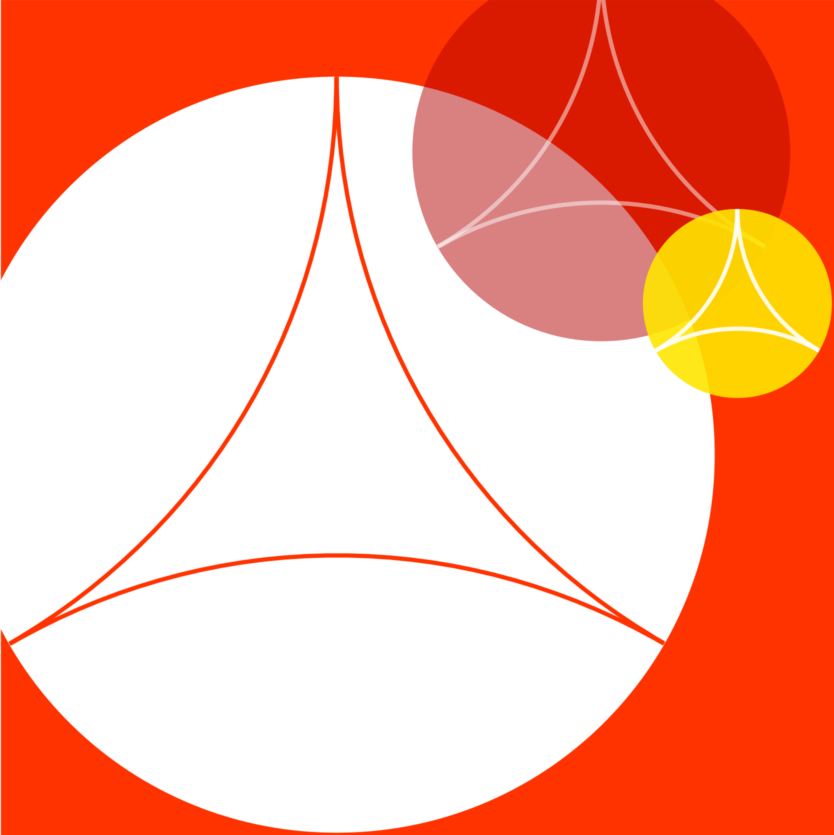

Clara Löh

Geometric group theory, an introduction

December 21, 2010 – 18:11

Preliminary version

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

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

NWF I – Mathematik
Universität Regensburg
93040 Regensburg
Germany

Contents

1 Introduction

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

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

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

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

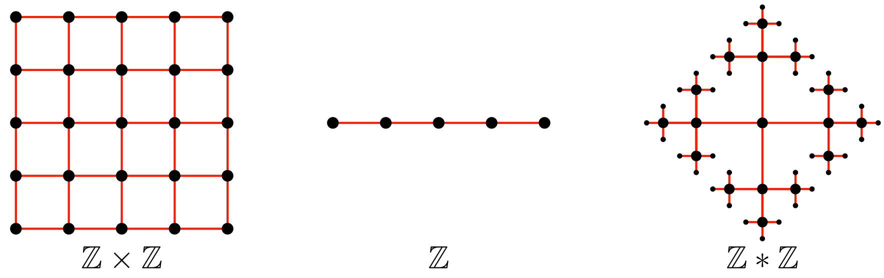

Figure 1.1: Basic examples of Cayley graphs

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

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

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

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

"Hic abundant leones." [5]
"A statement that holds for all finitely generated groups
has to be either trivial or wrong." [attributed to M. Gromov]

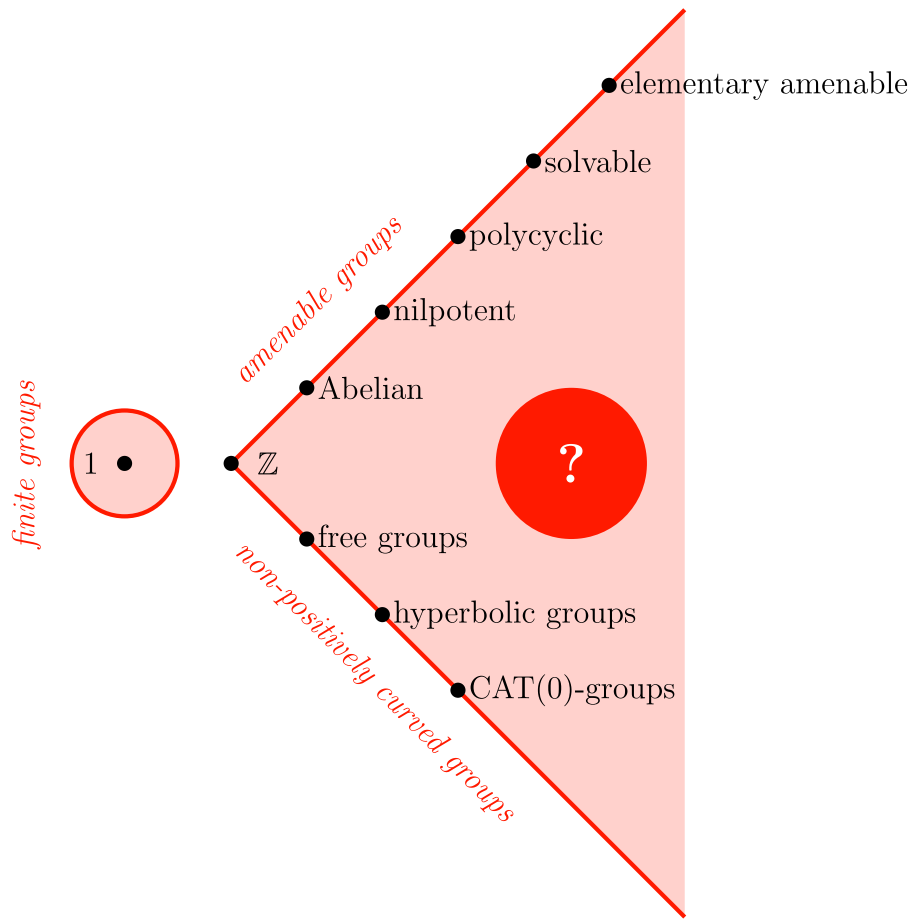

Figure 1.2: The universe of groups (simplified version of Bridson's universe of groups [5])

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

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

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

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

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

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

Literature. The standard resources for geometric group theory are:
– *Topics in Geometric group theory* by de la Harpe [12],
– *Metric spaces of non-positive curvature* by Bridson and Haefliger [7],
– *Trees* by Serre [19].

A short and comprehensible introduction into curvature in classical Riemannian geometry is given in the book *Riemannian manifolds. An introduction to curvature* by Lee [15].

Furthermore, I recommend to look at the overview articles by Bridson on geometric and combinatorial group theory [5, 6]. The original reference for modern large scale geometry of groups is the landmark paper *Hyperbolic groups* [11] by Gromov.

2 Generating groups

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

2.1 Review of the category of groups

2.1.1 Axiomatic description of groups

For the sake of completeness, we briefly recall the definition of a group; more information on basic properties of groups can be found in any textbook on algebra [14]. The category of groups has groups as objects and group homomorphisms as morhpisms.

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

- *Associativity.* For all $g_1, g_2, g_3 \in G$ we have

$$g_1 \cdot (g_2 \cdot g_3) = (g_1 \cdot g_2) \cdot g_3.$$

- *Existence of a neutral element.* There exists a *neutral element* $e \in G$ for "$\cdot$", i.e.,

$$\forall_{g \in G} \;\; e \cdot g = g = g \cdot e.$$

  (Notice that the neutral element is uniquely determined by this property.)
- *Existence of inverses.* For every $g \in G$ there exists an *inverse element* $g^{-1} \in G$ with respect to "$\cdot$", i.e.,

$$g \cdot g^{-1} = e = g^{-1} \cdot g.$$

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

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

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

- "The" trivial group; i.e., the group consisting only of a single element e and the composition $(e, e) \longmapsto e$. Clearly, every group contains "the" trivial group given by the neutral element as subgroup.
- The sets $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ are groups with respect to addition; moreover, $\mathbb{Z}$ is a subgroup of $\mathbb{Q}$, and $\mathbb{Q}$ is a subgroup of $\mathbb{R}$.
- The natural numbers $\mathbb{N}$ do not form a group with respect to addition (e.g., 1 does not have an additive inverse in $\mathbb{N}$); the rational numbers $\mathbb{Q}$ do not form a group with respect to multiplication (0 does not have a multiplicative inverse), but $\mathbb{Q} \setminus \{0\}$ is a group with respect to multiplication.
- Let X be a set. Then the set S_X of all bijections of type $X \longrightarrow X$ is a group with respect to composition of maps, the so-called *symmetric group over X*. If $n \in \mathbb{N}$, then we abbreviate $S_n := S_{\{1,\dots,n\}}$. In general, the group S_X is *not* Abelian.

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

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

- A map $\varphi\colon G \longrightarrow H$ is a *group homomorphism* if φ is compatible with the composition in G and H respectively, i.e., if

$$\varphi(g_1 \cdot g_2) = \varphi(g_1) \cdot \varphi(g_2)$$

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

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

- If H is a subgroup of a group G, then the inclusion $H \hookrightarrow G$ is a group homomorphism.

- Let $n \in \mathbb{Z}$. Then

$$\begin{aligned} \mathbb{Z} &\longrightarrow \mathbb{Z} \\ x &\longmapsto n \cdot x \end{aligned}$$

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

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

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

is the *kernel* of φ, and

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

is the *image* of φ.

Remark 2.1.7 (Isomorphisms via kernel/image)**.**

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 $\operatorname{im} \varphi = H$.

Proof. Exercise. □

2.1.2 Concrete groups – automorphism groups

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

Example 2.1.8 (Symmetric groups)**.** Let X be a set. Then the set of all bijections of type $X \longrightarrow X$ forms a group with respect to composition.

This example is generic in the following sense:

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

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

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

looking at $f_{g^{-1}}$ shows that f_g is a bijection. A straightforward computation shows that

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

is a group homomorphism and hence that G can be viewed as a subgroup of the symmetric group S_G over G. □

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

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

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

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

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

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

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

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

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

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

Exercise 2.1.15 (Groups and categories)**.**
1. Look up the definition of *category*, and *isomorphism in a category* and prove that automorphisms yield groups as described above.
2. What are the (natural) categories in the examples above?
3. Show that, conversely, every group arises as the isomorphism group of an object in some category.

Taking automorphism groups of geometric/algebraic objects is only one way to associate meaningful groups to interesting objects. Over time, many group-valued invariants have been developed in all fields of mathematics. For example:
- fundamental groups (in topology, algebraic geometry or operator algebra theory, ...)
- homology groups (in topology, algebra, algebraic geometry, operator algebra theory, ...)
- ...

2.1.3 Normal subgroups and quotients

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

Definition 2.1.16 (Normal subgroup)**.** Let G be a group. A subgroup N of G is *normal* if it is conjugation invariant, i.e., if

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

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

Example 2.1.17 (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 generated by the cycle given by $1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1$ is a normal subgroup.
- Kernels of group homomorphisms are normal in the domain group; conversely, every normal subgroup also is the kernel of a certain group homomorphism (namely of the canonical projection to the quotient (Proposition 2.1.18)).

Proposition 2.1.18 (Quotient group)**.** *Let G be a group, and let N be a subgroup.*

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

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

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

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

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

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

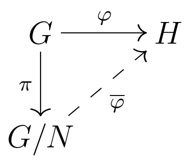

Proof. Ad 1. Suppose that N is normal in G. In this case, the composition map is well-defined (in the sense that the definition does not depend on the choice of the representatives of cosets) because: Let $g_1, g_2, \overline{g}_1, \overline{g}_2 \in G$ with

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

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

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

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

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

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

$g \cdot N = (g \cdot n) \cdot N$, and so (by well-definedness)

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

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

Ad 2. Exercise. □

Example 2.1.19 (Quotient groups)**.**

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

2.2 Groups via generators and relations

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

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

2.2.1 Generating sets of groups

We start by reviewing the concept of a generating set of a group; in geometric group theory, one usually is only interested in finitely generated groups.

Definition 2.2.1 (Generating set)**.**

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

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

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

Example 2.2.3 (Generating sets)**.**

- If G is a group, then G is a generating set of G.
- The trivial group is generated by the empty set.
- The set $\{1\}$ generates the additive group $\mathbb{Z}$; moreover, also, e.g., $\{2, 3\}$ is a generating set for $\mathbb{Z}$.
- Let X be a set. Then the symmetric group S_X is finitely generated if and only if X is finite (exercise).

2.2.2 Free groups

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

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

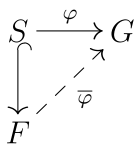

A group is *free* if it contains a free generating set.

Example 2.2.5 (Free groups)**.**

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

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

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

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

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

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

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

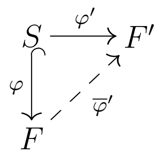

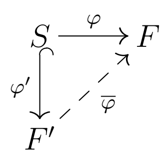

2. We now show that $\overline{\varphi} \circ \overline{\varphi}' = \mathrm{id}_F$ and $\overline{\varphi}' \circ \overline{\varphi} = \mathrm{id}_{F'}$ and hence that φ and φ' are isomorphisms: The composition $\overline{\varphi} \circ \overline{\varphi}'\colon F \longrightarrow F$ is a group homomorphism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

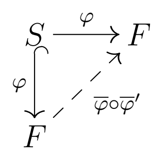

commutative. Moreover, also id_F is a group homomorphism fitting into this diagram. Because F is freely generated by S, the uniqueness part of the universal property thus tells us that these two homomorphisms have to coincide.

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

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

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

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

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

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

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

where $\sim$ is the equivalence relation generated by

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

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

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

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

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

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

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 $\sim$). This shows that inverses exist in $F(S)$.

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

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

inductively by

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

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

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

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

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

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

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 the constructive description as in the previous proof might be more appropriate than the other.

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

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

Proof. By construction, the statement holds for the free group generated by S constructed in the proof of Theorem 2.2.7; in view of the uniqueness result Proposition 2.2.6, we obtain that also the given free group F is generated by S. □

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

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

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

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

Caveat 2.2.11. While subspaces of vector spaces cannot have bigger dimension than the ambient space, free groups of rank 2 contain subgroups that are isomorphic to free groups of higher rank, even free subgroups of (countably) infinite rank.

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

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

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

2.2.3 Generators and relations

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

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

Remark 2.2.14 (Explicit description of generated normal subgroups)**.** Let G be a group and let $S \subset G$. Then the normal subgroup generated by S

in G always exists and can be described as follows:

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

Example 2.2.15.

- As all subgroups of Abelian groups are normal, we have $\langle S\rangle_G^{\triangleleft} = \langle S\rangle_G$ 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 $\langle\tau\rangle_{S_3}^{\triangleleft} = S_3$.

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

Definition 2.2.16 (Generators and relations)**.** Let S be a set, and let $R \subset (S \cup S^{-1})^*$ be a subset; let $F(S)$ be the free group generated by S. Then the group

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

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

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

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

Proposition 2.2.17 (Universal property of groups given by generators and relations)**.** *Let S be a set and let $R \subset (S \cup S^{-1})^*$. The group $\langle S \mid R\rangle$ generated by S with relations R together with the canonical map $\pi\colon S \longrightarrow F(S)/\langle R\rangle_{F(S)}^{\triangleleft} = \langle S \mid R\rangle$ has the following universal property: For any group G and any map $\varphi\colon S \longrightarrow G$ with the property that*

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

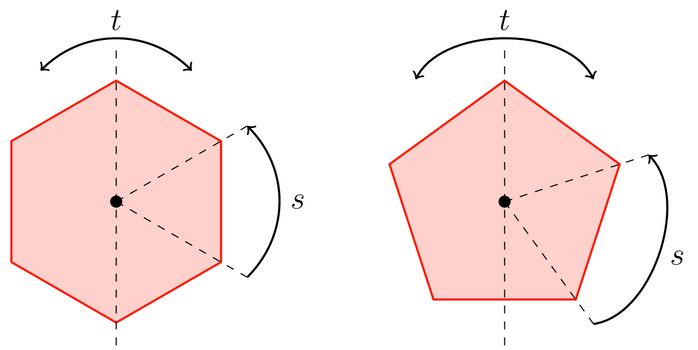

Figure 2.1: Generators of the dihedral groups D_6 and D_5

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

Proof. Exercise. □

Example 2.2.18 (Presentations of groups)**.**

- For all $n \in \mathbb{N}$, we have $\langle x \mid x^n\rangle \cong \mathbb{Z}/n$.
- We have $\langle x, y \mid xyx^{-1}y^{-1}\rangle \cong \mathbb{Z}^2$ (exercise).
- 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*:

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

  geometrically, s corresponds to a rotation about $2\pi/n$ around the centre of the regular n-gon X_n, and t corresponds to a reflection along a diameter passing through one of the vertices (Figure 2.1). One can show that the group $\operatorname{Isom}(X_n)$ contains exactly $2n$ elements, namely, $\mathrm{id}, s, \dots, s^{n-1}, t, t \cdot s, \dots, t \cdot s^{n-1}$.
- The group $G := \langle x, y \mid xyx^{-1} = y^2, yxy^{-1} = x^2\rangle$ is trivial because: Let $\overline{x} \in G$ and $\overline{y} \in G$ denote the images of x and y respectively under the canonical projection

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

By definition, in G we obtain

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

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

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

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

Caveat 2.2.19 (Word problem)**.** The problem to determine whether a group given by generators and relations is the trivial group is undecidable (in the sense of computability theory); i.e., there is no algorithmic procedure that, given generators and relations, can decide whether the corresponding group is trivial or not.

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

The undecidability of the word problem implies the undecidability of many other problems in pure mathematics.

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

Definition 2.2.20 (Finitely presented group)**.** A group G is *finitely presented* if there exists a finite generating set S and a finite set $R \subset (S \cup S^{-1})^*$ of relations such that $G \cong \langle S \mid R \rangle$.

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

Example 2.2.21 (A finitely generated group that is not finitely presented)**.** The group

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

is finitely generated, but not finitely presented [4]. Here, we used the *commutator* notation "$[x, y] := xyx^{-1}y^{-1}$." This group is an example of a so-called *lamplighter group* (see also Example 2.3.5).

While it might be difficult to prove that a specific group is not finitely presented (and such proofs usually require some input from algebraic topology), there is a non-constructive argument showing that there are finitely generated groups that are not finitely presented (Corollary 2.2.23):

Theorem 2.2.22 (Uncountably many finitely generated groups)**.** *There exist uncountably many isomorphism classes of groups generated by two elements.*

Before sketching de la Harpe's proof [12, Chapter III.C] of this theorem, we discuss an important consequence:

Corollary 2.2.23. *There exist uncountably many finitely generated groups that are* not *finitely presented.*

Proof. Notice that there are only countably many finite presentations of groups, and hence that there are only countably many isomorphism types of finitely presented groups. However, there are uncountably many finitely generated groups by Theorem 2.2.22. □

The proof of Theorem 2.2.22 consists of two steps:

1. We first show that there exists a group G generated by two elements that contains uncountably many different normal subgroups (Proposition 2.2.24).
2. We then show that G even has uncountably many quotient groups that are pairwise non-isomorphic (Proposition 2.2.25).

Proposition 2.2.24 (Uncountably many normal subgroups)**.** *There exists a group generated by two elements that has uncountably many normal subgroups.*

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

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

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

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

Proposition 2.2.25 (Uncountably many quotients)**.** *For a finitely generated group G the following are equivalent:*

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

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

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

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

2.3 New groups out of old

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

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

2.3.1 Products and extensions

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

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

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

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

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

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

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

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

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

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

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

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

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

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

is a split extension; here, $i\colon N \longrightarrow N \rtimes_\varphi$ 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 (exercise).

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

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

does not split because there is no non-trivial homomorphism from the torsion group $\mathbb{Z}/2\mathbb{Z}$ to $\mathbb{Z}$. One way to classify group extensions is to consider group cohomology [8, 16, Chapter IV,Chapter 1.4.4].

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

- Let $n \in \mathbb{N}_{\geq 3}$. Then the dihedral group D_n (see Example 2.2.18) is a semi-direct product

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

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

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

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

2.3.2 Free products and free amalgamated products

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

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

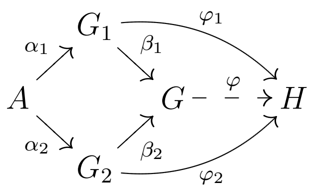

Such a free product with amalgamation is denoted by $G_1 *_A G_2$ (see Theorem 2.3.9 for existence and uniqueness).

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

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

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

- Free groups can also be viewed as free products of several copies of the additive group $\mathbb{Z}$; e.g., the free group of rank 2 is nothing but $\mathbb{Z} * \mathbb{Z}$ (which can be seen by comparing the respective universal properties and using uniqueness).

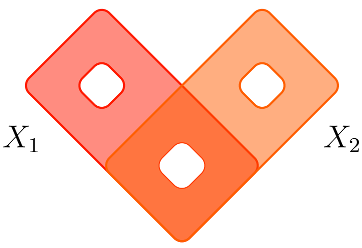

$$\pi_1(X_1 \cup_A X_2) \cong \pi_1(X_1) *_{\pi_1(A)} \pi_1(X_2)$$

Figure 2.2: The theorem of Seifert and van Kampen, schematically

- The infinite dihedral group $D_\infty = \operatorname{Isom}(\mathbb{Z})$ is isomorphic to the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$; for instance, reflection at 0 and reflection at 1/2 provide generators of D_∞ corresponding to the obvious generators of $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z}$.
- The matrix group $\operatorname{SL}(2, \mathbb{Z})$ is isomorphic to the free amalgamated product $\mathbb{Z}/6\mathbb{Z} *_{\mathbb{Z}/2\mathbb{Z}} \mathbb{Z}/4\mathbb{Z}$ [**?**].
- Free amalgamated products occur naturally in topology: By the theorem of Seifert and van Kampen, the fundamental group of a space glued together of several components is a free amalgamated product of the fundamental groups of the components over the fundamental group of the intersection (the two subspaces and their intersection have to be non-empty and path-connected) [17, Chapter IV] (see Figure 2.2).

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

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

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

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

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

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

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

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

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

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

Let H be a group and let $\varphi_1 \colon G_1 \longrightarrow H$ and $\varphi_2 \colon G_2 \longrightarrow H$ be homomorphisms with $\varphi_1 \circ \alpha_1 = \varphi_2 \circ \alpha_2$. We define a homomorphism $\varphi \colon G \longrightarrow H$ using the universal property of groups given by generators and relations (Proposition 2.2.17): The map on the set of all words in the generators $\{x_g \mid g \in G\} \sqcup \{x_g \mid g \in G\}$ and their formal inverses induced by the map

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

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

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

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

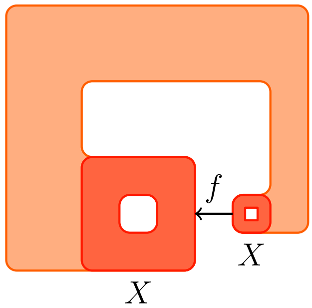

$$\pi_1(\text{mapping torus of } f) \cong \pi_1(X)*_{\pi(f)}$$

Figure 2.3: The fundamental group of a mapping torus, schematically

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

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

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

where

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

In other words, using an HNN-extension, we can force two given subgroups to be conjugate; iterating this construction leads to interesting examples of groups [**?**]. HNN-extensions are named after G. Higman, B.H. Neumann, and H. Neumann who were the first to systematically study such groups. Topologically, HNN-extensions arise naturally as fundamental groups of mapping tori of maps that are injective on the level of fundamental groups [**?**] (see Figure 2.3).

Remark 2.3.11 (Outlook – ends of groups)**.** The class of (non-trivial) free amalgamated products and of (non-trivial) HNN-extensions plays an important rôle in geometric group theory, more precisely, they are the key objects in Stallings's classification of groups with infinitely many ends [20].

3 Groups $\longrightarrow$ geometry, I: Cayley graphs

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

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

We start by reviewing some basic notation from graph theory (Section 3.1); more information on graph theory can be found in various textbooks [10, 13, 9].

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

3.1 Review of graph notation

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

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

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

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

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

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

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

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

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

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

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

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

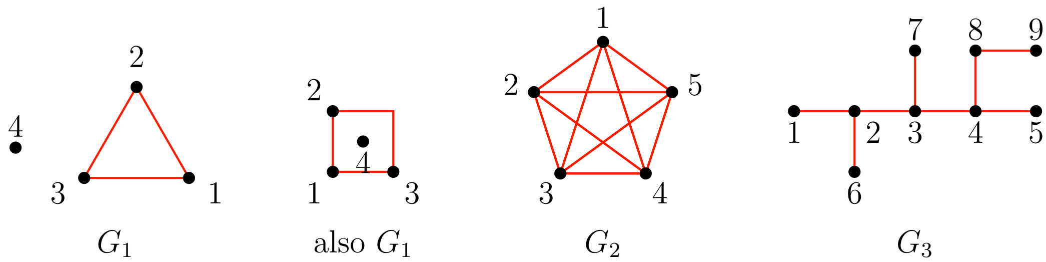

Figure 3.1: Some graphs

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

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

The problem to decide whether two given graphs are isomorphic or not is a difficult problem – in the case of finite graphs, this problem is a so-called NP-complete problem [**?**].

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

- Let $n \in \mathbb{N} \cup \{\infty\}$. A *path in G of length n* is a sequence $v_0, \dots, v_n$ of different vertices $v_0, \dots, v_n \in V$ with the property that $\{v_j, v_{j+1}\} \in E$ holds for all $j \in \{0, \dots, n-1\}$; if $n < \infty$, then we say that this path *connects the vertices v_0 and v_n*.
- The graph G is called *connected* if any two of its vertices can be connected by a path in G.
- Let $n \in \mathbb{N}_{>2}$. A *cycle in G of length n* is a sequence $v_0, \dots, v_{n-1}$ of different vertices $v_0, \dots, v_{n-1} \in V$ with $\{v_{n-1}, v_0\} \in E$ and moreover $\{v_j, v_{j+1}\} \in E$ for all $j \in \{0, \dots, n-2\}$.

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

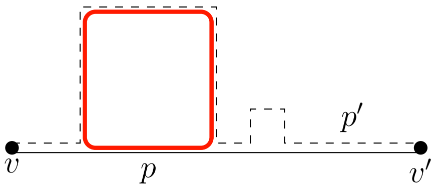

Figure 3.2: Constructing a cycle (orange) out of two different paths.

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

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

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

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

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

3.2 Cayley graphs

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

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

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

$$\bigl\{\{g, g \cdot s\} \bigm| g \in G,\ s \in S \setminus \{e\}\bigr\}.$$

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

Example 3.2.2 (Cayley graphs)**.**

- The Cayley graphs of the additive group $\mathbb{Z}$ with respect to the generating sets $\{1\}$ and $\{2, 3\}$ respectively are illustrated in Figure 3.3. Notice that when looking at these two graphs "from far away" they seem to have the same global structure, namely they look like the real line; in more technical terms, these graphs are quasi-isometric with respect to the corresponding word metrics – a concept that we will study thoroughly in later chapters (Chapter 5 ff.).
- The Cayley graph of the additive group $\mathbb{Z}^2$ with respect to the generating set $\{(1, 0), (0, 1)\}$ looks like the integer lattice in $\mathbb{R}^2$, see Figure 3.4; when viewed from far away, this Cayley graph looks like the Euclidean plane.
- The Cayley graph of the cyclic group $\mathbb{Z}/6\mathbb{Z}$ with respect to the generating set $\{[1]\}$ looks like a cycle graph (Figure 3.5).
- We now consider the symmetric group S_3. Let τ be the transposition exchanging 1 and 2, and let σ be the cycle $1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1$; the

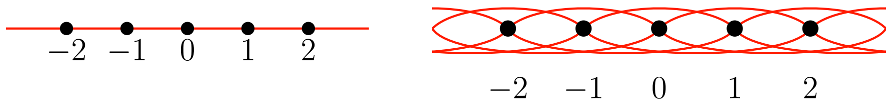

$\mathrm{Cay}\big(\mathbb{Z}, \{1\}\big)$ $\mathrm{Cay}\big(\mathbb{Z}, \{2,3\}\big)$

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

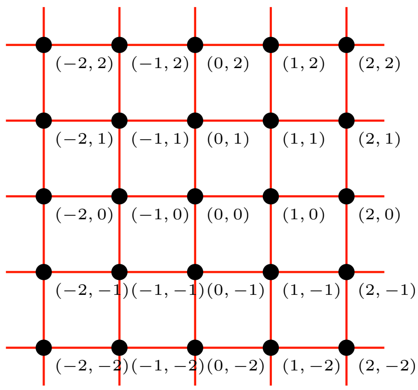

Figure 3.4: The Cayley graph $\mathrm{Cay}\big(\mathbb{Z}^2, \{(1,0),(0,1)\}\big)$

Cayley graph of S_3 with respect to the generating system $\{\tau, \sigma\}$ is depicted in Figure 3.5.
Notice that the Cayley graph $\mathrm{Cay}(S_3, S_3)$ is a complete graph on six vertices; similarly $\mathrm{Cay}(\mathbb{Z}/6\mathbb{Z}, \mathbb{Z}/6\mathbb{Z})$ is a complete graph on six vertices. In particular, we see that non-isomorphic groups may have isomorphic Cayley graphs with respect to certain generating systems.

– The Cayley graph of a free group with respect to a free generating set is a tree (see Theorem 3.3.1 below).

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

1. Cayley graphs are connected as every vertex g can be reached from the vertex of the neutral element by walking along the edges corresponding to a presentation of g in terms of the given generators.

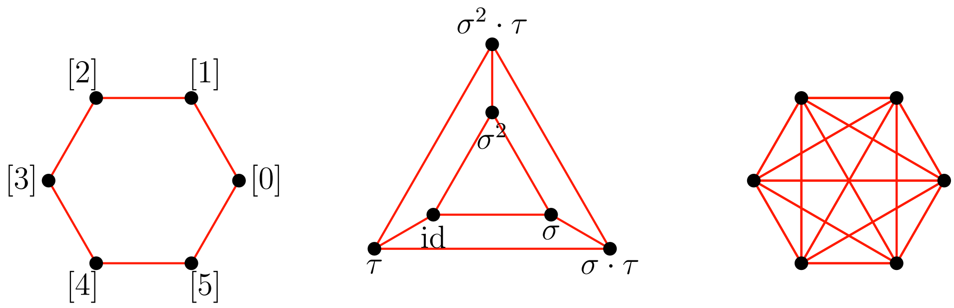

$\mathrm{Cay}\big(\mathbb{Z}/6\mathbb{Z}, \{[1]\}\big)$ $\mathrm{Cay}\big(S_3, \{\tau, \sigma\}\big)$ $\mathrm{Cay}(S_3, S_3)$ $\cong \mathrm{Cay}(\mathbb{Z}/6\mathbb{Z}, \mathbb{Z}/6\mathbb{Z})$

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

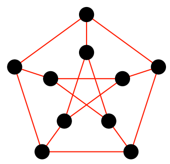

Figure 3.6: The Petersen graph

2. Cayley graphs are regular in the sense that every vertex has the same number $|(S \cup S^{-1}) \setminus \{e\}|$ of neighbours.
3. A Cayley graph is locally finite if and only if the generating set is finite; a graph is said to be *locally finite* if every vertex has only finitely many neighbours.

Exercise 3.2.4 (Petersen graph)**.** Show that the Petersen graph (depicted in Figure 3.6), even though it is a highly regular graph, is *not* the Cayley graph of any group.

Remark 3.2.5 (Cayley complexes, classifying spaces)**.** There are higher dimensional analogues of Cayley graphs in topology: Associated with a presentation of a group, there is the so-called Cayley complex [7, Chapter 8A], which is a 2-dimensional object. More generally, every group admits a so-called classifying space, a space whose fundamental group is the given group, and whose higher dimensional homotopy groups are trivial [?]. These spaces allow to model group theory in topology and play an important rôle in the study of group cohomology [8, 16].

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

3.3 Cayley graphs of free groups

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

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

The converse is *not* true in general:

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

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

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

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

While it might be intuitively clear that free generating set do not lead to any cycles in the corresponding Cayley graphs and vice versa, a formal proof requires the description of free groups in terms of reduced words (Section 3.3.1).

3.3.1 Free groups and reduced words

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

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

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

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

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

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

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

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

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

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

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

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

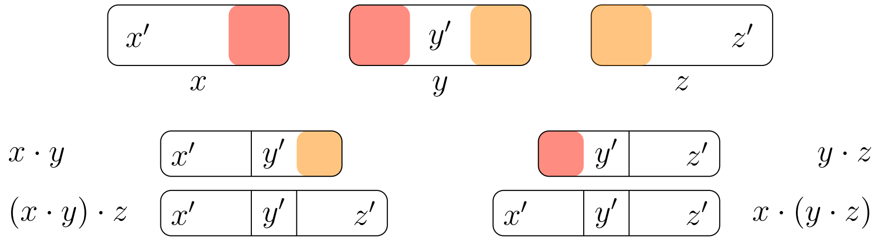

Figure 3.7: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do not interfere

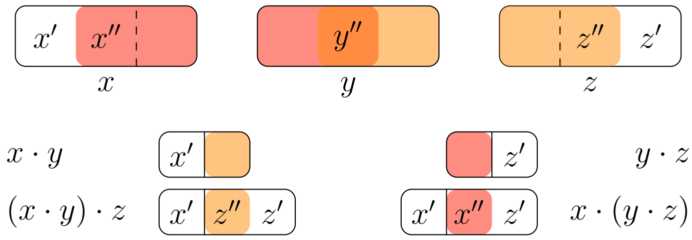

Figure 3.8: Associativity of the composition in $F_{\text{red}}(S)$; if the reduction areas of the outer elements do interfere

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

Thus it remains to prove that this composition is associative (which is the ugly part of this construction): Instead of giving a formal proof involving lots of indices, we sketch the argument graphically (Figures 3.7 and 3.8): Let x, y, $z \in F_{\text{red}}(S)$; we want to show that $(x \cdot y) \cdot z = x \cdot (y \cdot z)$. By definition, when composing two reduced words, we have to remove the maximal reduction area where the two words meet.

– If the reduction areas of x, y and y, z have no intersection in y, then clearly $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ (Figure 3.7).
– If the reduction areas of x, y and y, z have a non-trivial intersection y''

in y, then the equality $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ follows by carefully inspecting the reduction areas in x and z and the neighbouring regions, as indicated in Figure 3.8; notice that because of the overlap in y'', we know that x'' and z'' coincide (they both are the inverse of y'').

Ad. 2. We show that S is a free generating set of $F_{\text{red}}(S)$ by verifying that the universal property is satsified: 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 φ to the set $(S \cup \overline{S})^*$ of all words). Clearly, $\overline{\varphi}|_S = \varphi$; because S generates $F_{\text{red}}(S)$ it follows that $\overline{\varphi}$ is the only such homomorphism. Hence, $F_{\text{red}}(S)$ is freely generated by S. □

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

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

Corollary 3.3.7. *The word problem in free groups is solvable – we just need to consider and compare reduced words.*

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

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

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

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

Similarly, one can also describe free amalgamated products and HNN-extensions by suitable classes of reduced words [19, Chapter I].

3.3.2 Free groups → trees

Proof of Theorem 3.3.1. Suppose the group F is freely generated by S. By Proposition 3.3.5, the group F is isomorphic to $F_{\text{red}}(S)$ via an isomorphism

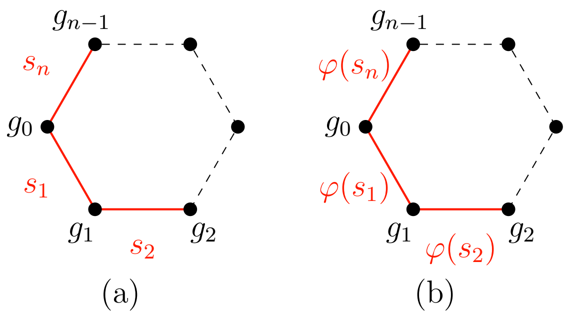

Figure 3.9: Cycles lead to reduced words, and vice versa

that is the identity on S; without loss of generality we can therefore assume that F is $F_{\text{red}}(S)$.

We now show that then 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 $s_n := g_0 \cdot g_{n-1}{}^{-1} \in S \cup S^{-1}$ (Figure 3.9 (a)). Because the vertices are distinct, the word $s_0 \dots s_{n-1}$ is reduced; on the other hand, we obtain

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

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

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

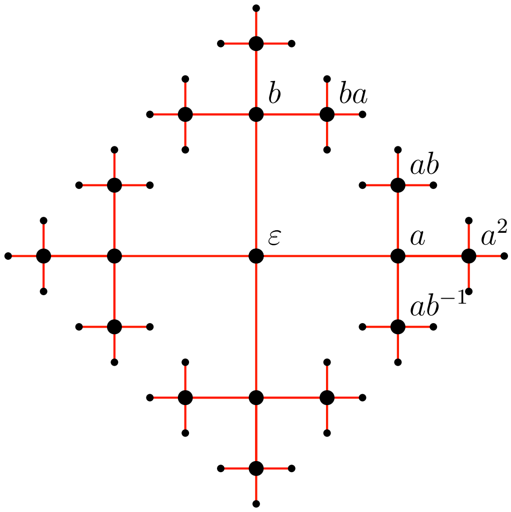

Figure 3.10: Cayley graph of the free group of rank 2 with respect to a free generating set $\{a, b\}$

3.3.3 Trees $\longrightarrow$ free groups

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

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

- Because $\varphi|_S = \mathrm{id}_S$ is injective, it follows that $n > 1$.
- If $n = 2$, then it would follow that

$$e = \varphi(s_1 \cdot s_2) = \varphi(s_1) \cdot \varphi(s_2) = s_1 \cdot s_2$$

in G, contradicting that $s_1 \dots s_n$ is reduced and that $s \cdot t \neq e$ for all s, $t \in S$.

– If $n \geq 3$, we consider the sequence $g_0, \dots, g_{n-1}$ of elements of G given inductively by $g_0 := e$ and

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

for all $j \in \{0, \dots, n-2\}$ (Figure 3.9 (b)). The sequence $g_0, \dots, g_{n-1}$ is a cycle in $\mathrm{Cay}(G, S)$ because by minimality of the word $s_1 \dots s_n$, the elements $g_0, \dots, g_{n-1}$ are all distinct; moreover, $\mathrm{Cay}(G, S)$ contains the edges $\{g_0, g_1\}, \dots, \{g_{n-2}, g_{n-1}\}$, and the edge

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

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

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

4

Groups $\longrightarrow$ geometry, II: Group actions

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

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

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

4.1 Review of group actions

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

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

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

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

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

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

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

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

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

  In other words: group actions are a concept generalising automorphism/symmetry groups.
- Let G be a group and let X be a set. If $\varrho\colon G \longrightarrow \mathrm{Aut}_{\mathsf{Set}}(X)$ is an action of G on X by bijections, then we also use the notation

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

  for $g \in G$ and $x \in X$, and we can view ϱ as a map $G \times X \longrightarrow X$.

More generally, we also use this notation whenever the group G actions on an object in a category, where morphisms are maps of sets and the composition of morphisms is nothing but composition of maps. This applies for example to
  - actions by isometries on a metric space,
  - actions by homeomorphisms on a topological space,
  - ...
- Further examples of group actions are actions of groups on a topological space by homotopy equivalences or actions on a metric space by quasi-isometries (see Chapter 5); notice however, that in these cases, automorphisms are equivalence classes of maps of sets and composition of morphisms is done by composing representatives of the corresponding equivalence classes.
- An action of a group on a vector space by linear isomorphisms is called a *representation* of the group in question.

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

4.1.1 Free actions

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

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

$$g \cdot x \neq x$$

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

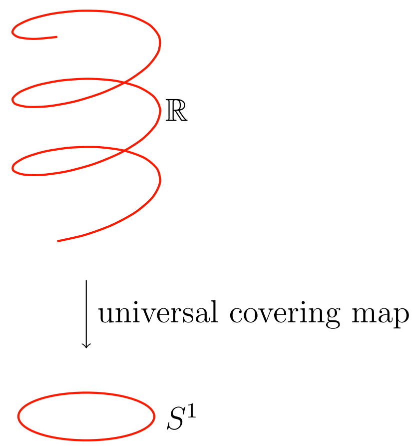

Figure 4.1: Universal covering of S^1

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

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

is a free action of G on itself by bijections.

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

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

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

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

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

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).

E.g., the fundamental group of S^1 is isomorphic to $\mathbb{Z}$, the universal covering of S^1 is the exponential map $\mathbb{R} \longrightarrow S^1$, and the action of the fundamental group of S^1 on $\mathbb{R}$ is given by translation (Figure 4.1).

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

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

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

Example 4.1.9 (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:

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

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

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

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

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

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

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

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

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

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

4.1.2 Orbits

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

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

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

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

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

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

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

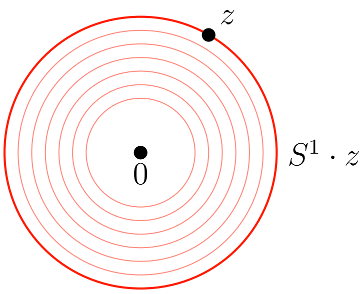

Figure 4.2: Orbits of the rotation action of S^1 on $\mathbb{C}$

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

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

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

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

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

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

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

- Let $t \in G$ be the reflection along the diagonal passing through $(0,0)$ and $(1,1)$. Then
$$Q^t = \bigl\{(x,x) \bigm| x \in [0,1]\bigr\}.$$
- Let $s \in G$ be rotation around $2\pi/4$. Then
$$Q^s = \{(1/2, 1/2)\}.$$
- The orbit of $(0,0)$ are all four vertices of Q, and the stabiliser of $(0,0)$ is $G_{(0,0)} = \{\mathrm{id}_Q, t\}$.
- The stabiliser of $(1/3,0)$ is the trivial group.
- The stabiliser of $(1/2,1/2)$ is $G_{(1/2,1/2)} = G$, so $(1/2,1/2)$ is a global fixed point of this action.

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

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

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

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

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

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

Ad 2. This equality is proved by double counting: More precisely, we consider the set
$$F := \{(g, x) \mid g \in G, x \in X, g \cdot x = x\} \subset G \times X.$$

By definition of stabiliser groups and fixed sets, we obtain
$$\sum_{x \in X} |G_x| = |F| = \sum_{g \in G} |X^g|.$$

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

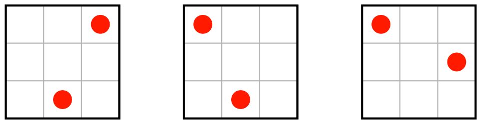

Figure 4.3: The same punchcard seen from different sides/angles

the first part, we obtain

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

□

4.1.3 Application: Counting via group actions

Group actions can be used to solve certain counting problems. A standard example from algebra is the proof of the Sylow theorems in finite group theory: in this proof, the conjugation action of a group on the set of certain subgroups is considered [2, Chapter 20].

Another class of examples arises in combinatorics:

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

In terms of group actions, this question can be reformulated as follows:

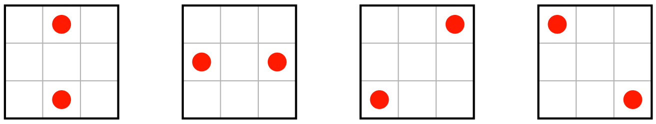

Figure 4.4: The punchcards fixed by rotation around π

We consider the set of all configurations

$$\{(x,y) \mid x,y \in \{0,1,2\} \times \{0,1,2\},\ x \neq y\}$$

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

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

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

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

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

different punchcards.

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

4.2 Free groups and actions on trees

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

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

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

Applying the universal property of F with respect to the free generating system S to maps $S \longrightarrow \mathbb{Z}$ it is easily seen that S cannot contain any element of order 2; therefore, the left translation action of F on $\operatorname{Cay}(F, S)$ is free by Proposition 4.1.10. □

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

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

4.2.1 Spanning trees

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

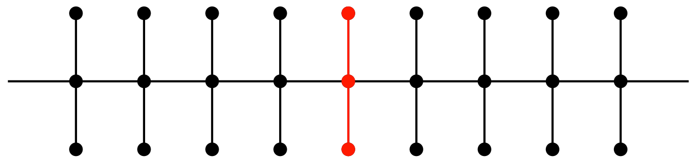

Figure 4.5: A spanning tree (red) for a shift action of $\mathbb{Z}$

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

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

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

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

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

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

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

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

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

Let v be a vertex such that none of the vertices of the orbit $G \cdot v$ is in T, and such that some neighbour u of v is in T. Then clearly, adding v and the edge $\{u, v\}$ to T produces a tree in T_G, which contains T as a proper subgraph. This contradicts the maximality of T. Hence, T is a spanning tree for the G-action on X. □

4.2.2 Completing the proof

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

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

The idea is to think about the graph obtained from T by contracting T' and all its copies $g \cdot T'$ for $g \in G$ each to a single vertex (Figure 4.6 shows this in the situation of Example 4.2.3). Hence, the candidates for a generating set come from the edges joining these new vertices: An edge

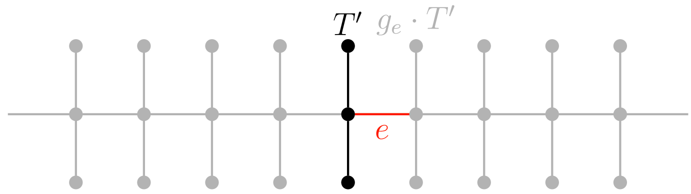

Figure 4.7: An essential edge (red) for the shift action of $\mathbb{Z}$ (Example 4.2.3)

of T is called *essential* if it does not belong to T', but if one of the vertices of the edge in question belongs to T' (then the other vertex cannot belong to T' as well, by uniqueness of paths in trees (Proposition 3.1.9)).

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

We define

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

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

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

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

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

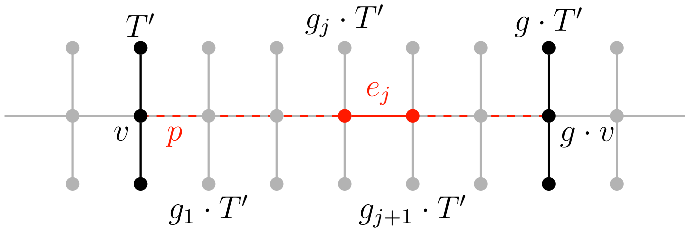

Figure 4.8: The set $\widetilde{S}$ generates G

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

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

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

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

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

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

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

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

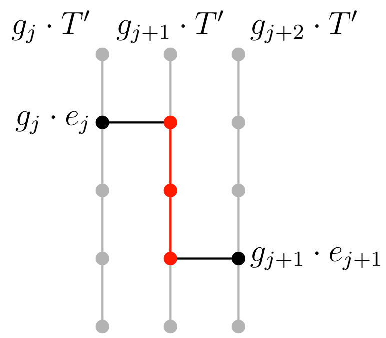

Figure 4.9: Cycles in $\mathrm{Cay}(G, \widetilde{S})$ lead to cycles in T by connecting translates of essential edges in the corresponding translates of T' (red path)

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

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

Remark 4.2.5 (Topological proof)**.** Let G be a group acting freely on a tree; then G also acts freely, continuously and properly discontinuously on the realisation X of this tree, which is contractible and homeomorphic to a one-dimensional CW-complex (such that the G-action is cellular). 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.

Remark 4.2.6 (Characterising free products with amalgamation)**.** Similarly, also free products with amalgamations and HNN-extensions can be characterised geometrically by suitable actions on graphs [19]; generalisations of this type eventually lead to complexes of groups and Bass-Serre theory [19].

4.3 Application: Subgroups of free groups are free

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Proof. Exercise. □

Corollary 4.3.5. *Finite index subgroups of finitely generated groups are finitely generated.*

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

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

We will later see an alternative proof of Corollary 4.3.5 via the Švarc-Milnor lemma.

Corollary 4.3.6 (Subgroups of free products)**.** *Let G and H be finite groups. Then all torsion-free subgroups of the free product $G * H$ are free groups.*

Sketch of proof. 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 vertices is

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

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

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

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

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

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

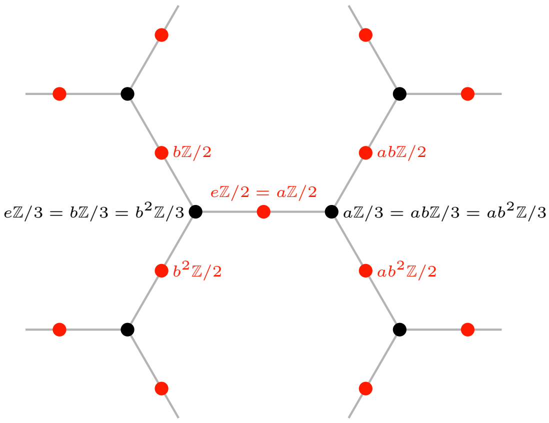

Figure 4.10: The tree for the free product $\mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z} \cong \langle a, b \,|\, a^2, b^3 \rangle$

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. Therefore, any torsion-free subgroup of $G * H$ acts freely on the tree X. Applying Theorem 4.2.1 finishes the proof. □

A similar technique as in the previous proof shows that all torsion-free subgroups of $\mathrm{SL}(2, \mathbb{Q}_p)$ are free [19].

4.4 The ping-pong lemma

The following criterion for freeness via suitable actions is due to Klein [**?**]:

Theorem 4.4.1 (Ping-pong lemma)**.** *Let G be a group, generated by two elements a and b of infinite order. Suppose there is a G-action on a set X*

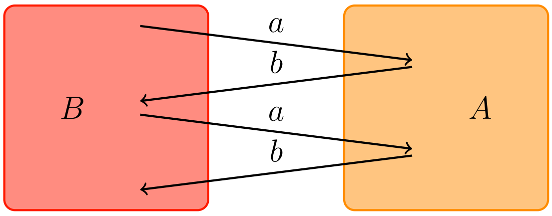

Figure 4.11: The ping-pong lemma

such that there are non-empty subsets A, $B \subset X$ with B not included in A and such that for all $n \in \mathbb{Z} \setminus \{0\}$ we have

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

Then G is freely generated by $\{a, b\}$.

Proof. It suffices to find an isomorphism $F_{\text{red}}(\{a, b\}) \cong G$ that extends the identity on $\{a, b\}$. By the universal property of the free group $F_{\text{red}}(\{a, b\})$ there is a group homomorphism $\varphi\colon F_{\text{red}}(\{a, b\}) \longrightarrow G$ extending the identity on $\{a, b\}$. 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 $w \in F_{\text{red}}(\{a, b\}) \setminus \{\varepsilon\}$ 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 a, i.e., we can write $w = a^{n_0} b^{m_1} a^{n_1} \dots b^{m_k} a^{n_k}$ for some $k \in \mathbb{N}$ and certain $n_0, \dots, n_k, m_1, \dots, m_k \in \mathbb{Z} \setminus \{0\}$. Then (ping-pong! – see Figure 4.11)

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

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

2. The word w starts and ends with non-trivial powers of b. Then awa^{-1} is a reduced word starting and ending in non-trivial powers of a. So

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

contradicting what we already proved for the first case.

3. The word w starts with a non-trivial power of a and ends with a non-trivial power of b, say $w = a^n w' b^m$ with $n.m \in \mathbb{Z} \setminus \{0\}$ and w' a reduced word not starting with a non-trivial power of a and not ending in a non-trivial power of b. Let $r \in \mathbb{Z} \setminus \{0, -n\}$. Then $a^r w a^{-r} = a^{r+n} w' b^m a^r$ starts and ends with a non-trivial power of a and

$$e = \varphi(a^r w a^{-r}),$$

contradicting what we already proved for the first case.

4. The word w starts with a non-trivial power of b and ends with a non-trivial power of a. Then the inverse of w falls into the third case and $\varphi(w^{-1}) = e$, which cannot be.

Therefore, φ is injective, and so $\varphi\colon F_{\text{red}}(\{a,b\}) \longrightarrow G$ is an isomorphism extending the identity on $\{a,b\}$, as was to be shown. □

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

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

then $G \cong G_1 * G_2$.

4.5 Application: Free subgroups of matrix groups

Using the ping-pong lemma we can establish the freeness of certain matrix groups.

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

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

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

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

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

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

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

and

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

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

Further examples are given in de la Harpe's book [12, Chapter II.B]; in particular, it can be shown that the group of homeomorphisms $\mathbb{R} \longrightarrow \mathbb{R}$ contains a free group of rank 2.

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

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

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

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

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

5.1 Quasi-isometry types of metric spaces

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

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

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

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

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

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

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

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

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

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

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

  In other words, an isometry is nothing but a bijective isometric embedding.

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

Remark 5.1.3. Clearly, every isometric embedding is injective, and every isometry is a homeomorphism with respect to the topologies induced by the metrics.

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

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

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

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

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

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

In other words, a bilipschitz embedding is nothing but a bijective bilipschitz embedding.

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

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

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

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

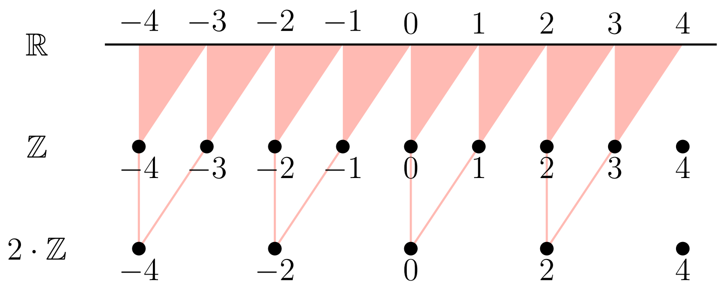

Figure 5.1: The metric spaces $\mathbb{R}$, $\mathbb{Z}$, and $2 \cdot \mathbb{Z}$ are quasi-isometric

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

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

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

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

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

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

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

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

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

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

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

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

Caveat 5.1.8. In particular, we see that

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

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

- All non-empty metric spaces of finite diameter are quasi-isometric; the *diameter* of a metric space (X, d) is
$$\operatorname{diam} X := \sup_{x,y \in X} d(x, y).$$
- Conversely, if a space is quasi-isometric to a space of finite diamter, then it has finite diameter as well. So the metric space $\mathbb{Z}$ (with the metric induced from $\mathbb{R}$) is *not* quasi-isometric to a non-empty metric space of finite diameter.
- The metric spaces $\mathbb{R}$ and $\mathbb{R}^2$ (with respect to the Euclidean metric) are *not* quasi-isometric (exercise).

Exercise 5.1.10.
1. Are the metric spaces $\mathbb{Z}$ and $\mathbb{N}$ with the metrics induced from the standard metric on $\mathbb{R}$ quasi-isometric?
2. Are the metric spaces $\mathbb{Z}$ and $\{n^3 \mid n \in \mathbb{Z}\}$ with the metrics induced from the standard metric on $\mathbb{R}$ quasi-isometric?

Proposition 5.1.11 (Alternative characterisation of quasi-isometries)**.** *A map $f\colon X \longrightarrow Y$ between metric spaces (X, d_X) and (Y, d_Y) is a quasi-isometry if and only if it is a quasi-isometric embedding with quasi-dense image; a map $f\colon X \longrightarrow Y$ is said to have* quasi-dense image *if there is a constant $c \in \mathbb{R}_{>0}$ such that*

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

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

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

in particular, f has quasi-dense image.

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

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

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

By the axiom of choice, we can define a map

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

where we choose for every $y \in Y$ an element x_y with $d_Y(f(x_y), y) \leq c$.

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

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

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

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

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

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

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

and

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

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

As quasi-isometries are not bijective in general, some care has to be taken when defining quasi-isometry groups of metric spaces:

Proposition 5.1.12 (Quasi-isometry groups)**.**

1. *The composition of two quasi-isometric embeddings is a quasi-isometric embedding and the composition of two quasi-isometries is a quasi-isometry.*
2. *Let (X, d) be a metric space. Let $\widetilde{\mathrm{QI}}(X)$ be the set of all quasi-isometries of type $X \longrightarrow X$, and let $\mathrm{QI}(X) := \widetilde{\mathrm{QI}}(X)/\sim$ be the*

set of quasi-isometries of X modulo finite distance. Then the composition

$$\begin{aligned} \mathrm{QI}(X) \times \mathrm{QI}(X) &\longrightarrow \mathrm{QI}(X) \\ \bigl([f],[g]\bigr) &\longmapsto [f \circ g] \end{aligned}$$

is well-defined, and $\mathrm{QI}(X)$ *is a group with respect to this composition. The group* $\mathrm{QI}(X)$ *is the* quasi-isometry group of X.

Proof. Exercise. □

The first part also shows that quasi-isometry behaves like an equivalence relation on the class of all metric spaces. Moreover, we can define a category whose objects are metric spaces and whose morphisms are quasi-isometric embeddings modulo finite distance; the automorphism groups in this category then coincide with the quasi-isometry groups defined above.

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

Example 5.1.13 (Quasi-isometry groups)**.**

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

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

together with many rather large and non-commutative groups [18].

5.2 Quasi-isometry types of groups

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

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

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

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

Definition 5.2.2 (Word metric)**.** Let G be a group and let $S \subset G$ be a generating set. The *word metric* d_S *on* G *with respect to* S is the metric on G associated with the Cayley graph $\operatorname{Cay}(G, S)$. In other words,

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

for all $g, h \in G$.

Example 5.2.3 (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 neglegible when looking at the group from far away:

Proposition 5.2.4. *Let G be a finitely generated group, and let S and S' be two finite generating sets of G.*

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

Proof. The second part directly follows from the first part because the composition of bilipischitz equivalences is a bilipschitz equivalence, and the composition of quasi-isometries is a quasi-isometry (Proposition 5.1.12).

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

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

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

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

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

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

For infinite generating sets the first part of the above proposition does not hold in general; for example taking $\mathbb{Z}$ as a generating set for $\mathbb{Z}$ leads to the space $(\mathbb{Z}, d_{\mathbb{Z}})$ of finite diameter, while $(\mathbb{Z}, d_{\{1\}})$ does *not* have finite diameter (Example 5.2.3).

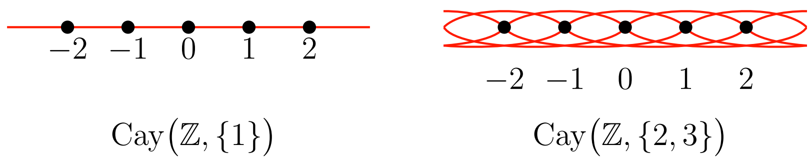

Figure 5.2: Cayley graphs of $\mathbb{Z}$ having the same large scale geometry

Definition 5.2.5 (Quasi-isometry type of finitely generated groups)**.** Let G be a finitely generated group.

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

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

Example 5.2.6. If $n \in \mathbb{N}$, then $\mathbb{Z}^n$ is quasi-isometric to Euclidean space $\mathbb{R}^n$ because the inclusion $\mathbb{Z}^n \hookrightarrow \mathbb{R}^n$ is a quasi-isometric embedding with quasi-dense image. In this sense, Cayley graphs of $\mathbb{Z}^n$ (with respect to finite generating sets) resemble the geometry of $\mathbb{R}^n$ (Figure 5.3).

At this point it might be more natural to consider bilipschitz equivalence of groups as a good geometric equivalence of finitely generated groups; however, later we will see why considering quasi-isometry types of groups is more appropriate.

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:

Exercise 5.2.7 (Quasi-isometry vs. bilipschitz equivalence)**.**

1. Show that a bijective quasi-isometry between finitely generated groups (with respect to the word metric of certain finite generating sets) is a bilipschitz equivalence.

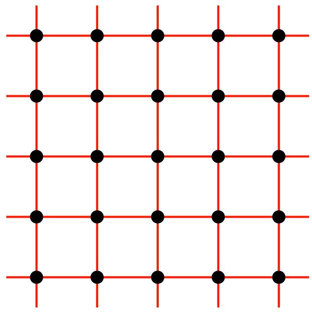

Figure 5.3: The Cayley graph $\operatorname{Cay}\bigl(\mathbb{Z}^2, \{(1,0),(0,1)\}\bigr)$ resembles the geometry of the Euclidean plane $\mathbb{R}^2$

2. Does this also hold in general? I.e., are all bijective quasi-isometries between general metric spaces necessarily bilipschitz equivalences?

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

5.2.1 First examples

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

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

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

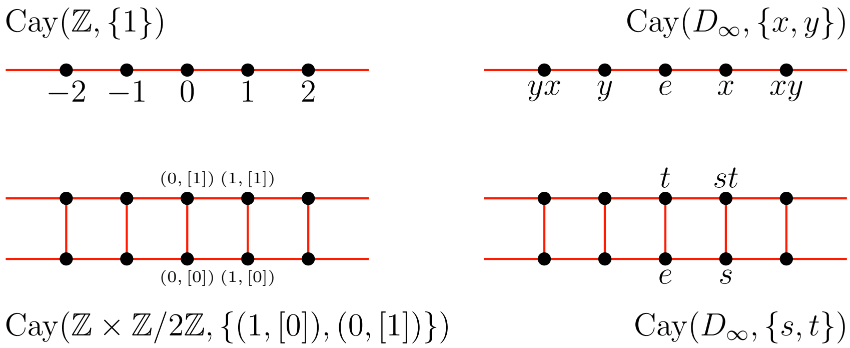

Figure 5.4: The groups $\mathbb{Z}$, $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ and D_∞ are quasi-isometric

translation action of G, it follows that all balls in (G, d_S) of finite radius are finite.

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

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

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

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

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

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

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

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

Exercise 5.2.11 (Groups not quasi-isometric to $\mathbb{Z}$)**.** Let $n \in \mathbb{N}_{\geq 2}$.

1. Show that the groups $\mathbb{Z}$ and $\mathbb{Z}^n$ are *not* quasi-isometric. In particular, $\mathbb{R}$ is not quasi-isometric to $\mathbb{R}^n$ with the Euclidean metric (because these spaces are quasi-isometric to $\mathbb{Z}$ and $\mathbb{Z}^n$ respectively).
2. Show that the group $\mathbb{Z}$ is *not* quasi-isometric to a free group of rank n.

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

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

In other words, the property of being virtually $\mathbb{Z}$ is a geometric property of groups.

We will give a proof of this result later when we have more tools at hand.

5.3 The Švarc-Milnor lemma

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

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

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

5.3.1 Quasi-geodesics and quasi-geodesic spaces

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

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

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

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

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

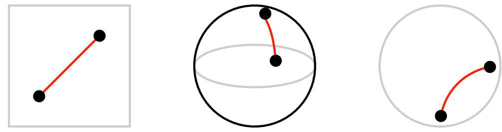

Figure 5.5: Geodesic spaces (and example geodesics (red))

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

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

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

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

Clearly, any geodesic space is also quasi-geodesic (namely, $(1, 0)$-quasi-geodesic); the converse does not hold in general:

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

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

5.3.2 The Švarc-Milnor lemma

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

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

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

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

Then the following holds:

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

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

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

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

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

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

and $t_n := L$, as well as

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

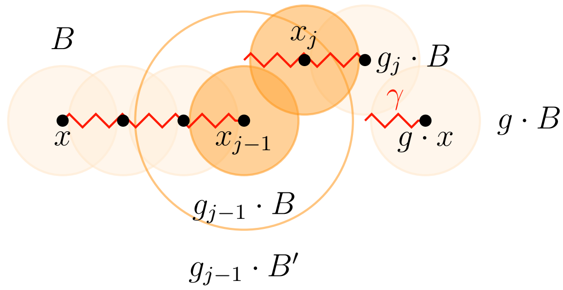

Figure 5.6: Covering a quasi-geodesic by translates of B

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

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

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

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

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

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

In particular,

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

lies in the subgroup generated by S, as desired.

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

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

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

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

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

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

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

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

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

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

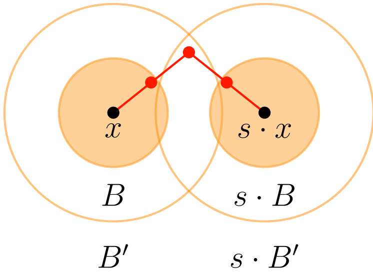

Figure 5.7: If $s \in S$, then $d(x, s \cdot x) \leq 2 \cdot (\operatorname{diam} B + 2 \cdot b)$

(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. □

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

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

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

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

is a quasi-isometry.

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

- A metric space is *proper* if all balls of finite radius are compact with respect to the topology induced by the metric.
  Hence, proper metric spaces are locally compact.
- An action $G \times X \longrightarrow X$ of a group G on a topological space X (e.g., with the topology coming from a metric on X) is *proper* if for all compact sets $B \subset X$ the set $\{g \in G \mid g \cdot B \cap B \neq \emptyset\}$ is finite.
  Example 5.3.8 (Proper actions)**.**
  - The translation action of $\mathbb{Z}$ on $\mathbb{R}$ is proper (with respect to the standard topology on $\mathbb{R}$).
  - More generally, the action by deck transformations of the fundamental group of a locally compact path-connected topological space on its universal covering is proper [**?**].
  - Clearly, all stabiliser groups of a proper action are finite. The converse is *not* necessarily true: For example, the action of $\mathbb{Z}$ on the circle S^1 given by rotation around an irrational angle is free but not proper (because $\mathbb{Z}$ is infinite and S^1 is compact).
- An action $G \times X \longrightarrow X$ of a group G on a topological space X is *cocompact* if the quotient space $G \setminus X$ with respect to the quotient topology is compact.
  Example 5.3.9 (Cocompact actions)**.**
  - The translation action of $\mathbb{Z}$ on $\mathbb{R}$ is cocompact (with respect to the standard topology on $\mathbb{R}$), because the quotient is homeomorphic to the circle S^1, which is compact.
  - More generally, the action by deck transformations of the fundamental group of a compact path-connected topological space X on its universal covering is cocompact because the quotient is homeomorphic to X [**?**].
  - 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.

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

A topological argument shows that because the action of G on X is proper and cocompact and because X is locally compact there exists a closed subspace $B \subset X$ of finite diameter such that $\pi(B) = X$ (e.g., a suitable finite union of closed balls [?]); here, $\pi\colon X \longrightarrow G \setminus X$ is the canonical projection onto the quotient space. In particular, $\bigcup_{g \in G} g \cdot B = X$ and

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

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

Hence, we can apply the Švarc-Milnor lemma. □

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

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

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

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

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

Proof. Let G be a finitely generated group, and let H be a subgroup of finite index. If S is a finite generating set of G, then the left translation action of H on (G, d_S) is an isometric action satisfying the conditions of the Švarc-Milnor lemma (Proposition 5.3.6): The space (G, d_S) is $(1,1)$-quasi-geodesic. Moreover, we let $B \subset G$ be a finite set of representatives of $H \setminus G$ (in particular, the diameter of B is finite). Then $H \cdot B = G$, and the set $B' := B_2(B)$ is finite as well, and so the set

$$\{h \in H \mid h \cdot B' \cap B' \neq \emptyset\}$$

is finite.

Therefore, H is finitely generated and the inclusion $H \hookrightarrow G$ is a quasi-isometry (with respect to any word metrics on H and G coming from finite generating sets). □

Pursuing this line of thought consequently leads to the notion of (weak) commensurability of groups:

Definition 5.3.11 ((Weak) commensurability)**.**

- Two groups G and H are *commensurable* if they contain finite index subgroups $G' \subset G$ and $H' \subset H$ with $G' \cong H'$.
- More generally, two groups G and H are *weakly commensurable* if they contain finite index subgroups $G' \subset G$ and $H' \subset H$ satisfying the following condition: There are normal subgroups N of G' and M of H' respectively such that the quotient groups G'/N and H'/M are isomorphic.

In fact, both commensurability and weak commensurability are equivalence relations on the class of groups (exercise).

Corollary 5.3.12 (Weak commensurability and quasi-isometry)**.** *Let G be a group.*

1. *Let G' be a finite index subgroup of G. Then G' is finitely generated if and only if G is finitely generated. If these groups are finitely generated, then $G \sim_{\text{QI}} G'$.*
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 $G \sim_{\text{QI}} G/N$.*

In particular, if G is finitely generated, then any group weakly commensurable to G is finitely generated and quasi-isometric to G.

Proof. Ad 1. In view of Corollary 5.3.10, it suffices to show that G is finitely generated if G' is; but clearly combining a finite generating set of G' with a finite set of representatives of the G'-cosets in G yields a finite generating set of G.

Ad 2. If G is finitely generated, then so is the quotient G/N; conversely, if G/N is finitely generated, then combining lifts with respect to the canonical projection $G \longrightarrow G/N$ of a finite generating set of G/N with the finite set N gives a finite generating set of G.

Let G and G/N be finitely generated, and let S be a finite generating set of G/N. Then the (pre-)composition of the left translation action of G/N on $(G/N, d_S)$ with the canonical projection $G \longrightarrow G/N$ gives an isometric action of G on G/N that satisfies the conditions of the Švarc-Milnor lemma (Proposition 5.3.6). Therefore, we obtain $G \sim_{\text{QI}} G/N$. □

Example 5.3.13.

- Let $n \in \mathbb{N}_{\geq 2}$. Then the free group of rank 2 contains a free group of rank n as finite index subgroup, and hence these groups are commensurable; in particular, all free groups of finite rank bigger than 1 are quasi-isometric.
- The subgroup of $\mathrm{SL}(2,\mathbb{Z})$ generated by the two matrices

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$$

  is free of rank 2; moreover, one can show that this subgroup has finite index in $\mathrm{SL}(2,\mathbb{Z})$ (exercise). Thus, $\mathrm{SL}(2,\mathbb{Z})$ is commensurable to a free group of rank 2 and so 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).
- There exist groups that are weakly commensurable but not commensurable [12].

Caveat 5.3.14. However, not all quasi-isometric groups are commensurable. 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) (Section **??**), but using the Euler characteristic one can show that these groups cannot be commensurable; moreover, because these groups are torsion-free, they also are not weakly commensurable.

As second topic, we look at applications of the Švarc-Milnor lemma in algebraic topology/Riemannian geometry via fundamental groups:

Corollary 5.3.15 (Fundamental groups and quasi-isometry)**.** *Let M be a closed (i.e., compact and without boundary) connected Riemannian manifold, and let $\widetilde{M}$ be its Riemannian universal covering manifold. Then the fundamental group $\pi_1(M)$ is finitely generated and for every $x \in \widetilde{M}$, the map*

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

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

Sketch of proof. Standard arguments from Riemannian geometry and topology show that in this case $\widetilde{M}$ is a proper geodesic metric space and that the action of $\pi_1(M)$ on $\widetilde{M}$ is isometric, proper, and cocompact (the quotient being the compact space M) [**?**].

Applying the topological version of the Švarc-Milnor lemma (Corollary 5.3.7) finishes the proof. □

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

Definition 5.3.16 (Flat manifold, hyperbolic manifold)**.**
- A Riemannian manifold is called *flat* if its Riemannian universal covering is isometric to the Euclidean space of the same dimension.
- A Riemannaian manifold is called *hyperbolic* if its Riemannian universal covering is isometric to the hyperbolic space of the same dimension.

Example 5.3.17 (Surfaces)**.** Oriented closed connected surfaces are determined up to homeomorphism/diffeomorphism by their genus (i.e., the number of "handles", see Figure 5.8) [**?**].
- 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.

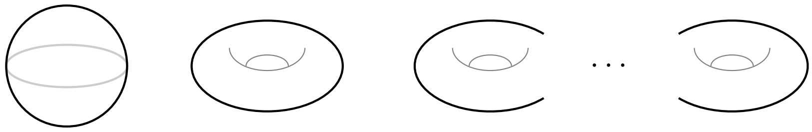

Figure 5.8: Oriented closed connected surfaces

- The oriented surface of genus 1 is the torus of dimension 2, which has fundamental group isomorphic to $\mathbb{Z}^2$. One can show that the torus admits a flat Riemannian metric [**?**].
- Oriented surfaces of genus $g \geq 2$ have fundamental group isomorphic to
$$\Big\langle a_1, \dots, a_g, b_1, \dots, b_g \;\Big|\; \prod_{j=1}^{g} [a_j, b_j] \Big\rangle$$
and one can show that these surfaces admit hyperbolic Riemannian metrics [**?**].

Being flat is the same as having vanishing sectional curvature and being hyperbolic is the same as having constant sectional curvature -1 [**?**].

Corollary 5.3.18 ((Non-)Existence of flat/hyperbolic structures)**.**
- *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 $\mathbb{Z}^n$. 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.*
- *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 $\mathbb{H}^n$. In other words: If the fundamental group of a closed connected smooth n-manifold is not quasi-isometric to $\mathbb{H}^n$, then this manifold does not admit a hyperbolic Riemannian metric.*

So, classifying finitely generated groups up to quasi-isometry and studying the quasi-geometry of finitely generated groups gives insights into geometry and topology of smooth/Riemannian manifolds.

5.4 The dynamic criterion for quasi-isometry

The Švarc-Milnor lemma translates an action of a group into a quasi-isometry of the group in question to the metric space acted upon. Similarly, we can also use certain actions to compare two groups with each other:

Definition 5.4.1 (Set-theoretic coupling)**.** Let G and H be groups. A *set-theoretic coupling for G and H* is a non-empty set X together with a left action of G on X and a right action 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 subset $F_G(h) \subset G$ with $K \cdot h \subset F_G(h) \cdot K$.

A *right action* of a group H on a set X is a map $X \times H \longrightarrow H$ such that $(x \cdot h) \cdot h' = x \cdot (h \cdot h')$ holds for all $x \in X$ and all $h, h' \in H$. In other words, a right action is the same as an antihomomorphism $H \longrightarrow S_X$.

Example 5.4.2. Let X be a group and let $G \subset X$ and $H \subset X$ be subgroups of finite index. Then the left action

$$\begin{aligned} G \times X &\longrightarrow X \\ (g, x) &\longmapsto g \cdot x \end{aligned}$$

of G on X and the right action

$$\begin{aligned} X \times H &\longrightarrow X \\ (x, h) &\longmapsto x \cdot h \end{aligned}$$

of H on X commute with each other (because multiplication in the group X is associative). The set X together with these actions is a set-theoretic coupling – a suitable subset $K \subset X$ can for example be obtained by taking the union of finite sets of representatives for G-cosets in X and H-cosets in X respectively.

Proposition 5.4.3 (Quasi-isometry and set-theoretic couplings)**.** *Let G and H be two finitely generated groups that admit a set-theoretic coupling. Then $G \sim_{\text{QI}} H$.*

Proof. Let X be a set-theoretic coupling space for G and H with the corresponding commuting actions by G and H; in the following, we will use the notation from Definition 5.4.1. We prove $G \sim_{\text{QI}} H$ by writing down a candidate for a quasi-isometry $G \longrightarrow H$ and by then verifying 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, because: Let $h \in H$. Using $G \cdot K = X$ we find a $g \in G$ with $x \cdot h \in g \cdot K$; because the actions of G and H commute with each other, it follows that $g^{-1} \cdot x \in K \cdot h^{-1}$. On the other hand, also $g^{-1} \cdot x \in K \cdot f(g)^{-1}$, by definition of f. In particular, $K \cdot h^{-1} \cap K \cdot f(g)^{-1} \neq \emptyset$, and so $h^{-1} \cdot f(g) \in F_H$. Therefore,

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

which is finite (the set F_H is finite by assumption) and independent of h; hence, f has quasi-dense image.

The map f is a quasi-isometric embedding, because: Notice that the sets

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

are finite by assumption. Let $g,\ g' \in G$.

– We first give an upper bound of $d_T(f(g), f(g'))$ in terms of $d_S(g, g')$: More precisely, we will show that

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

To this end let $n := d_S(g, g')$. As first step we show that the intersection $K \cdot f(g)^{-1} \cdot f(g') \cap K \cdot F_H(S)^n$ is non-empty: On the one hand,

$$g^{-1} \cdot x \cdot f(g') \in K \cdot f(g)^{-1} \cdot f(g')$$

by construction of f. On the other hand, because $d_S(g, g') = n$ we can write $g^{-1} \cdot g' = s_1 \cdot \dots \cdot s_n$ for certain $s_1, \dots, s_n \in S \cup S^{-1}$, and thus

$$\begin{aligned} g^{-1} \cdot x \cdot f(g') &= g^{-1} \cdot g' \cdot g'^{-1} \cdot x \cdot f(g') \\ &\in g^{-1} \cdot g' \cdot K \cdot f(g')^{-1} \cdot f(g') \\ &= g^{-1} \cdot g' \cdot K \\ &= s_1 \cdot \dots \cdot s_{n-1} \cdot s_n \cdot K \\ &\subset s_1 \cdot \dots \cdot s_{n-1} \cdot K \cdot F_H(S) \\ &\vdots \\ &\subset K \cdot F_H(S)^n. \end{aligned}$$

In particular, $K \cdot f(g)^{-1} \cdot f(g') \cap K \cdot F_H(S)^n \neq \emptyset$. Notice that in all these computations we used heavily that the actions of G and H on X commute with each other.

Using the definition of F_H, we see that

$$f(g)^{-1} \cdot f(g') \in F_H \cdot F_H(S)^n;$$

in particular, we obtain (via the triangle inequality)

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

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

$$d_S(g, g') \leq D_S F_G(T) \cdot d_T\big(f(g), f(g')\big) + D_S F_G,$$

which gives the desired lower bound. □

Remark 5.4.4 (Cocycles)**.** The construction of the map f in the proof above is an instance of a more general principle associating interesting maps with actions. Namely, suitable actions lead to so-called *cocycles* (which is an algebraic object); considering cocycles up to an appropriate equivalence relation ("being a coboundary") then gives rise to so-called *cohomology groups*. In this way, certain aspects of group actions on a space can be translated into an algebraic theory.

Moreover, we need not assume that both groups are finitely generated as being finitely generated is preserved by set-theoretic couplings:

Exercise 5.4.5 (Set-theoretic couplings and finite generation)**.** Let G and H be groups that admit a set-theoretic coupling. Show that if G is finitely generated, then so is H.

The converse of Proposition 5.4.3 also holds: whenever two finitely generated groups are quasi-isometric, then there exists a coupling (even a topological coupling) between them:

Definition 5.4.6 (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.

Theorem 5.4.7 (Dynamic criterion for quasi-isometry)**.** *Let G and H be finitely generated groups. Then the following are equivalent:*

1. *There is a topological coupling for G and H.*
2. *There is a set-theoretic coupling for G and H.*
3. *The groups G and H are quasi-isometric.*

Sketch of proof. Ad "1 $\Longrightarrow$ 2". Let G and H be finitely generated groups that admit a topological coupling, i.e., there is a non-empty locally compact space X together with a proper cocompact action from G on the left and from H on the right such that these two actions commute with each other. We show that such a topological coupling gives rise to a set-theoretic coupling:

A standard argument from topology shows that in this situation there is a compact subset $K \subset X$ such that $G \cdot K = X = K \cdot H$. Because the actions of G and H on X are proper, the sets

$$\{g \in G \mid g \cdot K \cap K \neq \emptyset\}, \qquad \{h \in H \mid K \cdot h \cap K \neq \emptyset\}$$

are finite; moreover, properness of the actions, compactness of the set K as well as the locally compactness of X also give us that for every $g \in G$ there is a finite set $F_H(g) \subset H$ satisfying $g \cdot K \subset K \cdot F_H(g)$, and similarly for elements of H. Hence, we obtain a set-theoretic coupling for G and H; in particular, G and H are quasi-isometric by Proposition 5.4.3.

Ad "2 $\Longrightarrow$ 3". This was proved in Proposition 5.4.3.

Ad "3 $\Longrightarrow$ 1". Suppose that the finitely generated groups G and H are quasi-isometric. We now sketch how this leads to a topological coupling of G and H:

Let $S \subset G$ and $T \subset H$ be finite generating sets of G and H respectively. As first step, we show that there is a finite group F and a constant $C \in \mathbb{R}_{>0}$

such that the set

$$X := \Big\{ f\colon G \longrightarrow H \times F \;\Big|\; f \text{ has } C\text{-dense image in } H \times F, \text{ and}$$
$$\forall_{g,g' \in G}\ \frac{1}{C} \cdot d_S(g,g') \leq d_{T\times F}\big(f(g), f(g')\big) \leq C \cdot d_S(g,g') \Big\}$$

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

$$\forall_{g,g' \in G}\ \frac{1}{c} \cdot d_S(g,g') - c \leq d_T\big(f(g), f(g')\big) \leq c \cdot d_S(g,g') + c.$$

In particular, if g, $g' \in G$ with $f(g) = f(g')$, then $d_S(g,g') \leq c^2$. Let F be a finite group that has more elements than the d_S-ball of radius c^2 in G (around the neutral element). Then out of f we can construct an *injective* quasi-isometry $\overline{f}\colon G \longrightarrow H \times F$. Let $\overline{c} \in \mathbb{R}_{>0}$ be chosen in such a way that $\overline{f}$ is a $(\overline{c},\overline{c})$-quasi-isometric embedding with $\overline{c}$-dense image. Because $\overline{f}$ is injective, then $\overline{f}$ satisfies a $\max(2\cdot\overline{c}, \overline{c}^2 + \overline{c})$-bilipschitz estimate (this follows as in Exercise 5.2.7 from the fact that any two different elements of a finitely generated group have distance at least 1 in any word metric). Hence, F and $C := \max(2\cdot\overline{c}, \overline{c}^2 + \overline{c})$ have the desired property that the corresponding set X is non-empty.

We now consider the following left G-action and right H-action on X:

$$\begin{aligned} G \times X &\longrightarrow X \\ (g,f) &\longmapsto \big(x \mapsto f(g^{-1}\cdot x)\big) \end{aligned}$$

$$\begin{aligned} H \times X &\longrightarrow X \\ (f,h) &\longmapsto \big(x \mapsto f(x)\cdot(h,e)\big) \end{aligned}$$

Clearly, these two actions commute with each other.

Furthermore, we can 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 [**?**], the space X is locally compact with respect to this topology. A straightforward computation (also using the Arzelá-Ascoli theorem) shows that the actions of G and H on X are proper and cocompact [**?**]. □

5.4.1 Applications of the dynamic criterion

A topological version of subgroups of finite index are uniform lattices; the dynamic criterion shows that finitely generated uniform lattices in the same ambient locally compact group are quasi-isometric (Corollary 5.4.9).

Definition 5.4.8 (Uniform lattices)**.** Let G be a locally compact topological group. A *uniform (or cocompact) lattice in* G is a discrete subgroup Γ of G such that the left translation action of Γ on G is cocompact.

Recall that a *topological group* is a group G that in addition is a topological space such that the composition $G \times G \longrightarrow G$ in the group and the map $G \longrightarrow G$ given by taking inverses are continuous (on $G \times G$ we take the product topology). A subgroup Γ of a topological group G is *discrete* if there is an open neighbourhood U of the neutral element e in G such that $U \cap \Gamma = \{e\}$.

A topological space X is called *locally compact* if for every $x \in X$ and every open neighbourhood $U \subset X$ of x there exists a compact neighbourhood $K \subset X$ of x with $K \subset U$. For example, a metric space is locally compact if and only if it is proper.

Corollary 5.4.9 (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 two finitely generated uniform lattices in G. Then the left action

$$\begin{aligned} \Gamma \times G &\longrightarrow G \\ (\gamma, g) &\longmapsto \gamma \cdot g \end{aligned}$$

of Γ on G and the right action

$$\begin{aligned} G \times \Lambda &\longrightarrow G \\ (g, \lambda) &\longmapsto g \cdot \lambda \end{aligned}$$

of Λ on G are continuous (because G is a topological group) and commute with each other, and these actions are cocompact and proper. Hence, Γ and Λ are quasi-isometric by the dynamic criterion (Theorem 5.4.7). $\square$

Example 5.4.10 (Uniform lattices)**.**

- Let $n \in \mathbb{N}$. Then $\mathbb{Z}^n$ is a discrete subgroup of the locally compact topological group $\mathbb{R}^n$, and $\mathbb{Z}^n \setminus \mathbb{R}^n$ is compact (namely, the n-torus); hence, $\mathbb{Z}^n$ is a uniform lattice in $\mathbb{R}^n$.
- The subgroup $\mathbb{Q} \subset \mathbb{R}$ is *not* discrete in $\mathbb{R}$.
- Because the quotient $\mathbb{Z} \times \{0\} \setminus \mathbb{R}^2$ is not compact, $\mathbb{Z} \times \{0\}$ is *not* a uniform lattice in $\mathbb{R}^2$. Notice that the above corollary would not hold in general without requiring that the lattices are uniform: the group $\mathbb{Z}$ is *not* quasi-isometric to $\mathbb{R}^2$.
- Let $H_\mathbb{R}$ be the *real Heisenberg group*, and let H be the *Heisenberg group*, i.e.,

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

  Then $H_\mathbb{R}$ is a locally compact topological group (with respect to the topology given by convergence of all matrix coefficients), and H is a finitely generated uniform lattice in $H_\mathbb{R}$ (exercise).
  So, any finitely generated group that is not quasi-isometric to H cannot be a uniform lattice in $H_\mathbb{R}$; for example, we will see later 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; 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 $\mathrm{Isom}(\widetilde{M})$ [**?**].

5.5 Preview: Quasi-isometry invariants and geometric properties

The central classification problem of geometric group theory is to classify all finitely generated groups up to quasi-isometry. As we have seen in the previous sections, knowing that certain groups are *not* quasi-isometric leads to interesting consequences in group theory, topology, and geometry.

5.5.1 Quasi-isometry invariants

While a complete classification of finitely generated groups up to quasi-isometry is far out of reach, partial results can be obtained. A general principle to obtain partial classification results in many mathematical fields is to construct suitable invariants:

Definition 5.5.1 (Quasi-isometry invariants)**.** Let V be a set. A *quasi-isometry invariant with values in* V is a map I from the class of all finitely generated groups to V such that

$$I(G) = I(H)$$

holds for all finitely generated groups G and H with $G \sim_{\mathrm{QI}} H$.

Proposition 5.5.2 (Using quasi-isometry invariants)**.** *Let V be a set, and let I be a quasi-isometry invariant with values in V, and let G and H be finitely generated groups with $I(G) \neq I(H)$. Then G and H are* not *quasi-isometric.*

Proof. *Assume* for a contradiction that G and H are quasi-isometric. Because I is a quasi-isometry invariant, this implies $I(G) = I(H)$, which contradicts the assumption $I(G) \neq I(H)$. Hence, G and H cannot be quasi-isometric. □

So, the more quasi-isometry invariants we can find, the more finitely generated groups we can distinguish up to quasi-isometry.

Caveat 5.5.3. If I is a quasi-isometry invariant of finitely generated groups, and G and H are finitely generated groups with $I(G) = I(H)$, then in general we *cannot* deduce that G and H are quasi-isometric, as the example of the trivial invariant shows (see below).

Some basic examples of quasi-isometry invariants are the following:

Example 5.5.4 (Quasi-isometry invariants)**.**

- *The trivial invariant.* Let V be a set containing exactly one element, and let I be the map associating with every finitely generated group this one element. Then clearly I is a quasi-isometry invariant – however, I does not contain any interesting information.
- *Finiteness.* Let $V := \{0, 1\}$, and let I be the map that sends all finite groups to 0 and all finitely generated infinite groups to 1. Then I is a quasi-isometry invariant, because a finitely generated group is quasi-isometric to a finite group if and only if it is finite (Example 5.2.9).
- *Rank of free groups.* Let $V := \mathbb{N}$, and let I be the map from the class of all finitely generated free groups to V that associates with a finitely generated free group its rank. Then I is *not* a quasi-isometry invariant on the class of all finitely generated free groups, because free groups of rank 2 and rank 3 are quasi-isometric.

In order to obtain more interesting classification results we need further quasi-isometry invariants. In the following chapters, we will, for instance, study the growth of groups, and ends and boundaries of groups (i.e., geometry at infinity).

5.5.2 Geometric properties of groups and rigidity

Moreover, it is common to use the following term:

Definition 5.5.5 (Geometric property of groups)**.** Let P be a property of finitely generated groups (i.e., any finitely generated group either has P or does not have P). We say that P is a *geometric property of groups*, in case the following holds for all finitely generated groups G and H: If G has P and H is quasi-isometric to G, then also H has P.

Example 5.5.6 (Geometric properties)**.**

- Being finite is a geometric property of groups (Example 5.2.9).
- Being Abelian is *not* a geometric property of groups: For example, the trivial group and the symmetric group S_3 are quasi-isometric (because they are both finite), but the trivial group is Abelian and S_3 is not Abelian.

Surprisingly, there are many interesting (many of them purely algebraic!) properties of groups that are geometric:

- Being virtually infinite cyclic is a geometric property.
- More generally, for every $n \in \mathbb{N}$ the property of being virtually $\mathbb{Z}^n$ is geometric.
- Being finitely generated and virtually free is a geometric property.
- Being finitely generated and virtually nilpotent is a geometric property of groups.
- Being finitely presented is a geometric property of groups.

Proving that these properties are geometric is far from easy; some of the techniques and invariants needed to prove such statements are explained in later chapters.

That a certain algebraic property of groups turns out to be geometric is an instance of the *rigidity* phenomenon; so, for example, the fact that being virtually infinite cyclic is a geometric property can also be formulated as the group $\mathbb{Z}$ being *quasi-isometrically rigid*.

Conversely, in the following chapters, we will also study geometrically defined properties of finitely generated groups such as hyperbolicity and amenability (both of which are geometric properties of groups), and we will investigate how the geometry of these groups affects their algebraic structure.

6 Growth types of groups

The first quasi-isometry invariant we discuss is the growth type. Basically, we measure the "volume"/size of balls in a given finitely generated group and study the asymptotic behaviour when the radius tends to infinity.

We will start by introducing growth functions for finitely generated groups (with respect to finite generating system); while these growth functions depend on the chose finite generating set, a straightforward calculation shows that growth functions for different finite generating sets only differ by a small amount, and more generally that growth functions of quasi-isometric groups are asymptotically equivalent. This leads to the notion of growth type of a finitely generated group.

The quasi-isometry invariance of the growth type allows us to show for many groups that they are not quasi-isometric.

Surprisingly, having polynomial growth is a rather strong constraint for finitely generated groups: By a celebrated theorem of Gromov, all finitely generated groups of polynomial growth are virtually nilpotent(!). We will briefly sketch an elementary approach to this theorem due to Kleiner, Tao, and Shalom.

6.1 Growth functions of finitely generated groups

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

Definition 6.1.1 (Growth function)**.** Let G be a finitely generated group and let $S \subset G$ be a finite generating set of G. Then

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

is the *growth function of G with respect to S*; here, $B_r^{G,S}(e)$ denotes the (closed) ball of radius r around e with respect to the word metric d_S on G.

Notice that this definition makes sense because balls for word metrics with respect to finite generating sets are finite (Remark 5.2.8).

Example 6.1.2 (Growth functions of groups)**.**

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

$$\begin{aligned} \beta_{\mathbb{Z},\{1\}} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto 2 \cdot r + 1. \end{aligned}$$

  On the other hand, the growth function of $\mathbb{Z}$ with respect to the generating set $\{2, 3\}$ is given by

$$\begin{aligned} \beta_{\mathbb{Z},\{2,3\}} \colon \mathbb{N} &\longrightarrow \mathbb{N} \\ r &\longmapsto \begin{cases} 1 & \text{if } r = 0 \\ 5 & \text{if } r = 1 \\ 6 \cdot r + 1 & \text{if } r > 1 \end{cases} \end{aligned}$$

  (exercise). So, in general, growth functions of different generating sets are different.

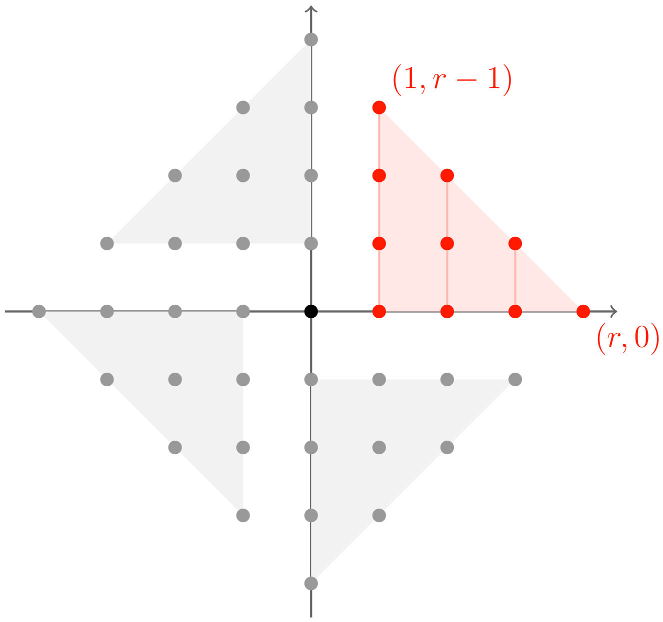

Figure 6.1: An r-ball in $\mathrm{Cay}\big(\mathbb{Z}^2, \{(1,0),(0,1)\}\big)$

– The growth function of $\mathbb{Z}^2$ with respect to the standard generating set $S := \{(1,0),(0,1)\}$ is quadratic (see Figure 6.1):

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

– More generally, if $n \in \mathbb{N}$, then the growth functions of $\mathbb{Z}^n$ grows like a polynomial of degree n.
– The growth function of the Heisenberg group

$$H = \left\{ \begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \middle|\, x,\ y,\ z \in \mathbb{Z} \right\} \cong \big\langle x, y, z \bigm| [x,y],[y,z],[x,y] = z \big\rangle$$

with respect to the generating set $\{x, y, z\}$ grows like a polynomial of degree 4 (exercise).

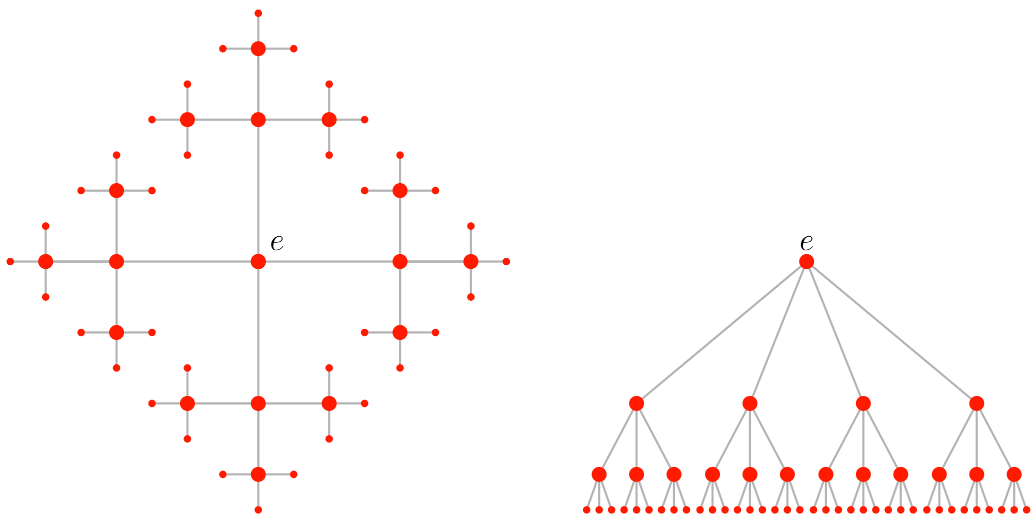

Figure 6.2: A 3-ball in $\operatorname{Cay}\bigl(F(\{a,b\}),\{a,b\}\bigr)$, drawn in two ways

– The growth function of a free group F of finite rank $n \geq 2$ with respect to a free generating set S is exponential (see Figure 6.2):

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

Proposition 6.1.3 (Basic properties of growth functions)**.** *Let G be a finitely generated group, and let $S \subset G$ be a finite generating set.*

1. Sub-multiplicativity. *For all $r,\ r' \in \mathbb{N}$ we have*
$$\beta_{G,S}(r + r') \leq \beta_{G,S}(r) \cdot \beta_{G,S}(r').$$
2. A general lower bound. *Let G be infinite. Then $\beta_{G,S}$ is strictly increasing; in particular, $\beta_{G,S}(r) \geq r$ for all $r \in \mathbb{N}$.*
3. A general upper bound. *For all $r \in \mathbb{N}$ we have*
$$\beta_{G,S}(r) \leq \beta_{F(S),S}(r) = 1 + \frac{|S|}{|S|-1} \cdot \bigl((2 \cdot |S| - 1)^r - 1\bigr).$$

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

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

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

for all $r \in \mathbb{N}$. The growth function $\beta_{F(S),S}$ is calculated in Example 6.1.2. □

6.2 Growth types of groups

As we have seen, different finite generating sets in general lead to different growth functions; however, one might suspect already that growth functions coming from different generating sets only differ by some "finite" terms. We therefore introduce the following notion of equivalence for growth functions:

6.2.1 Growth types

Definition 6.2.1 (Quasi-equivalence of (generalised) growth functions)**.**

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

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

  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-domination defines a partial order on the set of all generalised growth functions and that quasi-equivalence indeed defines an equivalence relation.

Example 6.2.2 (Generalised growth functions)**.**

– *Monomials.* If $a \in \mathbb{R}_{\geq 0}$, then

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ x &\longmapsto x^a \end{aligned}$$

is a generalised growth function.
For all a, $a' \in \mathbb{R}_{\geq 0}$ we have

$$(x \mapsto x^a) \prec (x \mapsto x^{a'}) \iff a \leq a',$$

because: If $a \leq a'$, then for all $r \in \mathbb{R}_{\geq 0}$

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

and so $(x \mapsto x^a) \prec (x \mapsto x^{a'})$.
Conversely, if $a > a'$, then for all c, $b \in \mathbb{R}_{>0}$ we have

$$\lim_{r \to \infty} \frac{r^a}{c \cdot (c \cdot r + b)^{a'}} = \infty;$$

thus, for all $c, b \in \mathbb{R}_{>0}$ there is $r \in \mathbb{R}_{\geq 0}$ such that $r^a \geq c \cdot (c \cdot r + b)^{a'} + b$, and so $(x \mapsto x^a) \not\prec (x \mapsto x^{a'})$.
In particular, $(x \mapsto x^a) \sim (a \mapsto x^{a'})$ if and only if $a = a'$.

– *Exponential functions.* If $a \in \mathbb{R}_{>1}$, then

$$\begin{aligned} \mathbb{R}_{\geq 0} &\longrightarrow \mathbb{R}_{\geq 0} \\ x &\longmapsto a^x \end{aligned}$$

is a generalised growth function. A straightforward calculation shows that

$$(x \mapsto a^x) \sim (x \mapsto a'^x)$$

holds for all $a,\ a' \in \mathbb{R}_{>1}$, as well as

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

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

Example 6.2.3 (Growth functions yield generalised growth functions)**.** Let G be a finitely generated group, and let $S \subset G$ be a finite generating set. Then the function

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

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

If G and H are finitely generated groups with finite generating sets S and T respectively, then we say that the growth function $\beta_{G,S}$ is *quasi-dominated by/quasi-equivalent to* the growth function $\beta_{H,T}$ if the associated generalised growth functions are quasi-dominated by/quasi-equivalent to each other.

Notice that then $\beta_{G,S}$ is quasi-dominated by $\beta_{H,T}$ if and only if there exist $c,\ b \in \mathbb{N}$ such that

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

6.2.2 Growth types and quasi-isometry

We will now show that growth functions of different finite generating sets are quasi-equivalent; more generally, we show that the quasi-equivalence class of growth functions of finite generating sets is a quasi-isometry invariant:

Proposition 6.2.4 (Growth functions and quasi-isometry)**.** *Let G and H be finitely generated groups, and let $S \subset G$ and $T \subset H$ be finite generating sets of G and H respectively.*

1. *If there exists a quasi-isometric embedding $(G, d_S) \longrightarrow (H, d_T)$, then*

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

2. *In particular, if G and H are quasi-isometric, then the growth functions $\beta_{G,S}$ and $\beta_{H,T}$ are quasi-equivalent.*

Proof. The second part follows directly from the first one (and the definition of quasi-isometry and quasi-equivalence of generalised growth functions).

For the first part, let $f\colon G \longrightarrow H$ be a quasi-isometric embedding; hence, there is a $c \in \mathbb{R}_{>0}$ such that

$$\forall_{g,g' \in G} \quad \frac{1}{c} \cdot d_S(g,g') - c \leq d_T\bigl(f(g), f(g')\bigr) \leq c \cdot d_S(g,g') + c.$$

We write $e' := f(e)$, and let $r \in \mathbb{N}$. Using the estimates above we obtain the following:

– If $g \in B_r^{G,S}(e)$, then $d_T(f(g), e') \leq c \cdot d_S(g,e) + c \leq c \cdot r + c$, and thus

$$f\bigl(B_r^{G,S}(e)\bigr) \subset B_{c \cdot r + c}^{H,T}(e').$$

– For all g, $g' \in G$ with $f(g) = f(g')$, we have

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

Because the metric d_T on H is invariant under left translation, it follows that

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

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

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

Definition 6.2.5 (Growth types of finitely generated groups)**.** Let G be a finitely generated group. The *growth type of G* is the (common) quasi-equivalence class of all growth functions of G with respect to finite generating sets of G.

Corollary 6.2.6 (Quasi-isometry invariance of the growth type)**.** *By Proposition 6.2.4, the growth type of finitely generated groups is a quasi-isometry invariant, i.e., quasi-isometric finitely generated groups have the same growth type.*

In other words: Finitely generated groups having different growth types cannot be quasi-isometric.

Example 6.2.7 (Growth types)**.**
- If $n \in \mathbb{N}$, then $\mathbb{Z}^n$ has the growth type of $(x \mapsto x^n)$.
- The Heisenberg group has the growth type of $(x \mapsto x^4)$.
- Non-Abelian free groups of finite rank have the growth type of the exponential function $(x \mapsto e^x)$.

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

Definition 6.2.8 (Exponential growth, polynomial growth, intermediate growth)**.** Let G be a finitely generated group.
- 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 any) finite generating set S of G there is an $a \in \mathbb{R}_{>0}$ such that $\beta_{G,S} \prec (x \mapsto x^a)$.
- The group G is of *intermediate growth*, if it is neither of exponential nor of polynomial growth.

Recall that growth functions of finitely generated groups grow at most exponentially (Proposition 6.1.3), and that polynomials and exponential functions are not quasi-equivalent; hence the term "intermediate growth" does make sense and a group cannot have exponential and polynomial growth at the same time.

Moreover, it follows from Proposition 6.2.4 and Example 6.2.2 that having exponential growth/polynomial growth/intermediate growth respectively is a geometric property of groups.

Example 6.2.9 (Distinguishing quasi-isometry types of basic groups)**.**
- We can recover the rank of free Abelian groups from their quasi-isometry type, namely: For all m, $n \in \mathbb{N}$ we have

$$\mathbb{Z}^m \sim_{\mathrm{QI}} \mathbb{Z}^n \iff m = n;$$

this follows by combining Example 6.2.2, Example 6.2.7, and the above corollary.

– Similarly, we obtain for the Heisenberg group H that $H \not\sim_{\mathrm{QI}} \mathbb{Z}^3$, which might be surprising because H fits into a short exact sequence

$$1 \longrightarrow \mathbb{Z} \longrightarrow H \longrightarrow \mathbb{Z}^2 \longrightarrow 1$$

of groups!

– Let F be a non-Abelian free group of finite rank, and let $n \in \mathbb{N}$. Because F grows exponentially, but $\mathbb{Z}^n$ has polynomial growth, we obtain

$$F \not\sim_{\mathrm{QI}} \mathbb{Z}^n.$$

Corollary 6.2.10 (Growth of subgroups)**.** *Let G be a finitely generated group, and let H be a finitely generated subgroup of G. If T is a finite generating set of H, and S is a finite generating set of G, then*

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

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

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

and so $B_r^{H,T}(e) \subset B_r^{G,S}(e)$. In particular,

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

and thus $\beta_{H,T} \prec \beta_{G,S'}$. Moreover, we know that (G, d_S) and $(G, d_{S'})$ are quasi-isometric, and hence the growth functions $\beta_{G,S'}$ and $\beta_{G,S}$ are quasi-equivalent by Proposition 6.2.4. Therefore, we obtain $\beta_{H,T} \prec \beta_{G,S}$, as desired. □

Example 6.2.11 (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.

Grigorchuk was the first to show that there indeed exist groups that have intermediate growth [**?**, 12]:

Theorem 6.2.12 (Existence of groups of intermediate growth)**.** *There exists a finitely generated group of intermediate growth.*

This group constructed by Grigorchuk can be described via automorphisms of trees or as a so-called automatic group.

Furthermore, the group constructed by Grigorchuk also has several other interesting properties [12, Chapter VIII]; for example, it is a finitely generated infinite torsion group, and it is commensurable to the direct product with itself.

Bibliography

[1] Martin Aigner. *Diskrete Mathematik*, fifth edition, Vieweg Verlag, 2004. Cited on page: 58

[2] Margaret A. Armstrong. *Groups and Symmetry*, Undergraduate Texts in Mathematics, Springer, 1988. Cited on page: 49, 58

[3] Tim Austin. Rational group ring elements with kernels having irrational dimension, preprint, available online at arXiv:0909.2360 [math.GR], 2009. Cited on page: 27

[4] Gilbert Baumslag. Wreath products and finitely presented groups, *Math. Z.*, 75, pp. 22–28, 1961. Cited on page: 25

[5] Martin R. Bridson. Non-positive curvature and complexity for finitely presented groups, volume 2 of *International Congress of Mathematicians*, pp. 961–987, EMS, 2006. Cited on page: 2, 3, 5

[6] Martin R. Bridson. Geometric and combinatorial group theory, in *The Princeton Companion to Mathematics*, pp. 431–448, Princeton University Press, 2008. Cited on page: 5

[7] Martin R. Bridson, André Haefliger. *Metric Spaces of Non-positive Curvature*, Band 319 der *Grundlehren der Mathematischen Wissenschaften*, Springer, 1999. Cited on page: 5, 41

[8] Kenneth S. Brown. *Cohomology of Groups*, volume 87 of *Graduate Texts in Mathematics*, Springer, 1982. Cited on page: 30, 42

[9] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein. *Introduction to algorithms*, third edition. MIT Press, 2009. Cited on page: 35

[10] Reinhard Diestel. *Graph theory*, third edition, volume 173 of *Graduate Texts in Mathematics*, Springer, 2005. Cited on page: 35

[11] Misha Gromov. Hyperbolic groups. in *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pp. 75–263, Springer, 1987. Cited on page: 5

[12] Pierre de la Harpe. *Topics in Geometric Group Theory*, Chicago University Press, 2000. Cited on page: 5, 26, 73, 98, 122, 123

[13] John M. Harris, Jeffrey L. Hirst, Michael J. Mossinghoff. *Combinatorics and Graph Theory*, second edition, *Undergraduate Texts in Mathematics, Springer*, 2008. Cited on page: 35

[14] Serge Lang. *Algebra*, Revised third edition. Volume 211 of *Graduate Texts in Mathematics*, Springer, 2002. Cited on page: 8

[15] John M. Lee. *Riemannian Manifolds. An Introduction to Curvature.* Volume 176 of *Graduate Texts in Mathematics*, Springer, 1997. Cited on page: 5

[16] Clara Löh. *Group Cohomology and Bounded Cohomology. An Introduction for Topologists.* Lecture notes, 2010. Available online at http://www.mathematik.uni-regensburg.de/loeh/teachin/topologie3_ws0910 Cited on page: 30, 42

[17] William S. Massey. *A Basic Course in Algebraic Topology.* Volume 127 of *Graduate Texts in Mathematics*, Springer, 1991. Cited on page: 32

[18] Parameswaran Sankaran. On homeomorphisms and quasi-isometries of the real line, *Proc. Amer. Math. Soc.*, 134, no. 7, pp. 1875–1880, 2006. Cited on page: 82

[19] Jean-Pierre Serre. *Trees*, corrected 2nd printing of the 1980 English translation, *Springer Monographs in Mathematics*, Springer, 2003. Cited on page: 5, 45, 65, 69

[20] John R. Stallings. On torsion-free groups with infinitely many ends, *Annals of Mathematics*, 88, pp. 312–334, 1968. Cited on page: 34