. *Trees*, corrected second printing of the 1980 English translation, *Springer Monographs in Mathematics*, Springer, 2003.

Related Fields

- Riccardo Benedetti, Carlo Petronio. *Lectures on Hyperbolic Geometry*, *Universitext*, Springer, 1992.
- Allen Hatcher. *Algebraic Topology*, Cambridge University Press, 2002. http://www.math.cornell.edu/~hatcher/AT/ATpage.html
- John M. Lee. *Introduction to Riemannian manifolds*, second edition, Graduate Texts in Mathematics, 176, Springer, 2018.
- Clara Löh. *Group Cohomology*, lecture notes, Universität Regensburg, 2019.
  https://loeh.app.ur.de/teaching/grouphom_ss19/lecture_notes.pdf
- Clara Löh. *Algebraic Topology*, lecture notes, Universität Regensburg, 2022.
  https://loeh.app.ur.de/teaching/grouphom_ss19/lecture_notes.pdf
- Clara Löh. *Ergodic theoretic methods in group homology. A minicourse on L^2-Betti numbers in group theory*, SpringerBriefs in Mathematics, Springer, 2020.
- William S. Massey. *Algebraic Topology: an Introduction*, seventh edition, Springer, 1989.
- Peter May. *A Concise Course in Algebraic Topology*, University of Chicago Press, 1999.

0

Introduction

This course provides an introduction to geometric group theory. 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

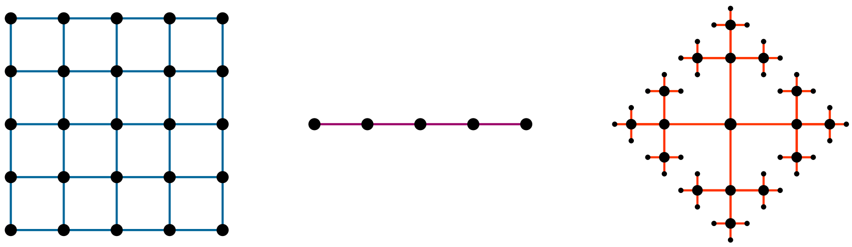

Figure 0.1.: Basic examples of Cayley graphs

(such as the Euclidean plane or the hyperbolic plane) or an action from classical geometric theories.

2. We take geometric invariants and apply these to the geometric objects obtained by the first step. This allows us 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.

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 0.1).

More generally, in (large scale) geometric group-theoretic terms, the universe of (finitely generated) groups roughly unfolds as depicted in Figure 0.2. 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 in the inner regions of the universe:

“Hic abundant leones.” [19]
“A statement that holds for all finitely generated groups
has to be either trivial or wrong.” [attributed to M. Gromov]

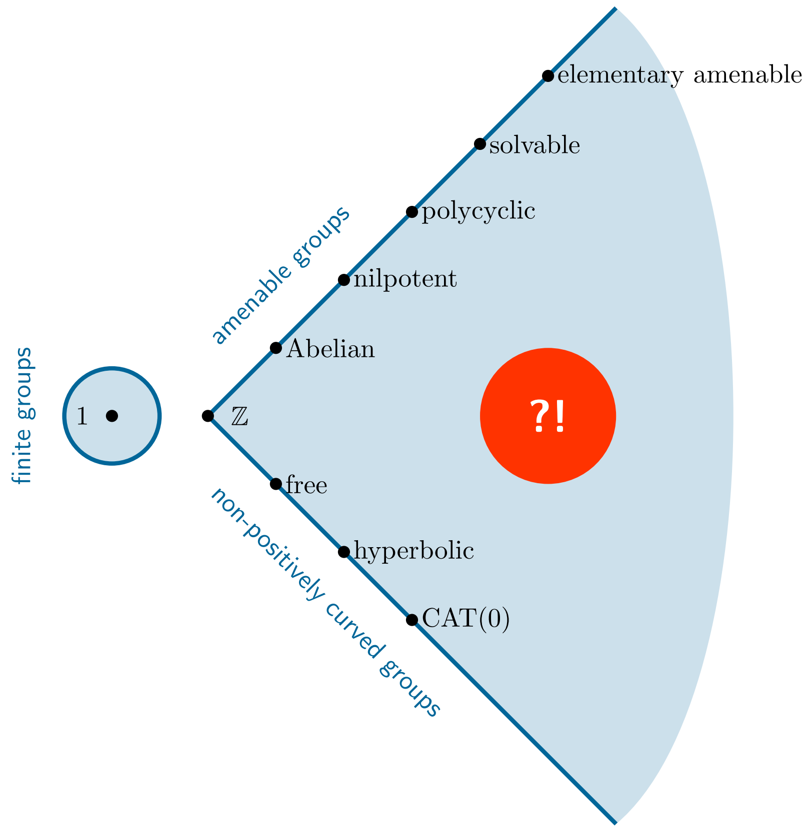

Figure 0.2.: The universe of groups (simplified version of Bridson's universe of groups [19])

Why 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:

- *Recognising that certain matrix groups are free groups;* there is a geometric criterion, the *ping-pong lemma*, that allows us to deduce freeness of a group by looking at a suitable action (not necessarily on a tree).
- *The construction of finite regular graphs with certain connectivity properties;* for example, one can use groups to construct finite regular graphs of high girth or families of expanders.
- *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 similar behaviour to certain matrix groups).

Overview of this Course

The goal of this course is to explain the basic terminology of geometric group theory, the standard proof techniques, and how these concepts can be applied.

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 us to generate interesting groups (Chapter 1).

Then we will introduce one of the main combinatorial objects in geometric group theory, Cayley graphs, and review some basic notions concerning actions of groups (Chapter 2–3). 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 4).

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

- growth types (Chapter 5),
- hyperbolicity (Chapter 6),
- geometry at infinity [63, Chapter 8].

In particular, we will discuss several of the applications discussed above.

Literature exercise. Where in the math library (including electronic resources) can you find books on geometric group theory, group theory, metric geometry, and large-scale geometry?

Convention. The set $\mathbb{N}$ of natural numbers contains 0. All rings are unital and associative. Usually, we assume manifolds to be non-empty (but we might not always mention this explicitly).

1

Generating groups

As the main characters in geometric group theory are groups, we start by reviewing some concepts and examples from group theory. In particular, we will present basic construction principles that allow us 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.

Overview of this chapter.

1.1 Review of the category of groups

1.1.1 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 [54, 91]. The category of groups has groups as objects and group homomorphisms as morphisms.

Definition 1.1.1 (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.,
$$\forall_{g \in G} \quad e \cdot g = g = g \cdot e.$$
(This property uniquely determines the neutral element; check!)
- *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; check!)

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 1.1.2 (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 1.1.3 (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 1.1.4 (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; check!)

- 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 1.1.5 (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 1.1.6 (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 1.1.7 (Isomorphisms via kernel/image)**.** It is a simple exercise in algebra to verify the following (check!):

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

1.1.2 Concrete groups: automorphism groups

The concept, and hence the axiomatisation, of groups originally developed 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 1.1.8 (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 1.1.9 (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 1.1.10 (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 1.1.11 (Isometry groups/Symmetry groups)**.** Let X be a metric space (basic notions for metric spaces are recalled in Chapter 4.1). The set $\operatorname{Isom}(X)$ of all isometries of type $X \longrightarrow X$ forms a group with respect to composition (a subgroup of the symmetric group S_X). For example, in this way the dihedral groups naturally occur as symmetry groups of regular polygons (Example 1.2.20).

Example 1.1.12 (Matrix groups)**.** Let k be a commutative ring with unit, and let V be a k-module. Then the set $\operatorname{Aut}(V)$ of all k-linear isomorphisms $V \longrightarrow V$ forms a group with respect to composition. In particular, the set $\operatorname{GL}(n, k)$ of invertible $n \times n$-matrices over k is a group (with respect to matrix multiplication) for every $n \in \mathbb{N}$. Similarly, $\operatorname{SL}(n, k)$, the subgroup of invertible matrices of determinant 1, is also a group.

Example 1.1.13 (Galois groups)**.** Let $K \subset L$ be a Galois extension of fields. Then the set

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

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

Example 1.1.14 (Deck transformation groups)**.** Let $\pi \colon X \longrightarrow Y$ be a covering map of topological spaces. Then the set

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

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

In 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 $\operatorname{Aut}_C(X)$ 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 1.1.15 (Category)**.** A *category* C consists of the following data:

- 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 [101].)
- 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.

This data is required 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

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

  (The morphism id_X is uniquely determined by this property; check! 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 1.1.16. 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 1.1.17 (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

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

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 1.1.18 (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 1.1.19), the *automorphism group of X in C*.

Proposition 1.1.19 (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

$$\begin{aligned} \circ\colon \mathrm{Mor}_C(X,X) \times \mathrm{Mor}_C(X,X) &\longrightarrow \mathrm{Mor}_C(X,X) \\ (g,h) &\longmapsto g \cdot h. \end{aligned}$$

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

We will now illustrate these terms in more concrete examples:

Example 1.1.20 (Set theory)**.** The category Set of sets consists of:

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

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

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

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 $\mathrm{Aut}_{\mathsf{Set}}(X)$.

Example 1.1.21 (Algebra)**.** The category Group of groups consists of:

- Objects: Let $\mathrm{Ob}(\mathsf{Group})$ be the class of all groups.
- Morphisms: For groups G and H we let $\mathrm{Mor}_{\mathsf{Group}}(G, H)$ 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 1.1.22 (Geometry of isometric embeddings)**.** The category $\mathsf{Met}_{\mathsf{isom}}$ of metric spaces and isometric embeddings consists of:

- Objects: Let $\mathrm{Ob}(\mathsf{Met}_{\mathsf{isom}})$ 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 1.1.23 (Topology)**.** The category Top of topological spaces consists of:

- Objects: Let $\mathrm{Ob}(\mathsf{Top})$ 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, ...)
- ...

1.1.3 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 every vector space by every subspace again canonically 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 1.1.24 (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 $N \lhd G$.

Example 1.1.25 (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 (check!); conversely, every normal subgroup is also the kernel of a certain group homomorphism, namely of the canonical projection to the quotient (Proposition 1.1.26).

Proposition 1.1.26 (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* $\overline{\varphi} \circ \pi = \varphi$*:*

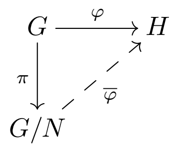

Proof. Ad 1. Suppose that N is normal in G. In this case, the composition map is well-defined (in the sense that the definition does not depend on the choice of the representatives of cosets): Let $g_1,\ g_2,\ g_1',\ g_2' \in G$ with

$$g_1 \cdot N = g_1' \cdot N, \quad \text{and} \quad g_2 \cdot N = g_2' \cdot N.$$

In particular, there are $n_1,\ n_2 \in N$ with $g_1' = g_1 \cdot n_1$ and $g_2' = g_2 \cdot n_2$. Thus we obtain

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

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 is indeed a group with respect to this composition follows easily from the fact that the group axioms are satisfied in G (check!).

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 (check!). □

Example 1.1.27 (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 $\{z \mid z \in \mathbb{C},\ |z| = 1\} \subset \mathbb{C} \setminus \{0\}$ (check!).
- 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 1.1.28 (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

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

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 $\mathrm{Aut}(G)$ (check!), 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 [106]. Curiously, one has for all $n \in \mathbb{N}$ that [79]

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

1.2 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, if we are 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.

1.2.1 Generating sets of groups

We start by reviewing the concept of a generating set of a group; in geometric group theory, we usually are only interested in finitely generated groups (for reasons that will become clear in Chapter 4).

Definition 1.2.1 (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 1.2.2 (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 (check!):

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

Example 1.2.3 (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 not 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).

1.2.2 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 generating sets. 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 1.2.4 (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:

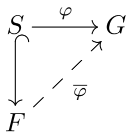

A group is *free* if it contains a free generating set.

Example 1.2.5 (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 $\{3\}$ (check!); 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).

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 1.2.7) that for every set there indeed exists a group freely generated by the given set.

Proposition 1.2.6 (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: Namely, we consider two objects that have the universal property in question. We then proceed as follows:

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

Proof of Proposition 1.2.6. Let F and F' be two groups freely generated by S. We denote the inclusion of S into F and F' by φ and φ' respectively.

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

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

   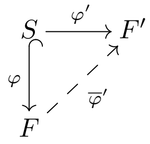

   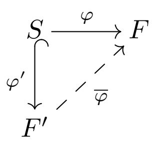

2. We now show that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$ and $\overline{\varphi}' \circ \overline{\varphi} = \mathrm{id}_{F'}$ and hence that φ and φ' are isomorphisms: The composition $\overline{\varphi} \circ \overline{\varphi}' \colon F \longrightarrow F$ is a group homomorphism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

   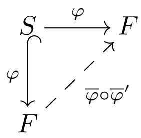

   commutative. Moreover, id_F is also a group homomorphism fitting into this diagram. Because F is freely generated by S, the uniqueness part of the universal property thus tells us that these two homomorphisms have to coincide, i.e., that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$. 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 1.2.7 (Free groups, existence)**.** *Let S be a set. Then there exists a group freely generated by S. (By Proposition 1.2.6, 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., $\widehat{\cdot} : S \longrightarrow \widehat{S}$ 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 the 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 the second step we define

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

where $\sim$ is the equivalence relation generated by

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

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 $\cdot : F(S) \times F(S) \longrightarrow F(S)$ 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

$$I(sx) := I(x)\widehat{s},$$
$$I(\widehat{s}x) := I(x)s$$

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 $i(S)$ via i with 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 2.3.1).

We conclude by collecting some properties of free generating sets in free groups: First of all, free groups are indeed generated (in the sense of Definition 1.2.1) by every free generating set (Corollary 1.2.8); secondly, free generating sets are generating sets of minimal size (Proposition 1.2.9); moreover, finitely generated groups can be characterised as the quotients of finitely generated free groups (Corollary 1.2.12).

Corollary 1.2.8. *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 1.2.7. In view of the uniqueness result Proposition 1.2.6, we find an isomorphism $F(S) \cong F$ that is the identity on S. Hence, it follows that the given free group F is also generated by S. □

Proposition 1.2.9 (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 or through the abelianisation (Exercise). The second part is a consequence of the first part. □

Definition 1.2.10 (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 1.2.11. 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 [77, Chapter VI.8] or via actions on trees (Chapter 3.2.3).

Corollary 1.2.12. *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 1.2.8, 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 $\operatorname{im} \pi = G$. □

1.2.3 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 1.2.13 (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 1.2.14 (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 (check!):

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

Example 1.2.15 (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 1.2.16. 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 1.2.17 (Generators and relations)**.** Let S be a set, let $R \subset (S \cup S^{-1})^*$ be a subset (alternatively: let $R \subset F(S)$); let $F(S)$ be the free group generated by S. Then the group

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

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

If G is a group with $G \cong \langle S \mid R \rangle$, then the pair (S, R) is a *presentation of G*; it is common to abuse notation and also use the symbol $\langle S \mid R \rangle$ to denote the presentation (S, R).

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

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

Proposition 1.2.18 (Universal property of generators and relations). *Let S be a set and let $R \subset (S \cup S^{-1})^*$. The group $\langle S \mid R \rangle$ generated by S with relations R together with the canonical map $\pi \colon S \longrightarrow F(S)/\langle R \rangle^{\triangleleft}_{F(S)} = \langle S \mid R \rangle$ has the following universal property: For every group G and every map $\varphi \colon S \longrightarrow G$ with the property that*

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

there exists precisely one group homomorphism $\overline{\varphi} \colon \langle S \mid R \rangle \longrightarrow G$ such that $\overline{\varphi} \circ \pi = \varphi$; here, $\varphi^* \colon (S \cup S^{-1})^* \longrightarrow G$ is the canonical extension of φ to words over $S \cup S^{-1}$ (as described in the proof of Theorem 1.2.7). Moreover, $\langle S \mid R \rangle$ (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 1.2.4) and of the universal property of quotient groups (Proposition 1.1.26) (check!). ☐

Example 1.2.19 (Presentations of groups).

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

Example 1.2.20 (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*:

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

This can be seen as follows: Let $\sigma \in \operatorname{Isom}(X_n)$ be the rotation about $2\pi/n$ around the centre of the regular n-gon X_n and let τ be the reflection along one of the diameters passing through one of the vertices (Figure 1.1). 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 1.2.18) 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

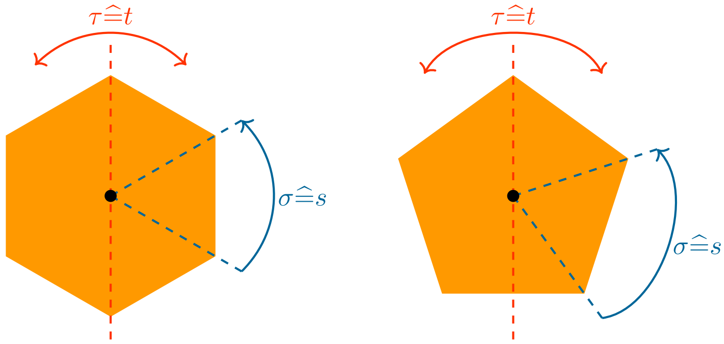

Figure 1.1.: Generators of the dihedral groups D_6 and D_5

by other means: Using the fact that isometries of X_n map (neighbouring) vertices to (neighbouring) vertices (check!), 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 1.2.21 (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

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

(this is discussed in the seminar on Thompson groups). 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 [24].

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 [24]. Moreover, it can be shown that F does not contain subgroups isomorphic to F_2 [24]. However, the question whether F belongs to the class of so-called amenable groups [63, Chapter 9] is a long-standing open problem in geometric group theory with an interesting history [93]: "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 1.2.22 (Baumslag–Solitar groups)**.** For $m, n \in \mathbb{N}_{>0}$ the *Baumslag–Solitar group* $\operatorname{BS}(m,n)$ is defined via the presentation

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

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

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

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

Example 1.2.23 (Complicated trivial group)**.** The group

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

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 1.2.24 (Word problem). The problem of determining 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., one can prove that there is no algorithmic procedure that, given generators and relations, can decide whether the corresponding group is trivial or not [91, 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 6.4 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 [76], and there are far-reaching consequences for the global shape of moduli spaces [107].

1.2.4 Finitely presented groups

Finitely presented groups are groups that admit a finite presentation. We briefly discuss the difference between finite generation and finite presentation.

Definition 1.2.25 (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 finitely presented groups can also be rather complicated.

Outlook 1.2.26 (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. For example, this implies that compact connected manifolds have a finitely presented fundamental group. Conversely, every finitely presented group is the fundamental group of a finite CW-complex (Outlook 2.2.4) 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 1.2.27 (A finitely generated group that is not finitely presented). The group

$$\bigl\langle s, t \bigm| \{[s, t^n s t^{-n}] \mid n \in \mathbb{N}_{>0}\} \bigr\rangle$$

[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.

is finitely generated, but not finitely presented (Exercise). This group is an example of a *lamplighter group* (see also Example 1.3.5).

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 1.2.29):

Theorem 1.2.28 (Uncountably many finitely generated groups). *There exist uncountably many isomorphism classes of groups generated by two elements.*

Before sketching Hall's proof [43, Theorem 7][44, Chapter III.C] of this theorem, we discuss an important consequence:

Corollary 1.2.29. *There are uncountably many isomorphism classes of finitely generated groups that are* not *finitely presented.*

Proof. On the one hand, up to renaming there are only countably many finite presentations of groups; hence, there are only countably many isomorphism types of finitely presented groups.

On the other hand, there are uncountably many finitely generated groups by Theorem 1.2.28. ☐

The proof of Theorem 1.2.28 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 1.2.30).
2. We then show that G even has uncountably many quotient groups that are pairwise non-isomorphic (Proposition 1.2.31).

Proposition 1.2.30 (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

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

Let C be the subgroup of G generated by the set $\{[s, t^n s t^{-n}] \mid n \in \mathbb{Z}\}$. All elements of C are invariant under conjugation with s by the first part of the relations, and they are invariant under conjugation with t by the second part of the relations; thus, C is central in G. Moreover, using the so-called calculus of commutators, it can be shown that C contains the additive group $\bigoplus_{\mathbb{N}} \mathbb{Z}$ [43, p. 434f][73, Corollary 5.12]. □

Proposition 1.2.31 (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 1.2.32 (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 version of a question of Atiyah in the negative) [2].

1.3 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.

1.3.1 Products and extensions

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

Definition 1.3.1 (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:

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

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 via the projections in natural 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 1.3.2 (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 1.3.3 (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

$$(N \rtimes_\varphi Q) \times (N \rtimes_\varphi Q) \longrightarrow (N \rtimes_\varphi Q)$$
$$\bigl((n_1, q_1), (n_2, q_2)\bigr) \longmapsto \bigl(n_1 \cdot \varphi(q_1)(n_2), q_1 \cdot q_2\bigr)$$

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 1.3.4 (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 $\varphi\colon Q \longrightarrow \operatorname{Aut}(N)$ 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

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

are well-defined mutually inverse group homomorphisms, where

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

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 [22, Chapter IV][65, Chapter 1.5.2].

Example 1.3.5 (Semi-direct product groups)**.**

- If N and Q are groups and $\varphi\colon Q \longrightarrow \operatorname{Aut}(N)$ is the trivial homomorphism, then the identity map (on the level of sets) yields an isomorphism $N \rtimes_\varphi Q \cong N \times Q$.
- Let $n \in \mathbb{N}_{\geq 3}$. Then the dihedral group $D_n = \langle s, t \mid s^n, t^2, tst^{-1} = s^{-1}\rangle$ (see Example 1.2.20) 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 (check!). Similarly, the infinite dihedral group $D_\infty = \langle s, t \mid t^2, tst^{-1} = s^{-1}\rangle \cong \operatorname{Isom}(\mathbb{Z})$ can also be written as a semi-direct product of $\mathbb{Z}/2$ by $\mathbb{Z}$ with respect to multiplication by -1 (check!).

- Let G be a group. Then the *lamplighter group over* G is the semi-direct product group $\left(\prod_{\mathbb{Z}} G\right) \rtimes_\varphi \mathbb{Z}$, 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 $\left(\prod_H G\right) \rtimes_\varphi H$, 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 versions of lamplighter and wreath product groups using $\bigoplus_H G$ instead of $\prod_H G$.

1.3.2 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 1.3.6 (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:

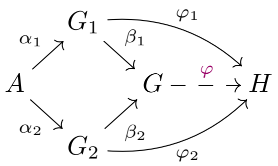

For every group H and all group homomorphisms $\varphi_1\colon G_1 \longrightarrow H$ and $\varphi_2\colon G_2 \longrightarrow H$ with $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$ there is exactly one homomorphism $\varphi\colon G \longrightarrow H$ with $\varphi \circ \beta_1 = \varphi_1$ and $\varphi \circ \beta_2 = \varphi_2$. Such a pushout is denoted by $G_1 *_A G_2$ (see Theorem 1.3.9 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 are both 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*)*.

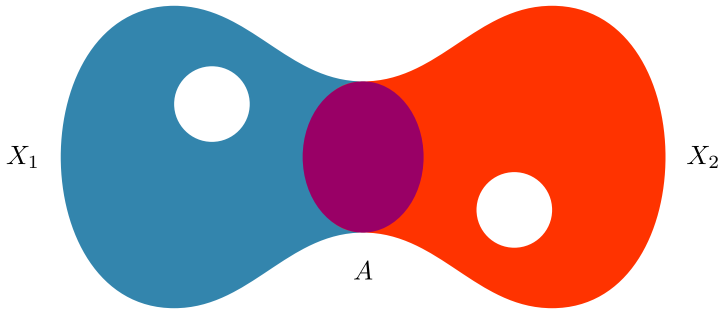

$$\pi_1(X_1 \cup_A X_2) \cong \pi_1(X_1) *_{\pi_1(A)} \pi_1(X_2)$$

Figure 1.2.: The theorem of Seifert and van Kampen, schematically

Caveat 1.3.7. 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 1.3.8 (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; check!).
- The infinite dihedral group $D_\infty \cong \operatorname{Isom}(\mathbb{Z})$ (Example 1.3.5) is isomorphic to the free product $\mathbb{Z}/2 * \mathbb{Z}/2$ (Exercise).
- The matrix group $\operatorname{SL}(2, \mathbb{Z})$ is isomorphic to the amalgamated free product $\mathbb{Z}/6 *_{\mathbb{Z}/2} \mathbb{Z}/4$ [96, Example I.4.2] (Outlook 3.4.3).
- 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 and the interiors of the two parts should cover the whole space) [77, Chapter IV][67, Chapter 2.2.2] (see Figure 1.2).

Theorem 1.3.9 (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 proof that free groups are uniquely determined up to canonical isomorphism by the universal property of free groups (Proposition 1.2.6; check!).

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. To minimise confusion, we introduce new letters x_g for $g \in G_1 \sqcup G_2$. Let

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

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

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

the relations R_{G_j} ensure that β_j is indeed 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 1.2.18): The map on the set of all words in the generators $\{x_g \mid g \in G_1 \sqcup G_2\}$ and their formal inverses induced by the map

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

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

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

As (the image of) $S := \{x_g \mid g \in G_1 \sqcup G_2\}$ 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 $\varphi \circ \beta_2 = \varphi_2$. □

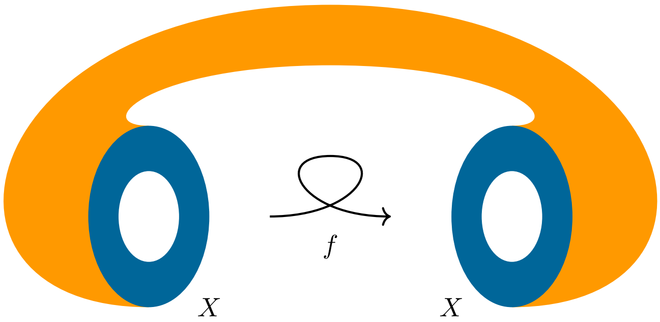

$$\pi_1(\text{mapping torus of } f) \cong \pi_1(X)*_{\pi_1(f)}$$

Figure 1.3.: The fundamental group of a mapping torus, schematically

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 1.3.10 (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 [72, Chapter IV][91, 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 1.3.12). 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 [72, p. 180] (see Figure 1.3).

Outlook 1.3.11 (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' classification of groups with infinitely many ends [102], and they are the starting point of Bass–Serre theory [96] (Outlook 3.2.6), which is concerned with actions of groups on trees (Outlook 3.2.6). Moreover, free groups, free products, amalgamated free products, and HNN-extensions can be understood in very concrete terms via suitable normal forms (Outlook 2.3.8).

Remark 1.3.12 (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:

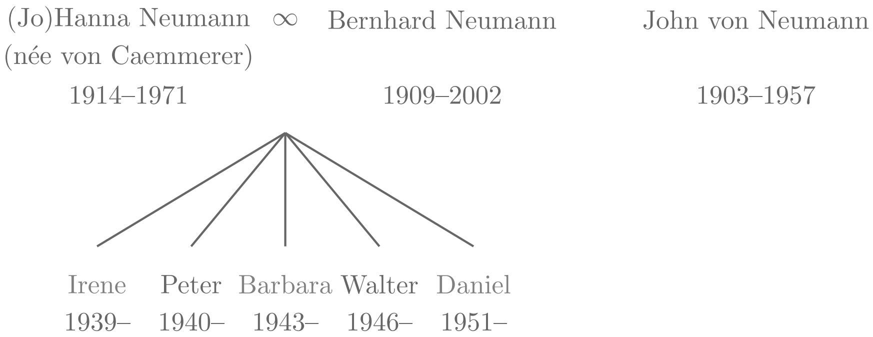

The contributions to geometric group theory of the (von) Neumanns are too numerous to be listed here [4, 3, 87]; for example:

- 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 1.3.10).
- The Hanna Neumann conjecture on ranks of certain subgroups of free groups (Exercise).
- The joint article by Bernhard, Hanna, and Peter Neumann [83].
- The von Neumann conjecture on the relation between non-amenability and free subgroups [63, Remark 9.1.12].

2

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 the 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 the second step, one adds a metric structure to Cayley graphs via word metrics. We will study this step in Chapter 4.

We start by reviewing some basic notation from graph theory (Chapter 2.1). We will then introduce Cayley graphs and discuss basic examples of Cayley graphs (Chapter 2.2); 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 2.3).

Overview of this chapter.

2.1 Review of graph notation

We start by reviewing basic terminology from graph theory; more information can be found in the literature [28, 46, 25]. In the following, we will always consider undirected, simple graphs without loops:

Definition 2.1.1 (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 2.1.2 (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 2.1.3 (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 2.1; however, pictures that may appear to be different 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 2.1):

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

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

Example 2.1.4 (Complete graphs). Let $n, m \in \mathbb{N}$. The graph

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

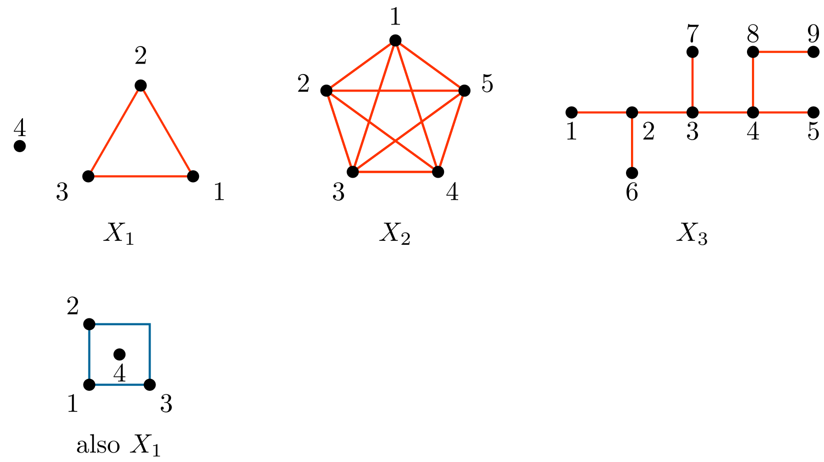

Figure 2.1.: Some graphs

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

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

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

Definition 2.1.5 (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 $f \colon V \longrightarrow V'$ such that for all $v, w \in V$ we have $\{v,w\} \in E$ if and only if $\big\{f(v), f(w)\big\} \in E'$. Thus, isomorphic graphs only differ in the labels of their vertices.[1]

The problem of deciding 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 [51].

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

Definition 2.1.6 (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 $\{v_j, v_{j+1}\} \in E$

[1] Of course, this notion of graph isomorphism can also be obtained as the 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.

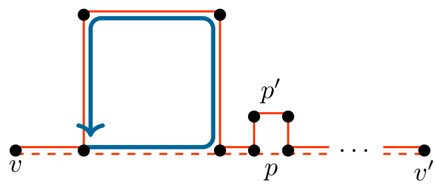

Figure 2.2.: Constructing a cycle (blue) out of two different paths.

holds for all $j \in \{0, \dots, n-1\}$; if $n < \infty$, then we say that this path *connects the vertices* v_0 *and* v_n.

- The graph X is called *connected* if every pair 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 $\{v_{n-1}, v_0\} \in E$.

Example 2.1.7. In Example 2.1.3, 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 not paths in X_3. In X_1, the sequence $1, 2, 3$ is a cycle.

Definition 2.1.8 (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 2.1.9 (Trees)**.** The graph X_3 in Example 2.1.3 is a tree, while X_1 and X_2 are not.

Proposition 2.1.10 (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 2.2), contradicting the fact that X is a tree. Hence, every pair of vertices of X can be connected by exactly one path in X. □

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

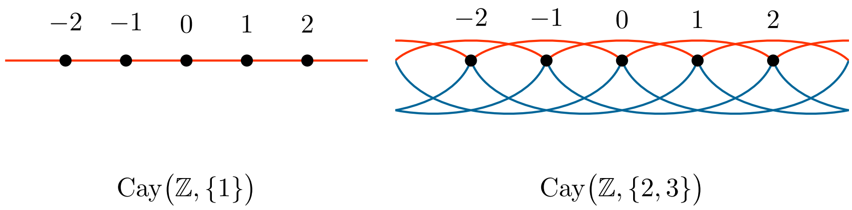

Figure 2.3.: Cayley graphs of the additive group $\mathbb{Z}$

Definition 2.1.11 (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 3.2.1 in order to characterise free groups.

2.2 Cayley graphs

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

Definition 2.2.1 (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 $\mathrm{Cay}(G, S)$ whose

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

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

That is, two vertices in a Cayley graph are adjacent if and only if they differ by right multiplication by an element (or its inverse) 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 2.2.2 (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 2.3.

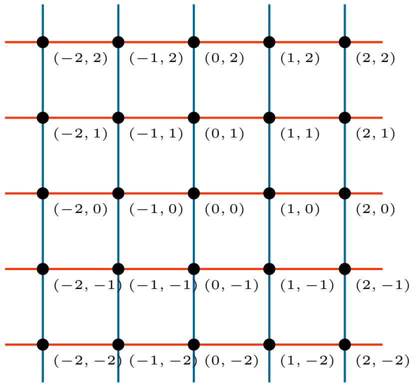

Figure 2.4.: The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$

When looking at these two graphs "from far away" they seem to have the same global structure, namely they look like the real line; in more technical terms, these graphs are quasi-isometric with respect to the corresponding word metrics – a concept that we will introduce and study thoroughly in later chapters (Chapters 4–**??**).

- 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 2.4; 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 2.5).
- 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 2.5.

  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 Cayley graph of a free group with respect to a free generating set is a tree (see Theorem 2.3.1 below).

Remark 2.2.3 (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.

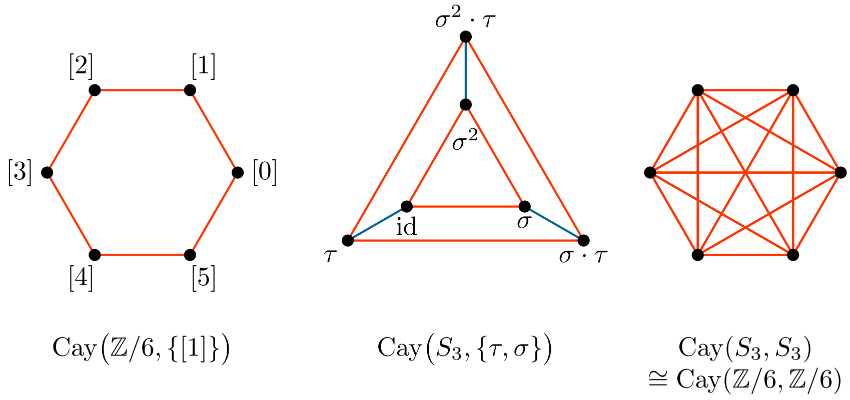

Figure 2.5.: Cayley graphs of $\mathbb{Z}/6$ and S_3

2. Cayley graphs are regular in the sense that every vertex has the same degree $|(S \cup S^{-1}) \setminus \{e\}|$.
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 finite degree.

So far, we have 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 3), and from the point of view of large scale geometry, by introducing metric structures on Cayley graphs (Chapter 4).

Outlook 2.2.4 (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* [20, 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.

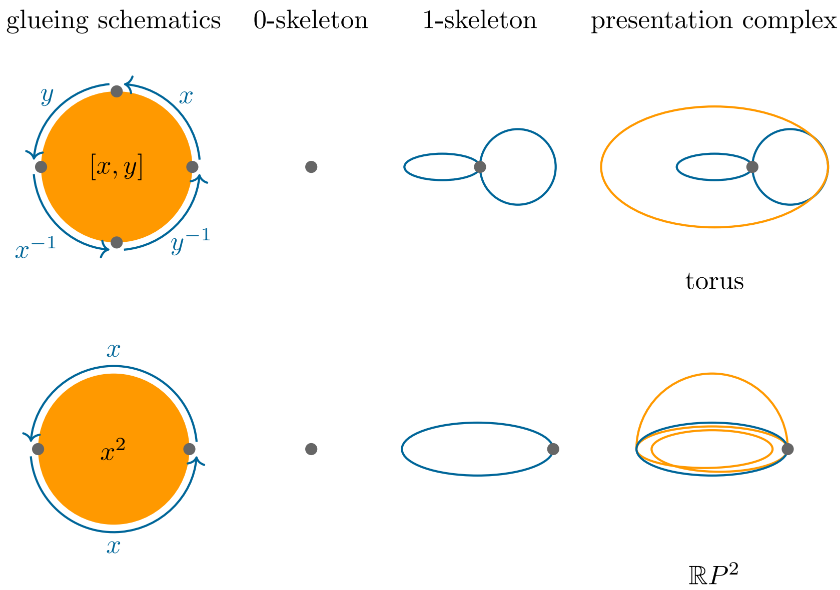

Figure 2.6.: Examples of presentation complexes

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 2.6).

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 [47, 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 us to model group theory (both groups and homomorphisms) in topology. Classifying spaces play an important role in the study of group cohomology [22, 65]. 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 $\langle S \mid R\rangle$ is almost the Cayley graph $\operatorname{Cay}(\langle S \mid R\rangle, S)$ (in the case of generators of order 2 some modifications might be necessary) [27, Chapter 2.2].

2.3 Cayley graphs of free groups

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

Theorem 2.3.1 (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 2.3.2 (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 essentially the only types of things that can go wrong:

Theorem 2.3.3 (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 2.3.1). More generally, any explicit and complete description of the Cayley graph of a group G with respect to a generating set S essentially requires us to solve the word problem of G with respect to S.

2.3.1 Free groups and reduced words

We constructed the free group $F(S)$ generated by S by taking the set of all words in elements of S and their formal inverses and then taking the quotient by the cancellation relation (proof of Theorem 1.2.7). While this construction is technically clean and simple, it has the disadvantage that getting hold of the precise nature of the 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 2.3.4 (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.

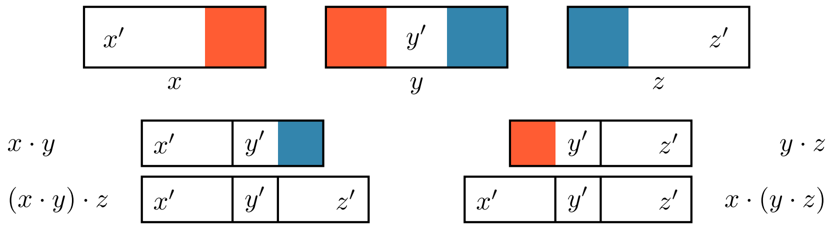

Figure 2.7.: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do not interfere

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

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

  holds for all $j \in \{1, \dots, n-1\}$. (In particular, ε is reduced.)

- We write $F_{\text{red}}(S)$ for the set of all reduced words in $(S \cup \widehat{S})^*$.

Proposition 2.3.5 (Free groups via reduced words)**.** *Let S be a set.*

1. *The set $F_{\text{red}}(S)$ of reduced words over $S \cup \widehat{S}$ forms a group with respect to the composition $F_{\text{red}}(S) \times F_{\text{red}}(S) \longrightarrow F_{\text{red}}(S)$ given by*

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

   where $s_1 \dots s_n$ and $s_{n+1} \dots s_m$ are in $F_{\text{red}}(S)$ (with $s_1, \dots, s_m \in S \cup \widehat{S}$), and

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

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

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

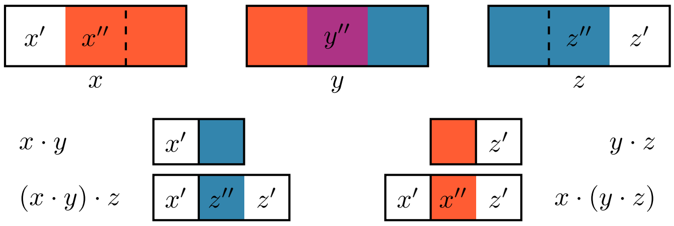

Figure 2.8.: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do interfere

lots of indices, we explain the argument graphically (Figures 2.7 and 2.8): 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 2.7).
- 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 2.8; 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 2.3.6 (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 2.3.7 (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 2.3.5 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 2.3.8 (Reduced words in free products etc.). Using the same method of proof as in Proposition 2.3.5, 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 [96, Chapter I][91, 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

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

   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

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

   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.

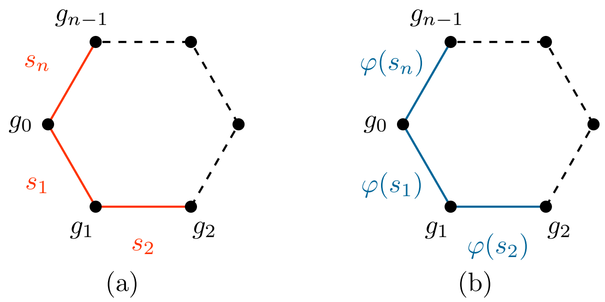

Figure 2.9.: Cycles lead to reduced words, and vice versa

2.3.2 Free groups → trees

Proof of Theorem 2.3.1. Suppose the group F is freely generated by S. By Proposition 2.3.5, the group F is isomorphic to $F_{\text{red}}(S)$ via an isomorphism that is the identity on S; without loss of generality we can therefore assume that F is $F_{\text{red}}(S)$.

We show that the Cayley graph $\text{Cay}(F, S)$ is a tree: Because S generates F, the graph $\text{Cay}(F, S)$ is connected. *Assume* for a contradiction that $\text{Cay}(F, S)$ 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 $s_n := {g_{n-1}}^{-1} \cdot g_0 \in S \cup S^{-1}$ (Figure 2.9 (a)). Because the vertices are distinct, the word $s_1 \dots s_n$ is reduced; on the other hand, we obtain

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

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

Example 2.3.9 (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 $\text{Cay}\big(F(S), \{a, b\}\big)$ is a regular tree whose vertices have exactly four neighbours (see Figure 2.10).

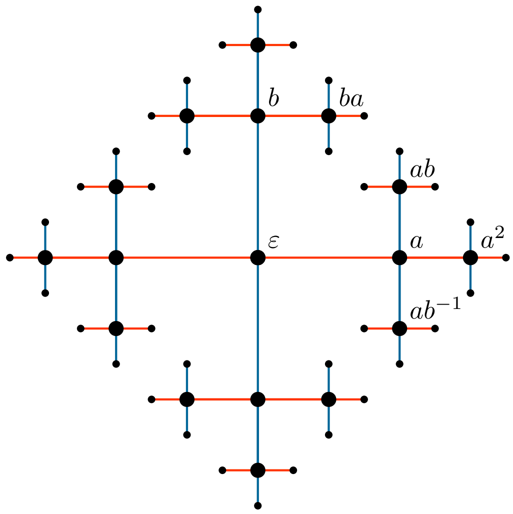

Figure 2.10.: Cayley graph of the free group of rank 2 with respect to a free generating set $\{a, b\}$

2.3.3 Trees $\to$ free groups

Proof of Theorem 2.3.3. 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 2.3.5, 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 $\varphi|_S = \operatorname{id}_S$ 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 2.9 (b)). The sequence $g_0, \dots, g_{n-1}$ is a cycle in $\operatorname{Cay}(G, S)$ because by minimality of the word $s_1 \dots s_n$, the elements $g_0, \dots, g_{n-1}$ are all distinct; moreover, $\operatorname{Cay}(G, S)$ contains the edges $\{g_0, g_1\}, \dots, \{g_{n-2}, g_{n-1}\}$, and the edge

$$\begin{aligned} \{g_{n-1}, g_0\} &= \{s_1 \cdot s_2 \cdot \dots \cdot s_{n-1}, e\} \\ &= \{s_1 \cdot s_2 \cdot \dots \cdot s_{n-1}, s_1 \cdot s_2 \cdot \dots \cdot s_{n-1} \cdot s_n\}. \end{aligned}$$

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

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

3

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 3.1).

As we have seen, free groups can be characterised combinatorially as the groups admitting trees as Cayley graphs (Chapter 2.3). In Chapter 3.2, 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 3.2.3.)

Another group action tool that helps to recognise that a given group is free is the ping-pong lemma (Chapter 3.3); this is particularly useful when proving that certain matrix groups are free – which also is a purely algebraic statement (Chapter 3.4).

Overview of this chapter.

3.1 Review of group actions

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

Definition 3.1.1 (Group action)**.** Let G be a group, let C be a category, and let X be an object in C. An *action of* G *on* X *in the category* C is a group homomorphism $G \longrightarrow \mathrm{Aut}_C(X)$. In other words, a group action of G on X consists of a family $(f_g)_{g\in G}$ of automorphisms of X such that

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

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

Example 3.1.2 (Group actions, generic examples)**.**

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

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

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

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

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

- Let G be a group and let X be a set. If $\varrho\colon G \longrightarrow \mathrm{Aut}_{\mathrm{Set}}(X)$ 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 $\varrho\colon G \longrightarrow \mathrm{Aut}_{\mathrm{Set}}(X)$ is a map, then ϱ is a G-action on X if and only if (check!)

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

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

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

We also use this notation whenever the group G acts on an object in a category whose objects are sets (with additional structure), whose morphisms are (structure preserving) maps of sets, and where 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 4); 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 [1].

3.1.1 Free actions

The relation between groups and geometric objects acted upon is particularly strong if the group action is a 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 3.1.3 (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 3.1.4 (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 3.1.5 (Rotations on the circle). Let $S^1 := \{z \in \mathbb{C} \mid |z| = 1\}$ be the unit circle in $\mathbb{C}$, and let $\alpha \in \mathbb{R}$. Then the rotation action

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

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

Example 3.1.6 (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 3.1.7 (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 [77, 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) [77, 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 3.1).
- 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 2.10), 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:

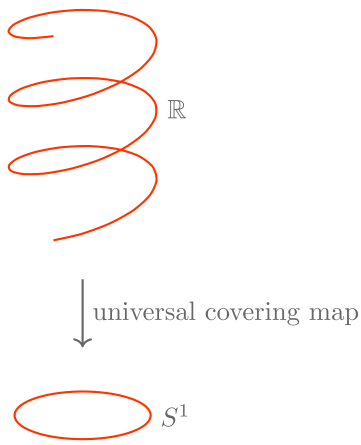

Figure 3.1.: Universal covering of S^1

Definition 3.1.8 (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 3.1.9 (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:

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

this map is indeed well-defined and a group homomorphism (check!).

Proposition 3.1.10 (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.

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 $\operatorname{Cay}(G, S)$. □

3.1.2 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 3.1.11 (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".

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 the orbit space usually also 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 3.1.12 (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 3.2). The quotient of $\mathbb{C}$ by this action can be identified with $\mathbb{R}_{\geq 0}$ (via the absolute value).

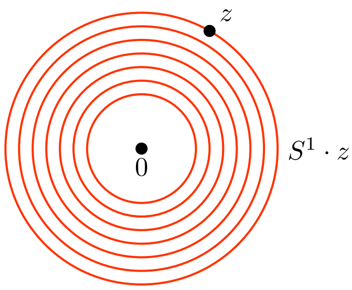

Figure 3.2.: Orbits of the rotation action of S^1 on $\mathbb{C}$

Example 3.1.13 (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 [77, Chapter V].

Definition 3.1.14 (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 is indeed a group (a subgroup of G; check!).

- 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 3.1.15 (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 canonically acts on Q by isometries and we know that $G \cong D_4$ (Example 1.2.20).

- 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 of Q by $2\pi/4$. 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 3.1.16 (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). □

Proposition 3.1.17 (Counting orbits)**.** *Let G be a group acting on a set X.*

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

$$\begin{aligned} A_x \colon G/G_x &\longrightarrow G \cdot x \\ g \cdot G_x &\longmapsto g \cdot x \end{aligned}$$

is well-defined and bijective. Here, G/G_x denotes the set of all G_x-cosets in G, i.e., $G/G_x = \{g \cdot G_x \mid g \in G\}$.

2. *Moreover, the number of distinct orbits equals the average number of points fixed by a group element: If G and X are finite, then*

$$|G \setminus X| = \frac{1}{|G|} \cdot \sum_{g \in G} |X^g|.$$

Proof. Ad 1. We start by showing that A_x is well-defined, i.e., that the 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

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

□

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 [1, 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 [91, Chapter 11]

3.1.3 Transitive actions

Transitive actions on "connected spaces" yield generating sets through "close neighbours". A first instance of this general principle is Proposition 3.1.19. A metric version of this principle is the Švarc–Milnor lemma (Chapter 4.4).

Definition 3.1.18 (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 $\mathrm{Cay}(G, S)$ is transitive (and free). In fact, this property can be used to characterise Cayley graphs in terms of actions on graphs:

Proposition 3.1.19 (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 the set 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 $\mathrm{Cay}(G, S)$ is isomorphic to X.

Proof. As the 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

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

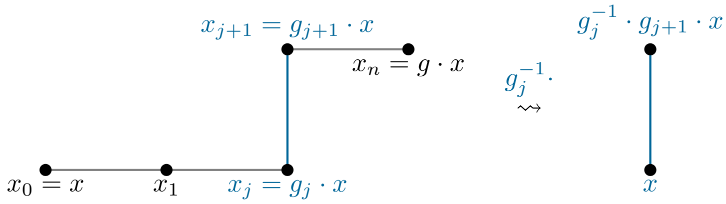

Figure 3.3.: From paths to words, using a transitive action

is bijective (Proposition 3.1.17). As the second step, we show that the set S is indeed 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 3.3). 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, $\{g_j^{-1} \cdot x_j, g_j^{-1} \cdot x_{j+1}\}$ is also an edge of X. By construction,

$$\begin{aligned} g_j^{-1} \cdot x_j &= \bigl(\varphi^{-1}(x_j)\bigr)^{-1} \cdot x_j = \bigl(\varphi^{-1}(x_j)\bigr)^{-1} \cdot \varphi^{-1}(x_j) \cdot x = x, \\ g_j^{-1} \cdot x_{j+1} &= g_j^{-1} \cdot g_{j+1} \cdot x = s_j \cdot x. \end{aligned}$$

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

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

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. □

3.2 Free groups and actions on trees

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 3.1.8).

Theorem 3.2.1 (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 3.2.1, 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 (Theorem 2.3.1). 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 2.3.5) 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 (Proposition 3.1.10). □

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 from T and the G-action on T we can construct – by contracting certain subtrees – a tree that is a Cayley graph of G for a suitable generating set and such that the assumptions of Theorem 2.3.3 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.

3.2.1 Spanning trees for group actions

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

Definition 3.2.2 (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 3.2.3 (Spanning trees). We consider the action of $\mathbb{Z}$ by "horizontal" shifting on the (infinite) tree depicted in Figure 3.4. Then the red subgraph is a spanning tree for this action.

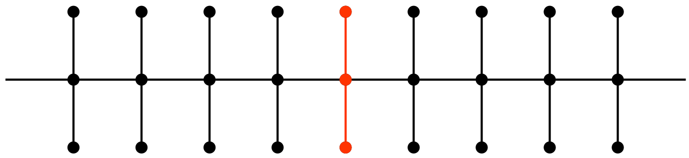

Figure 3.4.: A spanning tree (red) for a shift action of $\mathbb{Z}$

Theorem 3.2.4 (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 first show that T_G contains an element T that is maximal with respect to the subtree relation, using Zorn's lemma: 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; check!). 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 vertex of T. We 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 at v' and ending in v, then $g \cdot p'$ is a path starting at 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 eventually produces 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 that contains T as a proper subgraph. This contradicts the maximality of T. Hence, T is a spanning tree for the G-action on X. ☐

Figure 3.5.: Contracting the spanning tree and all its translates to vertices (red), in the situation of Example 3.2.3

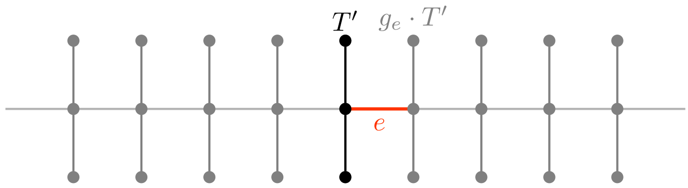

Figure 3.6.: An essential edge (red) for the shift action of $\mathbb{Z}$ (Example 3.2.3)

3.2.2 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 3.2.1, part II. Let G be a group acting freely on a tree T by graph automorphisms. By Theorem 3.2.4 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 3.5 shows this in the situation of Example 3.2.3); 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 3.1.19. However, we prefer to proceed directly in the original tree T:

As in Proposition 3.1.19, 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' and if one of the vertices of the edge in question belongs to T' (then the other vertex cannot belong to T' as well, by uniqueness of paths in trees (Proposition 2.1.10)).

As the 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

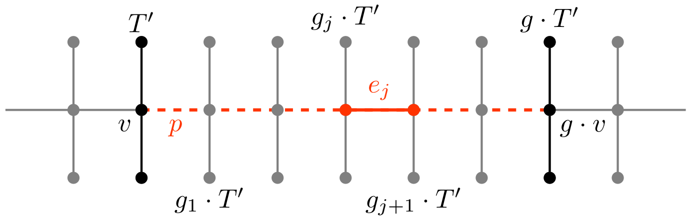

Figure 3.7.: The set $\widetilde{S}$ generates G

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

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 3.1.16).
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 $g^{-1} = g_{g^{-1} \cdot e}$ is also 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

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

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 3.7).

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

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

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

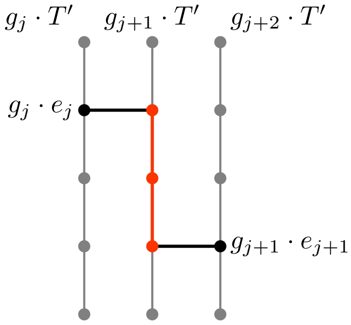

Figure 3.8.: Cycles in $\operatorname{Cay}(G, \widetilde{S})$ lead to cycles in T by connecting translates of essential edges in the corresponding translates of T' (red path)

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

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 2.3.3 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

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

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 3.8). 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 3.2.5 (Topological proof)**.** Let G be a group acting freely on a tree; then G also acts freely, continuously, cellularly, and properly discontinuously on the CW-realisation X of this tree, which is contractible. Covering theory shows that the quotient space $G \setminus X$ is homeomorphic to a one-dimensional CW-complex and that $G \cong \pi_1(G \setminus X)$. The Seifert-van Kampen theorem then yields that G is a free group [77, Chapter VI].

Outlook 3.2.6 (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 [96, Chapter 5][37, 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.

3.2.3 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 3.2.7 (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, the (sub)group G also acts freely on this non-empty tree. Therefore, G is a free group by Theorem 3.2.1. □

Example 3.2.8. 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 3.2.7). Because $\mathbb{Z}^2$ is *not* free (Exercise), we obtain that $H \not\cong \mathbb{Z}^2$. We will see a vast, geometric, generalisation of this fact in Corollary 6.5.15.

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 3.2.9 (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, the left translation action of the subgroup G on T is also free (and so G is free).

Looking at the proof of Theorem 3.2.1 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). 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 3.2.10 (Topological proof of the Nielsen–Schreier theorem)**.** A topological version of the proof of the Nielsen–Schreier theorem can be given via covering theory [77, 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, 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 3.2.11. *If F is a free group of rank at least 2, and $n \in \mathbb{N}$, then there is a subgroup of F that is free of finite rank at least n.*

Proof. Using a surjective homomorphism $F \longrightarrow \mathbb{Z}$, one can apply the quantitative version of the Nielsen–Schreier theorem to obtain free subgroups of large rank (Exercise). □

Corollary 3.2.12. *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 $[F(S) : H'] = [G : H]$ (check!).

By Corollary 3.2.9, the group H' is finitely generated; thus, the image $H = \pi(H')$ is also finitely generated. □

We will later see an alternative proof of Corollary 3.2.12 via the Švarc–Milnor lemma (Corollary 4.4.5).

Corollary 3.2.13 (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.*

Recall that a group is *torsion-free* if it does not contain non-trivial elements of finite order.

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 3.9 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 2.3.8) 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

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

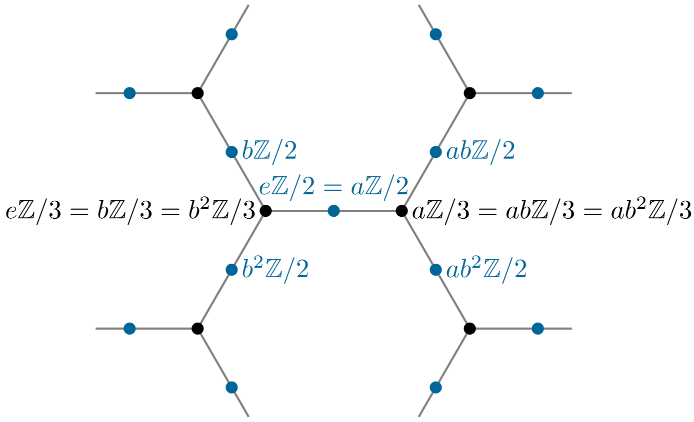

Figure 3.9.: The tree for the free product $\mathbb{Z}/2 * \mathbb{Z}/3 \cong \langle a, b \,|\, a^2, b^3 \rangle$

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 3.2.1 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 [96].

3.3 The ping-pong lemma

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

Theorem 3.3.1 (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*

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

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

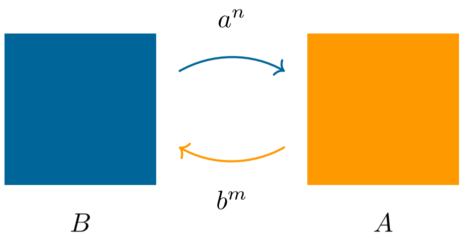

Figure 3.10.: The ping-pong lemma

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 3.10)

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

   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 3.3.2 (Ping-pong lemma for free products)**.** Similarly, using the description of free products in terms of reduced words (Outlook 2.3.8), one can show the following [44, 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

$$\forall_{g\in G_1\setminus\{e\}} \quad g \cdot X_2 \subset X_1 \quad \text{and} \quad \forall_{g\in G_2\setminus\{e\}} \quad g\cdot X_1 \subset X_2,$$

then $G \cong G_1 * G_2$.

The ping-pong lemma is a standard tool to establish that certain matrix groups are free (Chapter 3.4). Further examples are given in de la Harpe's book [44, 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.

3.4 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 3.4.1), which has applications in graph theory (Chapter 3.4.2); finally, we will briefly discuss the Tits alternative (Chapter 3.4.3).

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

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

Example 3.4.1 (A free subgroup of $\mathrm{SL}(2,\mathbb{Z})$)**.** Let a, $b \in \mathrm{SL}(2,\mathbb{Z})$ be given by

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

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

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

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

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

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

and

$$\begin{aligned} |x + 2 \cdot n \cdot y| &\geq |2 \cdot n \cdot y| - |x| \geq 2 \cdot |y| - |x| && \text{(because } |n| \geq 1\text{)} \\ &> 2 \cdot |y| - |y| && \text{(because } (x,y)^{\mathsf{T}} \in B\text{)} \\ &= |y|; \end{aligned}$$

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

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

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 $[\mathrm{SL}(2,\mathbb{Z}) : G] = 6$. In particular,*

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

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 when 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 $\bigl|\mathrm{SL}(2,\mathbb{Z}/2)\bigr| = 6$.

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 the converse inclusion $F' \subset F$ also 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).

Outlook 3.4.3 ($\mathrm{SL}(2, \mathbb{Z})$ as amalgamated free product)**.** The discussion above can be extended to prove the following fact [96, 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

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

3.4.2 Application: Regular graphs of large girth

We will now discuss a graph-theoretic application of Example 3.4.1, namely the construction of regular graphs with few vertices and large girth.

Definition 3.4.4 (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 3.4.5 (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 [11] of the existence of finite graphs of large girth and large chromatic number, which shows that colouring graphs is indeed a global problem. A first, constructive, step in this direction is Mycielski's iterated graph construction [82]. Another construction problem of this type is to exhibit regular graphs of large girth with "few" vertices.

Margulis [75] solved this problem, using Cayley graphs; for simplicity, we only treat the case of 4-regular graphs.

Theorem 3.4.6 (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 $d := 1 + \sqrt{2}$.

This result is asymptotically optimal in the sense that the girth grows at most logarithmically in the number of vertices. Moreover, the proof by Margulis is constructive.

Proof. We construct the desired graphs explicitly: Let

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

be the matrices in $\mathrm{SL}(2,\mathbb{Z})$ of Example 3.4.1. 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

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

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 (as $N \geq 5$) 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 $F_{\text{red}}(\alpha, \beta) \longrightarrow F$ 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 the girth $g(X_N)$, we may assume that the lengths $m, n \in \mathbb{N}$ of w and v respectively satisfy

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

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 a $c' \in M_{2\times 2}(\mathbb{Z})$ 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 and the Euclidean norm; Remark 3.4.7)

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

the last inequality follows because the entries of c' are integral and $c' \neq 0$. On the other hand, computing the eigenvalues of $a^{\mathsf{T}} \cdot a$ shows

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

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 $g(X_N)$. □

Remark 3.4.7 (Operator norm)**.** Let $n \in \mathbb{N}_{>0}$ and let $A \in M_{n\times n}(\mathbb{R})$. Then the *operator norm* of A (with respect to the Euclidean norm $\|\cdot\|_2$ on $\mathbb{R}^n$) is defined as

$$\|A\| := \sup\Bigl\{ \frac{\|A \cdot x\|_2}{\|x\|_2} \Bigm| x \in \mathbb{R}^n \setminus \{0\} \Bigr\} \in \mathbb{R}_{\geq 0}.$$

Straightforward estimates show that (check!):

- We have $\|A\| \geq \max\{|A_{jk}| \mid j, k \in \{1, \dots, n\}\}$.
- We have $\|A\| \leq \sqrt{\lambda}$, where λ is the biggest eigenvalue of $A^{\mathsf{T}} \cdot A$.

3.4.3 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. But non-Abelian free groups appear frequently as building blocks in linear groups; more precisely, J. Tits [105] discovered the following:

Theorem 3.4.8 (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 5.3.1; the definition of solvability is recalled in Definition 5.3.3.

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 [31]. The proof of the Tits alternative consists of the following components:

- The two alternatives exclude each other.
- Recognition of free groups via the ping-pong lemma.
- Eigenvalue analysis to set up the ping-pong lemma.

In Example 3.4.1 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 example:

Example 3.4.9 (Another free linear group). Let $\lambda \in \mathbb{C}$. We consider the matrices

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

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

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

satisfy the condition of the ping-pong lemma (Theorem 3.3.1) provided that $|\lambda|$ is large enough and $\varepsilon \in \mathbb{R}_{>0}$ is small enough (Exercise). 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 when 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. At this point, we have not only changed 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 the absolute irreducibility of the G-action, we can pass to a finite extension of k' so that a' is diagonalisable and the G-action is still 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 3.4.9. □

A quantitative version of the Tits alternative was recently established by Breuillard [15].

4

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 $\operatorname{Cay}(G, S)$ 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 the 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 4.4).

Overview of this chapter.

4.1 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 4.1.12.

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

Definition 4.1.1 (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 4.1.2 (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

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

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

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

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

Remark 4.1.3. 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 4.1.4 (Bilipschitz equivalence)**.** Let $f\colon X \longrightarrow Y$ be a map between metric spaces (X, d_X) and (Y, d_Y).

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

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

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

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

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

Remark 4.1.5. 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.

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

Definition 4.1.6 (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.

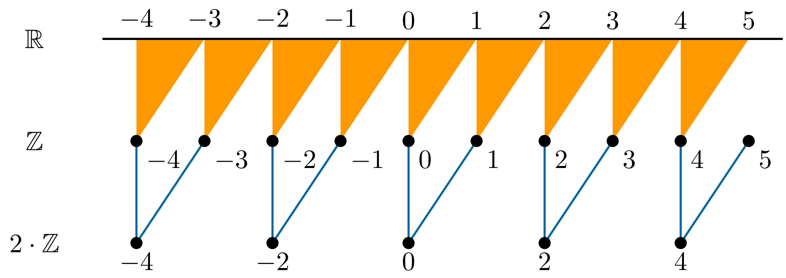

Figure 4.1.: The metric spaces $\mathbb{R}$, $\mathbb{Z}$, and $2 \cdot \mathbb{Z}$ are quasi-isometric

- 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 4.1.7 (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:

We consider $\mathbb{R}$ as a metric space with respect to the standard 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 4.1).

The inclusions $2 \cdot \mathbb{Z} \hookrightarrow \mathbb{Z}$ and $\mathbb{Z} \hookrightarrow \mathbb{R}$ are quasi-isometric embeddings but not 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). In particular, $\mathbb{Z} \sim_{\mathrm{QI}} \mathbb{R}$.

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 4.1.8. 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 4.1.9 (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 (check!). 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 (Example 5.2.8).

Proposition 4.1.10 (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*

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

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

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

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

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

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)

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

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

Moreover, g is also 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.) □

A formalisation of Proposition 4.1.10 in the proof assistant Lean is given in Appendix A.1.

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

Proposition 4.1.11 (Inheritance properties of quasi-isometric embeddings)**.**

1. *Every map at finite distance from a quasi-isometric embedding is a quasi-isometric embedding.*

2. *Every map at finite distance from 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 $g \circ f$ and $g \circ f'$ also have finite distance from each other.*
4. *Compositions of quasi-isometric [bilipschitz] embeddings are quasi-isometric [bilipschitz] embeddings.*
5. *Compositions of quasi-isometries [bilipschitz equivalences] are quasi-isometries [bilipschitz equivalences].*

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

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

Remark 4.1.12 (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 4.1.11). 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

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

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

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

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

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

Definition 4.1.13 (Quasi-isometry group). Let X be a metric space. Then the *quasi-isometry group of* X is defined by

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

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 us 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 4.1.14 (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

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

  together with many rather large and non-commutative groups [92].

4.2 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 allows us to associate geometric types to finitely generated groups that do not depend on the choice of finite generating sets.

Definition 4.2.1 (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 (check!), the *metric on V associated with X*.

Remark 4.2.2. 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 (check!).

Definition 4.2.3 (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,

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

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

Example 4.2.4 (Word metrics on $\mathbb{Z}$)**.** The word metric on $\mathbb{Z}$ corresponding to the generating set $\{1\}$ coincides with the metric on $\mathbb{Z}$ induced from the standard metric on $\mathbb{R}$. On the other hand, in the word metric on $\mathbb{Z}$ corresponding to the generating set $\mathbb{Z}$, all group elements have distance 1 from every other group element.

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

Proposition 4.2.5. *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 4.1.11).

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

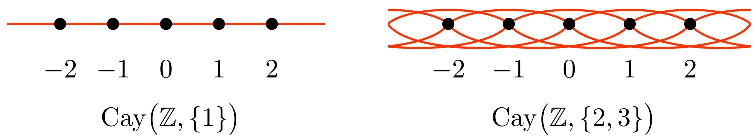

Figure 4.2.: Cayley graphs of $\mathbb{Z}$ with the same large scale geometry

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

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

Example 4.2.6 (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 4.2.

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 4.2.4).

Definition 4.2.7 (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 4.2.8 ($\mathbb{Z}^n \sim_{\mathrm{QI}} \mathbb{R}^n$)**.** 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

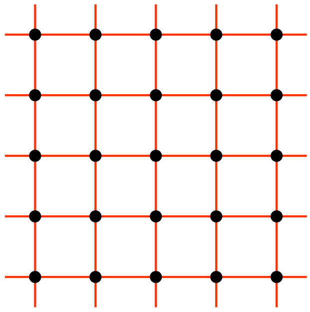

Figure 4.3.: The Cayley graph $\mathrm{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ resembles the geometry of the Euclidean plane $\mathbb{R}^2$

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 4.3).

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 4.4).

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 4.2.9 (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). □

However, not all infinite finitely generated groups that are quasi-isometric are bilipschitz equivalent [63, Chapter 9.4].

4.2.1 First examples

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

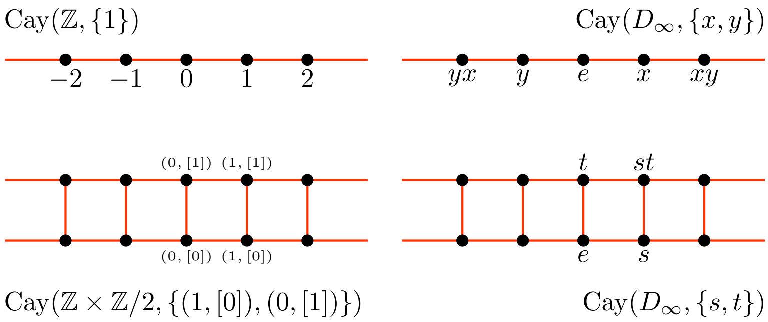

Figure 4.4.: The groups $\mathbb{Z}$, $\mathbb{Z} \times \mathbb{Z}/2$, and D_∞ are quasi-isometric

Remark 4.2.10 (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 4.2.11 (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 4.2.10), 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 0.2).

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

Example 4.2.12 (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 4.4):

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

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

The Cayley graph $\mathrm{Cay}(D_\infty, \{x, y\})$ is isomorphic to $\mathrm{Cay}(\mathbb{Z}, \{1\})$; in particular, D_∞ and $\mathbb{Z}$ are bilipschitz equivalent. On the other hand, the Cayley graph $\mathrm{Cay}(D_\infty, \{s, t\})$ is isomorphic to $\mathrm{Cay}(\mathbb{Z}\times\mathbb{Z}/2, \{(1, [0]), (0, [1])\})$; 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 $\mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}/2$ are also bilipschitz equivalent.

Caveat 4.2.13 (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 [63, 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 [62]. More generally, a similar classification also applies to finitely generated nilpotent groups [109].

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

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

Theorem 4.2.14 (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 5.3.6, Corollary 6.5.8).

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

4.3 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.

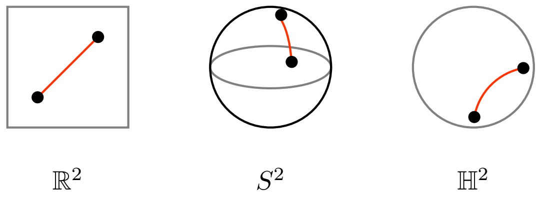

Figure 4.5.: Geodesic spaces and example geodesics

4.3.1 (Quasi-)Geodesic spaces

Definition 4.3.1 (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 4.3.2 (Geodesic spaces)**.** The following statements are illustrated in Figure 4.5.

- 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) [69, Proposition 2.2.6]. 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).

- 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 [63, Appendix A.3]. In the Poincaré disk model, geodesics are parts of circles that intersect the boundary circle orthogonally.

Caveat 4.3.3. 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 [68, Chapter 3].

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 4.3.4 (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 4.3.5 (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; check!).
- 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).

4.3.2 Geodesification via geometric realisation of graphs

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

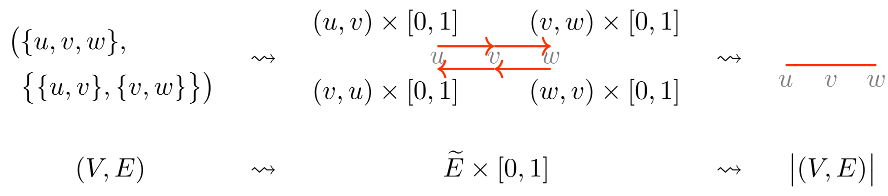

Figure 4.6.: Geometric realisation of graphs

(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 4.6):

Definition 4.3.6 (Geometric realisation of graphs)**.** Let $X = (V, E)$ be a connected graph. The *geometric realisation of* X is the metric space

$$\big(|X|, d_{|X|}\big)$$

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,

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

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 4.7)

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

Figure 4.7.: Parametrisations describing the same points in the geometric realisation

- 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

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

for all $[((u,v),t)]$, $[((u',v'),t')] \in |X|$, where d_X denotes the metric on V induced from the graph structure (Definition 4.2.1).

Clearly, this construction can be extended to a functor from the category of graphs to the category of metric spaces (and distance-non-increasing maps; check!). 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 3.1.8 and Proposition 3.1.16 (check!).

Example 4.3.7 (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 4.8).
- The geometric realisation of $\mathrm{Cay}(\mathbb{Z},\{1\})$ is isometric to the real line $\mathbb{R}$ with the standard metric (check!).
- 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 $\mathbb{R}^2$ (check!).

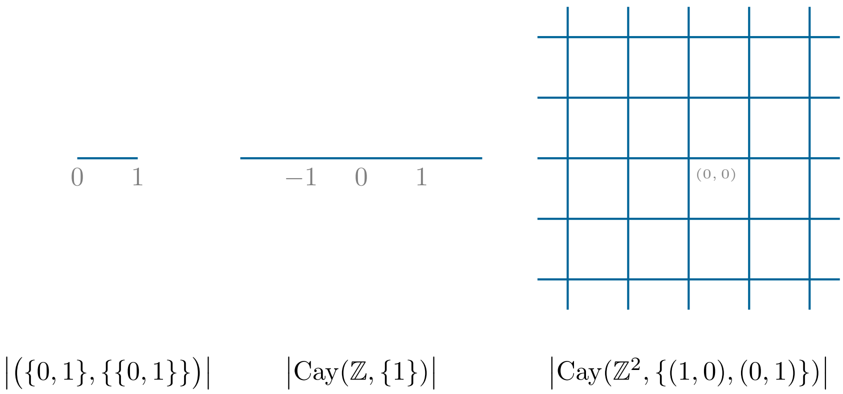

Figure 4.8.: Geometric realisations (Example 4.3.7)

Proposition 4.3.8 (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 (check!). □

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

Proposition 4.3.9 (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, $Y \sim_{\text{QI}} |Y|$ (Proposition 4.3.8). Hence, X and $|Y|$ are quasi-isometric. Moreover, $|Y|$ is a geodesic metric space by Proposition 4.3.8. □

However, it is not always desirable to replace the original space by a geodesic space, because some control is lost during this replacement. In particular, 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.

4.4 The Švarc–Milnor lemma

Why should we be interested in understanding what 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 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 4.4.1 (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*

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

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).

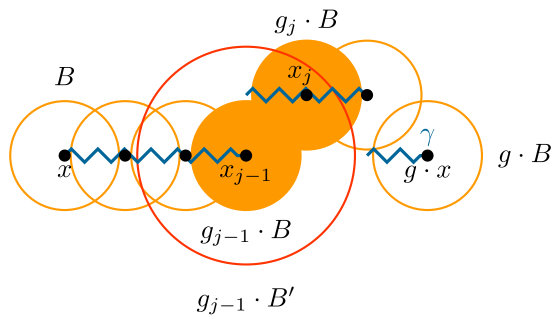

Figure 4.9.: Covering a quasi-geodesic by translates of B

Proof. The set S generates G: The argument follows the transitivity principle used in the proof of Proposition 3.1.19. Let $g \in G$. We show that $g \in \langle S \rangle_G$ by using a suitable quasi-geodesic and following translates of B along this quasi-geodesic (Figure 4.9): 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 at x and ending in $g \cdot x$. We look at close enough points on this quasi-geodesic and "connect the dots":

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

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

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

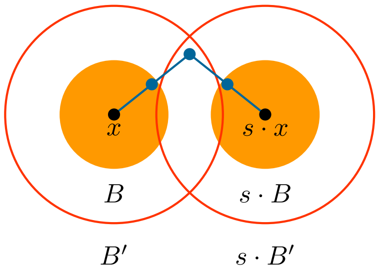

Figure 4.10.: If $s \in S$, then $d(x, s \cdot x) \leq 2 \cdot (\operatorname{diam} B + 2 \cdot b)$

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

lies in the subgroup generated by S, as desired.

Implicitly, we worked with the graph $Y = (V, E)$ with $V = \{g \cdot B \mid g \in G\}$ and $E = \{\{g \cdot B, h \cdot B\} \mid g \cdot B' \cap h \cdot B' \neq \emptyset\}$ and the left translation action by G on Y by graph automorphisms. The argument given above shows that Y is connected and a variation of the proof of Proposition 3.1.19 then establishes that S is a generating set of G.

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, 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 $d_S(e, g)$: As above, let $\gamma \colon [0, L] \longrightarrow X$ be 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

$$\begin{aligned} d\bigl(\varphi(e),\varphi(g)\bigr) &= d(x, g\cdot x) = d\bigl(\gamma(0),\gamma(L)\bigr) \\ &\geq \frac{1}{c}\cdot L - b && (\gamma \text{ is a QI embedding}) \\ &\geq \frac{1}{c}\cdot \frac{b\cdot(n-1)}{c} - b && (\text{by definition of } n) \\ &= \frac{b}{c^2}\cdot n - \frac{b}{c^2} - b \\ &\geq \frac{b}{c^2}\cdot d_S(e,g) - \frac{b}{c^2} - b && (\text{proof of the first part}). \end{aligned}$$

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 4.10) we obtain

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

(Recall that $\operatorname{diam} B$ is assumed to be finite.)

Because

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

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

The proof of the Švarc–Milnor lemma only gives 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 4.4.2 (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 4.4.3 (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 [77, 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 by 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 4.4.4 (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 3.1.13).
  - The (horizontal) translation action of $\mathbb{Z}$ on $\mathbb{R}^2$ is *not* cocompact (with respect to the standard topology on $\mathbb{R}^2$), because the quotient is homeomorphic to the infinite cylinder $S^1 \times \mathbb{R}$, which is not compact.

– The action of $\mathrm{SL}(2,\mathbb{Z})$ by Möbius transformations, i.e., via

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

on the upper half plane $H := \{z \in \mathbb{C} \mid \operatorname{Re} z > 0\}$ is *not* cocompact [63, Exercise 5.E.20].

Proof of Corollary 4.4.2. 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 4.4.1), 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 4.4.1). □

4.4.1 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 4.4.5. *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 4.4.1): 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 4.4.6 ((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 (Exercisesbut).

Corollary 4.4.7 (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 4.4.5, it suffices to show that G is finitely generated if G' is; 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 4.4.1). Therefore, we obtain $G \sim_{\mathrm{QI}} G/N$. □

Example 4.4.8 (Commensurability)**.**

- Let $n \in \mathbb{N}_{\geq 2}$. Then the free group of rank 2 contains a free group of rank n as a finite index subgroup (Exercise), 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 3.4.1) and has index 12 in $\mathrm{SL}(2,\mathbb{Z})$ (Proposition 3.4.2). 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 4.4.9. Not all quasi-isometric groups are commensurable [44, 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 [63, 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 [22]. Hence, commensurable groups G and G' (that admit sufficiently finite models of classifying spaces) satisfy

$$\chi(G) = 0 \Longleftrightarrow \chi(G') = 0.$$

However, the inheritance properties of the Euler characteristic [22] 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 are also not weakly commensurable.

There also exist groups that are weakly commensurable but not commensurable [44, III.18(xi)].

4.4.2 Application: Geometric structures on manifolds

As a second example, we look at applications of the Švarc–Milnor lemma in algebraic topology and Riemannian geometry via fundamental groups; for introductions to Riemannian geometry, we refer to the literature [58, 68].

Corollary 4.4.10 (Fundamental groups and quasi-isometry)**.** *Let M be a closed (i.e., compact and without boundary) connected non-empty Riemannian manifold, and let $\widetilde{M}$ be its Riemannian universal covering manifold. Then the fundamental group $\pi_1(M)$ is finitely generated and for every $x \in \widetilde{M}$, the map*

$$\begin{aligned}\pi_1(M) &\longrightarrow \widetilde{M}\\ g &\longmapsto g \cdot x\end{aligned}$$

given by the action of the fundamental group $\pi_1(M)$ on $\widetilde{M}$ via deck transformations is a quasi-isometry. Here, M and $\widetilde{M}$ are equipped with the metrics induced from their Riemannian metrics.

Sketch of proof. Standard arguments from Riemannian geometry and topology show that in this case $\widetilde{M}$ is a proper geodesic metric space and that the action of $\pi_1(M)$ on $\widetilde{M}$ is isometric, proper, and cocompact (the quotient being the compact space M). Applying the topological version of the Švarc–Milnor lemma (Corollary 4.4.2) finishes the proof. □

We give a sample application of this consequence of the Švarc–Milnor lemma to Riemannian geometry:

Definition 4.4.11 (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.

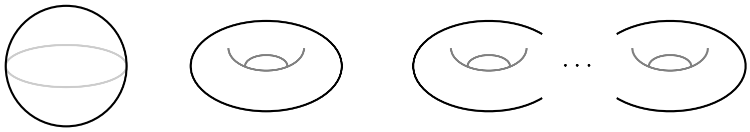

Figure 4.11.: Oriented closed connected surfaces

- 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 [58, Chapter 11].

Example 4.4.12 (Surfaces). Oriented closed connected surfaces are determined up to homeomorphism/diffeomorphism by their genus (i.e., the number of "handles", see Figure 4.11) [77, 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
$$\Bigl\langle a_1, \dots, a_g, b_1, \dots, b_g \Bigm| \prod_{j=1}^{g} [a_j, b_j] \Bigr\rangle$$
and one can show that these surfaces admit hyperbolic Riemannian metrics [10, Chapter B.1, B.3][103].

Corollary 4.4.13 ((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 4.4.10. □

Moreover, by the Bonnet–Myers theorem, closed connected Riemannian manifolds of positive sectional curvature have finite fundamental group [58, 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.

4.5 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 4.5.1 (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. $G \cdot K = X = K \cdot H$.

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 $K \cdot h \subset F_G(h) \cdot K$.

[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 $H \longrightarrow S_X$.

Example 4.5.2 (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 4.5.3 (Quasi-isometry and set-theoretic couplings)**.** *Let G and H be two finitely generated groups that admit a set-theoretic coupling. Then*

$$G \sim_{\mathrm{QI}} H.$$

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 4.5.1. We prove $G \sim_{\mathrm{QI}} H$ by writing down a candidate for a quasi-isometry $G \longrightarrow H$ and by then verifying (similarly 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,

$$d_T\big(h, f(g)\big) \leq D_T F_H,$$

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

$$F_H(S) := \bigcup_{s \in S \cup S^{-1}} F_H(s) \quad \text{and} \quad F_G(T) := \bigcup_{t \in T \cup T^{-1}} F_G(t)$$

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

  $$d_T\big(f(g), f(g')\big) \leq D_T F_H(S) \cdot d_S(g, g') + D_T F_H.$$

  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 heavily used the fact 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)

  $$\begin{aligned} d_T\big(f(g), f(g')\big) &= d_T\big(e, f(g)^{-1} \cdot f(g')\big) \\ &\leq D_T(F_H \cdot F_H(S)^n) \\ &\leq n \cdot D_T F_H(S) + D_T F_H \\ &= d_S(g, g') \cdot D_T F_H(S) + D_T F_H, \end{aligned}$$

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 4.5.4 (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* [65]. 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 [38, 98, 95, 59].

Moreover, we do not need to assume that both groups are finitely generated as being finitely generated is preserved by set-theoretic couplings (Exercise).

The converse of Proposition 4.5.3 also holds: whenever two finitely generated groups are quasi-isometric, then there exists a coupling (even a topological coupling) between them:

Definition 4.5.5 (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 neighbourhood $K \subset X$ of x with $K \subset U$. For example, proper metric spaces are locally compact.

We can now formulate Gromov's dynamic criterion for quasi-isometry:

Theorem 4.5.6 (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 are several *different* notions of local compactness in the literature!

2. *There is a set-theoretic coupling for G and H.*

3. *The groups G and H are quasi-isometric.*

Proof. Ad "2 $\Longrightarrow$ 3". This was proved in Proposition 4.5.3.

Ad "3 $\Longrightarrow$ 1". Suppose that the finitely generated groups G and H are quasi-isometric. We 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 a 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, the map $\overline{f}$ satisfies a $\max(2 \cdot \overline{c}, \overline{c}^2 + \overline{c})$-bilipschitz estimate (this follows as for bijective quasi-isometries 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:

$$\begin{aligned} G \times X &\longrightarrow X \\ (g, f) &\longmapsto \bigl(x \mapsto f(g^{-1} \cdot x)\bigr) \\[1ex] X \times H &\longrightarrow X \\ (f, h) &\longmapsto \bigl(x \mapsto f(x) \cdot (h, e)\bigr) \end{aligned}$$

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 [49, 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 [94].

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

$$\{g \in G \mid g \cdot K \cap K \neq \emptyset\} \quad \text{and} \quad \{h \in H \mid K \cdot h \cap K \neq \emptyset\}$$

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. □

Outlook 4.5.7 (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 [78, Theorem 3.2][59].

4.5.1 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 4.5.9).

Definition 4.5.8 (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 4.5.9 (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 4.5.6). $\square$

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 4.5.10 (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 $\mathbb{R}^2$ (as we will see later).
- Let $H_{\mathbb{R}}$ be the *real Heisenberg group*, and let H be the *Heisenberg group*, i.e.,

  $$H_{\mathbb{R}} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \, x, y, z \in \mathbb{R} \right\}, \ H := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle| \, x, y, z \in \mathbb{Z} \right\}.$$

  Then $H_{\mathbb{R}}$ is a locally compact topological group (with respect to the topology given by convergence of all matrix coefficients), and H is a finitely generated uniform lattice in $H_{\mathbb{R}}$ (Exercise).

  So, a 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 5 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); 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 $\mathrm{Isom}(\widetilde{M})$ [94, Theorem 2.35].

Outlook 4.5.11 (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 [36, 66].

4.6 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.

4.6.1 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 4.6.1 (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 4.6.2 (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 4.6.3. 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 4.6.4 (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.
- *The perfect invariant.* Let V be the set(!) of quasi-isometry types of finitely generated groups and let I be the map associating with every finitely generated group its quasi-isometry type. Then I contains perfect information on which finitely generated groups are quasi-isometric, but I is *not* directly computable.
- *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 4.2.11).
- *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 4.4.8).

In order to obtain interesting classification results we need further quasi-isometry invariants and strategies for (partial) computation. In the following chapters, we will, for instance, study

- the growth of groups (Chapter 5),
- hyperbolicity (Chapter 6),

- ends of groups (i.e., geometry at infinity) [63, Chapter 8].

Another interesting example is amenability [63, Chapter 9].

Caveat 4.6.5. 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 1.2.28 and it can, for example, be proved by producing uncountably many different growth types of groups [39] or via small cancellation theory and the geometry of loops in Cayley graphs [14].

4.6.2 Geometric properties of groups and rigidity

In geometric group theory, it is common to use the following term:

Definition 4.6.6 (Geometric property of groups)**.** Let P be a property of finitely generated groups (i.e., every finitely generated group either satisfies P or does not satisfy P; more formally, P is a subclass of the class of finitely generated groups). We say that P is a *geometric property of groups* if the following holds for all finitely generated groups G and H: If G satisfies P and H is quasi-isometric to G, then H also satisfies P (i.e., if "having property P" is a quasi-isometry invariant).

Example 4.6.7 (Geometric properties)**.**

- Being finite is a geometric property of groups (Example 4.2.11).
- 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 [31]:

- Being virtually[3] infinite cyclic is a geometric property (Chapter 5.3).
- More generally, for every $n \in \mathbb{N}$ the property of being virtually $\mathbb{Z}^n$ is geometric (Chapter 5.3).
- Being finitely generated and virtually free is a geometric property [110, 31].
- Being finitely generated and virtually nilpotent is a geometric property of groups (Chapter 5.3).

[3] Let P be a property of groups. A group is *virtually* P if it contains a finite index subgroup that has property P.

- Being finitely presented is a geometric property of groups [20, Proposition I.8.24].

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 6) and we will investigate how the geometry of these groups affects their algebraic structure.

4.6.3 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:

Functors, by definition, satisfy a fundamental invariance principle:

Proposition 4.6.8 (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 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.

Definition 4.6.9 (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 5.2.9) [63, Chapter 8].

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 [63, 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 5).

Moreover, there is a general principle turning functors from algebraic topology into quasi-isometry invariants, based on the coarsening construction of Higson and Roe [48, 90, 84]; a more general and more conceptual approach was recently developed by Bunke and Engel [23].

5

Growth types of groups

The first quasi-isometry invariant we discuss in detail is the growth type. We essentially 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 5.3. In contrast, in Chapter 5.4 we will briefly summarise some aspects of exponential growth.

Overview of this chapter.

5.1 Growth functions of finitely generated groups

We start by introducing growth functions of groups with respect to finite generating sets:

Definition 5.1.1 (Growth function). Let G be a finitely generated group and let $S \subset G$ be a finite generating set of G. Then

$$\begin{aligned} \beta_{G,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto \left|B_r^{G,S}(e)\right| \end{aligned}$$

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 4.2.10).

Example 5.1.2 (Growth functions of groups).

- The growth function of the additive group $\mathbb{Z}$ with respect to the generating set $\{1\}$ is clearly 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 5.1):

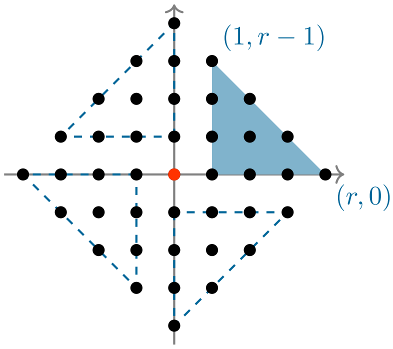

Figure 5.1.: The r-ball in $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ around $(0,0)$

$$\begin{aligned} \beta_{\mathbb{Z}^2,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longrightarrow 1 + 4 \cdot \sum_{j=1}^{r} (r+1-j) = 2 \cdot r^2 + 2 \cdot r + 1. \end{aligned}$$

- More generally, if $n \in \mathbb{N}$, then the growth functions of $\mathbb{Z}^n$ grow like a polynomial of degree n (Exercise).
- 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).
- 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 5.2):

$$\begin{aligned} \beta_{F,S} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto 1 + 2 \cdot n \cdot \sum_{j=0}^{r-1} (2 \cdot n - 1)^j = 1 + \frac{n}{n-1} \cdot \bigl((2 \cdot n - 1)^r - 1\bigr). \end{aligned}$$

Proposition 5.1.3 (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*

$$\beta_{G,S}(r + r') \leq \beta_{G,S}(r) \cdot \beta_{G,S}(r').$$

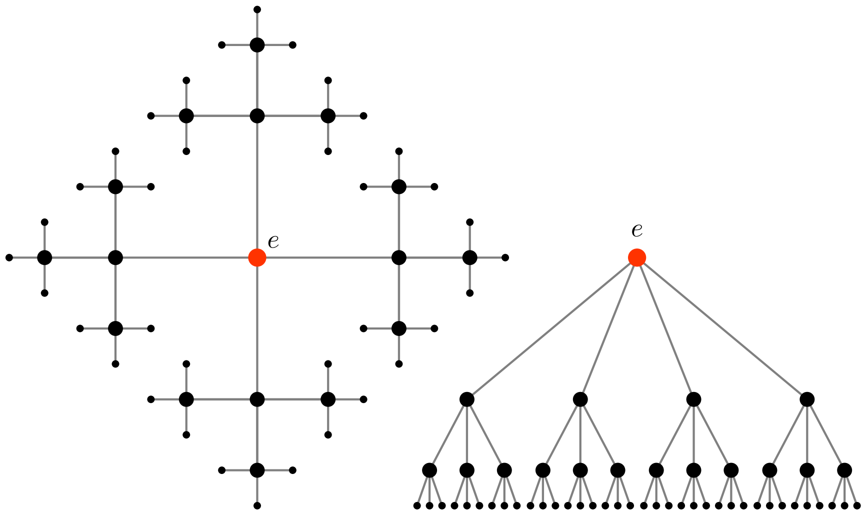

Figure 5.2.: The 3-ball in $\operatorname{Cay}\bigl(F(\{a,b\}),\{a,b\}\bigr)$ around e, drawn in two ways

2. A general lower bound. *Let G be infinite. Then $\beta_{G,S}$ is strictly increasing; in particular, $\beta_{G,S}(r) \geq r$ for all $r \in \mathbb{N}$.*

3. A general upper bound. *For all $r \in \mathbb{N}$ we have*

$$\beta_{G,S}(r) \leq \beta_{F(S),S}(r).$$

Proof. Ad 1./2. This follows easily from the definition of the word metric d_S on G (Exercise).

Ad 3. The homomorphism $\varphi\colon F(S) \longrightarrow G$ characterised by $\varphi|_S = \mathrm{id}_S$ is contracting with respect to the word metrics given by S on $F(S)$ and G respectively. Moreover, φ is surjective. Therefore, we obtain

$$\beta_{G,S}(r) = \bigl|B_r^{G,S}(e)\bigr| = \bigl|\varphi(B_r^{F(S),S}(e))\bigr| \leq \bigl|B_r^{F(S),S}(e)\bigr| = \beta_{F(S),S}(r)$$

for all $r \in \mathbb{N}$. The growth function $\beta_{F(S),S}$ is calculated in Example 5.1.2. □

5.2 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.

5.2.1 Growth types

We therefore introduce the following notion of equivalence for growth functions:

Definition 5.2.1 (Quasi-equivalence of (generalised) growth functions).

- A *generalised growth function* is a function of type $\mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ that is increasing.
- Let f, $g\colon \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ be generalised growth functions. We say that g *quasi-dominates* f if there exist $c, b \in \mathbb{R}_{>0}$ such that

  $$\forall_{r \in \mathbb{R}_{\geq 0}} \quad f(r) \leq c \cdot g(c \cdot r + b) + b.$$

  If g quasi-dominates f, then we write $f \prec g$.
- Two generalised growth functions f, $g\colon \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}$ are *quasi-equivalent* if both $f \prec g$ and $g \prec f$; if f and g are quasi-equivalent, then we write $f \sim g$.

A straightforward computation shows that quasi-equivalence is an equivalence relation on the set of all generalised growth functions (check!). Quasi-domination then induces a partial order on the set of equivalence classes; however, this partial order is *not* total (check!).

Example 5.2.2 (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 $r \in \mathbb{R}_{\geq 0}$

  $$r^a \leq r^{a'} + 1 \leq (r+1)^{a'} + 1,$$

  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 an $r \in \mathbb{R}_{\geq 0}$ 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 $a = a'$.

Moreover: If $p \in \mathbb{R}[X]$ is a polynomial all of whose coefficients are non-negative, then $p \sim (x \mapsto x^{\deg p})$.

- *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

$$(x \mapsto a^x) \succ (x \mapsto x^{a'}) \quad \text{and} \quad (x \mapsto a^x) \not\prec (x \mapsto x^{a'})$$

for all $a \in \mathbb{R}_{>1}$ and all $a' \in \mathbb{R}_{\geq 0}$ (check!). 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 $a \in \mathbb{R}_{>1}$, $a' \in \mathbb{R}_{\geq 0}$ (check!).

Example 5.2.3 (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

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ r &\longmapsto \beta_{G,S}(\lceil r \rceil) \end{aligned}$$

associated with the growth function $\beta_{G,S} \colon \mathbb{N} \longrightarrow \mathbb{N}$ is indeed a generalised growth function (which is also sub-multiplicative; check!).

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

$$\forall_{r \in \mathbb{N}} \quad \beta_{G,S}(r) \leq c \cdot \beta_{H,T}(c \cdot r + b) + b.$$

5.2.2 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 5.2.4 (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
$$d_S(g,g') \leq c \cdot \bigl(d_T(f(g), f(g')) + c\bigr) = c^2.$$

Using these two properties and that word metrics are invariant under left translation, it follows that

$$\begin{aligned} \beta_{G,S}(r) &\leq \bigl|B_{c^2}^{G,S}(e)\bigr| \cdot \bigl|B_{c \cdot r + c}^{H,T}(e')\bigr| \\ &= \bigl|B_{c^2}^{G,S}(e)\bigr| \cdot \bigl|B_{c \cdot r + c}^{H,T}(e)\bigr| \\ &= \beta_{G,S}(c^2) \cdot \beta_{H,T}(c \cdot r + c). \end{aligned}$$

This shows that $\beta_{G,S} \prec \beta_{H,T}$ (the term $\beta_{G,S}(c^2)$ does not depend on the radius r). □

Proposition 5.2.4 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 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 (check!).

In particular, we can define the growth type of finitely generated groups:

Definition 5.2.5 (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 5.1.3 and Example 5.1.2), and that polynomials and exponential functions are not quasi-equivalent (Example 5.2.2); 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 5.2.4 and Example 5.2.2 that having exponential growth/polynomial growth/intermediate growth respectively is a geometric property of groups. More generally:

Corollary 5.2.6 (Quasi-isometry invariance of the growth type)**.** *By Proposition 5.2.4, 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 5.2.7 (Growth types)**.** From Example 5.1.2 we obtain:

- If $n \in \mathbb{N}$, then $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$ (Exercise).
- The Heisenberg group has the growth type of $(x \mapsto x^4)$ (Exercise).
- Non-Abelian free groups of finite rank have the growth type of the exponential function $(x \mapsto e^x)$.

The groups $\mathbb{Z}^n$ and the Heisenberg group hence have polynomial growth, while non-Abelian free groups have exponential growth.

Example 5.2.8 (Quasi-isometry classification of Abelian groups). In analogy with the topological invariance of dimension, the following holds: We can recover the rank of free Abelian groups from their quasi-isometry type: 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 5.2.2, Example 5.2.7, and Corollary 5.2.6. 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 $\operatorname{rk}_{\mathbb{Z}} A := r$. Hence, combining the above observation with Corollary 4.4.5, we obtain for all finitely generated Abelian groups A and A' the equivalence

$$A \sim_{\mathrm{QI}} A' \iff \operatorname{rk}_{\mathbb{Z}} A = \operatorname{rk}_{\mathbb{Z}} A'.$$

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 [71].

Example 5.2.9 (Distinguishing quasi-isometry types of basic groups).

- We obtain for the Heisenberg group H that $H \not\sim_{\mathrm{QI}} \mathbb{Z}^3$ (Example 5.2.7), 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

$$F \not\sim_{\mathrm{QI}} \mathbb{Z}^n \quad \text{and} \quad F \not\sim_{\mathrm{QI}} H.$$

Grigorchuk was the first to show that there do indeed exist groups that have intermediate growth [39][44, Chapter VIII]:

Theorem 5.2.10 (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, which can be described via automorphisms of trees or as an automatic group. Furthermore, this group also has several other interesting properties [44, Chapter VIII]; for example, it is a finitely generated infinite torsion group and it is commensurable to the direct product with itself [63, Exercises 4.E.31ff, Exercises 6.E.11–6.E.13].

Proposition 5.2.11 (Growth of subgroups). *Let G be a finitely generated group and let H be a finitely generated subgroup of G. If T is a finite generating set of H, and S is a finite generating set of G, then*

$$\beta_{H,T} \prec \beta_{G,S}.$$

Proof. Let $S' := S \cup T$; then S' is a finite generating set of G. Let $r \in \mathbb{N}$; then for all $h \in B_r^{H,T}(e)$ we have

$$d_{S'}(h,e) \leq d_T(h,e) \leq r,$$

and so $B_r^{H,T}(e) \subset B_r^{G,S'}(e)$. In particular,

$$\beta_{H,T}(r) \leq \beta_{G,S'}(r),$$

and thus $\beta_{H,T} \prec \beta_{G,S'}$. Moreover, we know that (G, d_S) and $(G, d_{S'})$ are quasi-isometric, and hence the growth functions $\beta_{G,S'}$ and $\beta_{G,S}$ are quasi-equivalent by Proposition 5.2.4. Therefore, we obtain $\beta_{H,T} \prec \beta_{G,S}$. □

Example 5.2.12 (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; e.g., there exist solvable groups of exponential growth (Exercise).

Caveat 5.2.13 (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 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.

5.2.3 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

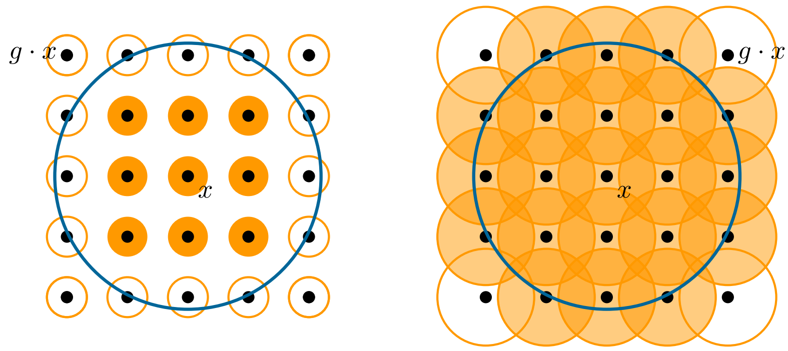

(a) (b)

Figure 5.3.: Packing balls into balls, and covering balls by balls

the growth type of the group acting and the growth type of the metric space acted upon. One instance of this type of result is the following [80]:

Proposition 5.2.14 (Švarc–Milnor lemma for growth types)**.** *Let M be a non-empty 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 $\widetilde{M}$ (not of 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 4.4.10). 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 Riemannian metric on $\widetilde{M}$, and "$\operatorname{vol}_{\widetilde{M}}$" denotes the Riemannian volume with respect to the Riemannian metric on $\widetilde{M}$.

Sketch of proof. By the Švarc–Milnor lemma (Corollary 4.4.10), 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 us to translate radii for balls in $\pi_1(M)$ into radii for balls in $\widetilde{M}$, and vice versa. Therefore, it only remains to translate counting of points in the orbit $\pi_1(M) \cdot x$ in a ball to the Riemannian volume of such a ball.

- 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 5.3 (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 5.3 (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 [44, Proposition VI.36]. □

In particular, we obtain the following obstruction for the existence of maps of non-zero degree:

Corollary 5.2.15 (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*
$$\bigl(r \mapsto \operatorname{vol}_{\widetilde{M}} B^{\widetilde{M}}_r(x)\bigr) \succ \bigl(r \mapsto \operatorname{vol}_{\widetilde{N}} B^{\widetilde{N}}_r(y)\bigr)$$
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 [29, Chapter VIII.4].

Proof. In view of Proposition 5.2.14 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 $\operatorname{im} \pi_1(p) = G$ [77, Theorem V.10.2 and V.4.2]; in particular, $\overline{N}$ is also an orientable 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 $p \circ \overline{f} = f$ [77, Theorem V.5.1]:

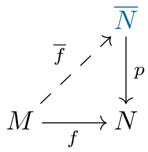

Looking at the induced diagram in singular homology with integral coefficients in top degree $n := \dim M = \dim N = \dim \overline{N}$ shows $H_n(\overline{N};\mathbb{Z}) \not\cong 0$:

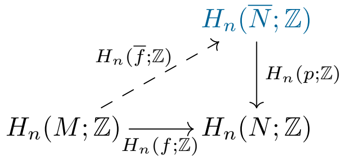

Namely, $H_n(M;\mathbb{Z}) \cong \mathbb{Z} \cong H_n(N;\mathbb{Z})$ and $H_n(\overline{f};\mathbb{Z})$ corresponds to multiplication by the mapping degree $\deg f$. In particular, $\overline{N}$ is compact [29, Corollary VIII.3.4]. On the other hand, $|\deg p|$ coincides with the number of sheets of the covering p [29, Proposition VIII.4.7], which in turn equals the index of $\operatorname{im} \pi_1(p)$ in $\pi_1(N)$ [77, p. 133]. Hence,

$$\bigl[\pi_1(N) : G\bigr] = \bigl[\pi_1(N) : \operatorname{im} \pi_1(p)\bigr] = |\deg p| < \infty,$$

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 5.2.16 (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 [89, Chapter 3.4][68, Example 1.5.22]. □

In Chapter 6, we will discuss a concept of negative curvature for finitely generated groups, leading to generalisations of Corollary 5.2.16. Alternatively, Corollary 5.2.16 can also be obtained via simplicial volume [60, 61].

5.3 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 [41] was subsequently simplified by van den Dries and Wilkie [30, 74]; alternative proofs have been given by Kleiner [50], Shalom and Tao [99], Ozawa [88], and Breuillard, Green, and Tao [16]. A complete proof is also given in the textbook by Druţu and Kapovich [31].

Theorem 5.3.1 (Gromov's polynomial growth theorem). *Finitely generated groups have polynomial growth if and only if they are virtually nilpotent.*

In Chapter 5.3.1 and 5.3.2 we briefly discuss nilpotent groups and their growth properties; in Chapter 5.3.3, 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.

5.3.1 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 5.3.2 ((Virtually) nilpotent group).

- Let G be a group. For $n \in \mathbb{N}$ we inductively define $C_{(n)}(G)$ by

  $$C_{(0)}(G) := G \quad \text{and} \quad \forall_{n\in\mathbb{N}} \quad C_{(n+1)}(G) := \bigl[G, C_{(n)}(G)\bigr].$$

  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 5.3.3 ((Virtually) solvable group).

- Let G be a group. For $n \in \mathbb{N}$ we inductively define $G^{(n)}$ by

  $$G^{(0)} := G \quad \text{and} \quad \forall_{n\in\mathbb{N}} \quad G^{(n+1)} := [G^{(n)}, G^{(n)}].$$

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 [54, Chapter VI.7][64, Kapitel 3.5.2].

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 5.3.4 (Nilpotent/solvable groups).

- Free groups of rank at least 2 are *not* virtually solvable (Exercise).
- 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 (check!)
$$C_{(1)}(H) = [H, H] \cong \langle z\rangle_H,$$
and hence $C_{(2)}(H) = [H, C_{(1)}(H)] \cong [H, \langle z\rangle_H] = \{e\}$ (check!).
- 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 for $n \in \mathbb{N}_{\geq 5}$ are simple and non-Abelian and so not even solvable [54, 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 solvable but not virtually nilpotent (this can be shown by a direct computation).

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 0.2):

Proposition 5.3.5 (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 (check!). □

Except for the last statement on centrality, the analogous statements also hold for the derived series instead of the lower central series (check!).

5.3.2 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 5.3.6 (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 5.3.5). Moreover, it can be shown that it is finitely generated (because G is finitely generated [111, Lemma 3.7][73, Theorem 5.4]). Therefore, $\operatorname{rk}_{\mathbb{Z}} C_{(j)}(G)/C_{(j+1)}(G)$ is well-defined.

The proof of Theorem 5.3.6 proceeds by induction over the nilpotence degree n and uses a suitable normal form of group elements in terms of the lower central series [8, 31]; the arguments from the computation of the growth rate of the Heisenberg group (Exercise) give a first impression of how this inductive proof works. We refrain from going into the details for the general case.

Example 5.3.7. For the Heisenberg group H, one can compute directly that it has polynomial growth of degree 4 (Exercise). On the other hand, we know that $C_{(2)}(H) = \{e\}$ and obtain

$$\sum_{j=0}^{2-1} (j+1) \cdot \operatorname{rk}_{\mathbb{Z}} C_{(j)}(H)/C_{(j+1)}(H) = 1 \cdot \operatorname{rk}_{\mathbb{Z}} \mathbb{Z}^2 + 2 \cdot \operatorname{rk}_{z} \mathbb{Z} = 4,$$

as predicted by Theorem 5.3.6.

Caveat 5.3.8 (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 5.3.1) and the fact that there exist solvable groups that are not virtually nilpotent (Example 5.3.4); moreover, it can also be shown by elementary calculations that the group $\mathbb{Z}^2 \rtimes_A \mathbb{Z}$ given by the action of the matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}$$

on $\mathbb{Z}^2$ has exponential growth (Exercise).

Wolf [111] and Milnor [80] used algebraic means (similar to the calculations in Caveat 5.3.8) to prove the following predecessor of the polynomial growth theorem:

Theorem 5.3.9 (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.

5.3.3 Polynomial growth implies virtual nilpotence

We sketch Gromov's argument that finitely generated groups of polynomial growth are virtually nilpotent, mainly following the exposition by van den Dries and Wilkie [30]:

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 of Gromov [41]:

Theorem 5.3.10. *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's 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 5.3.10. We will now briefly sketch Gromov's argument:

Sketch of proof of Theorem 5.3.10. 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 [41]. Using the solution of Hilbert's fifth problem [81, 104], 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 3.4.3)) allow us to construct a surjective homomorphism from a finite index subgroup of G' to $\mathbb{Z}$.

A detailed proof is given in the paper by van den Dries and Wilkie [30]. Gromov's considerations of the sequence $(G, 1/n \cdot d_S)_{n \in \mathbb{N}}$ are a precursor of asymptotic cones [32, 31]. □

In view of Theorem 5.3.10 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 5.3.11 (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 can be shown by an elementary argument (Exercise).

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

$$\forall_{r \in \mathbb{N}} \quad \beta_{G,S}(r) \leq c \cdot r^d.$$

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 (check!), 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 5.3.11, 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 [30, Lemma 2.1]:

Lemma 5.3.12. *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 G is also 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 a subgroup of H) of finite index in K (Exercise) 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 5.3.9 lets us deduce that G is indeed virtually nilpotent, as desired. This finishes the proof sketch of Gromov's polynomial growth theorem.

5.3.4 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 4.6.6. 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 5.3.13. *Being virtually nilpotent is a geometric property of finitely generated groups.*

Proof. In view of Gromov's polynomial growth theorem (Theorem 5.3.1), 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 5.2.6). □

5.3.5 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 5.3.14 (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 5.3.1), 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 5.3.6 (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 5.3.15 (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 5.2.14). Therefore, the previous corollary implies integrality of the growth exponent. □

5.3.6 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 5.3.16 (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 5.3.1). Let $H \subset G$ be a nilpotent subgroup of finite index; so H

has linear growth as well. By Bass' theorem on the growth rate of nilpotent groups (Theorem 5.3.6) 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

$$1 = \operatorname{rk}_{\mathbb{Z}} C_{(0)}(H)/C_{(1)}(H) \quad \text{and} \quad \forall_{j \in \mathbb{N}_{\geq 1}} \quad 0 = \operatorname{rk}_{\mathbb{Z}} C_{(j)}(H)/C_{(j+1)}(H).$$

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 $H = C_{(0)}(H)$ is also virtually $\mathbb{Z}$ (check!). In particular, G is virtually $\mathbb{Z}$. □

We will see more elementary proofs of the quasi-isometry rigidity of $\mathbb{Z}$ in Chapter 6 (Corollary 6.5.8).

More generally, a similar argument yields quasi-isometry rigidity of higher-dimensional Abelian groups [21, Theorem 5.8]:

Corollary 5.3.17 (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 [20, 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 [98, 26]. On the other hand, a full quasi-isometry classification of virtually nilpotent groups is out of reach.

5.3.7 Application: Expanding maps of manifolds

We conclude the discussion of polynomial growth with Gromov's geometric application [41, Geometric corollary on p. 55] of the polynomial growth theorem (Theorem 5.3.1) to infra-nil-endomorphisms:

Corollary 5.3.18. *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

$$\forall_{x,x' \in X} \quad x \neq x' \Longrightarrow d_Y\big(f(x), f(x')\big) > d_X(x, x').$$

A map $f \colon X \longrightarrow Y$ is *expanding* if every point of X has a neighbourhood U such that the restriction $f|_U \colon U \longrightarrow Y$ is expanding. As Riemannian manifolds can be viewed as metric spaces, we obtain a notion of expanding maps of Riemannian manifolds.

As a simple example, let us consider the n-dimensional torus $\mathbb{Z}^n \setminus \mathbb{R}^n$. A straightforward calculation shows that a linear map $f \colon \mathbb{R}^n \longrightarrow \mathbb{R}^n$ with $f(\mathbb{Z}^n) \subset \mathbb{Z}^n$ induces a self-map $\mathbb{Z}^n \setminus \mathbb{R}^n \longrightarrow \mathbb{Z}^n \setminus \mathbb{R}^n$ and that this self-map is expanding if and only if all complex eigenvalues of f have absolute value bigger than 1.

A *nil-manifold* is a compact Riemannian manifold that can be obtained as a quotient $\Gamma \setminus N$, where N is a simply connected nilpotent Lie group and $\Gamma \subset N$ is a cocompact lattice. More generally, an *infra-nil-manifold* is a compact Riemannian manifold that can be obtained as a quotient $\Gamma \setminus N$, where N is a simply connected nilpotent Lie group and Γ is a subgroup of the group of all isometries of N generated by left translations of N and all automorphisms of N. Clearly, all nil-manifolds are also infra-nil-manifolds, and it can be shown that every infra-nil-manifold is finitely covered by a nil-manifold.

Let $\Gamma \setminus N$ be such an infra-nil-manifold. An expanding *infra-nil-endomorphism* is an expanding map $\Gamma \setminus N \longrightarrow \Gamma \setminus N$ that is induced by an expanding automorphism $N \longrightarrow N$ of the Lie group N.

For example, all tori and the quotient $H \setminus H_{\mathbb{R}}$ of the Heisenberg group $H_{\mathbb{R}}$ with real coefficients by the Heisenberg group H are nil-manifolds (and so also infra-nil-manifolds). The expanding maps on tori mentioned above are examples of expanding infra-nil-endomorphisms.

Two self-maps $f \colon X \longrightarrow X$ and $g \colon Y \longrightarrow Y$ between topological spaces are *topologically conjugate* if there exists a homeomorphism $h \colon X \longrightarrow Y$ with $h \circ f = g \circ h$, i.e., which fits into a commutative diagram:

$$\begin{array}{ccc} X & \xrightarrow{\;f\;} & X \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle h} \\ Y & \xrightarrow[\;g\;]{} & Y \end{array}$$

Sketch of proof of Corollary 5.3.18. By a theorem of Franks [35], 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 5.2.14), the fundamental group $\pi_1(M)$ is finitely generated and has polynomial growth as well.

In view of the polynomial growth theorem (Theorem 5.3.1), we obtain that $\pi_1(M)$ is virtually nilpotent.

By a result of Shub [100, 35], this implies that every expanding self-map of M is topologically conjugate to an infra-nil-endomorphism. □

5.4 Groups of uniform exponential growth

In contrast to Chapter 5.3, we will now focus on groups of exponential growth. We will first introduce a stronger version of exponential growth (Chapter 5.4.1). We will then discuss the interesting relation between exponential growth and number theory (Chapter 5.4.2–5.4.3), as discovered by Breuillard [15, 17].

5.4.1 Uniform exponential growth

The rate of exponential growth can be measured as follows.

Proposition 5.4.1 (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 5.1.3). 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 5.4.2 (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 5.1.2 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).

Definition 5.4.3 (Uniform exponential growth)**.** A finitely generated group G has *uniform exponential growth* if

$$\inf\bigl\{\varrho_{G,S} \bigm| S \subset G \text{ is a finite generating set}\bigr\} > 1.$$

Example 5.4.4 (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 3.2.7) 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 5.4.2.

It seems to be an open problem to decide whether having uniform exponential growth is a quasi-isometry invariant or not [45].

It is known that finitely generated solvable groups [85] and finitely generated linear groups [34] 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 [108].

5.4.2 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 5.4.5 (Breuillard's growth conjecture [15]). For every $d \in \mathbb{N}$ there exists an $\varepsilon(d) \in \mathbb{R}_{>0}$ 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 $\varrho_{\langle S \rangle_{\mathrm{GL}(d,K)}, S} > 1 + \varepsilon(d)$.

While the growth conjecture is open in general, partial results are known:

Theorem 5.4.6 (Growth gap [15]). *For every $d \in \mathbb{N}$ there exists an $\varepsilon(d) \in \mathbb{R}_{>0}$ 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 [63, Chapter 6.4.3].

5.4.3 Application: The Lehmer conjecture

We digress briefly to 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 [12].

Numerical experiments support the following gap phenomenon:

Conjecture 5.4.7 (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 [15, 17].

Theorem 5.4.8 (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 5.4.9. 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

$$S(\alpha) := \bigl\{e, A(\alpha), A(\alpha)^{-1}, B, B^{-1}\bigr\} \subset \mathrm{GL}(2,\overline{\mathbb{Q}}).$$

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). 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. Hence, despite the growth gap theorem (Theorem 5.4.6), 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 [17], which are far beyond the scope of this course.

6

Hyperbolic groups

In the universe of groups (Figure 0.2), 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 surfaces (Chapter 6.1). We will then introduce Gromov's extension of the notion of negative curvature to large scale geometry via slim triangles (Chapter 6.2). In particular, this leads to a notion of negatively curved finitely generated groups: hyperbolic groups (Chapter 6.3).

The hyperbolicity condition for groups has far-reaching algebraic consequences: The word problem is solvable for hyperbolic groups (Chapter 6.4) and elements of infinite order in hyperbolic groups are well-behaved (Chapter 6.5).

Overview of this chapter.

6.1 Classical curvature, intuitively

The 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 and they usually rely on a notion of second derivatives (and their effect on curves, triangles, angles, volumes, ...).

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 [58].

The *Gaussian curvature* of a smooth surface S embedded into $\mathbb{R}^3$ at a point $x \in S$ is the product of the two principal curvatures at x, i.e., the product of the supremum and the infimum of the signed curvature at x of smooth curves in S passing through x. While the principal curvatures are not intrinsic invariants of a surface (i.e., they are in general *not* invariant under isometries) [58, p. 6], the Gaussian curvatures of a surface are intrinsic (*Theorema Egregium* [58, Chapter 8]).

Example 6.1.1. The following examples are illustrated in Figure 6.1.

- 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 [63, Appendix A.3]. For example, one can calculate via the Theorema Egregium that the hyperbolic plane has everywhere negative Gaussian curvature [63, Theorem A.3.29].

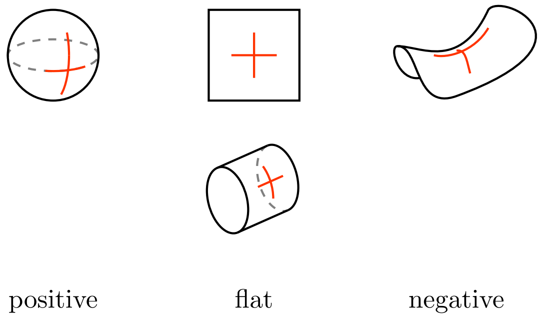

Figure 6.1.: Examples of Gaussian curvatures of surfaces

Outlook 6.1.2 (Curvature of 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 at 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 is in fact an intrinsic invariant of Riemannian manifolds and can be described analytically in terms of tensors on M [58, 68]. Taking suitable averages of sectional curvatures leads to the weaker curvature notions of *Ricci curvature* and *scalar curvature*, respectively.

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. 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. To this end, we look at geodesic triangles in surfaces in $\mathbb{R}^3$ (Figure 6.2), i.e., at triangles in surfaces whose sides are geodesics in the surface in question.

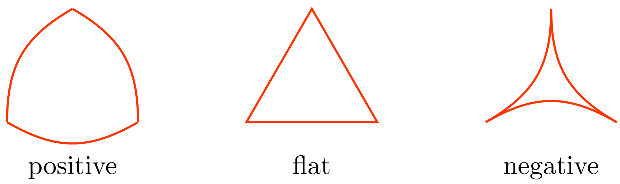

Figure 6.2.: Geodesic triangles in surfaces

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 [63, Theorem A.3.27].

6.2 (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 [42]. 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 6.2.1). As the next step, we translate this notion to quasi-geometry (Chapter 6.2.2). Finally, in Chapter 6.2.3, we establish the quasi-isometry invariance of hyperbolicity.

6.2.1 Hyperbolic spaces

Taking slim geodesic triangles as the defining property for negative curvature leads to the notion of (Gromov) hyperbolic spaces:

Definition 6.2.1 (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 6.3)

$$\begin{aligned} \operatorname{im}\gamma_0 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_1 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_1 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_2\bigr),\\ \operatorname{im}\gamma_2 &\subset B_\delta^{X,d}\bigl(\operatorname{im}\gamma_0 \cup \operatorname{im}\gamma_1\bigr). \end{aligned}$$

  Here, for $\gamma \colon [0, L] \longrightarrow X$ we use the abbreviation $\operatorname{im}\gamma := \gamma([0, L])$, and for $A \subset X$ we write $B_\delta^{X,d}(A) := \bigl\{x \in X \bigm| \exists_{a \in A} \quad d(x, a) \leq \delta\bigr\}$.

Definition 6.2.2 (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 6.2.3 (Hyperbolic spaces)**.**

- Every geodesic metric space X of finite diameter is $\operatorname{diam}(X)$-hyperbolic.

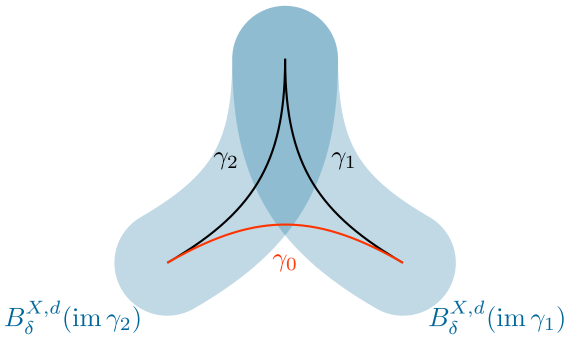

Figure 6.3.: A δ-slim triangle

- 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 6.4).
- The hyperbolic plane $\mathbb{H}^2$ is a hyperbolic metric space in the sense of Definition 6.2.2 [63, 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 6.2.2 [20, Chapter II.1.A, Proposition III.H.1.2].
- If T is a tree, then the geometric realisation $|T|$ of T (Chapter 4.3.2) is 0-hyperbolic because all geodesic triangles in $|T|$ are degenerate tripods (see also Proposition 6.2.17 below); in a sense, hyperbolic spaces can be viewed as thickenings of metric trees [63, Exercise 7.E.9].

Caveat 6.2.4 (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 in general is only a quasi-geodesic triangle and not a geodesic triangle. However, we will see in Corollary 6.2.13 below that hyperbolicity is indeed a quasi-isometry invariant among geodesic spaces.

6.2.2 Quasi-hyperbolic spaces

We translate the definition of hyperbolicity to quasi-geometry; so instead of geodesics we consider quasi-geodesics. This results in a notion of quasi-

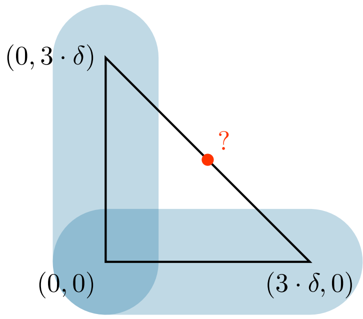

Figure 6.4.: The Euclidean plane $\mathbb{R}^2$ is *not* hyperbolic

hyperbolicity for quasi-geodesic spaces. On the one hand, it will be immediate from the definition that quasi-hyperbolicity is indeed a quasi-isometry invariant (Proposition 6.2.9); on the other hand, we will relate hyperbolicity of geodesic spaces to quasi-hyperbolicity (Theorem 6.2.10), which shows that hyperbolicity is also a quasi-isometry invariant in the class of geodesic spaces (Corollary 6.2.13).

Definition 6.2.5 (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 6.5)

$$\begin{aligned} \operatorname{im} \gamma_0 &\subset B_\delta^{X,d}\bigl(\operatorname{im} \gamma_1 \cup \operatorname{im} \gamma_2\bigr), \\ \operatorname{im} \gamma_1 &\subset B_\delta^{X,d}\bigl(\operatorname{im} \gamma_0 \cup \operatorname{im} \gamma_2\bigr), \\ \operatorname{im} \gamma_2 &\subset B_\delta^{X,d}\bigl(\operatorname{im} \gamma_0 \cup \operatorname{im} \gamma_1\bigr). \end{aligned}$$

Definition 6.2.6 (Quasi-hyperbolic space)**.** Let X be a metric space.

- Let c, $b \in \mathbb{R}_{>0}$, $\delta \in \mathbb{R}_{\geq 0}$. We say that X is (c, b, δ)-*quasi-hyperbolic* if X is (c, b)-quasi-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.

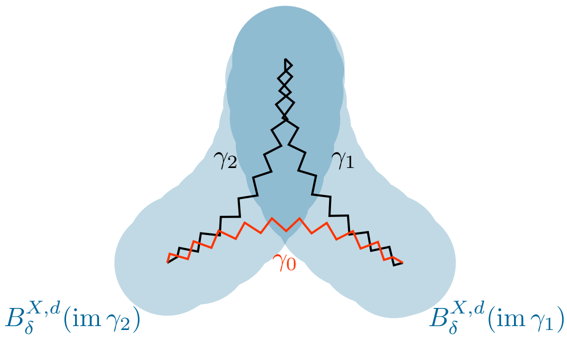

Figure 6.5.: A quasi-slim quasi-geodesic triangle

- The space X is *quasi-hyperbolic* if there exist c, $b \in \mathbb{R}_{>0}$ such that X is (c, b)-quasi-hyperbolic.

Example 6.2.7. All metric spaces of finite diameter are quasi-hyperbolic.

Caveat 6.2.8. 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 6.2.13 below simplifies this task considerably if 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 6.2.9 (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 X is also 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 X is also 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.

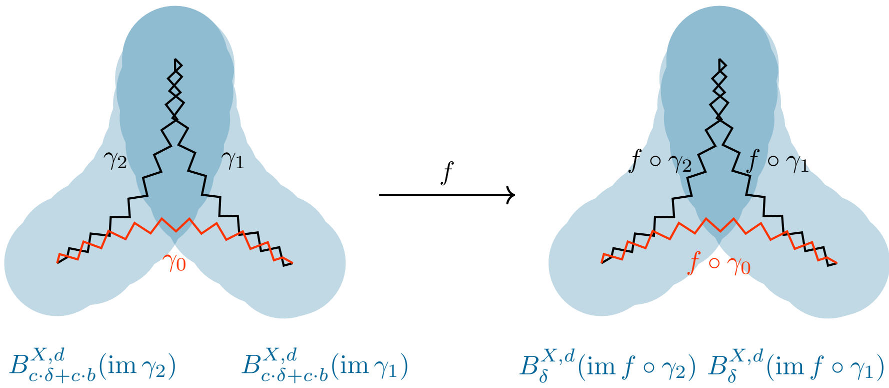

Figure 6.6.: Using a quasi-isometric embedding to translate quasi-geodesic triangles and neighbourhoods back and forth

We start by proving the first part. Let Y be quasi-geodesic and suppose that $f\colon X \longrightarrow Y$ is a quasi-isometric embedding; let $c \in \mathbb{R}_{\geq 0}$ be so large that Y is (c,c)-quasi-geodesic and such that f is a (c,c)-quasi-isometric embedding with c-dense image. Furthermore, let $x, x' \in X$. Then there is a (c,c)-quasi-geodesic $\gamma\colon [0,L] \longrightarrow Y$ joining $f(x)$ and $f(x')$. Using the axiom of choice and the fact that f has c-dense image, we can find a map

$$\widetilde{\gamma}\colon [0,L] \longrightarrow X$$

such that $\widetilde{\gamma}(0) = x$, $\widetilde{\gamma}(L) = x'$, and

$$d_Y\bigl(f \circ \widetilde{\gamma}(t), \gamma(t)\bigr) \leq c$$

for all $t \in [0,L]$. The same arguments as in the proof of Proposition 4.1.10 show that $\widetilde{\gamma}$ is a $(c, \max(3\cdot c^2, 3))$-quasi-geodesic joining x and x'. Hence, X is $(c, \max(3\cdot c^2, 3))$-quasi-geodesic.

As for the second part, suppose that Y is quasi-hyperbolic, that X is quasi-geodesic, and that $f\colon X \longrightarrow Y$ is a quasi-isometric embedding. Hence, there are c, $b \in \mathbb{R}_{>0}$ such that Y is (c,b)-quasi-hyperbolic, such that X is (c,b)-quasi-geodesic, and such that f is a (c,b)-quasi-isometric 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 6.6 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 (by the built-in freedom of quasi-hyperbolicity!). 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. □

6.2.3 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 6.2.13). To this end we first compare hyperbolicity and quasi-hyperbolicity on geodesic spaces (Theorem 6.2.10); then we apply quasi-isometry invariance of quasi-hyperbolicity.

Theorem 6.2.10 (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 are indeed 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 6.2.11 (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*

$$\operatorname{im}\gamma' \subset B^{X,d}_{\Delta}(\operatorname{im}\gamma) \quad \textit{and} \quad \operatorname{im}\gamma \subset B^{X,d}_{\Delta}(\operatorname{im}\gamma').$$

Caveat 6.2.12. In general, the stability theorem for quasi-geodesics does *not* hold in non-hyperbolic spaces: For example, the logarithmic spiral (Figure 6.7)

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}^2 \\ t &\longmapsto t \cdot \bigl(\sin(\ln(1+t)), \cos(\ln(1+t))\bigr) \end{aligned}$$

[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.

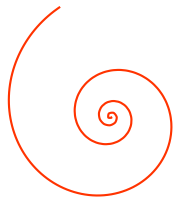

Figure 6.7.: The logarithmic spiral

is a quasi-isometric embedding with respect to the standard metrics on $\mathbb{R}$ and $\mathbb{R}^2$ (Exercise), but this quasi-geodesic ray does not have bounded distance from any geodesic ray (Exercise). 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 for geodesic spaces:

Proof of Theoren 6.2.10. 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 6.2.11) for the constants c, b, δ. We show 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 6.8). In particular, $(\gamma_0', \gamma_1', \gamma_2')$ is a geodesic triangle in X, and by the stability theorem (Theorem 6.2.11)

$$\operatorname{im} \gamma_j' \subset B_\Delta^{X,d}(\operatorname{im} \gamma_j) \quad \text{and} \quad \operatorname{im} \gamma_j \subset B_\Delta^{X,d}(\operatorname{im} \gamma_j')$$

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. $\square$

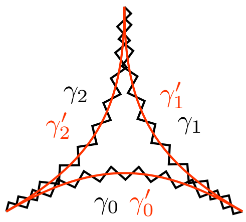

Figure 6.8.: Approximating quasi-geodesic triangles by geodesic triangles

Corollary 6.2.13 (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 6.2.10, the space Y is also quasi-hyperbolic. Therefore, X is quasi-hyperbolic as well (Proposition 6.2.9).

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 6.2.9. In particular, Y is hyperbolic.

Ad 3. This follows easily from the previous parts and Theorem 6.2.10. ☐

It remains to prove the stability theorem (Theorem 6.2.11). In order to do so, we need two facts about approximating curves by geodesics:

Lemma 6.2.14 (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:*

$$L_X(\gamma) := \sup\biggl\{\sum_{j=0}^{k-1} d\bigl(\gamma(t_j), \gamma(t_{j+1})\bigr) \biggm| k \in \mathbb{N},\ t_0, \dots, t_k \in [0, L],$$
$$t_0 \leq t_1 \leq \dots \leq t_k\biggr\} \in \mathbb{R}_{\geq 0} \cup \{\infty\}.$$

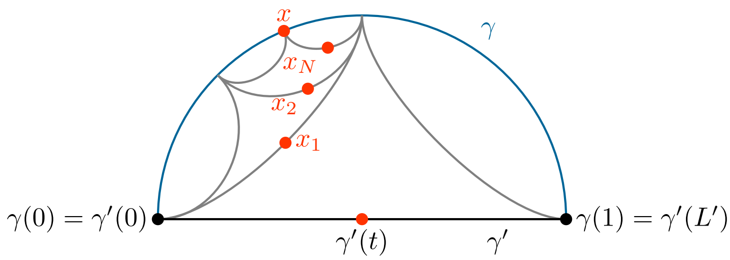

Figure 6.9.: Distance from geodesics to curves in hyperbolic spaces

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 (check!). 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 6.9). In particular, there is an $x \in \operatorname{im}\gamma$ with

$$\begin{aligned} d\bigl(\gamma'(t), x\bigr) &\leq d\bigl(\gamma'(t), x_N\bigr) + d(x_N, x) && \text{(triangle inequality)} \\ &\leq \delta \cdot N + \frac{L_X(\gamma)}{2^{N+1}} && \text{(by construction of } x_1, \dots, x_N) \\ &\leq \delta \cdot \bigl|\log_2(L_X(\gamma))\bigr| + 1, && \text{(by construction of } N) \end{aligned}$$

as claimed □

Lemma 6.2.15 (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, $\operatorname{im}\gamma' \subset B^{X,d}_{c+b}(\operatorname{im}\gamma)$ and $\operatorname{im}\gamma \subset B^{X,d}_{c+b}(\operatorname{im}\gamma')$.*

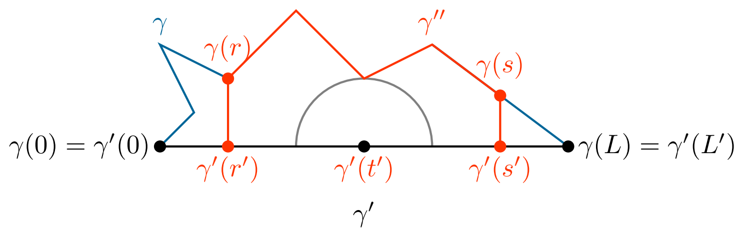

Figure 6.10.: In hyperbolic spaces, geodesics are close to quasi-geodesics

Proof. Let $I := \big([0,L) \cap \mathbb{Z}\big) \cup \{L\}$. As a first step, we set $\gamma'|_I := \gamma|_I$. We then geodesically connect the dots: We 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 (check!). $\square$

Proof of Theorem 6.2.11. Let γ be a quasi-geodesic and let γ' be a geodesic as in the statement of the stability theorem (Theorem 6.2.11). In view of Lemma 6.2.15, 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 the said lemma, i.e.,

$$L_X(\gamma|_{[r,s]}) \leq c \cdot d\big(\gamma(r), \gamma(s)\big) + b$$

for all r, $s \in [0, L]$ with $r \leq s$.

As a *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

$$\Delta := \sup\big\{d(\gamma'(t'), \operatorname{im} \gamma) \bigm| t' \in [0, L']\big\};$$

as γ is continuous, a topological argument shows that there is a $t' \in [0, L]$ at which this supremum is attained (check!). We deduce an upper bound for Δ (see Figure 6.10 for an illustration of the notation):

Let

$$r' := \max(0, t' - 2 \cdot \Delta) \quad \text{and} \quad s' := \min(L', t' + 2 \cdot \Delta);$$

by construction of Δ, there exist r, $s \in [0, L]$ with

$$d\big(\gamma(r), \gamma'(r')\big) \leq \Delta \quad \text{and} \quad d\big(\gamma(s), \gamma'(s')\big) \leq \Delta.$$

We consider the curve γ'' in X that starts at $\gamma'(r')$, then follows a geodesic to $\gamma(r)$, then follows γ until $\gamma(s)$, and finally follows a geodesic to $\gamma'(s')$. From Lemma 6.2.14 and the construction of γ'' we obtain that

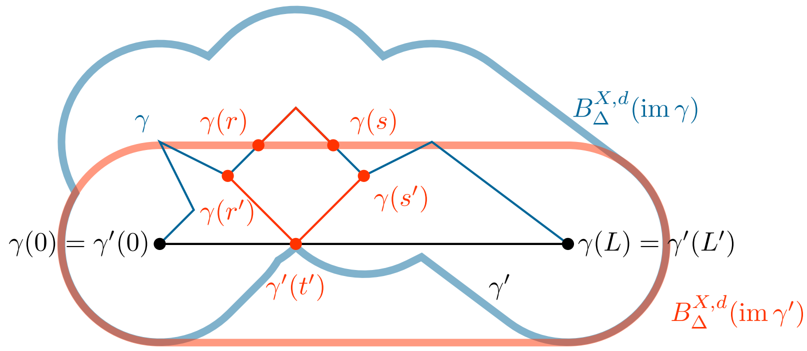

Figure 6.11.: In hyperbolic spaces, quasi-geodesics are close to geodesics

$$\Delta \leq d\bigl(\gamma'(t'), \operatorname{im}\gamma''\bigr) \leq \delta \cdot \bigl|\log_2 L_X(\gamma'')\bigr| + 1;$$

moreover, because γ satisfies the length estimate from Lemma 6.2.15 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 a *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 $B^{X,d}_\Delta(\operatorname{im}\gamma')$. To this end, let r, $s \in [0, L]$ be such that $[r, s]$ is (with respect to inclusion) a maximal interval with

$$\gamma((r, s)) \subset X \setminus B^{X,d}_\Delta(\operatorname{im}\gamma').$$

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 $\operatorname{im}\gamma' \subset B^{X,d}_\Delta(\operatorname{im}\gamma)$, and so $\operatorname{im}\gamma' \subset B^{X,d}_\Delta(\operatorname{im}\gamma|_{[0,r]} \cup \operatorname{im}\gamma|_{[s,L]})$. Because γ is continuous and the interval $[0, L']$ is connected, there is a $t' \in [0, L']$ such that there are $r' \in [0, r]$ and $s' \in [s, L]$ with

$$d\bigl(\gamma'(t'), \gamma(r')\bigr) \leq \Delta \quad \text{and} \quad d\bigl(\gamma'(t'), \gamma(s')\bigr) \leq \Delta$$

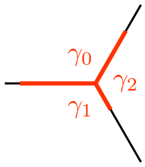

Figure 6.12.: Geodesic triangles (red) in trees are tripods

(see also Figure 6.11). 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. □

6.2.4 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 than 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 6.2.16 (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 6.2.13, 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 4.3.8). □

Proposition 6.2.17 (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 6.12, which implies 0-hyperbolicity. □

6.3 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 [42]:

Definition 6.3.1 (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 6.2.16, we can also verify hyperbolicity of a finitely generated group 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 6.3.2 (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 6.2.9) and the fact that Cayley graphs of groups are quasi-geodesic. □

Example 6.3.3 (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 6.2.17.
- In particular, $\mathrm{SL}(2, \mathbb{Z})$ is hyperbolic, because $\mathrm{SL}(2, \mathbb{Z})$ is quasi-isometric to a free group of rank 2 (Example 4.4.8).

- Let M be a compact Riemannian manifold of negative sectional curvature (e.g., a hyperbolic manifold in the sense of Definition 4.4.11). Then the fundamental group $\pi_1(M)$ is hyperbolic, because by the Švarc–Milnor lemma (Corollary 4.4.10) $\pi_1(M)$ is quasi-isometric to the Riemannian universal covering of M, which is hyperbolic (Example 6.2.3).

  In particular, the fundamental groups of oriented closed connected surfaces of genus at least 2 are hyperbolic (Example 4.4.12).

- 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 6.2.3).

- We will see that the Heisenberg group is *not* hyperbolic (Example 6.5.16).

Caveat 6.3.4 (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* [86].

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 6.3.5. 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 [63, Appendix A.3] 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 [63, 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

$$\begin{aligned} d_H(z, A^n \cdot z) = d_{\mathbb{H}^2}(z, z+n) &= \operatorname{arcosh}\Bigl(1 + \frac{n^2}{2 \cdot (\operatorname{Im} z)^2}\Bigr) \\ &\in O\bigl(\operatorname{arcosh}(n^2)\bigr) = O\bigl(\ln(n^2 + \sqrt{n^2-1})\bigr) = O(\ln n). \end{aligned}$$

Geometrically, the points $\{A^n \cdot z \mid n \in \mathbb{Z}\}$ all lie 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.

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 6.4 and 6.5).

6.4 The word problem in hyperbolic groups

As a first algebraic consequence, we show that the geometric condition of being hyperbolic implies solvability of the word problem.

Definition 6.4.1 (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 [18, 13, 7].

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 [91, Chapter 12]:

Theorem 6.4.2. *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 – otherwise there would be a Turing machine that halts exactly on those Turing machines that do not halt on their own encoding. 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 1.2.24).

6.4.1 Application: "Solving" the word problem

Gromov observed that hyperbolic groups have solvable word problem [42], thereby generalising and unifying previous work in combinatorial group theory and on fundamental groups of negatively curved manifolds:

Theorem 6.4.3 (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.*

The proof of Theorem 6.4.3 relies on a basic idea due to Dehn:

Definition 6.4.4 (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 6.4.5. Looking at the characterisation of free groups in terms of reduced words shows that $\bigl\langle x, y \bigm| xx^{-1}\varepsilon, yy^{-1}\varepsilon, x^{-1}x\varepsilon, y^{-1}y\varepsilon\bigr\rangle$ is a Dehn presentation of the free group of rank 2. On the other hand, $\bigl\langle x, y \bigm| [x, y]\bigr\rangle$ is *not* a Dehn presentation for $\mathbb{Z}^2$.

The key property of Dehn presentations is the third one, as it allows us to replace words by *shorter* words that represent the same group element:

Proposition 6.4.6 (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).
  - If there is a $j \in \{1, \dots, n\}$ such that u_j is a subword of w, then we can write $w = w' u_j w''$ for certain words w', $w'' \in (S \cup S^{-1})^*$. Because $u_j v_j^{-1} \in R$, the words w and $w' v_j w''$ represent the same group element in $\langle S \mid R\rangle$; hence, w represents the trivial element of the group $\langle S \mid R\rangle$ if and only if the shorter word $w' v_j w''$ represents the trivial element of the group $\langle S \mid R\rangle$ (which we can check by applying this algorithm recursively to $w' v_j w''$).

Clearly, this algorithm terminates on all inputs from $(S \cup S^{-1})^*$ and decides whether the given word represents the trivial element in $\langle S \mid R\rangle$ or not. Hence, the word problem for $\langle S \mid R\rangle$ is solvable. ☐

In the setting of hyperbolic groups, short-cuts as required in the definition of Dehn presentations are enforced by negative curvature (see Lemma 6.4.8 below).

Theorem 6.4.7 (Dehn presentations and hyperbolic groups). *Let G be a hyperbolic group and let S be a finite generating set of G. Then there exists a finite set $R \subset (S \cup S^{-1})^*$ such that $\langle S \mid R\rangle$ is a Dehn presentation and $G \cong \langle S \mid R\rangle$.*

The proof of this theorem relies on the existence of short-cuts in cycles in hyperbolic groups (Figure 6.13), which then give rise to a nice set of relations:

Lemma 6.4.8 (Short-cuts in cycles in hyperbolic groups). *Let G be a hyperbolic group, let S be a finite generating set of G, and let $|\operatorname{Cay}(G,S)|$ be δ-hyperbolic with $\delta > 0$. If $\gamma\colon [0,n] \longrightarrow |\operatorname{Cay}(G,S)|$ is the geometric (piecewise linear) realisation of a graph-theoretic cycle in the Cayley graph $\operatorname{Cay}(G,S)$ of length $n > 0$, then there exist t, $t' \in [0,n]$ such that*

$$L_{|\operatorname{Cay}(G,S)|}(\gamma|_{[t,t']}) \leq 8 \cdot \delta$$

and such that the restriction $\gamma|_{[t,t']}$ is not *geodesic.*

A proof of this lemma will be given below.

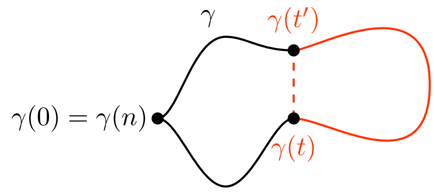

Figure 6.13.: Short-cuts in hyperbolic groups

Proof of Theorem 6.4.7. Because G is hyperbolic, there is a $\delta \in \mathbb{R}_{>0}$ such that $|\operatorname{Cay}(G,S)|$ is δ-hyperbolic. Let $D := \lceil 8 \cdot \delta \rceil + 2$, and let $\pi \colon F(S) \longrightarrow G$ be the canonical homomorphism. Modelling the short-cut lemma above in terms of group theory leads to the (finite) set

$$R := \bigl\{uv^{-1} \bigm| u, v \in (S \cup S^{-1})^*, |u| \leq D, |u| > d_S(e, \pi(u)),\\ \pi(v) = \pi(u), |v| = d_S(e, \pi(u))\bigr\};$$

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 2.1.6). Applying the short-cut lemma (Lemma 6.4.8) to the geometric realisation of this cycle in $|\operatorname{Cay}(G,S)|$ shows that in $(S \cup S^{-1})^*$ we can decompose

$$w = w'uw''$$

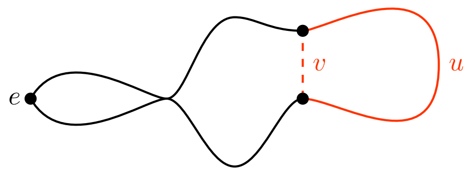

Figure 6.14.: Short-cuts in hyperbolic groups and Dehn presentations

into subwords w', u, w'' such that

$$d_S\bigl(e, \pi(u)\bigr) < |u| \leq D.$$

Let $v \in (S \cup S^{-1})^*$ be chosen in such a way that $\pi(v) = \pi(u)$ and $|v| = d_S(e, \pi(u)) < |u|$ (Figure 6.14). 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 6.4.9. *Every hyperbolic group admits a finite presentation.*

Proof. By Theorem 6.4.7, every hyperbolic group possesses a finite Dehn presentation, whence a finite presentation. □

In particular, we can now prove Theorem 6.4.3:

Proof of Theorem 6.4.3. In view of Theorem 6.4.7, every finite generating set S of a hyperbolic group can be extended to a Dehn presentation $\langle S \,|\, R \rangle$ of the group in question. By Proposition 6.4.6, the word problem for $\langle S \,|\, R \rangle$ 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 6.4.10 (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 can be shown by considering suitable geodesic quadrilaterals around points in $\operatorname{im}\gamma$ that have maximal distance to $\operatorname{im}\gamma'$ (Exercise). □

Proof of Lemma 6.4.8. 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 6.4.10, γ is $2\cdot\delta$-close to every geodesic starting at $\gamma(0)$ 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 $c \in \mathbb{R}_{>8\cdot\delta}$.

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 $\operatorname{Cay}(G,S)$, it follows that

$$L_{|\operatorname{Cay}(G,S)|}(\gamma|_{[t,t']}) = |t-t'| \leq 8\cdot\delta.$$ □

6.5 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 6.5.1) and that centralisers of elements of infinite order in hyperbolic groups are "small" (Chapter 6.5.2). Consequently, hyperbolic groups cannot contain $\mathbb{Z}^2$ as a subgroup. We will mostly follow the arguments of Bridson and Haefliger [20].

6.5.1 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. However, in the case of hyperbolic groups, the geometry forces the existence of elements of infinite order:

Theorem 6.5.1. *Every infinite hyperbolic group contains an element of infinite order.*

For the proof, we introduce the notion of cone types of group elements:

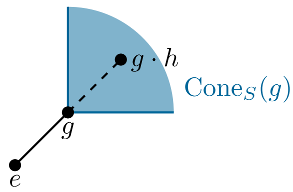

Figure 6.15.: Cone type, schematically; the drawing shows the intuitive visualisation $g \cdot \mathrm{Cone}_S(g)$ of the cone type instead of $\mathrm{Cone}_S(g)$.

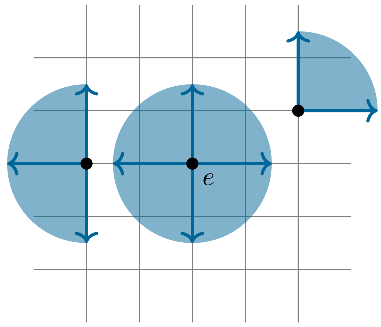

Figure 6.16.: Cone types of $\mathbb{Z}^2$

Definition 6.5.2 (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 6.15)

$$\mathrm{Cone}_S(g) := \bigl\{h \in G \bigm| d_S(e, g \cdot h) \geq d_S(e, g) + d_S(e, h)\bigr\}.$$

Example 6.5.3 (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 $\mathrm{Cone}_S(e) = F$, and

  $$\mathrm{Cone}_S(s) = \{w \mid w \text{ is a reduced word over } S \cup S^{-1} \text{ that does not start with } s^{-1}\}$$

  for all $s \in S \cup S^{-1}$ (check!).

- The group $\mathbb{Z}^2$ has only finitely many (nine) different cone types with respect to the generating set $\{(1,0),(0,1)\}$ (Figure 6.16; check!).

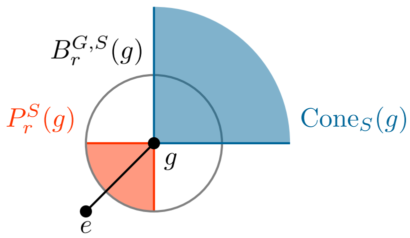

Figure 6.17.: The past of an element, schematically

Theorem 6.5.1 is proved by verifying that hyperbolic groups have only finitely many cone types (Proposition 6.5.4) and that infinite groups that have only finitely many cone types must contain an element of infinite order (Proposition 6.5.6).

Proposition 6.5.4 (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 6.17). Because G is hyperbolic there is a $\delta \in \mathbb{R}_{\geq 0}$ such that $|\operatorname{Cay}(G,S)|$ is δ-hyperbolic. Let

$$r := 2 \cdot \delta + 2.$$

We prove 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 $B_{d_S(e,h)-1}^{G,S}(e)$. 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

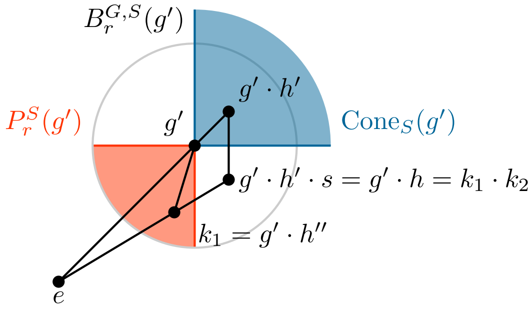

Figure 6.18.: The past of an element determines its cone type

$$d_S(e, g' \cdot h) < d_S(e, g') + d_S(e, h);$$

and we have $d_S(e, g' \cdot h) \geq d_S(e, g')$ (check!). Looking at a shortest path in $\mathrm{Cay}(G, S)$ from e to $g' \cdot h$, we find a decomposition of $g' \cdot h$ of the following form (see also Figure 6.18): 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

$$\begin{aligned} d_S(e, g' \cdot h) &= d_S(e, k_1) + d_S(e, k_2), \\ d_S(e, k_1) &= d_S(e, g'), \\ d_S(e, k_2) &\leq d_S(e, h) - 1 \end{aligned}$$

We 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)

$$d_S(e, g' \cdot h'') = d_S(e, k_1) \leq d_S(e, g').$$

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 at e, and that end distance 1 apart (namely

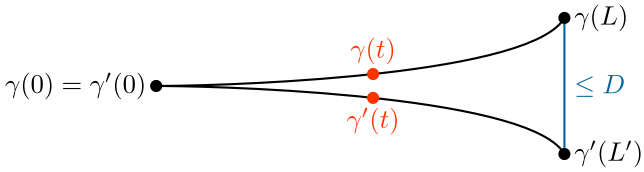

Figure 6.19.: Geodesics in hyperbolic spaces starting at the same point with close endpoints are uniformly close

in $g' \cdot h'$ and $g' \cdot h$ respectively) (see Lemma 6.5.5 below); such geodesics indeed exist because $h' \in \operatorname{Cone}_S(g')$ 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 $h \in \operatorname{Cone}_S(g)$, the choice of k_1 and k_2, as well as the fact that $h'' \in P_r^S(g)$, we obtain

$$\begin{aligned}
d_S(e,g) + d_S(e,h) &\leq d_S(e, g \cdot h) && \text{(because } h \in \operatorname{Cone}_S(g)\text{)}\\
&= d_S(e, g \cdot g'^{-1} \cdot g' \cdot h) && \text{(inserting } g'\text{)}\\
&= d_S(e, g \cdot g'^{-1} \cdot k_1 \cdot k_2) && \text{(choice of } k_1,\ k_2\text{)}\\
&\leq d_S(e, g \cdot h'') + d_S(e, k_2) && \text{(definition of } h''\text{)}\\
&\leq d_S(e,g) + d_S(e,h) - 1, && (h'' \in P_r(g) \text{ and choice of } k_2)
\end{aligned}$$

which is a contradiction. Therefore, $h \in \operatorname{Cone}_S(g')$, 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. Therefore, there are only finitely many cone types with respect to S in G, as claimed. □

Lemma 6.5.5 (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*

$$\gamma(0) = \gamma'(0) \quad \textit{and} \quad d\bigl(\gamma(L), \gamma'(L')\bigr) \leq D.$$

Then γ and γ' are uniformly $(2 \cdot \delta + D)$-close, i.e.,

$$\forall_{t \in [0, \min(L, L')]} \quad d\bigl(\gamma(t), \gamma'(t)\bigr) \leq 2 \cdot \delta + D \quad \textit{and} \quad |L - L'| \leq D.$$

Proof. This follows from a suitable application of the slim triangles condition and a case distinction (Exercise), see also Figure 6.19. □

Proposition 6.5.6 (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 6.5.1. 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 6.5.4), and so contains an element of infinite order (Proposition 6.5.6). □

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 6.5.7. *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 6.5.1). □

Corollary 6.5.8 (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 6.5.7, the group G contains an element g of infinite order. One then shows:

- The subgroup $\langle g \rangle_G$ generated by g is quasi-dense in G because: The image of $\langle g \rangle_G$ under a quasi-isometry $G \longrightarrow \mathbb{Z}$ contains infinitely many positive and infinitely many negative numbers (check!); moreover, because the distance between subsequent powers in $\langle g \rangle_G$ is uniformly bounded, this image in $\mathbb{Z}$ is quasi-dense (check!).
- Hence, $\langle g \rangle_G$ has finite index in G. □

6.5.2 Centralisers

Because (quasi-)isometrically embedded (geodesic) subspaces of hyperbolic spaces are hyperbolic (Proposition 6.2.9), 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 show that the, stronger, algebraic analogue also holds: A hyperbolic group cannot contain $\mathbb{Z}^2$ as a subgroup (Corollary 6.5.15).

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 6.5.9 (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 6.5.3 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 6.5.10 (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 6.5.11 (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 is always a subgroup of the ambient group (check!).

Example 6.5.12 (Centralisers).

- If G is Abelian, then $C_G(g) = G$ for all $g \in G$.
- If F is a free group and $g \in F \setminus \{e\}$ is not a proper power of an element of F, then $C_F(g) = \langle g \rangle_F \cong \mathbb{Z}$.
- If G and H are groups, then for all $h \in H$, we have

$$G \times \langle h \rangle_H \subset G \times C_H(h) = C_{G \times H}(e, h).$$

- If S is a generating set of a group G, then

$$\bigcap_{s \in S} C_G(s) = \bigcap_{g \in G} C_G(g) = \{h \in G \mid \forall_{g \in G} \quad h \cdot g = g \cdot h\}$$

  is the *centre of G* (check!).

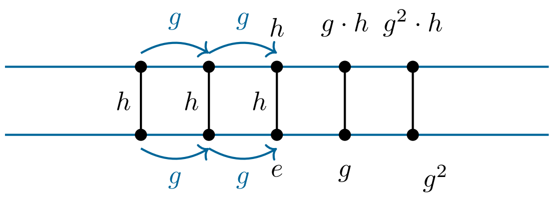

Figure 6.20.: Elements in the centraliser $C_G(g)$ lead to flat strips

Caveat 6.5.13. 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 5.2.13).

We now prove Theorem 6.5.10, using Theorem 6.5.9. This is a geometric argument involving invariant geodesics, inspired by Preissmann's theorem (and its proof) in Riemannian geometry [33, Theorem 10.2.2]:

In view of Theorem 6.5.9, 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 6.20). 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 first prove the following, slightly more general, statement about elements that quasi-commute with the given element g.

Lemma 6.5.14 (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

$$d_S\bigl(k, \langle g \rangle_G\bigr) \le \Delta.$$

Proof. We first make some of the constants explicit: By Theorem 6.5.9, there exists a constant $c \in \mathbb{R}_{\ge 1}$ such that the map

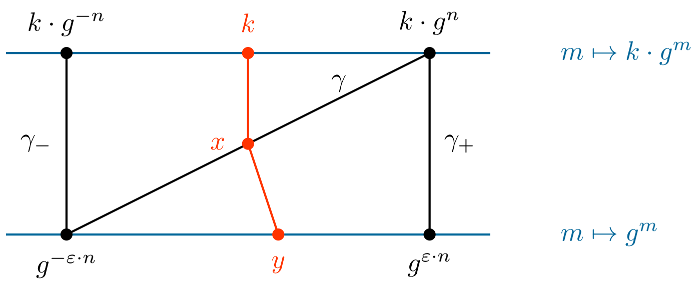

Figure 6.21.: Comparing the quasi-geodesic lines $n \mapsto k \cdot g^n$ and $n \mapsto g^n$

$$\begin{aligned} \mathbb{Z} &\longrightarrow G \\ n &\longmapsto g^n \end{aligned}$$

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 (check!). 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 6.5.9, 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 is illustrated in Figure 6.21. We now argue similarly as in Lemma 6.5.5 and Lemma 6.4.10:

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 (check!):

$$\begin{aligned} d_S(x, k) &\geq d_S(k \cdot g^{-n}, k) - d_S(x, k \cdot g^{-n}) \\ &\geq d_S(g^{-n}, e) - d_S(g^{-\varepsilon \cdot n}, k \cdot g^{-n}) - 2. \end{aligned}$$

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 6.5.10. Let S be a finite generating set and let $\Delta \in \mathbb{R}_{>0}$ be a constant for g as provided by Lemma 6.5.14. 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

$$d_S\bigl(h, \langle g \rangle_G\bigr) \leq \Delta.$$

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). □

Corollary 6.5.15. *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 (check!) that the centraliser $C_G(h)$ of h is virtually $\langle h \rangle_G$ (Theorem 6.5.10). □

Example 6.5.16 (Heisenberg group, $\mathrm{SL}(n, \mathbb{Z})$ and hyperbolicity)**.** By Corollary 6.5.15, 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.

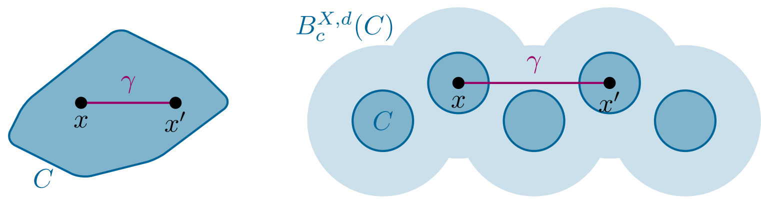

Figure 6.22.: Convexity and quasi-convexity, schematically

The proofs of Theorem 6.5.10 and Theorem 6.5.9 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 [63, Chapter 8.3]. For example, these techniques then also show that "generic" elements in hyperbolic groups fail to commute in the strongest possible way [63, Theorem 8.3.13].

Moreover, the question of whether group elements/isometries act by translation on (quasi-)geodesic lines is already present in the classical classification of isometries of the hyperbolic plane [10, Proposition A.5.14ff][63, Remark 7.5.17].

6.5.3 Quasi-convexity

In order to complete the proof of Theorem 6.5.10 we still need to provide a proof of Theorem 6.5.9, i.e., that elements of infinite order in hyperbolic groups give rise to quasi-geodesic lines. Also the proof of Theorem 6.5.9 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 6.22).

Definition 6.5.17 (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

$$\operatorname{im} \gamma \subset B_c^{X,d}(C).$$

One could also formulate a notion of quasi-convexity in quasi-geodesic spaces, but for simplicity we only consider the geodesic case.

Definition 6.5.18 (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 6.5.19 (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 the intersection $H \cap H'$ is also a quasi-convex subgroup of G with respect to S.*

Proof. Detailed proofs can be found in the book by Bridson and Haefliger [20, Proposition III.Γ.4.12, Proposition III.Γ.4.13]. The second part is non-trivial in the sense that the analogous statement for general quasi-convex subspaces is *not* true (Caveat 6.5.20).

We only sketch the main ideas: For the first part, let $c \in \mathbb{R}_{\geq 0}$ be chosen in such a way that H is c-quasi-convex in $|\operatorname{Cay}(G,S)|$. Then a connect-the-dots argument shows that H is generated by the (finite!) set

$$T := H \cap B^{G,S}_{2\cdot c+1}(e).$$

Moreover, straightforward estimates show that the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding with respect to d_T on H and d_S on G (check!).

For the second part, let $c \in \mathbb{R}_{\geq 0}$ be chosen large enough that H and H' both are c-quasi-convex in $|\operatorname{Cay}(G,S)|$. We set

$$C := \bigl|B^{G,S}_c(e)\bigr|^2$$

and show that $H \cap H'$ is C-quasi-convex in $|\operatorname{Cay}(G,S)|$: It suffices to check the quasi-convexity condition for $H \cap H'$ for geodesics that start in e and for points on the geodesic that lie in the vertex set G. Let $h \in H \cap H'$ and let $\gamma \colon [0,L] \longrightarrow |\operatorname{Cay}(G,S)|$ be a geodesic joining e and h and let $t \in [0,L]$ with $\gamma(t) \in G$. We then consider the set

$$\begin{aligned} M := \bigl\{ g \in G \bigm| & \gamma(t) \cdot g \in H \cap H', \text{ and for all } k \in G \\ & \text{that lie on a geodesic in } |\operatorname{Cay}(G,S)| \text{ from } e \text{ to } g \\ & \text{we have } d_S(\gamma(t)\cdot k, H) \leq c \text{ and } d_S(\gamma(t)\cdot k, H') \leq c \bigr\}. \end{aligned}$$

The set M is non-empty, because it contains the element $\gamma(t)^{-1} \cdot h$ (check!). Therefore, M contains a $d_S(e,\,\cdot\,)$-minimal element g. In the following, we show that g satisfies $d_S(e,g) \leq C$ (whence $d_S(\gamma(t), H \cap H') \leq C$, as was to be shown).

Assume for a contradiction that $d_S(e,g) > C$. Let η be a geodesic from e to g. By the pigeon-hole principle and the construction of M and C, there exist two vertices g_1 and g_2 on η with $g_1 \neq g_2$ that have the same "connecting legs" to H and H', respectively. Now a straightforward computation shows that "cutting out" the part of η from g_1 to g_2 leads to a group element g' with $d_S(e, g') < d_S(e, g)$ and $g' \in M$ (check!). This contradicts the minimality property of g. Therefore, $d_S(e, g) \leq C$. □

Caveat 6.5.20. 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

$$C := 2 \cdot \mathbb{Z} \quad \text{and} \quad C' := (2 \cdot \mathbb{Z} + 1) \cup \{n^2 \mid n \in \mathbb{Z}\}$$

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 6.5.21 (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 $|\operatorname{Cay}(G, S)|$ is δ-hyperbolic. Let $\gamma \colon [0, L] \longrightarrow |\operatorname{Cay}(G, S)|$ 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 $x := \gamma(t) \in G$ (otherwise, we pick a group element that is at most distance 1 away).

Clearly, $g \cdot \gamma \colon [0, L] \longrightarrow |\operatorname{Cay}(G, S)|$ is a geodesic starting at g and ending in $g \cdot h = h \cdot g$. Because $|\operatorname{Cay}(G, S)|$ 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 6.5.5 twice). In particular, we obtain

$$d_S(e, x^{-1} \cdot g \cdot x) = d_S(x, g \cdot x) = d_S\bigl(\gamma(t), g \cdot \gamma(t)\bigr) \leq c.$$

As the next step we show that there is a "small" element y satisfying

$$y^{-1} \cdot g \cdot y = x^{-1} \cdot g \cdot x.$$

To this end, we consider the following two geodesics in $|\operatorname{Cay}(G, S)|$: Let $\overline{\gamma} := \gamma|_{[0,t]} \colon [0, t] \longrightarrow |\operatorname{Cay}(G, S)|$; i.e., $\overline{\gamma}$ is a geodesic starting at e and ending in $\overline{h}$. Then $g \cdot \overline{\gamma}$ is a geodesic starting at 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)

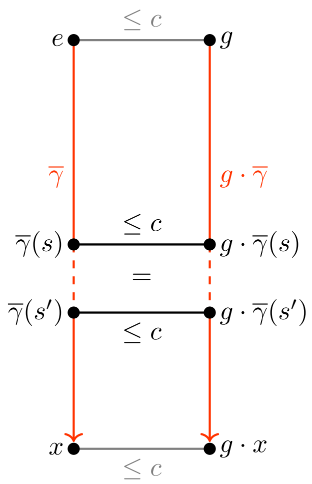

Figure 6.23.: Finding a shorter conjugating element

$$\overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) \in B_c^{G,S}(e)$$

for all $s \in [0, t]$ with $\overline{\gamma}(s) \in G$.

If $d_S(e, x) > |B_c^{G,S}(e)|$, then by the pigeon-hole principle and the fact that geodesics in $|\operatorname{Cay}(G, S)|$ basically follow a shortest path in $\operatorname{Cay}(G, S)$ there exist parameters s, $s' \in [0, t]$ such that $s < s'$ and $\overline{\gamma}(s)$, $\overline{\gamma}(s') \in G$, as well as

$$\overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) = \overline{\gamma}(s')^{-1} \cdot g \cdot \overline{\gamma}(s')$$

(see Figure 6.23). Then the element

$$y := \overline{\gamma}(s) \cdot \overline{\gamma}(s')^{-1} \cdot x$$

satisfies $d_S(e, y) < d_S(e, x)$ (because these elements are lined up on the same geodesic) and

$$\begin{aligned} y^{-1} \cdot g \cdot y &= x^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s)^{-1} \cdot g \cdot \overline{\gamma}(s) \cdot \overline{\gamma}(s')^{-1} \cdot x \\ &= x^{-1} \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot g \cdot \overline{\gamma}(s') \cdot \overline{\gamma}(s')^{-1} \cdot x \\ &= x^{-1} \cdot g \cdot x. \end{aligned}$$

Hence, inductively, we can find an element $y \in G$ with $d_S(e, y) \leq |B_c^{G,S}(e)|$ and

$$y^{-1} \cdot g \cdot y = x^{-1} \cdot g \cdot x.$$

The element $z := x \cdot y^{-1}$ now witnesses that $\gamma(t)$ is $|B_c^{G,S}(e)|$-close to $C_G(g)$: On the one hand we have

$$\begin{aligned} d_S\bigl(z, \gamma(t)\bigr) &= d_S(x \cdot y^{-1}, x) = d_S(y^{-1}, e) = d_S(e, y) \\ &\leq \bigl|B_c^{G,S}(e)\bigr|. \end{aligned}$$

On the other hand, $z \in C_G(g)$ because

$$\begin{aligned} z \cdot g &= x \cdot y^{-1} \cdot g \\ &= x \cdot y^{-1} \cdot g \cdot y \cdot y^{-1} \\ &= x \cdot x^{-1} \cdot g \cdot x \cdot y^{-1} && \text{(by construction of } y\text{)} \\ &= g \cdot z. \end{aligned}$$

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 6.5.9:

Proof of Theorem 6.5.9. We split the proof into two steps by factoring the inclusion $\langle G \rangle_g \hookrightarrow G$ into the following composition:

$$\langle G \rangle_g \hookrightarrow C\bigl(C_G(g)\bigr) \hookrightarrow G.$$

In view of Proposition 6.5.21, the centraliser $C_G(g)$ is a quasi-convex subgroup of G. In particular, $C_G(g)$ is finitely generated by Proposition 6.5.19, say by a finite generating set T. Then the intersection

$$\bigcap_{t \in T} C_G(t) = C\bigl(C_G(g)\bigr),$$

which is the centre of $C_G(g)$ (check!), is also 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 6.5.19). In particular, $C(C_G(g))$ is also a hyperbolic group (Proposition 6.3.2).

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 $\mathbb{Z}$ (check!). 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. □

6.5.4 Application: Products and negative curvature

In view of Corollary 6.5.15, most non-trivial products of finitely generated groups are *not* hyperbolic.

Corollary 6.5.22. *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 6.3.3); hence, we can apply Corollary 6.5.15 and rule out $\mathbb{Z}^2$ as a subgroup. □

Example 6.5.23 (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 6.5.16).

Outlook 6.5.24 (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 [40] (for simplicity, we only state the version where the fundamental group has trivial centre):

Theorem 6.5.25 (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 [52, 53]:

Theorem 6.5.26 (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.*

The End.[2]

[2] If you want to learn more about ends and how to get to infinity and beyond, you should read about "ends of groups and spaces" and about "Gromov boundaries"!

A

Appendix

Overview of this chapter.

A.1 Quasi-isometries in a proof assistant

We formalise basics on quasi-isometries in the proof assistant Lean. Our formalisation includes the definition of quasi-isometric embeddings and quasi-isometries, the characterisation of quasi-isometries as quasi-isometric embeddings with quasi-dense image (Proposition 4.1.10) as well the fact that $\mathbb{Z}$ is quasi-isometric to $\mathbb{R}$ via the inclusion map.

More information on Lean can be found in a variety of documentation and introduction material [5, 6, 55, 56, 57, 70].

```
/-
Copyright (c) 2022 Clara L"oh. All rights reserved.
Released under Apache 2.0 license as described in the file
    LICENSE.txt.
Author: Clara L"oh.
-/

import tactic                        -- standard proof tactics
import topology.metric_space.basic -- basics on metric spaces
import topology.instances.real      -- ℤ as metric space

open classical  -- we work in classical logic

/-
We define quasi-isometries as quasi-isometric embeddings
that admit a quasi-inverse quasi-isometric embedding.
We then prove that a quasi-isometric embedding is a
quasi-isometry if and only if it has quasi-dense image.
-/

/-
# Quasi-isometric embeddings and quasi-isometries
-/

-- quasi-isometric embeddings
def is_QIE_lower
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
:= ∀ x x' : X, dist (f x) (f x') ≥ 1/c * dist x x' - b

def is_QIE_upper
    {X Y : Type*} [metric_space X] [metric_space Y]
```

```
    (f : X → Y)
    (c b : ℝ)
:= ∀ x x' : X, dist (f x) (f x') ≤ c * dist x x' + b

def is_QIE'
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
:= is_QIE_upper f c b
 ∧ is_QIE_lower f c b

def is_QIE
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
:= ∃ c b : ℝ,
   c > 0
 ∧ b > 0
 ∧ is_QIE' f c b


-- finite distance
def has_fin_dist'
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f g : X → Y)
    (c : ℝ)
:= ∀ x : X, dist (f x) (g x) ≤ c

def has_fin_dist
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f g : X → Y)
:= ∃ c : ℝ,
   c > 0
 ∧ has_fin_dist' f g c

def are_quasi_inverse
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (g : Y → X)
:= has_fin_dist (g ∘ f) id
 ∧ has_fin_dist (f ∘ g) id

-- quasi-isometry
def is_QI
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
```

```
:= is_QIE f
 ∧ ∃ g : Y → X, is_QIE g
              ∧ are_quasi_inverse f g

/-
# Two lemmas on quasi-isometric embeddings
-/

-- rewriting the lower estimate for quasi-isometric embeddings
lemma QIE_lower_est
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
    (c_pos : c > 0)
    (f_is_QIE : is_QIE' f c b)
  : ∀ x x' : X, dist x x' ≤ c * dist (f x) (f x') + c * b
:=
begin
  have c_neq_0 : c ≠ 0,
       by exact ne_of_gt c_pos,
  have nonneg_c : 0 ≤ c,
       by exact le_of_lt c_pos,

  assume x x' : X,

  have lower_est : 1/c * dist x x' - b ≤ dist (f x) (f x'),
       by exact f_is_QIE.2 x x',

  calc dist x x'
         = c * 1/c * dist x x' - c * b + c * b
         : by simp[div_self,c_neq_0]
     ... = c * (1/c * dist x x' - b) + c * b
         : by ring
     ... ≤ c * dist (f x) (f x') + c * b
         : by {apply add_le_add_right,
               exact mul_le_mul_of_nonneg_left
                       lower_est nonneg_c},
end

-- Sometimes, it is convenient to be able to use
-- different constants for the upper/lower estimates
lemma QIE_from_different_constants
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
```

```
    (c1 b1 c2 b2: ℝ)
    (c1_pos : c1 > 0)
    (b1_pos : b1 > 0)
    (c2_pos : c2 > 0)
    (b2_pos : b2 > 0)
    (f_QIE_upper : is_QIE_upper f c1 b1)
    (f_QIE_lower : is_QIE_lower f c2 b2)
  : is_QIE f
:=
begin
  unfold is_QIE,
  unfold is_QIE',

  -- we increase the given constants suitably:
  let c := c1 + c2,
  let b := b1 + b2,

  use c,
  use b,

  -- preparation: basic estimates for the constants:
  have c_pos : c > 0,
       by exact add_pos c1_pos c2_pos,
  have nonneg_c : 0 ≤ c,
       by exact le_of_lt c_pos,
  have b_pos : b > 0,
       by exact add_pos b1_pos b2_pos,
  have c1_leq_c : c1 ≤ c,
       by simp[le_of_lt c2_pos],
  have b1_leq_b : b1 ≤ b,
       by simp[le_of_lt b2_pos],
  have b2_leq_b : -b2 ≥ -b,
       by simp[le_of_lt b1_pos],
  have c2_leq_c' : c2 ≤ c,
       by simp[le_of_lt c1_pos],
  have c2_leq_c : 1/c2 ≥ 1/c,
       by simp[c2_leq_c', inv_le_inv_of_le c2_pos],

  -- Now, the upper/lower estimates are basic calculations:
  have f_QIE_upper_cb : is_QIE_upper f c b, by
  begin
    unfold is_QIE_upper,
    assume x x' : X,
    calc dist (f x) (f x')
```

```
          ≤ c1 * dist x x' + b1
          : by exact f_QIE_upper x x'
      ... ≤ c1 * dist x x' + b
          : by exact add_le_add_left b1_leq_b (c1 * dist x x')
      ... ≤ c * dist x x' + b
          : by {apply add_le_add_right,
                exact mul_le_mul_of_nonneg_right
                        c1_leq_c dist_nonneg},
  end,

  have f_QIE_lower_cb : is_QIE_lower f c b, by
  begin
    unfold is_QIE_lower,
    assume x x' : X,
    calc dist (f x) (f x')
          ≥ 1/c2 * dist x x' - b2
          : by exact f_QIE_lower x x'
      ... ≥ 1/c2 * dist x x' - b
          : by exact add_le_add_left b2_leq_b _
      ... ≥ 1/c * dist x x' - b
          : by {apply add_le_add_right,
                exact mul_le_mul_of_nonneg_right
                        c2_leq_c dist_nonneg},
  end,

  show _,
      by exact ⟨ c_pos, b_pos,
                 ⟨ f_QIE_upper_cb, f_QIE_lower_cb ⟩⟩,
end


/-
# An alternative characterisation of quasi-isometries
-/

-- We show that a quasi-isometric embedding is a quasi-isometry
-- if and only if it has quasi-dense image.
-- We show the two implications separately:

def has_quasidense_image'
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c : ℝ)
:= ∀ y : Y, ∃ x : X, dist (f x) y ≤ c
```

```
def has_quasidense_image
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
:= ∃ c : ℝ,
   c > 0
 ∧ has_quasidense_image' f c

-- Quasi-isometries have quasi-dense image:
theorem QI_has_quasidense_image
     {X Y : Type*} [metric_space X] [metric_space Y]
     (f : X → Y)
     (f_is_QI : is_QI f)
   : has_quasidense_image f
:=
begin
  have ex_qinv : ∃ g : Y → X, is_QIE g
                          ∧ are_quasi_inverse f g,
       by exact f_is_QI.2,
  rcases ex_qinv with ⟨ g, ⟨ is_QIE_g, fg_qinv ⟩⟩,

  have fg_close_to_id : ∃ c : ℝ, c > 0
                            ∧ has_fin_dist' (f ∘ g) id c,
       by exact fg_qinv.2,
  rcases fg_close_to_id with ⟨ c, c_pos, fg_c_close_to_id ⟩,

  -- This constant c is a witness for quasi-density of im f:
  use c,
  split,
  show c > 0,
       by exact c_pos,

  show has_quasidense_image' f c, by
  begin
    unfold has_quasidense_image',
    assume y,
    let x := g y,
    use x,
    show dist (f x) y ≤ c, by
    calc dist (f x) y = dist (f (g y)) y : by simp
                  ... ≤ c                : by exact
    fg_c_close_to_id y,
  end,
end
```

```
-- Quasi-isometric embeddings with quasi-dense image are quasi-
    isometries:

-- Preparation:
-- Quasi-inverses of quasi-isometric embeddings
-- are quasi-isometric embeddings
lemma quasiinverse_of_QIE_is_QIE
     {X Y : Type*} [metric_space X] [metric_space Y]
     (f : X → Y)
     (g : Y → X)
     (f_is_QIE : is_QIE f)
     (fg_qinv : are_quasi_inverse f g)
   : is_QIE g
:=
begin
  -- We choose constants witnessing that
  -- f is a quasi-isometric embedding and that
  -- f and g are quasi-inverse to each other:
  rcases f_is_QIE with ⟨cf, bf, cf_pos, bf_pos,
                      f_is_QIE_upper, f_is_QIE_lower⟩,
  rcases fg_qinv with ⟨⟨c_gf, c_gf_pos, gf_close_to_id⟩,
                     ⟨c_fg, c_fg_pos, fg_close_to_id⟩⟩,
  have f_is_cfbf_QIE : is_QIE' f cf bf,
       by exact ⟨ f_is_QIE_upper, f_is_QIE_lower ⟩,

  -- We combine these constants appropriately:
  let c1 := cf,
  let b1 := cf * (2 * c_fg + bf),
  let c2 := cf,
  let b2 := 1/cf * (2 * c_fg + bf),
  have c1_pos : c1 > 0,
       by exact cf_pos,
  have b1_pos : b1 > 0,
       by {apply mul_pos cf_pos,
           apply add_pos _ bf_pos,
           simp[mul_pos _ c_fg_pos]},
  have c2_pos : c2 > 0,
       by exact cf_pos,
  have b2_pos : b2 > 0,
       by {apply mul_pos,
           apply one_div_pos.mpr cf_pos,
           apply add_pos _ bf_pos,
           simp[mul_pos _ c_fg_pos]},

  -- The upper estimate for g:
```

```
have g_is_QIE_upper : is_QIE_upper g c1 b1, by
begin
  unfold is_QIE_upper,
  assume y y' : Y,
  let x  := g y,
  let x' := g y',

  have dist_f_estimate : dist (f x) (f x') ≤ dist y y' + 2 *
  c_fg,
  by calc dist (f x) (f x')
             ≤ dist (f x) y + dist y (f x')
             : by simp[dist_triangle]
         ... ≤ dist (f x) y + dist y y' + dist y' (f x')
             : by {ring_nf,simp[dist_triangle y y' (f x')]}
         ... ≤ c_fg + dist y y' + dist y' (f x')
             : by {simp, exact fg_close_to_id y}
         ... ≤ c_fg + dist y y' + c_fg
             : by {simp,rw[dist_comm],exact fg_close_to_id y'}
         ... ≤ dist y y' + 2 *c_fg
             : by ring_nf,

  calc dist (g y) (g y')
         = dist x x'
         : by  refl
     ... ≤ cf * dist (f x) (f x') + cf * bf
         : by exact QIE_lower_est f cf bf cf_pos
                                f_is_cfbf_QIE x x'
     ... ≤ cf * (dist y y' + 2 * c_fg) + cf * bf
         : by simp[le_of_lt cf_pos,mul_le_mul_of_nonneg_left,
                   dist_f_estimate]
     ... = cf * dist y y' + cf * (2 * c_fg + bf)
         : by ring
     ... ≤ c1 * dist y y' + b1
         : by refl,
end,

-- The lower estimate for g:
have g_is_QIE_lower : is_QIE_lower g c2 b2, by
begin
  unfold is_QIE_lower,
  assume y y' : Y,
  let x  := g y,
  let x' := g y',
```

```
    have cf_times_claim :
         cf * dist (g y) (g y') ≥ dist y y' - cf * b2,
    by calc cf * dist (g y) (g y')
              = cf * dist x x'
              : by refl
          ... ≥ dist (f x) (f x') - bf
              : by {simp,exact f_is_QIE_upper x x'}
          ... ≥ dist (f x) y' - dist (f x') y' - bf
              : by simp[dist_triangle]
          ... ≥ dist y' y - dist (f x) y - dist (f x') y' - bf
              : by simp[dist_triangle_left y' y (f x)]
          ... ≥ dist y' y - c_fg - dist (f x') y' - bf
              : by {simp, exact fg_close_to_id y}
          ... ≥ dist y' y - c_fg - c_fg - bf
              : by {simp, exact fg_close_to_id y'}
          ... ≥ dist y y' - cf * (1/cf * (2 * c_fg + bf))
              : by {simp[ne_of_gt cf_pos,dist_comm], ring_nf}
          ... = dist y y' - cf * b2
              : by refl,

    have cf_inv_nonneg : 0 ≤ cf⁻¹,
         by simp[inv_nonneg.mpr (le_of_lt cf_pos)],

    calc dist (g y) (g y')
            = 1/cf * (cf * dist (g y) (g y'))
            : by simp[ne_of_gt cf_pos]
        ... ≥ 1/cf * (dist y y' - cf * b2)
            : by simp[mul_le_mul_of_nonneg_left cf_times_claim
                                                 cf_inv_nonneg]
        ... = 1/cf * dist y y' - 1/cf * cf * b2
            : by ring
        ... = 1/c2 * dist y y' - b2
            : by simp[ne_of_gt cf_pos],
  end,

  show is_QIE g,
       by exact QIE_from_different_constants g
                   c1 b1 c2 b2
                   c1_pos b1_pos c2_pos b2_pos
                   g_is_QIE_upper g_is_QIE_lower,
end

theorem QIE_with_quasidense_image_is_QI
     {X Y : Type*} [metric_space X] [metric_space Y]
```

```
      (f : X → Y)
      (f_is_QIE : is_QIE f)
      (f_qdense_im : has_quasidense_image f)
    : is_QI f
:=
begin
  -- We obtain a quasi-inverse from quasi-density of the image
  -- and the axiom of choice:
  rcases f_qdense_im with ⟨ c, c_pos, f_has_c_dense_im ⟩,
  rcases classical.axiom_of_choice f_has_c_dense_im
         with ⟨ g, fg_c_close_to_id ⟩,
  -- basic simplifications
  dsimp at g,
  dsimp at fg_c_close_to_id,

  -- This candidate indeed is quasi-inverse to f:
  have f_and_g_are_qinv : are_quasi_inverse f g, by
  begin
    -- By construction, f ∘ g has finite distance from id
    have fg_close_to_id : has_fin_dist (f ∘ g) id, by
    begin
      unfold has_fin_dist,
      unfold has_fin_dist',
      use c,
      split,
      show c > 0,
           by exact c_pos,

      assume y : Y,
      calc dist ((f ∘ g) y) y
             = dist (f (g y)) y : by refl
         ... ≤ c                : by exact fg_c_close_to_id y,
    end,

    -- Conversely, also g ∘ f has finite distance from id;
    have gf_close_to_id : has_fin_dist (g ∘ f) id, by
    begin
      unfold has_fin_dist,
      unfold has_fin_dist',
      -- we choose QIE-constants for f ...
      rcases f_is_QIE with ⟨ cf, bf, cf_pos, bf_pos,
                             f_is_cfbf_QIE ⟩,
      -- ... and construct a suitably large constant c':
      let c' := cf * c + cf * bf,
      use c',
```

```
      split,
      show c' > 0,
           by {apply add_pos,
               apply mul_pos cf_pos,
               exact c_pos,
               apply mul_pos cf_pos bf_pos},

      assume x,
      show dist ((g ∘ f) x) x ≤ c', by
      begin
        let x_fx := g (f x),
        calc dist ((g ∘ f) x) x
               = dist x_fx x
               : by refl
           ... ≤ cf * dist (f x_fx) (f x) + cf * bf
               : by exact QIE_lower_est f cf bf cf_pos
                                         f_is_cfbf_QIE x_fx x
           ... ≤ cf * c + cf * bf
               : by simp[fg_c_close_to_id (f x),
                         le_of_lt cf_pos,
                         mul_le_mul_of_nonneg_left]
           ... ≤ c'
               : by refl,
      end,
    end,

    show are_quasi_inverse f g,
         by exact ⟨ gf_close_to_id, fg_close_to_id ⟩,
  end,

  -- Hence, g is a quasi-isometric embedding:
  have g_is_QIE : is_QIE g,
       by exact quasiinverse_of_QIE_is_QIE f g
                  f_is_QIE f_and_g_are_qinv,

  -- We conclude that f is a quasi-isometry
  -- by putting everything together:
  show is_QI f,
       by exact ⟨ f_is_QIE,
                  begin
                    use g,
                    exact ⟨ g_is_QIE, f_and_g_are_qinv⟩
                  end⟩,
end
```

```
/-
# An example
-/

-- We use the quasi-density criterion to show that
-- the inclusion of ℤ into ℝ is a quasi-isometry
def i_ZR
   : ℤ → ℝ
:= λ x, x

lemma Z_into_R_is_QI
   : is_QI i_ZR
:=
begin
  apply QIE_with_quasidense_image_is_QI,

  show is_QIE i_ZR, by
  begin
    unfold is_QIE,
    use 1,
    use 1,

    have one_pos : (1:real) > 0,
         by simp,

    have i_ZR_is_QIE : is_QIE' i_ZR 1 1, by
    begin
      unfold is_QIE',

      have upper_estimate : ∀ x x' : ℤ,
           dist (i_ZR x) (i_ZR x') ≤ 1 * dist x x' + 1, by
      begin
        assume x x' : ℤ,
        simp[i_ZR],
      end,

      have lower_estimate : ∀ x x' : ℤ,
           dist (i_ZR x) (i_ZR x') ≥ 1/1 * dist x x' - 1, by
      begin
        assume x x' : ℤ,
        simp[i_ZR],
      end,

      show _,
```

```
          by {simp only[is_QIE_upper,is_QIE_lower],
              exact ⟨upper_estimate, lower_estimate ⟩},
    end,

    show _,
        by exact ⟨ one_pos, ⟨one_pos, i_ZR_is_QIE⟩⟩,
  end,

  show has_quasidense_image i_ZR, by
  begin
    unfold has_quasidense_image,
    use 1,

    have one_pos : (1:real) > 0,
         by simp,

    have qdense_im : has_quasidense_image' i_ZR 1, by
    begin
      unfold has_quasidense_image',

      assume y : ℝ,
      let x := int.floor y,
      use x,

      show dist (i_ZR x) y ≤ 1, by
      begin
        calc dist (i_ZR x) y
               = dist y (i_ZR x)
               : by exact dist_comm _ _
           ... = |y - ↑x|
               : by refl
           ... = y - ↑x
               : by simp[int.floor_le,int.fract_nonneg]
           ... ≤ 1
               : by simp[int.fract_lt_one,le_of_lt],
      end,
    end,

    show _,
         by exact ⟨one_pos, qdense_im⟩,
  end,
end
```

B

Exercise Sheets

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 1, April 26, 2022

Quick check A (isomorphic groups?).

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

Quick check B (automorphism groups). Describe Galois groups and (if you know some algebraic topology) deck transformation groups as automorphism groups in suitable categories!

Quick check C (unfree groups). Use the universal property of free groups to show that the groups $\mathbb{Z}/2022$ and $\mathbb{Z}^{2022}$ are *not* free.

Quick check D (isometry groups). Determine the isometry group of the following subset of $\mathbb{R}^2$ with respect to the Euclidean metric. How could one turn this argument into a rigorous proof? Generalise!

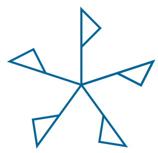

Quick check E (groups). What is your favourite example of a group? Do you know "exotic" examples of groups?

Exercise 1 (symmetric groups; 4 credits). Let X be an infinite set. Is the symmetric group S_X finitely generated? Justify your answer!

Exercise 2 (rank of free groups; 4 credits). Let S be a set, let F be the free group generated by S. Prove that if $T \subset F$ is a generating set of F, then $|T| \geq |S|$.
Hints. If you want, you may restrict to the case that S is finite.

Exercise 3 (the normal subgroup trick; 8 credits). Let G be a group.

1. Let $H, K \subset G$ be subgroups of finite index. Show that $H \cap K$ also has finite index in G.
2. Let $H \subset G$ be a subgroup and let $S \subset G$ 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 $N \subset H$.

Bonus problem (mapping class groups; 4 credits). Look up the term *mapping class group* (e.g., of manifolds, topological spaces); give a reference for the definition you found. Which formal similarity is there between this definition and the definition of outer automorphisms of groups?

Submission before May 3, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on May 2, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 2, May 3, 2022

Hints. Do not resort to dodgy arguments with generators and relations, but use universal properties whenever appropriate!

Quick check A (a presentation of $\mathbb{Z}^2$)**.** Show that $\langle x, y \mid xyx^{-1}y^{-1}\rangle \cong \mathbb{Z}^2$.

Quick check B (trivial groups?)**.** Which of the following groups are trivial?

1. $\langle x, y \mid xyx = yxy\rangle$
2. $\langle x, y \mid yx^{2022}y = x^{2021}, x^2y = x\rangle$
3. $\langle x, y \mid xy^{2022}x = yx^{2021}, x^{2022} = y^{2021}\rangle$
4. $\bigl\langle a_1, a_2, b_1, b_2 \bigm| [a_1, b_1][a_2, b_2]\bigr\rangle$

Quick check C (fundamental groups)**.** If you know fundamental groups: Which of the presentations above are related to the following picture?

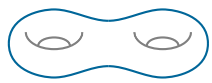

Exercise 1 (the infinite dihedral group; 4 credits)**.** The *infinite dihedral group* is defined as

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

Show that $D_\infty \cong \operatorname{Isom}(\mathbb{Z}, d)$, where d is the metric on $\mathbb{Z}$ induced by the absolute value on $\mathbb{R}$.

Exercise 2 (an alternative presentation of the infinite dihedral group; 4 credits)**.** Show that

$$D_\infty \cong \langle x, y \mid x^2, y^2\rangle$$

Hints. Try first to find geometric candidates for such generators!

Exercise 3 (finite normal generation of kernels; 8 credits)**.** Let $\varphi\colon G \longrightarrow H$ be a surjective group homomorphism, where G is finitely generated and H is finitely presented. Show that there exists a finite set $N \subset G$ with

$$\ker \varphi = \langle N\rangle^{\triangleleft}_G$$

Hints. Start with a finite generating set of G and then find a finite presentation of H that is related to this generating set. Structure your proof into steps.

Bonus problem (a finitely generated group without finite presentation; 4 credits)**.** We consider the group

$$G := \bigl\langle s, t \bigm| \{[s, t^n st^{-n}] \mid n \in \mathbb{N}_{>0}\}\bigr\rangle.$$

Show that G is *not* finitely presentable.

Hints. For $N \in \mathbb{N}_{>0}$, let $G_N := \langle s, t \mid \{[s, t^n st^{-n}] \mid n \in \{1, \dots, N\}\}\rangle$. Show that the homomorphism $\pi_N : G_N \longrightarrow G_{N+1}$ induced by the identity on $\{s, t\}$ is surjective but *not* injective. Homomorphisms to S_{2N+3} might help.

Submission before May 10, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on May 9, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 3, May 10, 2022

Quick check A (free products). Show that $\mathbb{Z} * \mathbb{Z}$ is free of rank 2.

Quick check B (HNN-extensions). Establish an appropriate universal property for HNN-extensions!

Quick check C (group sudoku). Play a round of *group sudoku* by Raphael Appenzeller: https://n.ethz.ch/~apraphae/sudokuformathematicians.html

Quick check D (Isomorphic graphs?). Which of the following graphs are isomorphic?

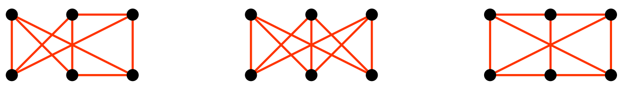

Exercise 1 (pushout groups; 4 credits). Let

$$\begin{array}{ccc} A & \xrightarrow{\alpha_1} & G_1 \\ {\scriptstyle \alpha_2}\downarrow & & \downarrow{\scriptstyle \beta_1} \\ G_2 & \xrightarrow[\beta_2]{} & G \end{array}$$

be a pushout diagram of groups. Are then β_1 and β_2 necessarily injective? Justify your answer by a proof or counterexample!

Exercise 2 (ascending HNN-extensions; 4 credits). Let G be a group and let $\vartheta \in \operatorname{Aut}(G)$. Show that $G*_\vartheta$ is isomorphic to a semi-direct product of $\mathbb{Z}$ with kernel G.

Exercise 3 (spanning trees; 8 credits). Use Zorn's lemma to show that every connected non-empty graph contains a spanning tree.
Hints. A *spanning tree* is a subgraph that is a tree and contains all vertices.

Bonus problem (social networks; 4 credits).

1. How can graphs be used to model social networks?

2. Give an example of how personal information could be inferred from knowledge of such a graph.

Submission before May 17, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on May 13(!), 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold | Sheet 4, May 17, 2022

Quick check A (Cayley graphs)**.** Pick you favourite group and your favourite generating set. Sketch the associated Cayley graph!

If your favourite group happens to be $\mathbb{Z}/3$, you should give D_{2022} (and a proper generating set) a try!

Quick check B (powers in free groups)**.** Let F be a free group of rank 2.

1. Which elements $g \in F$ satisfy $g^2 = e$?
2. Are there elements $g, h \in F$ with $g^{2022} = h^{2023}$?

Quick check C (actions on trees)**.**

1. Is every action of a free group on a tree free?
2. Can free groups act freely on graphs that are not trees?

Exercise 1 (Cayley graph? 4 credits)**.** Does there exist a group that has a Cayley graph with exactly 2022 vertices and exactly 2023 edges? Justify your answer!

Exercise 2 (isomorphic Cayley graphs; 4 credits)**.** Show that there exist finite generating sets S of $\mathbb{Z} \times \mathbb{Z}/2$ and T of D_∞ such that the graphs $\operatorname{Cay}(\mathbb{Z} \times \mathbb{Z}/2, S)$ and $\operatorname{Cay}(D_\infty, T)$ are isomorphic.
Hints. Do not only draw pictures, but give an actual proof!

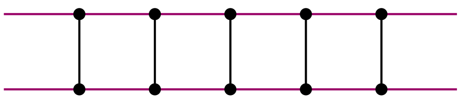

Exercise 3 (actions of finite groups on trees; 8 credits)**.** Prove (without using the characterisation of free groups in terms of free actions on trees) that every action of a finite group G on a non-empty tree T has a global fixed point (i.e., a vertex or an edge on which all group elements act trivially). Proceed in the following steps:

1. Why/How can one restrict to the case that T is finite?
2. A vertex of a tree of degree 1 is called a *leaf*. Show that removing the leaves of T produces a tree T' and that the G-action on T restricts to an action on T'.
3. Use the previous step to shrink the original tree and conclude.

Bonus problem (Women of Mathematics throughout Europe; 4 credits)**.** Pick four of the posters of the exhibition *Women of Mathematics throughout Europe* in the math building and for each of them list the name, the current affiliation (might not be the same as on the poster), the research field, and the title and full reference of one published paper of the corresponding researcher (the database https://mathscinet.ams.org might help).

Submission before May 24, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on May 23, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 5, May 24, 2022

Quick check A (faithful vs. free actions)**.** Let G be a group, let C be a category, and let X be an object of C. A group action $\varrho\colon G \longrightarrow \operatorname{Aut}_C(X)$ of G on X in C is *faithful* if ϱ is a monomorphism.

1. Is every free group action on a non-empty set faithful?
2. Is every faithful group action on a non-empty set free?

Quick check B (spanning trees)**.** The group $\mathbb{Z}/4$ acts on $\operatorname{Cay}(\mathbb{Z}^2, \{(1,0),(0,1)\})$ be rotation about $\pi/2$ around 0. Sketch a spanning tree for this group action! Is the spanning tree of your neighbour isomorphic to yours?

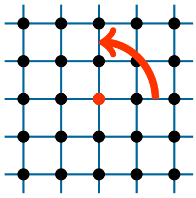

Quick check C (rank gradient)**.** Let F be a free group of finite rank. Compute the following number:

$$\inf\left\{ \frac{\operatorname{rk} H}{[F : H]} \;\middle|\; H \subset F \text{ is a subgroup of finite index}\right\}$$

Exercise 1 (counting in finite trees; 4 credits)**.** Let (V, E) be a finite non-empty tree. Show that

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

Exercise 2 (subgroups of large rank; 4 credits)**.** Let F be a finitely generated free group of rank at least 2. Prove: For each $n \in \mathbb{N}$, there exists a finitely generated free subgroup G of F of rank at least n.

Exercise 3 (characterisation of finite cyclic groups; 8 credits)**.** 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 C. Prove your claim!
Hints. What are "the" special Cayley graphs of finite cyclic groups?

Bonus problem (the Hanna Neumann conjecture; 4 credits)**.**

1. What is the statement of the Hanna Neumann conjecture?
2. Where was the Hanna Neumann conjecture originally formulated? Give the reference!
3. What is the statement of the strengthened Hanna Neumann conjecture?
4. Give two references that contain (different) proofs of the strengthened Hanna Neumann conjecture!

Submission before May 31, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on May 30, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 6, May 31, 2022

Quick check A (ping-pong?). 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

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

Is then G free of rank 2 ? Justify your answer!

Quick check B (unsolvability of free groups). Show that the free group of rank 2 is *not* solvable.

Quick check C (maps close to quasi-isometric embeddings). Show that every map at finite distance to a quasi-isometric embedding is a quasi-isometric embedding.

Exercise 1 (eigen-ping-pong; 4 credits). Let $\lambda \in \mathbb{C}$. We consider the matrices

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

in $\mathrm{GL}(2, \mathbb{C})$. Use the action of $\mathrm{GL}(2, \mathbb{C})$ on $\mathbb{C}^2$ by matrix multiplication and the sets

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

to show that $\langle \{a, b\} \rangle_{\mathrm{GL}(2,\mathbb{C})}$ is free of rank 2 provided $|\lambda|$ is "big" and ε is "small". Make "big" and "small" precise and illustrate the situation by suitable pictures (over $\mathbb{R}$).

Exercise 2 (free groups are residually finite; 8 credits). A group G is *residually finite* if for every $g \in G \setminus \{e\}$, there exists a finite group H and a group homomorphism $\varphi\colon G \longrightarrow H$ with $\varphi(g) \neq e$.

1. Show that the group $\mathrm{SL}(2, \mathbb{Z})$ is residually finite.
2. Conclude that free groups of rank 2 are residually finite.
3. Conclude that all free groups are residually finite.

Exercise 3 (quasi-isometry? 4 credits). Are the spaces $\mathbb{Z}$ and $\{n^3 \mid n \in \mathbb{Z}\}$ (with the standard metric on $\mathbb{R}$) quasi-isometric? Justify your answer!

Bonus problem (free rotation groups; 4 credits). Show that the special orthogonal group $\mathrm{SO}(3)$ 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 $\mathbb{Z}^3 \subset \mathbb{R}^3$.

Submission before June 8(!), 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on Thu, June 9(!), 16:15–17:45, H32.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 7, June 7, 2022

Quick check A (compositions of quasi-isometric embeddings). Show that compositions of quasi-isometric embeddings are quasi-isometric embeddings.

Quick check B (quasi-isometry types of finitely generated groups).

1. Is $\mathbb{Z}/2$ quasi-isometric to D_{2022} ?
2. Is $\mathbb{Z}/2$ quasi-isometric to D_∞ ?

Quick check C (free products and bilipschitz equivalence). Let G_1, G_2, H_1, H_2 be finitely generated groups, let G_1 be bilipschitz equivalent to H_1, and let G_2 be bilipschitz equivalent to H_2. Does this imply that $G_1 * G_2$ is bilipschitz equivalent to $H_1 * H_2$? Prove your claim!

Exercise 1 (free products and quasi-isometry; 4 credits). Let G_1, G_2, H_1, H_2 be finitely generated groups and let $G_1 \sim_{\mathrm{QI}} H_1$ and $G_2 \sim_{\mathrm{QI}} H_2$. Does this imply that $G_1 * G_2 \sim_{\mathrm{QI}} H_1 * H_2$? Prove your claim!

Exercise 2 (bijective quasi-isometries; 8 credits).

1. Let G and H be finitely generated groups and let $f \colon G \longrightarrow H$ be a bijective quasi-isometry. Prove that f is a bilipschitz equivalence.
2. Does the analogous statement also hold for general metric spaces instead of finitely generated groups? Prove your claim!

Exercise 3 (finite distance between maps; 4 credits). Let X and Y be metric spaces and let $f, g \colon X \longrightarrow Y$ be quasi-isometric embeddings. Show that the following are equivalent:

1. The maps f and g have finite distance to each other.
2. There exists a quasi-isometric embedding $h \colon X \times [0,1] \longrightarrow Y$ with
$$h(\,\cdot\,, 0) = f \quad \text{and} \quad h(\,\cdot\,, 1) = g.$$
Here, $X \times [0,1]$ carries the maximum metric of the given metric on X and the standard metric on $[0,1]$.

Bonus problem (isometries are quasi-isometries; 4 credits). Formalise the following definitions/statements/proofs in Lean:

1. definition: isometries of metric spaces;
2. lemma: every isometry is a quasi-isometry;
3. proof of this lemma.

Hints. Please submit a .lean source file. You can use the template quasiisometry_exercise.lean and try out your code in the Lean web interface: https://leanprover.github.io/live/latest/ For the proof: Split the proof into small steps and first write up a detailed(!) pen-and-paper proof. Abstraction helps to avoid doing the same thing multiple times.

Submission before June 14, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on June 13, 2022.

```
/-
Copyright (c) 2022 Clara L"oh. All rights reserved.
Released under Apache 2.0 license as described in the file LICENSE.txt.
Author: Clara L"oh.
-/

import tactic                       -- standard proof tactics
import topology.metric_space.basic -- basics on metric spaces

open classical  -- we work in classical logic

/-
We define quasi-isometries as quasi-isometric embeddings
that admit a quasi-inverse quasi-isometric embedding.

Similarly, we define isometries between metric spaces
and show that isometries are quasi-isometries.
-/

/-
# Quasi-isometric embeddings and quasi-isometries
-/

-- quasi-isometric embeddings
def is_QIE_lower
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
:= ∀ x x' : X, dist (f x) (f x') ≥ 1/c * dist x x' - b

def is_QIE_upper
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
:= ∀ x x' : X, dist (f x) (f x') ≤ c * dist x x' + b

def is_QIE'
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (c b : ℝ)
:= is_QIE_upper f c b
 ∧ is_QIE_lower f c b

def is_QIE
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
:= ∃ c : ℝ, ∃ b : ℝ,
   c > 0
 ∧ b > 0
 ∧ is_QIE' f c b

-- finite distance
def has_fin_dist'
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f g : X → Y)
    (c : ℝ)
:= ∀ x : X, dist (f x) (g x) ≤ c

def has_fin_dist
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f g : X → Y)
:= ∃ c : ℝ,
   c > 0
 ∧ has_fin_dist' f g c

def are_quasi_inverse
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
    (g : Y → X)
:= has_fin_dist (g ∘ f) id
 ∧ has_fin_dist (f ∘ g) id

-- quasi-isometry
def is_QI
    {X Y : Type*} [metric_space X] [metric_space Y]
    (f : X → Y)
:= is_QIE f
 ∧ ∃ g : Y → X, is_QIE g
             ∧ are_quasi_inverse f g

/-
# Isometric embeddings and isometries
-/

-- definition of isometries
def is_isometry
     {X Y : Type*} [metric_space X] [metric_space Y]
     (f : X → Y)
:= sorry
-- Exercise: complete this definition;
-- keeping the definition structurally close to the QI case
-- will make the proof of the theorem below easier

-- Exercise: it could be helpful to prove individual steps
-- of the main proof in separate lemmas

-- Every isometry is a quasi-isometry
theorem isometries_are_quasiisometries
-- Exercise: complete the statement of the theorem
:=
begin
  -- Exercise: complete the proof of the theorem
end
```

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 8, June 14, 2022

Quick check A (quasi-dense subgroups). Let G be a finitely generated group with finite generating set $S \subset G$ and let $H \subset G$ be a subgroup that is quasi-dense in G with respect to d_S. Show that H has finite index in G.

Quick check B (the maximum metric). Show that $(\mathbb{R}^2, d_\infty)$ is a geodesic metric space.

Quick check C (the punctured plane: quasi-geodesics). Let $\varepsilon \in \mathbb{R}_{>0}$. Show that $\mathbb{R}^2 \setminus \{0\}$ is $(1, \varepsilon)$-quasi-geodesic with respect to the standard metric.

Exercise 1 (the punctured plane: geodesics; 4 credits). Prove(!) that $\mathbb{R}^2 \setminus \{0\}$ is *not* geodesic with respect to the standard metric.

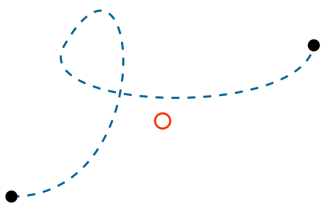

Exercise 2 (Švarc–Milnor lemma; 8 credits). For each of the following group actions name one of the conditions of the Švarc–Milnor lemma that is satisfied, and one that is not (and prove your claims).

1. The action of $\mathrm{SL}(2, \mathbb{Z})$ on $\mathbb{R}^2$ by matrix multiplication.
2. The action of $\mathbb{Z}$ on $X := \{(r^3, s) \mid r, s \in \mathbb{Z}\}$ (with the standard metric on $\mathbb{R}^2$) given by

$$\begin{aligned} \mathbb{Z} \times X &\longrightarrow X \\ \bigl(n, (r^3, s)\bigr) &\longmapsto (r^3, s + n). \end{aligned}$$

Exercise 3 (the Heisenberg group; 4 credits). Let $H_\mathbb{R}$ be the real Heisenberg group and let $H \subset H_\mathbb{R}$ be the Heisenberg group:

$$H_\mathbb{R} := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x, y, z \in \mathbb{R} \right\} \quad \text{and} \quad H := \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; x, y, z \in \mathbb{Z} \right\}.$$

We equip $H_\mathbb{R}$ with the topology given by convergence of all matrix coefficients. Show that $H \subset H_\mathbb{R}$ is discrete with respect to this topology and that the left translation action of H on $H_\mathbb{R}$ is cocompact (i.e., that the quotient $H \setminus H_\mathbb{R}$ is compact with respect to the quotient topology).

Bonus problem (Švarc–Milnor lemma via quasi-isometric actions; 4 credits). 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.

Submission before June 21, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on June 20, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 9, June 21, 2022

Quick check A (homomorphisms and quasi-isometry)**.** Characterise all group homomorphisms between finitely generated groups that are quasi-isometries!

Quick check B (a set-theoretic coupling of $\mathbb{Z}$ with $\mathbb{Z}$?)**.** Let $G := \mathbb{Z}$ and $H := \mathbb{Z}$. We consider the following left and right actions by G and H on $\mathbb{R}^2$, respectively:

$$\begin{aligned} G \times \mathbb{R}^2 &\longrightarrow \mathbb{R}^2 \\ \bigl(n, (x,y)\bigr) &\longmapsto (n+x, y) \\ \mathbb{R}^2 \times H &\longrightarrow \mathbb{R}^2 \\ \bigl((x,y), n\bigr) &\longmapsto (x, y+n) \end{aligned}$$

Is $\mathbb{R}^2$ together with these actions a set-theoretic coupling of G and H ?

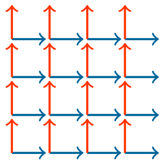

Quick check C (matrix groups)**.** Let $i\colon \mathbb{Z}[\sqrt{2}] \longrightarrow \mathbb{R}$ be the inclusion (a unital ring homomorphism). Does the group homomorphism $\mathrm{SL}(2, \mathbb{Z}[\sqrt{2}]) \longrightarrow \mathrm{SL}(2, \mathbb{R})$ induced by i have discrete image?

Exercise 1 (more matrix groups; 8 credits)**.** Let $\sigma_+, \sigma_- \colon \mathbb{Z}[\sqrt{2}] \longrightarrow \mathbb{R}$ be the two unital ring homomorphisms (what do they look like?!). Does the group homomorphism

$$\mathrm{SL}\bigl(2, \mathbb{Z}[\sqrt{2}]\bigr) \longrightarrow \mathrm{SL}(2, \mathbb{R}) \times \mathrm{SL}(2, \mathbb{R})$$

given by componentwise application of σ_+ and σ_-, respectively, have discrete image? Justify your answer!

Exercise 2 ((weak) commensurability; 4 credits)**.** Show that "commensurability" and "weak commensurability" are equivalence relations on the class of all groups.

Exercise 3 (set-theoretic couplings and finite generation; 4 credits)**.** Let G and H be groups that admit a set-theoretic coupling. Show that G is finitely generated if and only if H is finitely generated. Subdivide your proof into suitable intermediate steps.

Bonus problem (poetry; 4 credits)**.**

The fundamental lemma of GGT
urgently needs poetry:
State and prove the lemma by Milnor and Švarc,
while paying attention to the rhyming arts!

Submission before June 28, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on June 27, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 10, June 28, 2022

Quick check A (sub-multiplicativity of growth functions). Let G be a finitely generated group and let $S \subset G$ be a finite generating set. Show that then $\beta_{G,S}(r + r') \leq \beta_{G,S}(r) \cdot \beta_{G,S}(r')$ holds for all $r, r' \in \mathbb{N}$.

Quick check B (balls in free groups). Pick a (non-free!) finite generating set of a free group of rank 2 and illustrate the balls around the neutral element.

Quick check C (balls in $\mathbb{Z}^2$).

1. Is there a finite generating set $S \subset \mathbb{Z}^2$ with $\beta_{\mathbb{Z}^2,S}(42) = 2022$?

   Hints. Parity!

2. Is there a finite generating set $S \subset \mathbb{Z}^2$ that satisfies $B_r^{\mathbb{Z}^2,S}(e) = \{-r, \dots, r\}^2$ for all $r \in \mathbb{N}$?

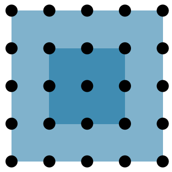

Exercise 1 (growth of $\mathbb{Z}^n$; 4 credits). Let $n \in \mathbb{N}$. Show that $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$. Illustrate your arguments!

Exercise 2 (growth functions of infinite groups; 4 credits). Let G be a finitely generated infinite group and let $S \subset G$ be a finite generating set. Show that $\beta_{G,S}(r) \geq r$ for all $r \in \mathbb{N}$. Illustrate your arguments!

Exercise 3 (geometric properties?; 8 credits). Which of the following properties of finitely generated groups are geometric? Justify your answers!

1. containing a generating set with at most 2022 elements;
2. being isomorphic to a subgroup of $\mathbb{Z}^{2022}$;
3. being isomorphic to a subgroup of $\mathrm{SL}(2, \mathbb{Z})$;
4. being isomorphic to a free product of two non-trivial groups;

Bonus problem (growth type of the Heisenberg group; 8 credits). Let $H \subset \mathrm{SL}(3, \mathbb{Z})$ be the Heisenberg group. Show that H has the growth type of $(x \mapsto x^4)$:

0. Show that $H \cong \langle x, y, z \mid [x, z], [y, z], [x, y] = z \rangle$.

In the following, 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.

Submission before July 5, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on July 4, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 11, July 5, 2022

Quick check A (polynomial growth?)**.** Which of the following groups have polynomial growth?

$$\mathbb{Z}^{2022}, \quad \mathbb{Z}^{2022} \times \mathbb{Z}/2, \quad \mathbb{Z}^{2022} * \mathbb{Z}/2, \quad \mathbb{Z}/2 * \mathbb{Z}/2, \quad \mathrm{SL}(2, \mathbb{Z})$$

Quick check B (non-polynomial growth)**.** Let G be a finitely generated infinite group that is quasi-isometric to $G \times G$. Show that G does *not* have polynomial growth.

Quick check C (exponential growth?)**.** Let G be a group that admits an epimorphism to a free group of rank 2. Does G have exponential growth?

Exercise 1 (exponential growth?!; 4 credits)**.** Does the group

$$\bigl\langle x_1, \dots, x_{2022} \bigm| [x_1, x_2] \cdot [x_3, x_4] \cdot \dots \cdot [x_{2021}, x_{2022}] \bigr\rangle$$

have exponential growth? Justify your answer!

Exercise 2 (solvable groups of exponential growth; 8 credits)**.** 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 $G_A = \mathbb{Z}^n \rtimes_A \mathbb{Z}$ 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.

1. Show that there exists an $x \in \mathbb{Z}^n$ with: 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.

   Hints. To simplify matters, you may restrict to the case that the "largest" eigenvalue of A is real.

2. Conclude that G_A has exponential growth.

Exercise 3 (finite generation and projections on $\mathbb{Z}$; 4 credits)**.** Let G be a finitely generated group with subexponential growth that admits a surjective homomorphism $\pi \colon G \longrightarrow \mathbb{Z}$. Show that then the kernel of π is also finitely generated.

Hints. Choose an element $g \in G$ with $\pi(g) = 1$. Show that there is a finite subset $S \subset \ker \pi$ such that $\{g\} \cup S$ generates G. For $s \in S$ and $n \in \mathbb{N}$, let $g_{n,s} := g^n \cdot s \cdot g^{-n} \in \ker \pi$. Prove that there is an $N \in \mathbb{N}$ such that $\ker \pi$ is generated by the (finite!) set $\{g_{n,s} \mid s \in S, n \in \{-N, \dots, N\}\}$ by considering elements of the form "$g_{n_0, s^{\varepsilon_0}} \cdot \dots \cdot g_{n_k, s^{\varepsilon_k}}$".

Bonus problem (die Hilbertschen Probleme; 4 credits)**.** 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*. Take a random number n between 1 and 23. Describe Hilbert's n-th problem and the status of its solution. Do not forget to cite sources properly!

Submission before July 12, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on July 11, 2022.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold Sheet 12, July 12, 2022

Quick check A (exponential growth rate). Let $H_3 \subset \mathrm{SL}(3, \mathbb{Z}/3)$ be the image of the Heisenberg group $H \subset \mathrm{SL}(3, \mathbb{Z})$ under the reduction $\mathbb{Z} \longrightarrow \mathbb{Z}/3$ and let $S \subset H_3$ be a finite generating set. What is the exponential growth rate of H_3 with respect to S ?

Quick check B (exponential growth rate in free groups). Let F be a free group of rank 2. Show that there exist finite generating sets $S, T \subset F$ with $\varrho_{F,S} \neq \varrho_{F,T}$.
Hints. There is a solution that requires no calculations.

Quick check C (2022-hyperbolicity). Is every hyperbolic metric space also a 2022-hyperbolic metric space?

Exercise 1 (groups of uniform exponential growth; 4 credits). Let G be a finitely generated group of uniform exponential growth. Does then also $G \times G$ have uniform exponential growth? Justify your answer!

Exercise 2 (groups and roots; 4 credits). Let $\alpha, \beta \in \mathbb{C}$. We consider

$$A(\alpha) := \begin{pmatrix} \alpha & 0 \\ 0 & 1 \end{pmatrix}, \quad B(\beta) := \begin{pmatrix} 1 & \beta \\ 0 & 1 \end{pmatrix} \in \mathrm{GL}(2, \mathbb{C}).$$

Prove that if α is *not* a root of unity, then the group $\langle A(\alpha), B(1) \rangle_{\mathrm{GL}(2,\mathbb{C})}$ is *not* virtually nilpotent.
Hints. Calculate $[A(\alpha^n), B(\beta)]$. Why does this help?

Exercise 3 (a weird quasi-geodesic ray in the Euclidean plane; 8 credits).

1. Show that the following map is a quasi-isometric embedding with respect to the standard metrics on $\mathbb{R}$ and $\mathbb{R}^2$, respectively:

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}^2 \\ t &\longmapsto t \cdot \big(\sin(\ln(1+t)), \cos(\ln(1+t))\big) \end{aligned}$$

2. Conclude that the stability theorem for quasi-geodesics does not hold in the Euclidean space $\mathbb{R}^2$.

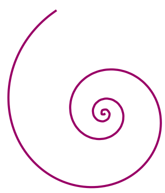

Bonus problem (Fekete's lemma; 4 credits). Translate the proof of Fekete's lemma given in the Lean mathlib library analysis.subadditive into a pen-and-paper version (all definitions, statements, and proofs). Moreover, explain how Fekete's lemma can be used to show that the exponential growth rate is well-defined.

Submission before July 19, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on July 18, 2022.

This is the last regular exercise sheet. Subsequent sheets will give bonus credits.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 13, July 19, 2022

Quick check A (summary). What is your favourite group? For this group,

1. determine the growth type;
2. decide whether it is hyperbolic or not;
3. decide whether it is quasi-isometric to $\mathbb{R}^3$ or not;
4. determine whether it is finitely presented; if so, give a finite presentation.

Quick check B (quadrilaterals). In a quadrilaterals, which triangles appear naturally? How could this help in hyperbolic spaces?

Exercise 1 (local geodesics in general spaces; 4 credits). Let X be a metric space and let $\gamma\colon [0,2022] \longrightarrow X$ be a 1-local geodesic. Is γ then necessarily a geodesic? Justify your answer!

Exercise 2 (geodesics in hyperbolic spaces starting at the same point; 4 credits). 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 $\gamma'(0) = \gamma(0)$ and $d(\gamma'(L'),\gamma(L)) \leq D$. Show that γ and γ' are uniformly $(2\cdot\delta + D)$-close, i.e.,

$$\forall_{t\in[0,\min(L,L')]} \quad d\bigl(\gamma(t),\gamma'(t)\bigr) \leq 2\cdot\delta + D \quad \text{and} \quad |L - L'| \leq D.$$

Illustrate your proof by suitable pictures!

Hints. Distinguish the different cases arising from δ-slimness!

Exercise 3 (local geodesics in hyperbolic spaces; 8 credits). Let X 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 and let $\gamma'\colon [0,L'] \longrightarrow X$ be a geodesic with $\gamma'(0) = \gamma(0)$ and $\gamma'(L') = \gamma(L)$. Prove that

$$\operatorname{im}\gamma \subset B_{2\cdot\delta}(\operatorname{im}\gamma').$$

Illustrate your proof by suitable pictures!

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. Exclude the "weird" cases by computation.

Bonus problem (real hyperbolicity; 4 credits). Crochet/knit a hyperbolic sphere or annulus! For submission: Take pictures of at least two intermediate stages and from at least two different perspectives of the completed model.

Hints. https://pi.math.cornell.edu/~dwh/papers/crochet/crochet.html

How ironic that marine flatworms are called *flat*worms!

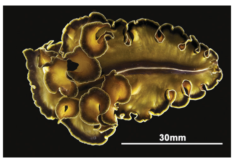

Pseudobiceros flowersi
https://commons.wikimedia.org/wiki/File:Pseudobiceros_flowersi_%2810.11646-zootaxa.4019.1.14%29_Figure_7_%28cropped%29.png (CC 3.0)

Submission before July 26, 2022, 8:30, via GRIPS (in English or German)

The Quick checks are not to be submitted and will not be graded; they will be solved and discussed in the exercise class on July 25, 2022.

All credits on this sheet count as bonus credits.

Geometric Group Theory: Exercises

Prof. Dr. C. Löh/M. Uschold — Sheet 14, July 26, 2022

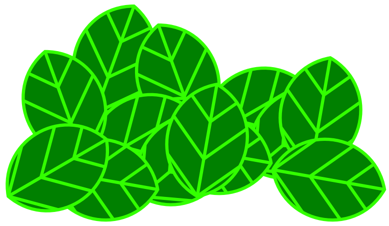

Commander Blorx was summoned to the Digit Council for interrogation. When asked for the result of $9 \cdot 6$, Blorx blacked out, forgot that the actual answer is 42, then decided to compute the result on the spot, but confused the German and English ordering when reading numbers and replied 45. The highest member of the Digit Council, Emperor Neun, thus was not amused and sentenced Blorx to solve the word problem in the infamous jungle on the planet Ji-Ji-Tee. Help Blorx! Beware: Multiple answers might be correct . . .

Problem 1 (the centipede)**.** Hidden under a small rock Blorx finds a centipede. According to the latest research, centipedes on Ji-Ji-Tee do not have 100 legs but are graphs of the form $\big(V, \{\{v,w\} \mid v,w \in V,\ \|v-w\|_1 = 1\}\big)$ on the vertex set $V := ((2 \cdot \mathbb{Z}) \times \mathbb{Z}) \cup (\mathbb{Z} \times \{0\}) \subset \mathbb{R}^2$.

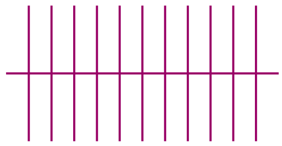

Which groups are quasi-isometric to the centipede?

$\mathbb{Z}$	OR
$\mathbb{Z}^2$	YOU
F_2	TH
There is no such group	FOR

Problem 2 (eek, spiders!)**.** Blorx journeys onward. Suddenly he realises that the little black sticks on the ground are moving and in fact are spider legs. Lots of them. The spiders are interlocked in an 8-regular tree. Which groups could have left such a Cayley mess?

$\mathrm{SL}(3,\mathbb{Z})$	GR
F_4	THR
$\mathbb{Z}^5$	IN
There is no such group	FRE

Problem 3 (the sloth)**.** After the previous encounters Blorx decides that it might be safer to travel further up through the trees. However, there the way is blocked by a particularly lazy sloth, which only travels a generator a day. In the tree

$$\langle \mathsf{A}, \mathsf{B}, \mathsf{C}, \mathsf{H}, \mathsf{N}, \mathsf{R} \mid \mathsf{NA}^2 = \mathsf{A}^3\mathsf{N}^2, \mathsf{AN} = \mathsf{A}^3\mathsf{N}^2\mathsf{A}^{-2}, \mathsf{RAN} = \mathsf{HAH}^{-1}\mathsf{A} \rangle,$$

will the sloth reach the BRANCH within three days (and thus let Blorx pass before he starves)?

Yes	HEC
No	UP

Problem 4 (the monkey puzzle tree)**.** The next trees that Blorx reaches are much more lively – being conquered by a monkey population. Can the group

$$\langle \mathsf{A}, \mathsf{C}, \mathsf{I}, \mathsf{N}, \mathsf{R}, \mathsf{U} \mid \mathsf{ARAUCARIA} = \mathsf{ARAUCANA} \rangle$$

of monkeys act freely on a non-empty tree?

Yes	UPO
No	TR

Problem 5 (more trees!)**.** Obviously, the monkeys need more space, i.e., more trees. Which of the following seeds are not suited for this purpose because they grow only polynomially?

$\mathbb{Z}/3 \times H$	EEN
$H \times H$	OTT
$H * H$	REE
$\mathrm{SL}(3, \mathbb{Z})$	INN
There is no such group in this list	ATT

Problem 6 (the mantis lord)**.** In any case, travelling through the trees didn't turn out as relaxing as Blorx hoped. He thus jumps into the water. There, he barely escapes the hyperbolicity-fuelled (seriously!) kick of a mantis shrimp. Which of the following groups are hyperbolic and could thus be used by Blorx for defense?

$\mathbb{Z}/2 * \mathbb{Z}/2$	ICG
$\mathbb{Z}^2 * \mathbb{Z}^2$	IS
$H * \mathrm{SL}(2, \mathbb{Z})$	RI
$F_2 * F_2$	RO
$F_2 * \mathrm{SL}(3, \mathbb{Z})$	NO
There is no such group in this list	EK

Problem 7 (action!)**.** The continuing problems on land, trees, and water cause a sense of discomfort and urgency in Blorx. He starts running. Quasi-geodesically. In which groups does *not* every element g of infinite order induce a quasi-geodesic line through $\mathbb{Z} \ni n \longmapsto g^n$?

$\mathbb{Z}/3$	LEM
$\mathbb{Z}^2$	AC
$H * \mathbb{Z}^2$	YCL
F_2	SO
There is no such group in this list	MA

Problem 8 (the temple)**.** Completely exhausted, Blorx reaches a temple. The temple bears a mysterious inscription:

$$\langle \mathsf{I}, \mathsf{F}, \mathsf{O}, \mathsf{R}, \mathsf{T}, \mathsf{U}, \mathsf{Y} \mid \mathsf{FOUR} = \mathsf{FIVE}, \mathsf{FIFTY} = \mathsf{FORTY} \rangle$$

Clearly, like every other jungle temple, this temple must have been built by aliens that are far more enlightened than Emperor Neun and the temple serves only as camouflage for an advanced rocket launchpad! To unlock the gate and hence reach the launchpad, Blorx needs to decide whether there is a group homomorphism from the inscription to $\mathbb{Z}$ that maps $\mathsf{FIFTYFOUR}$ to 45. Does such a homomorphism exist?

Yes	DER
No	DIE

Solution (clock-wise):

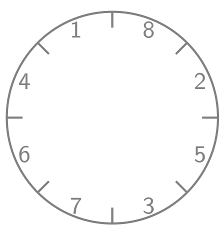

No submission! Enjoy the summer break!

C

General Information

Geometric Group Theory: Admin

Prof. Dr. C. Löh/M. Uschold April 2022

Homepage. Information and news concerning the lectures, exercise classes, office hours, literature, as well as the exercise sheets can be found on the course homepage and in GRIPS:

https://loeh.app.ur.de/teaching/ggt_ss22

https://elearning.uni-regensburg.de

Lectures. The lectures are on Tuesdays (8:30–10:00; M104) and on Fridays (8:30–10:00; M104). The exact start time will be discussed in the first lecture.

Basic lecture notes will be provided, containing an overview of the most important topics of the course. These lecture notes can be found on the course homepage and will be updated after each lecture. Please note that these lectures notes are not meant to replace attending the lectures or the exercise classes!

According to current plans (21.04.2022): This course will be taught on campus in person. On request, this could be turned into a hybrid format (with live zoom streaming). Please note that there will be no recordings of the lectures. The lectures are a precious opportunity for live interaction and I want to keep the atmosphere as casual and unintimidating as possible. For asynchronous self-study, lecture notes will be made available. Please send an email to Clara Löh in case there is a need for the hybrid option!

Exercises. Homework problems will be posted on Tuesdays (before 8:30) on the course homepage; submission is due one week later (before 8:30, via GRIPS).

Each exercise sheet contains regular exercises (16 credits in total) and more challenging bonus problems (4 credits each).

It is recommended to solve the exercises in small groups; however, solutions need to be written up individually (otherwise, no credits will be awarded). Solutions can be submitted alone or in teams of at most two participants; all participants must be able to present *all* solutions of their team.

The first exercise sheet will appear on Tuesday, April 26. The exercise classes start in the *second* week.

In addition, the exercise sheets will contain simple problems that will be solved and discussed during the exercise classes. These problems should ideally be easy enough to be solved within a few minutes. Solutions are not to be submitted and will not be graded.

Registration for the exercise classes. Please register for the exercise classes via GRIPS:

https://elearning.uni-regensburg.de

Please register before Wednesday, April 27, 2022, 10:00.

Credits/Exam. This course can be used as specified in the commented list of courses and in the module catalogue.

- *Studienleistung*: Successful participation in the exercise classes: 50% of the credits (of the regular exercises), presentation of solutions in class (twice).
- *Prüfungsleistung*: Oral exam (25 minutes), by individual appointment at the end of the lecture period/during the break.

You will have to register in FlexNow for the Studienleistung and the Prüfungsleistung (if applicable). Registration will open at the end of the lecture period.

Further information on formalities can be found at:

https://www.uni-regensburg.de/mathematik/fakultaet/studium/studierende/index.html

Contact.

- If you have questions regarding the organisation of the exercise classes or the exercises, please contact Matthias Uschold:

  matthias.uschold@ur.de

- If you have mathematical questions regarding the lectures, please contact Matthias Uschold or Clara Löh.
- If you have questions concerning your curriculum or the examination regulations, please contact the student counselling offices or the exam office:

http://www.uni-regensburg.de/mathematik/fakultaet/studium/ansprechpersonen/index.html

- In many cases, also the Fachschaft can help:

https://www-app.uni-regensburg.de/Studentisches/FS_MathePhysik/cmsms/

- Official information of the administration related to the COVID-19 pandemic can be found at:

  https://go.ur.de/corona

Bibliography

Please note that the bibliography will grow during the semester. Thus, also the numbers of the references will change!

[1] Mark A. Armstrong. *Groups and Symmetry*, Undergraduate Texts in Mathematics, Springer, 1988. Cited on page: 57, 63

[2] Tim Austin. Rational group ring elements with kernels having irrational dimension, *Proc. Lond. Math. Soc*, 107(6), pp. 1424–1448, 2013. Cited on page: 30

[3] Australian Academy of Science. *Bernhard Hermann Neumann (1909–2002)*, http://sciencearchive.org.au/fellows/memoirs/neumann-b.html Cited on page: 37

[4] Australian Academy of Science. *Professor Bernhard Neumann (1909–2002), Mathematician*, http://www.sciencearchive.org.au/scientists/interviews/n/bn.html Cited on page: 37

[5] J. Avigad, K. Buzzard, R.Y. Lewis, P. Massot. Mathematics in Lean. https://leanprover-community.github.io/mathematics_in_lean/ Cited on page: A.2

[6] A. Baanen, A. Bentkamp, J. Blanchette, J. Hölzl, J. Limperg. *The Hitchhiker's Guide to Logical Verification*, 2021 Standard Edition, 2021. https://github.com/blanchette/logical_verification_2021/raw/main/hitchhikers_guide.pdf Cited on page: A.2

[7] 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: 168

[8] Hyman Bass. The degree of polynomial growth of finitely generated nilpotent groups, *Proc. London Math. Soc.*, 25, pp. 603–614, 1972. Cited on page: 140

[9] Gilbert Baumslag, Donald Solitar. Some two-generator one-relator non-Hopfian groups, *Bull. Amer. Math. Soc.*, 68, pp. 199–201, 1962. Cited on page: 27

[10] Riccardo Benedetti, Carlo Petronio. *Lectures on Hyperbolic Geometry*, *Universitext*, Springer, 1992. Cited on page: 112, 184

[11] Béla Bollobás. *Random graphs*, second edition, volume 73 of *Cambridge Studies in Advanced Mathematics*, Cambridge University Press, 2001. Cited on page: 78

[12] Enrico Bombieri, Walter Gubler. *Heights in Diophantine Geometry*, volume 4 of *New Mathematical Monographs*, Cambridge University Press, 2007. Cited on page: 149

[13] George S. Boolos, John P. Burgess, Richard C. Jeffrey. *Computability and logic*, fifth edition, Cambridge University Press, 2007. Cited on page: 168

[14] Brian H. Bowditch. Continuously many quasiisometry classes of 2-generator groups, *Comment. Math. Helv.*, 73, pp. 232–236, 1998. Cited on page: 122

[15] 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: 83, 147, 149, 150

[16] 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: 138

[17] 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: 147, 150

[18] Douglas S. Bridges. *Computability. A mathematical sketchbook*, volume 146 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 168

[19] 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

[20] Martin R. Bridson, André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*, Springer, 1999. Cited on page: 45, 123, 145, 155, 173, 185

[21] 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: 145

[22] Kenneth S. Brown. *Cohomology of Groups*, volume 87 of *Graduate Texts in Mathematics*, Springer, 1982. Cited on page: 32, 46, 110

[23] Ulrich Bunke, Alexander Engel. Homotopy theory with bornological coarse spaces, preprint, arXiv:1607.03657 [math.AT], 2016. Cited on page: 124

[24] James W. Cannon, William J. Floyd, Walter R. Parry. Introductory notes on Richard Thompson's groups, *Enseign. Math. (2)*, 42, pp. 215–256, 1996. Cited on page: 26, 27

[25] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein. *Introduction to Algorithms*, third edition, MIT Press, 2009. Cited on page: 40

[26] 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: 145

[27] 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: 46

[28] Reinhard Diestel. *Graph theory*, third edition, volume 173 of *Graduate Texts in Mathematics*, Springer, 2005. Cited on page: 40

[29] Albrecht Dold. *Lectures on Algebraic Topology*, reprint of the 1972 edition, *Classics in Mathematics*, Springer, 1995. Cited on page: 136, 137

[30] 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: 138, 141, 142, 143

[31] Cornelia Druţu, Michael Kapovich. *Geometric Group Theory*, with an appendix by Bogdan Nica, volume 63 of *Colloquium Publications*, AMS, 2017. Cited on page: 81, 122, 138, 140, 142

[32] 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: 142

[33] Patrick B. Eberlein. *Geometry of Nonpositively Curved Manifolds*, University of Chicago Press, 1996. Cited on page: 181

[34] Alex Eskin, Shahar Mozes, Hee Oh. On uniform exponential growth for linear groups, *Invent. Math.*, 160(1), pp. 1–30, 2005. Cited on page: 148

[35] John Franks. Anosov diffeomorphisms, *1970 Global Analysis*, volume XIV of *Proc. Sympos. Pure Math. (Berkeley, 1968)*, pp. 61–93, AMS. Cited on page: 146, 147

[36] 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: 120

[37] Ross Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*, Springer, 2008. Cited on page: 70

[38] 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: 116

[39] 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: 122, 133

[40] 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: 189

[41] 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: 138, 141, 142, 145

[42] 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: 154, 166, 169

[43] Philip Hall. Finiteness conditions for soluble groups, *Proc. London Math. Soc.*, 4, pp. 419–436, 1954. Cited on page: 29, 30

[44] Pierre de la Harpe. *Topics in Geometric Group Theory*, Chicago University Press, 2000. Cited on page: 29, 75, 110, 111, 133, 136

[45] 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: 148

[46] John M. Harris, Jeffrey L. Hirst, Michael J. Mossinghoff. *Combinatorics and Graph Theory*, second edition, *Undergraduate Texts in Mathematics, Springer*, 2008. Cited on page: 40

[47] Allen Hatcher. *Algebraic Topology*, Cambridge University Press, 2002. http://www.math.cornell.edu/~hatcher/AT/ATpage.html Cited on page: 46

[48] 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: 124

[49] John L. Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*, Springer, 1975. Cited on page: 117

[50] 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: 138

[51] Johannes Köbler, Uwe Schöning, Jacobo Torán. *Graph Isomorphism Problem: The Structural Complexity*, Birkhäuser, 1993. Cited on page: 41

[52] Dieter Kotschick, Clara Löh. Fundamental classes not representable by products, *J. Lond. Math. Soc.*, 79(3), 545–561, 2009. Cited on page: 189

[53] Dieter Kotschick, Clara Löh. Groups not presentable by products, *Groups Geom. Dyn.*, 7(1), pp. 181–204, 2013. Cited on page: 189

[54] Serge Lang. *Algebra*, revised third edition, volume 211 of *Graduate Texts in Mathematics*, Springer, 2002. Cited on page: 8, 139

[55] Lean community. Learning Lean. https://leanprover-community.github.io/learn.html Cited on page: A.2

[56] Lean community. Get started with Lean. https://leanprover-community.github.io/get_started.html Cited on page: A.2

[57] Lean community. mathlib. https://leanprover-community.github.io/mathlib-overview.html Cited on page: A.2

[58] John M. Lee. *Riemannian Manifolds. An Introduction to Curvature*, volume 176 of *Graduate Texts in Mathematics*, Springer, 1997. Cited on page: 111, 112, 113, 152, 153

[59] Xin Li. Quasi-isometry, Kakutani equivalence, and applications to cohomology, preprint, arXiv:1604.07375 [math.GR], 2016 Cited on page: 116, 118

[60] Clara Löh. Simplicial volume, *Bull. Man. Atl.*, pp. 7–18, 2011. Cited on page: 137

[61] 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: 137

[62] Clara Löh. Which finitely generated Abelian groups admit isomorphic Cayley graphs? *Geom. Dedicata*, 164(1), pp. 97–111, 2013. Cited on page: 97

[63] Clara Löh. *Geometric Group Theory. An Introduction.* Universitext, Springer, 2017. Cited on page: 5, 27, 37, 95, 97, 98, 108, 110, 122, 124, 133, 149, 152, 153, 155, 167, 184

[64] Clara Löh. *Algebra*, lecture notes, 2018. Available online at https://loeh.app.ur.de/teaching/algebra_ws1718 Cited on page: 139

[65] Clara Löh. *Group Cohomology*, lecture notes, 2019. Available online at https://loeh.app.ur.de/teaching/grouphom_ss19 Cited on page: 32, 46, 116

[66] Clara Löh. *Ergodic Theory of Groups*, lecture notes, 2020. Available online at https://loeh.app.ur.de/teaching/erg_ss2020 Cited on page: 120

[67] Clara Löh. *Algebraic Topology*, lecture notes, 2019. Available online at https://loeh.app.ur.de/teaching/topologie1_ws2122 Cited on page: 34

[68] Clara Löh. *Differential Geometry*, lecture notes, 2021. Available online at https://loeh.app.ur.de/teaching/diffgeo1_ws2021 Cited on page: 98, 111, 137, 153

[69] Clara Löh. *Geometrie*, lecture notes, 2021. Available online at https://loeh.app.ur.de/teaching/geometrie_ss21 Cited on page: 98

[70] Clara Löh. *Exploring Formalisation. A Primer in Human-Readable Mathematics in Lean 3 with Examples from Simplicial Topology*, book project, 2022. https://loeh.app.ur.de/mapa Cited on page: A.2

[71] Clara Löh, Matthias Mann. Which finitely generated Abelian groups admit equal growth functions? preprint, arXiv: 1309.3381 [math.GR], 2013. Cited on page: 133

[72] Roger C. Lyndon, Paul E. Schupp. *Combinatorial Group Theory*, reprint of the 1977 edition, *Classics in Mathematics*, Springer, 2001. Cited on page: 36

[73] Wilhelm Magnus, Abraham Karrass, Donald Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*, Interscience Publishers, 1966. Cited on page: 30, 140

[74] Avinoam Mann. *How groups grow*, volume 395 of *London Math. Soc. Lecture Note Ser.*, Cambridge University Press, 2012. Cited on page: 138

[75] Grigorii A. Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1), pp. 71–78, 1982. Cited on page: 79

[76] 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: 28

[77] William S. Massey. *Algebraic Topology: An Introduction*, volume 56 of *Graduate Texts in Mathematics*, Springer, 1977. Cited on page: 23, 34, 58, 61, 69, 71, 107, 112, 136, 137

[78] Kostya Medynets, Roman Sauer, Andreas Thom. Cantor systems and quasi-isometry of groups, preprint, arXiv:1508.07578 [math.DS], 2015. Cited on page: 118

[79] Donald W. Miller. On a theorem of Hölder, *The American Mathematical Monthly*, 65(4), pp. 252–254, 1958. Cited on page: 17

[80] John Milnor. A note on curvature and the fundamental group, *Journal of Differential Geometry*, 2, pp. 1–7, 1968. Cited on page: 135, 141

[81] Deane Montgomery, Leo Zippin. Small subgroups of finite-dimensional groups, *Ann. of Math.*, 56, pp. 213–241, 1952. Cited on page: 142

[82] Jan Mycielski. Sur le coloriage des graphs. *Colloq. Math.*, 3, pp. 161–162, 1955. Cited on page: 78

[83] Bernhard H. Neumann, Hanna Neumann, Peter Neumann. Wreath products and varieties of groups, *Math. Z.*, 80(1), pp. 44–62, 1962. Cited on page: 37

[84] Piotr W. Nowak, Guoliang Yu. *Large Scale Geometry*, *EMS Textbooks in Mathematics*, EMS, 2012. Cited on page: 124

[85] Denis V. Osin. Algebraic entropy of elementary amenable groups *Geom. Dedicata*, 107, pp. 133–151, 2004. Cited on page: 148

[86] Denis V. Osin. *Relatively Hyperbolic Groups: Intrinsic Geometry, Algebraic Properties, and Algorithmic Problems*, volume 843 of *Memoirs of the AMS*, 2006. Cited on page: 167

[87] 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: 37

[88] Narutaka Ozawa. A functional analysis proof of Gromov's polynomial growth theorem, preprint, arXiv:1510.04223 [math.GR], 2015. Cited on page: 138

[89] John G. Ratcliffe. *Foundations of Hyperbolic Manifolds*, volume 149 of *Graduate Texts in Mathematics*, Springer, 1994. Cited on page: 137

[90] John Roe. *Index Theory, Coarse Geometry, and Topology of Manifolds*, volume 90 of *CBMS Regional Conf. Ser. in Math.*, AMS, 1996. Cited on page: 124

[91] Joseph J. Rotman. *An Introduction to the Theory of Groups*, fourth edition, volume 148 of *Graduate Texts in Mathematics*, Springer, 1999. Cited on page: 8, 28, 36, 50, 63, 168

[92] Parameswaran Sankaran. On homeomorphisms and quasi-isometries of the real line, *Proc. Amer. Math. Soc.*, 134(7), pp. 1875–1880, 2006. Cited on page: 92

[93] 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: 27

[94] 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: 118, 120

[95] Roman Sauer. Homological invariants and quasi-isometry, *Geom. Funct. Anal.*, 16(2), pp. 476–515, 2006. Cited on page: 116

[96] Jean-Pierre Serre. *Trees*, corrected second printing of the 1980 English translation, *Springer Monographs in Mathematics*, Springer, 2003. Cited on page: 34, 37, 50, 70, 73, 78

[97] Brandon Seward. Burnside's Problem, spanning trees and tilings, *Geom. Topol.*, 18(1), pp. 179–210, 2014. Cited on page:

[98] Yehuda Shalom. Harmonic analysis, cohomology, and the large-scale geometry of amenable groups, *Acta Math.*, 192(2), pp. 119-185, 2004. Cited on page: 116, 145

[99] 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: 138

[100] Michael Shub. Expanding maps. *1970 Global Analysis*, volume XIV of *Proc. Sympos. Pure Math. (Berkeley, 1968)*, pp. 273–276, AMS. Cited on page: 147

[101] Raymond M. Smullyan, Melvin Fitting. *Set Theory and the Continuum Problem*, Dover, 2010. Cited on page: 12

[102] John R. Stallings. On torsion-free groups with infinitely many ends, *Ann. of Math.*, 88, pp. 312–334, 1968. Cited on page: 37

[103] John Stillwell. *Geometry of Surfaces*, *Universitext*, Springer, 1992. Cited on page: 112

[104] Terence Tao. *Hilbert's Fifth Problem and Related Topics*, volume 153 of *Graduate Studies in Mathematics*, American Mathematical Society, 2014. Cited on page: 142

[105] Jacques Tits. Free subgroups in linear groups, *J. Algebra*, 20, pp. 250–270, 1972. Cited on page: 81

[106] 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: 17

[107] 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: 28

[108] John S. Wilson. On exponential growth and uniformly exponential growth for groups, *Invent. Math.*, 155(2), pp. 287–303, 2004. Cited on page: 148

[109] 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: 97

[110] Wolfgang Woess. Graphs and groups with tree-like properties, *J. Combin. Theory*, B 47(3), pp. 361–371, 1989. Cited on page: 122

[111] 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: 140, 141

Deutsch → English

A

abgeleitete Untergruppen/iterierte Kommutatoruntergruppen	derived subgroups	138
Alphabet	alphabet	20
amalgamiertes freies Produkt	amalgamated free product	33
auflösbar	solvable	138
äußere Automorphismengruppe	outer automorphism group	17
Automorphismus	automorphism	11

B

Bahn/Orbit	orbit	60
Baum	tree	42
benachbart	adjacent	40
Beweisassistent	proof assistant	A.2
Bild	image	10
Bilipschitzäquivalenz	bilipschitz equivalence	87
bipartiter Graph	bipartite graph	40

C

Cayleygraph	Cayley graph	43

D

dünn	slim	154
Darstellung	group representation	24
Diedergruppe	dihedral group	25

E

Eilenberg–MacLane-Raum	Eilenberg–MacLane space	45

endlich präsentiert	finitely presented	28
Erzeugendensystem	generating set	18
Erzeuger und Relationen	generators and relations	24

F

Fixpunktmenge	fixed set	61
flache Mannigfaltigkeit	flat manifold	111
freie Gruppe	free group	19
freie Operation	free action	57
freies Produkt	free product	33

G

Gaußkrümmung	Gaussian curvature	152
Geodäte/geodätisch	geodesic	98
geometrische Gruppentheorie	geometric group theory	1
geometrische Realisierung	geometric realisation	100
geschlossene Mannigfaltigkeit	closed manifold	111
Grad	degree	40
Graph	graph	40
Gruppe	group	8
Gruppenerweiterung	group extension	31
Gruppenoperation	group action	56
Gruppenpräsentation	group presentation	24

H

HNN-Erweiterung	HNN-extensions	36
Homomorphismus	homomorphism	9
hyperbolische Gruppe	hyperbolic group	151
hyperbolische Mannigfaltigkeit	hyperbolic manifold	111

I

Identitätsmorphismus	identity morphism	12
innerer Automorphismus	inner automorphism	17
invers/Inverses	inverse	8
Isometrie	isometry	11
Isometriegruppe	isometry group	11
Isomorphismus	isomorphism	9

K

Kante	edge	40
Kategorie	category	11
Kegel	cone	174
Kegeltyp	cone type	174
Kern	kernel	10
klassifizierender Raum	classifying space	45
Knoten	vertex	40
kommensurabel	commensurable	109
Kopplung	coupling	113
Krümmung	curvature	152
Kranzprodukt	wreath product	33

L

Lamplighter-Gruppe	lamplighter group	33

M

Morphismus	morphism	11

N

Nachbar	neighbour	40
neutrales Element	neutral element	8
nilpotent	nilpotent	138
normal erzeugte Untergruppe	normal subgroup, generated by	24
Normalteiler	normal subgroup	15

O

Objekt	object	11
Operatornorm	operator norm	80

P

Pfad	path	41
Präsentationskomplex	presentation complex	45
Pushout	pushout	33

Q

quasi-dominiert	quasi-dominates	129
Quasigeodäte/quasigeodätisch	quasi-geodesic	99
Quasiisometrie	quasi-isometry	87
quasiisometrische Einbettung	quasi-isometric embedding	87
Quotientengruppe	quotient group	15

R

Rang	rank	23
reduziertes Wort	reduced word	47
Relation	relation	24
Riccikrümmung	Ricci curvature	153

S

Schnittkrümmung	sectional curvature	153
semidirektes Produkt	semi-direct product	31
Skalarkrümmung	scalar curvature	153
Spalt/spaltend	split	31
Spannbaum	spanning tree	43
stabiler Buchstabe	stable letter	36
Standgruppe	stabiliser	61
symmetrische Gruppe	symmetric group	10

T

Taillenweite	girth	78

transitive Operation	transitive action	63
treu	faithful	B.6

U

Unentscheidbarkeit	undecidability	28
untere Zentralreihe	lower central series	138
Untergruppe	subgroup	8

V

virtuell ...	virtually ...	138
vollständiger Graph	complete graph	40

W

Wachstumsfunktion	growth function	126
Wachstumstyp	growth type	129
Wald	forest	42
Wort	word	20
Wortlänge	word length	93
Wortmetrik	word metric	93
Wortproblem	word problem	27

Z

Zentralisator	centraliser	180
Zentralreihe	central series	138
Zentrum	centre	180
Zykel	cycle	41

English → Deutsch

A

adjacent	benachbart	40
alphabet	Alphabet	20
amalgamated free product	amalgamiertes freies Produkt	33
automorphism	Automorphismus	11

B

bilipschitz equivalence	Bilipschitzäquivalenz	87
bipartite graph	bipartiter Graph	40

C

category	Kategorie	11
Cayley graph	Cayleygraph	43
central series	Zentralreihe	138
centraliser	Zentralisator	180
centre	Zentrum	180
classifying space	klassifizierender Raum	45
closed manifold	geschlossene Mannigfaltigkeit	111
commensurable	kommensurabel	109
complete graph	vollständiger Graph	40
cone	Kegel	174
cone type	Kegeltyp	174
coupling	Kopplung	113
curvature	Krümmung	152
cycle	Zykel	41

D

degree	Grad	40

derived subgroups	abgeleitete Untergruppen/iterierte Kommutatoruntergruppen	138
dihedral group	Diedergruppe	25

E

edge	Kante	40
Eilenberg–MacLane space	Eilenberg–MacLane-Raum	45

F

faithful	treu	B.6
finitely presented	endlich präsentiert	28
fixed set	Fixpunktmenge	61
flat manifold	flache Mannigfaltigkeit	111
forest	Wald	42
free action	freie Operation	57
free group	freie Gruppe	19
free product	freies Produkt	33

G

Gaussian curvature	Gaußkrümmung	152
generating set	Erzeugendensystem	18
generators and relations	Erzeuger und Relationen	24
geodesic	Geodäte/geodätisch	98
geometric group theory	geometrische Gruppentheorie	1
geometric realisation	geometrische Realisierung	100
girth	Taillenweite	78
graph	Graph	40
group	Gruppe	8
group action	Gruppenoperation	56
group extension	Gruppenerweiterung	31
group presentation	Gruppenpräsentation	24
group representation	Darstellung	24
growth function	Wachstumsfunktion	126
growth type	Wachstumstyp	129

H

HNN-extensions	HNN-Erweiterung	36
homomorphism	Homomorphismus	9
hyperbolic group	hyperbolische Gruppe	151
hyperbolic manifold	hyperbolische Mannigfaltigkeit	111

I

identity morphism	Identitätsmorphismus	12
image	Bild	10
inner automorphism	innerer Automorphismus	17
inverse	invers/Inverses	8
isometry	Isometrie	11
isometry group	Isometriegruppe	11
isomorphism	Isomorphismus	9

Symbols

Symbols

A

B

C

Index

D

E

F

H

S

T