Matthias Blank, Werner Thumann

Geometric Group Theory II

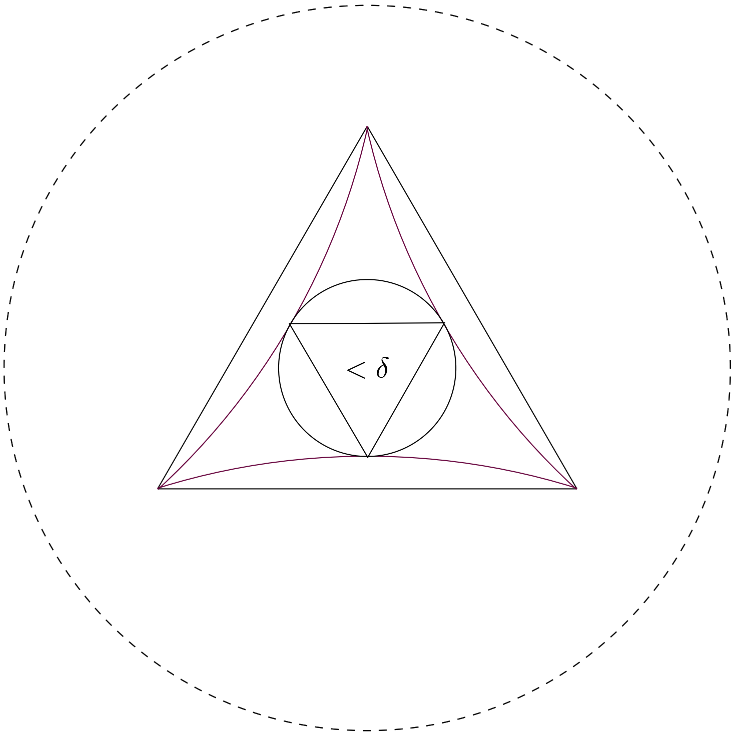

Lecture Notes
University of Regensburg
Summer Semester 2015

Contents

List of Figures

CHAPTER 1

The Geometry of Metric Spaces

Last semester, we have studied, in particular, two examples of metric spaces coming from geometric group theory, that is finitely generated groups with a word metric and Cayley graphs associated to such groups. In this chapter, we will explore the geometry of metric spaces in more detail.

1.1 Basics of Metric Geometry

1.1.1 Basic definitions

We begin by recalling some basic concepts and fixing some notations:

Definition 1.1.1. Let X be a set. A map $d\colon X \times X \longrightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is called a *pseudo-metric*, if the following holds

(i) For all $x \in X$, we have $d(x,x) = 0$.

(ii) The map d is symmetric, i.e.

$$\forall_{x,y \in X} \quad d(x,y) = d(y,x).$$

(iii) The triangle inequality holds, i.e.:

$$\forall_{x,y,z \in X} \quad d(x,z) \leq d(x,y) + d(y,z).$$

A pseudo-metric d is called a *generalised metric* if the following holds:

(i') For all $x, y \in X$, we have $d(x,y) = 0$ if and only if $x = y$.

A generalised metric d with image in $\mathbb{R}_{\geq 0}$ is called a *metric*. The pair (X,d) is then called a *metric space*.

Notation 1.1.2. *Let (X,d) be a pseudo-metric space and $x \in X$.*

(i) For any $r \in \mathbb{R}_{>0}$, we write $B(x,r) := \{y \in X \mid d(x,y) < r\}$ for the open r-ball around x in X.

(ii) Similarly, for any $r \in \mathbb{R}_{\geq 0}$, we write $\bar{B}(x,r) := \{y \in X \mid d(x,y) \leq r\}$ for the closed r-ball around x in X. Note that in general, $\bar{B}(x,r)$ can be strictly larger than the closure $\overline{B(x,r)}$.

Example 1.1.3 (One nice and one boring metric space)**.**

(i) For $n \in \mathbb{N}$, we write $\mathbb{E}^n$ to denote the standard Euclidean metric space $(\mathbb{R}^n, d_2^n)$ where

$$d_2^n(x,y) = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2}$$

for $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$ and $y = (y_1, \ldots, y_n) \in \mathbb{R}^n$.

(ii) For any set X, there is a metric on X given by setting for all $x, y \in X$

$$d(x,y) := \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y, \end{cases}$$

called the *discrete metric* on X.

General metric spaces can be rather savage and we will in general restrict our attention to well-behaved classes. For example, we will often consider metric spaces in which, like in Euclidean space, every closed bounded set is compact:

Definition 1.1.4 (Locally compact and proper)**.** Let (X,d) be a metric space.

(i) We call (X,d) *locally compact*, if for every $x \in X$, there exists an $r \in \mathbb{R}_{>0}$, such that $\bar{B}(x,r)$ is compact.

(ii) We call (X,d) *proper*, if for every $x \in X$ and every $r \in \mathbb{R}_{>0}$ the closed r-ball $\bar{B}(x,r)$ is compact.

Remark 1.1.5. Clearly, every proper metric space is locally compact and complete. The converse does not hold, however. For example, every set with the discrete metric is locally compact and complete, but an infinite set with the discrete metric is clearly not proper.

1.1.2 Lengths, Geodesics and Angles

In this section, we will discuss three elementary geometric notions, lengths of paths, geodesics and angles, in the metric setting.

We begin by recalling the definition of the length of a path:

Definition 1.1.6 (Length of paths)**.**

(i) Let X be a topological space. A continuous map $I \longrightarrow X$, where $I \subset \mathbb{R}$ is a non-empty compact interval in $\mathbb{R}$, is called a *path* in X.

(ii) For $a, b \in \mathbb{R}$ with $a \leq b$, we call a tuple $(t_0, \dots, t_{n+1})$ in $\mathbb{R}$ *a partition of* $[a, b]$ if $t_0 = a$ and $t_{n+1} = b$ and $t_i \leq t_{i+1}$ for all $i \in \{0, \dots, n\}$.

(iii) Let (X, d) be a metric space and $c\colon [a, b] \longrightarrow X$ a path in X. We call

$$l(c) := \sup\Big\{\sum_{i=0}^{n} d(c(t_{i+1}), c(t_i)) \;\Big|\; (t_0, \dots, t_{n+1}) \text{ a partition of } [a, b]\Big\},$$

the length of c.

(iv) We call a path $c\colon [a, b] \longrightarrow X$ *rectifiable* if $l(c)$ is finite.

Caveat 1.1.7. Not every path is rectifiable and there are path-connected metric spaces that do not admit rectifiable paths between every pair of points.

Proposition 1.1.8 (Elementary properties of length)**.** *Let (X, d) be a metric space.*

(i) Let $c\colon [a, b] \longrightarrow X$ be a path in X. We have $l(c) \geq d(c(a), c(b))$. Furthermore, we have $l(c) = 0$ if and only if c is constant.

(ii) Let $c_1\colon [a_1, b_1] \longrightarrow X$ and $c_2\colon [a_2, b_2] \longrightarrow X$ be two paths in X such that $c_1(b_1) = c_2(a_2)$. Let $c_1 * c_2$ denote the concatenation of c_1 and c_2, i.e. the path $c_1 * c_2\colon [0, b_2 + b_1 - a_2 - a_1] \longrightarrow X$ given by

$$t \longmapsto \begin{cases} c_1(t + a_1) & \text{if } t \in [0, b_1 - a_1] \\ c_2(t + a_1 + a_2 - b_1) & \text{if } t \in [b_1 - a_1, b_2 + b_1 - a_2 - a_1]. \end{cases}$$

Then $l(c_1 * c_2) = l(c_1) + l(c_2)$.

(iii) Let $c\colon [a, b] \longrightarrow X$ be a rectifiable path. Let $\bar{c}\colon [-b, -a] \longrightarrow X$ be the inverse path of c, i.e. the path given by

$$t \longmapsto c(-t).$$

Then $l(\bar{c}) = l(c)$.

(iv) Let $c\colon [a, b] \longrightarrow X$ be a rectifiable path. Then the arc length function of c

$$\begin{aligned} \lambda_c\colon [a, b] &\longrightarrow [0, l(c)] \\ t &\longmapsto l\big(c|_{[a,t]}\big) \end{aligned}$$

is continuous and weakly monotonic.

(v) *Let* $c\colon [a,b] \longrightarrow X$ *be a rectifiable path. Then there exists a unique rectifiable path* $\tilde{c}\colon [0, l(c)] \longrightarrow X$*, s.t.* $\tilde{c} \circ \lambda_c = c$ *and* $l(\tilde{c}|_{[0,t]}) = t$*.*

(vi) *Let* $(c_n\colon [a,b] \longrightarrow X)_{n\in\mathbb{N}}$ *be a sequence of paths converging uniformly to a path* $c\colon [a,b] \longrightarrow X$*. If* c *is rectifiable, then* $l(c) \leq \liminf_{n\to\infty} l(c_n)$*.*

Proof. Parts (i),(ii),(iii) and (v) are easy and left as an exercise.

(iv) It is clear that λ_c is weakly monotonic using property (ii). Fix $\varepsilon \in \mathbb{R}_{>0}$. Since $[a,b]$ is compact, c is uniformly continuous. Hence, there exists $\delta \in \mathbb{R}_{>0}$ such that for all $t, t' \in [a,b]$ with $d(t,t') < \delta$, we have $d(c(t), c(t')) < \varepsilon/2$. Let $(t_0, \dots, t_{n+1})$ be a partition of $[a,b]$, such that

(a) $\forall_{i\in\{0,\dots,n\}} \quad d(t_i, t_{i+1}) < \delta$.
(b) $\sum_{i=0}^n d(c(t_{i+1}), c(t_i)) > l(c) - \varepsilon/2$.

Applying part (i) and (ii), we have

$$l(c) = \sum_{i=0}^n l(c|_{[t_i,t_{i+1}]}) \geq \sum_{i=0}^n d(c(t_{i+1}), c(t_i)) > l(c) - \varepsilon/2.$$

So, for all $i \in \{0, \dots, n\}$, we have $|l(c|_{[t_i,t_{i+1}]}) - d(c(t_{i+1}), c(t_i))| < \varepsilon/2$ and therefore also $l(c|_{[t_i,t_{i+1}]}) < \varepsilon$. Thus, for all $t < t' \in [a,b]$ with $d(t,t') < \min_i d(t_{i+1}, t_i)$, we have

$$|l(c|_{[a,t]}) - l(c|_{[a,t']})| = l(c|_{[t,t']}) < 2 \cdot \varepsilon.$$

Hence λ_c is continuous.

(vi) Fix $\varepsilon \in \mathbb{R}_{>0}$ and choose again a partition $(t_0, \dots, t_{n+1})$ of $[a,b]$ such that $\sum_{i=0}^n d(c(t_{i+1}), c(t_i)) > l(c) - \varepsilon/2$. Pick $N \in \mathbb{N}$, such that for all $m \in \mathbb{N}_{>N}$ and $t \in [a,b]$ we have

$$d(c(t), c_m(t)) < \frac{\varepsilon}{4 \cdot (n+1)}.$$

Then we have for all $m \in \mathbb{N}_{>N}$

$$\begin{aligned} l(c) &< \sum_{i=0}^n d(c(t_{i+1}), c(t_i)) + \frac{\varepsilon}{2} \\ &< \sum_{i=0}^n \left(d(c_m(t_{i+1}), c_m(t_i)) + 2 \cdot \frac{\varepsilon}{4 \cdot (n+1)} \right) + \frac{\varepsilon}{2} \\ &= l(c_m) + \varepsilon. \end{aligned}$$

□

For technical reasons, it will sometimes be useful to consider only paths that are parametrised in a simple fashion:

Definition 1.1.9. Let (X, d) be a metric space. We call a rectifiable path $c\colon [a, b] \longrightarrow X$ *parametrised proportional to arc length* if the arc length function λ_c is linear.

Last semester, we have studied a particularly well-behaved type of paths, geodesics [14]. We repeat the basic definitions:

Definition 1.1.10. Let (X, d) be a metric space.

(i) A *geodesic* or *geodesic path* (of *length* $L \in \mathbb{R}_{\geq 0}$) in X is an isometric embedding $\gamma\colon [0, L] \longrightarrow X$. We call $\gamma(0)$ the start point of γ and $\gamma(L)$ the endpoint of γ. Furthermore, we call the image $\gamma([0, L])$ *a geodesic segment in X between $\gamma(0)$ and $\gamma(L)$*.

(ii) Similarly, we call an isometric embedding $\mathbb{R}_{\geq 0} \longrightarrow X$ a *geodesic ray in X* and an isometric embedding $\mathbb{R} \longrightarrow X$ a *geodesic line in X*.

(iii) We call (X, d) a *geodesic space* if for each pair $x, y \in X$ there is a geodesic joining x and y in X. For $r \in \mathbb{R}_{>0}$, we call (X, d) *an r-geodesic space* if for each $x, y \in X$ with $d(x, y) < r$ there is a geodesic segment joining x and y in X.

(iv) We call X *uniquely geodesic* if for each pair $x, y \in X$, there is a unique geodesic from x to y. In this case we also write $[x, y]$ to denote the unique geodesic segment between x and y.

(v) We call X *locally uniquely geodesic* if for all $x \in X$ exists $r > 0$ such that for each $a, b \in B(x, r)$ there is a unique geodesic from a to b in X and this geodesic lies in $B(x, r)$.

(vi) We call a subset $C \subset X$ *convex* if for each $x, y \in C$ every geodesic segment joining x and y in X is contained in C.

(vii) We call a map $\gamma\colon [0, L] \longrightarrow X$ a *local geodesic* if for each $t \in [0, L]$ there is a $\varepsilon \in \mathbb{R}_{>0}$, such that $\gamma|_{(t-\varepsilon,t+\varepsilon)\cap[0,L]}$ is geodesic.

The existence of geodesics will be crucial for us to study non-trivial geometric properties of metric spaces. We have already encountered a very important class of geodesic spaces:

Example 1.1.11. Let (V, E) be a connected graph. Then the geometric realisation of (V, E) is a geodesic space [14, A.3.3]. In particular, this is true for the realisation of Cayley graphs of finitely generated groups. In general, these spaces are not uniquely geodesic. (Which ones are?)

Another simple example of geodesic spaces are normed spaces:

Example 1.1.12. Let $(V, \|\cdot\|)$ be a normed real vector space and $d_{\|\cdot\|}$ the associated metric. Then $(V, d_{\|\cdot\|})$ is a geodesic space. More concretely, for each pair $u, v \in V$ with $u \neq v$, a geodesic from u to v of length $L := \|u - v\|$ is given by

$$\begin{aligned} [0, L] &\longrightarrow V \\ t &\longmapsto \left(1 - \frac{t}{L}\right) \cdot u + \frac{t}{L} \cdot v. \end{aligned}$$

In general, these spaces are not uniquely geodesic, but if V is a Hilbert space, then $(V, d_{\|\cdot\|})$ is indeed uniquely geodesic. For instance, for each $n \in \mathbb{N}$, the space $(\mathbb{R}^n, d_2^n)$ is uniquely geodesic, while the spaces $(\mathbb{R}^2, d_1^2)$ and $(\mathbb{R}^2, d_\infty^2)$ are not uniquely geodesic. Exercise!

Next, we will see that we can define the angle between a pair of geodesics in an arbitrary metric space, generalising the notion of angle in Euclidean geometry. We will do this by comparing certain distances with distances in Euclidean space via so called comparison triangles, a technique that will be crucial in the second chapter.

Definition 1.1.13. Let (X, d) be a metric space and (a, b, c) a triple of points in X. A *comparison triangle in* $\mathbb{E}^2$ is a geodesic triangle $\Delta(\bar{a}, \bar{b}, \bar{c})$ in $\mathbb{E}^2$ satisfying

$$\begin{aligned} d_{\mathbb{E}^2}(\bar{a}, \bar{b}) &= d(a, b) \\ d_{\mathbb{E}^2}(\bar{a}, \bar{c}) &= d(a, c) \\ d_{\mathbb{E}^2}(\bar{b}, \bar{c}) &= d(b, c). \end{aligned}$$

By the triangle inequality, for each triple (a, b, c) in X, there always exists a comparison triangle in $\mathbb{E}^2$, and such a comparison triangle is uniquely determined up to isometry by (a, b, c). Hence we write $\overline{\Delta}(a, b, c)$ to denote (any choice of such) a comparison triangle.

Definition 1.1.14 (Angles)**.** Let (X, d) be a metric space.

(i) Let (a, b, c) be a triple of points in X and $\overline{\Delta}(a, b, c)$ a comparison triangle. We write $\overline{\angle}_a(b, c)$ to denote the Euclidean angle in $\overline{\Delta}(a, b, c)$ at a.

(ii) Let $c\colon [0, l] \longrightarrow X$ and $c'\colon [0, l'] \longrightarrow X$ be two geodesics in X starting from the same point $p := c(0) = c'(0)$. We call

$$\begin{aligned} \angle(c, c') &:= \limsup_{t, t' \longrightarrow 0} \overline{\angle}_{c(0)}(c(t), c'(t')) \\ &:= \lim_{n \to \infty} \sup \left\{ \overline{\angle}_{c(0)}(c(t), c'(t')) \mid t, t' \in (0, 1/n) \right\} \end{aligned}$$

the *angle between* c *and* c'.

Caveat 1.1.15. One has to take the limit superior in the above definition since the corresponding limit does not exist in general. For nice spaces the limit *does* always exist, however.

Remark 1.1.16 (Elementary properties of angles)**.** Let (X, d) be a metric space.

(i) By the definition, the angle between two geodesics g and g' in X only depends on the *germs* of g and g', i.e. if h and h' are two geodesics in X and there is an $\varepsilon \in \mathbb{R}_{>0}$ such that $g|_{[0,\varepsilon]} = h|_{[0,\varepsilon]}$ and $g'|_{[0,\varepsilon]} = h'|_{[0,\varepsilon]}$, then $\angle(g, g') = \angle(h, h')$.

(ii) The turnaround angle is π, i.e., if g and h are two geodesics in X starting in a common point $p \in X$, such that $\bar{g} * h$ is a geodesic in X, then $\angle(g, h) = \pi$.

Caveat 1.1.17. It is possible that the angle between two distinct germs is zero. Consider for instance the space $(\mathbb{R}^2, d^2_\infty)$. For each $n \in \mathbb{N}_{>1}$, the following map is a geodesic

$$\begin{aligned} c_n \colon [0, 1/n] &\longrightarrow \mathbb{R}^2 \\ t &\longmapsto (t, (t \cdot (1-t))^n), \end{aligned}$$

and the family $(c_n)_{n \in \mathbb{N}_{>1}}$ has pairwise disjoint germs. But for any $n, m \in \mathbb{N}_{>1}$, we have $\angle(c_n, c_m) = 0$.

Proposition 1.1.18 (Triangle inequality for angles)**.** *Let (X, d) be a metric space and c, c', c'' three geodesics in X starting in a common point in X. Then*

$$\angle(c, c'') \leq \angle(c, c') + \angle(c', c'').$$

Proof. The proof is a simple comparison argument. Assume that the claim does not hold. Then there exists an $\varepsilon \in \mathbb{R}_{>0}$ such that

$$\angle(c, c'') > \angle(c, c') + \angle(c', c'') + 3 \cdot \varepsilon. \tag{1}$$

By the definition of the angle, there exists a $\delta \in (0, \min(l(c), l(c'), l(c'')))$ such that the following holds:

$$\forall_{t,t' \in (0,\delta]} \quad \overline{\angle}_{c(0)}(c(t), c'(t')) < \angle(c, c') + \varepsilon \tag{2}$$

$$\forall_{t',t'' \in (0,\delta]} \quad \overline{\angle}_{c(0)}(c'(t'), c''(t'')) < \angle(c', c'') + \varepsilon. \tag{3}$$

Furthermore, there exist $t, t'' \in (0, \delta)$ such that

$$\overline{\angle}_{c(0)}(c(t), c''(t'')) > \angle(c, c'') - \varepsilon. \tag{4}$$

Choose $\alpha \in [0, \pi]$ such that

$$\overline{\angle}_{c(0)}(c(t), c''(t'')) > \alpha > \angle(c, c'') - \varepsilon. \tag{5}$$

Let $\Delta(0, x, x'')$ be a triangle in $\mathbb{E}^2$, such that $d(0, x) = t$, $d(0, x'') = t''$ and that the angle at 0 is α. Since by (1) and (5), we have

$$\alpha > \angle(c, c') + \angle(c', c'') + 2 \cdot \varepsilon,$$

there exists an $x' \in [x, x'']$ such that

$$\angle([0, x], [0, x']) > \angle(c, c') + \varepsilon \tag{6}$$

$$\angle([0, x'], [0, x'']) > \angle(c', c'') + \varepsilon. \tag{7}$$

Since $d(0, x), d(0, x'') \in (0, \delta)$, also $t' := d(0, x') \in (0, \delta)$. Hence by (2),(3),(6) and (7), we have

$$\overline{\angle}_{c(0)}(c(t), c'(t')) < \angle([0, x], [0, x'])$$

$$\overline{\angle}_{c(0)}(c'(t'), c''(t'')) < \angle([0, x'], [0, x'']).$$

Therefore we have

$$\begin{aligned} d(c(t), c'(t')) + d(c'(t'), c''(t'')) &< d(x, x') + d(x', x'') \\ &= d(x, x'') \\ &< (d(c(t), c''(t'')), \end{aligned}$$

contradiction. □

Definition 1.1.19 (Space of directions)**.** Let (X, d) be a metric space and $p \in X$ a point. Let Σ be the set of all non-trivial geodesics in X issuing from p. Define an equivalence relation $\sim$ on Σ by setting $c \sim c'$ if and only if $\angle(c, c') = 0$ (see also Caveat 1.1.17). Write $S_p(X) := \Sigma / \sim$ for the set of equivalence classes. By Proposition 1.1.18, $\angle$ induces a metric on $S_p(X)$ and we call $S_p(X)$ equipped with this metric the *space of directions at p*.

Remark 1.1.20. If X is a Riemannian manifold, then $S_p(X)$ is isometric to the unit sphere in the tangent space $T_p(X)$.

1.1.3 Length Spaces

Remark 1.1.21 (Taking a plane vs. digging through the core)**.** Consider the two-sphere $S^2 \subset \mathbb{R}^2$, endowed with the metric induced by $d_{\mathbb{E}^2}$. It turns out that this is not a very useful metric to do geometry with. Consider for instance two antipodal points Taipeh and Asunción in S^2. Then the distance between Taipeh and Asunción is 2, while the length of the shortest rectifiable path between Taipeh and Asunción is π. In particular, S^2 together with the induced metric is not geodesic.

Definition 1.1.22 (Length spaces)**.** Let (X, d) be a metric space.

(i) We call the function

$$\begin{aligned} \bar{d}\colon X \times X &\longrightarrow \mathbb{R} \cup \{\infty\} \\ (x, y) &\longmapsto \inf\{l(c) \mid c \text{ a rectifiable path between } x \text{ and } y \text{ in } X\} \end{aligned}$$

the *length metric on X associated to d.*

(ii) We call d simply a *length metric* if $d = \bar{d}$. In this case, we call (X, d) a *length space*.

(iii) For any subset $Y \subset X$, we call $\overline{d|_{Y\times Y}}$ the *induced length metric* on Y.

Proposition 1.1.23 (Elementary properties of the length metric)**.** *Let (X, d) be a metric space.*

(i) The length metric $\bar{d}$ is indeed a generalised metric and $\bar{d} \geq d$.

(ii) If $c\colon [a, b] \longrightarrow X$ is a rectifiable path in (X, d), then it is a rectifiable path in $(X, \bar{d})$ and it has the same length in (X, d) and $(X, \bar{d})$.

(iii) We have $\bar{\bar{d}} = \bar{d}$.

Proof. Property (i) is immediate from the definition and property (iii) follows from property (ii). Regarding (ii), let $c\colon [a, b] \longrightarrow X$ be a rectifiable path in (X, d). We have to show first, that c is continuous also with respect to $\bar{d}$. This follows since

$$\bar{d}(c(t), c(t')) \leq l(c|_{[t,t']}) = |\lambda_c(t') - \lambda_c(t)|,$$

and λ_c is continuous by Proposition 1.1.8 (iv). Because $\bar{d} \geq d$, we also have $l^{\bar{d}}(c) \geq l^d(c)$. Furthermore, we have

$$\begin{aligned} l^{\bar{d}}(c) &= \sup\Big\{\sum_{i=0}^{n} \bar{d}(c(t_{i+1}), c(t_i)) \;\Big|\; (t_0, \dots, t_{n+1}) \text{ a partition of } [a, b]\Big\} \\ &\leq \sup\Big\{\sum_{i=0}^{n} l^d\left(c|_{[t_i,t_{i+1}]}\right) \;\Big|\; (t_0, \dots, t_{n+1}) \text{ a partition of } [a, b]\Big\} \\ &= l^d(c). \end{aligned}$$

□

Example 1.1.24 (Tits star)**.** For $n \in \mathbb{N}_{>1}$ define a metric on $\mathbb{R}^n$ by setting

$$\begin{aligned} d\colon \mathbb{R}^n \times \mathbb{R}^n &\longrightarrow \mathbb{R} \\ (x, y) &\longmapsto \begin{cases} |r_x - r_y| + \min\{r_x, r_y\} \cdot \|s_x - s_y\|_2^{1/2} & \text{for } \min\{r_x, r_y\} > 0 \\ |r_x - r_y| & \text{else.} \end{cases} \end{aligned}$$

Here, we use polar coordinates, i.e. for each $x \in \mathbb{R}^n_{\neq 0}$ we set $r_x = \|x\|_2$ and $s_x = x/\|x\|_2$. This metric gives the usual topology on $\mathbb{R}^n$, but the metric is badly behaved with respect to lengths: The shortest rectifiable path between any two points $x, y \in \mathbb{R}^d$ with respect to d is given by going straight from x to 0 and then from 0 to y. In particular, the associated length space $(\mathbb{R}^n, \bar{d})$ has the topology of a disjoint union of rays, one for each point in S^{n-1}, glued together in their starting points. Exercise!

Definition 1.1.25. Let (X, d) and (X', d') be two metric spaces. A family $(f_n : X \longrightarrow X')_{n \in \mathbb{N}}$ is called *equicontinuous* if for each $x \in X$ and each $\varepsilon \in \mathbb{R}_{>0}$ there is a $\delta \in \mathbb{R}_{>0}$ such that for all $n \in \mathbb{N}$

$$\forall_{x' \in B(x,\delta)} \quad d(f_n(x), f_n(x')) < \varepsilon.$$

Definition 1.1.26. A metric space (X, d) is called *separable* if X contains a countable dense subset.

Example 1.1.27.

(i) The real line $\mathbb{R}$ is separable.

(ii) Every compact metric space is separable.

Theorem 1.1.28 (Arzelà-Ascoli)**.** *Let (X, d) be a separable metric space and (X', d') a compact metric space. Let $(f_n : X \longrightarrow X')_{n \in \mathbb{N}}$ be an equicontinuous family. Then $(f_n)_{n \in \mathbb{N}}$ has a subsequence that converges uniformly on compact sets to a continuous function $f : X \longrightarrow X'$.*

Proof. See [3, Lemma I.3.10]. □

Lemma 1.1.29 (Diagonal trick)**.** Let (X, d) be a metric space and let $(y_n^m)_{n,m \in \mathbb{N}}$ be a sequence in X. Assume that for each $m \in \mathbb{N}$ the following holds: Each subsequence of $(y_n^m)_{n \in \mathbb{N}}$ contains a convergent subsequence. Then there is a strictly increasing sequence $(n_k)_{k \in \mathbb{N}}$ in $\mathbb{N}$, such that $(y_{n_k}^m)_{k \in \mathbb{N}}$ converges in X for all $m \in \mathbb{N}$.

Proof. Let $(y_{n_k^1}^1)_{k \in \mathbb{N}}$ be a convergent subsequence of $(y_n^1)_{n \in \mathbb{N}}$. Let $(y_{n_k^2}^2)_{k \in \mathbb{N}}$ be a convergent subsequence of $(y_{n_k^1}^2)_{k \in \mathbb{N}}$, and so on. Set for all $k \in \mathbb{N}$

$$n_k := n_k^k.$$

Then $(y_{n_k}^m)_{k \in \mathbb{N}}$ is a convergent sequence for all $m \in \mathbb{N}$. □

The following theorem is due to Hopf and Rinow for surfaces and Cohn and Voss for the general result. It is often also called Hopf-Rinow Theorem.

Theorem 1.1.30 (Hopf-Rinow-Cohn-Voss Theorem)**.** *Let (X, d) be a length space. Then the following are equivalent*

(i) The space (X, d) is complete and locally compact.

(ii) The space (X, d) is proper.

If one of these holds, (X, d) is a geodesic space.

Proof. We first show that (i) implies (ii), the converse is obvious. Fix a point $x \in X$ and set

$$\Sigma := \{r \in \mathbb{R}_{\geq 0} \mid \bar{B}(x, r) \text{ is compact}\}.$$

We have to show that $\Sigma = \mathbb{R}_{\geq 0}$. Clearly, Σ is a non-empty subset of the connected set $\mathbb{R}_{\geq 0}$, so it suffices to show that it is open and closed. Note also that for $r \in \Sigma$, also $[0, r] \subset \Sigma$.

- Σ is open: This is just a compactness argument: Since X is locally compact, Σ contains a neighbourhood of 0. Pick any $r \in \Sigma \cap \mathbb{R}_{>0}$. Since $\bar{B}(x, r)$ is compact and X is locally compact, there are $x_1, \ldots, x_n \in X$ and $r_1, \ldots, r_n \in \mathbb{R}_{>0}$ such that $\bar{B}(x, r) \subset \bigcup_{i=1}^n B(x_i, r_i)$ and $\bar{B}(x_i, r_i)$ is compact for all $i \in \{0, \ldots, n\}$. Since $\bar{B}(x, r)$ is compact and $X \setminus \bigcup_{i=1}^n B(x_i, r_i)$ is closed, and the two sets are disjoint, there is $\delta \in \mathbb{R}_{>0}$ such that

  $$\inf\Big\{d(a, b) \;\Big|\; a \in \bar{B}(x, r),\ b \in X \setminus \bigcup_{i=1}^n B(x_i, r_i)\Big\} > 2 \cdot \delta.$$

  Therefore $\bar{B}(x, r+\delta)$ is a closed subset of the compact set $\bigcup_{i=1}^n \bar{B}(x_i, r_i)$, and thus $r + \delta \in \Sigma$.

- Σ is closed: It suffices to show that for any $r \in \mathbb{R}_{>0}$ with $[0, r) \subset \Sigma$, also $r \in \Sigma$. Fix such an $r \in \mathbb{R}_{>0}$. Let $(x_n)_{n \in \mathbb{N}}$ be a sequence in $\bar{B}(x, r)$. We have to show that this sequence has a convergent subsequence. We can assume that $\lim_{n \to \infty} d(x_n, x) = r$, otherwise we could find a subsequence of $(x_n)_{n \in \mathbb{N}}$ contained in a compact set $\bar{B}(x, r')$ for some $r' \in [0, r)$. First, we choose a sequence $(y_n^m)_{n,m \in \mathbb{N}}$ such that for all $n, m \in \mathbb{N}$

  $$d(y_n^m, x) < r - \frac{1}{2 \cdot (m+1)} \quad \text{and} \quad d(y_n^m, x_n) < \frac{1}{m+1}.$$

  We can for instance choose for each $n, m \in \mathbb{N}$ a rectifiable path in X between x and x_n of length smaller than $r + 1/(2 + 2m)$ and pick y_n^m on this path accordingly. Since for each $m \in \mathbb{N}$, we have that $(y_n^m)_{n \in \mathbb{N}}$ is contained in the compact set $\bar{B}(x, r - 1/(2+2m))$, by Lemma 1.1.29 we can find a subsequence $(n_k)_{k \in \mathbb{N}}$ in $\mathbb{N}$ such that $(y_{n_k}^m)_{k \in \mathbb{N}}$ converges for all $m \in \mathbb{N}$. Then also $(x_{n_k})_{k \in \mathbb{N}}$ converges, since for all $k, l, m \in \mathbb{N}$

  $$d(x_{n_k}, x_{n_l}) \leq d(x_{n_k}, y_{n_k}^m) + d(y_{n_k}^m, y_{n_l}^m) + d(y_{n_l}^m, x_{n_l}) \leq \frac{2}{m+1} + d(y_{n_k}^m, y_{n_l}^m).$$

We now prove that X is geodesic if X is proper. Let $x, y \in X$ be two points. Choose a sequence of rectifiable paths $(c_n\colon [0,1] \longrightarrow X)_{n\in\mathbb{N}}$ parameterised proportional to arc length such that for each $n \in \mathbb{N}$ we have $l(c_n) \leq d(x,y) + 1/(n+1)$. The family $(c_n)_{n\in\mathbb{N}}$ is equicontinuous, since for all $t < t' \in [0,1]$ and all $n \in \mathbb{N}$

$$d(c_n(t), c_n(t')) \leq l(c_n|_{[t,t']}) \leq l(c_n|_{[t,t']}) \cdot \frac{d(x,y)+1}{l(c_n)} = |t'-t| \cdot (d(x,y)+1).$$

Since for each $n \in \mathbb{N}$, the image of c_n is contained in the compact set $\bar{B}(x, d(x,y)+1)$, by Arzelà-Ascoli, Theorem 1.1.28, there is a subsequence $(c_{n_k})_{k\in\mathbb{N}}$ converging uniformly to a path $c\colon [0,1] \longrightarrow X$ between x and y. By Proposition 1.1.8 (vi), we have

$$d(x,y) \leq l(c) \leq \liminf_{n\to\infty} l(c_n) \leq d(x,y).$$

Hence, up to reparametrisation, c is a geodesic between x and y. □

Remark 1.1.31. Given a Riemannian manifold, one can assign to each piecewise differentiable path its *Riemannian length*. Analogously to the definition of the length metric of a metric space, one can use the Riemannian length to define an actual metric on a Riemannian manifold. It is not difficult to see that this metric is indeed a length metric in the metric sense. Thus, using our version of the Hopf-Rinow theorem, we have:

Corollary 1.1.32. Every connected, complete Riemannian manifold is a geodesic metric space.

We refer to the book of John Lee [13] for more details about the Hopf-Rinow thereom in the context of Riemannian manifolds.

1.2 Model Spaces of Constant Curvature

One of our goals in this coarse is to introduce a concept of curvature for metric spaces, or more precisely of upper bounds for the curvature of a metric space. To do this, we will now give a definition of three families of spaces, Euclidean n-space, the n-Sphere and the hyperbolic n-space, that model the situation of constant curvature 0,1 and -1 respectively.

1.2.1 The Euclidean Space

In this section, we recall the definition of Euclidean space and its most elementary properties. This section will be well-known to the reader, but we repeat it here to stress the comparison with the other models.

Definition 1.2.1 (Euclidean n-space)**.** Let $n \in \mathbb{N}$ be a number.

(i) We call the map

$$\begin{aligned}\langle\cdot,\cdot\rangle\colon \mathbb{R}^n\times\mathbb{R}^n &\longrightarrow \mathbb{R}\\ (x,y) &\longmapsto \sum_{i=0}^{n-1} x_i\cdot y_i,\end{aligned}$$

the Euclidean scalar product on $\mathbb{R}^n$, where $x = (x_0,\dots,x_{n-1})$ and $y = (y_0,\dots,y_{n-1})$.

(ii) We write $d^n_{\mathbb{E}} := d^n_2$ for the metric on $\mathbb{R}^n$ associated to the Euclidean scalar product, i.e. we set

$$\begin{aligned}d^n_{\mathbb{E}}\colon \mathbb{R}^n\times\mathbb{R}^n &\longrightarrow \mathbb{R}_{\geq 0}\\ (x,y) &\longmapsto \left(\sum_{i=0}^{n-1} |x_i-y_i|^2\right)^{1/2}.\end{aligned}$$

Furthermore, we call the metric space $\mathbb{E}^n := (\mathbb{R}^n, d^n_{\mathbb{E}})$ the *Euclidean n-space*.

Definition 1.2.2 (Basic concepts in Euclidean geometry)**.** Fix $n\in\mathbb{N}_{>0}$.

(i) For each $x,y\in\mathbb{E}^n$ with $x\neq y$, the *straight line segment* between x and y is

$$[x,y] = \big\{(1-t)\cdot x + t\cdot y \;\big|\; t\in[0,1]\big\}.$$

(ii) For each tripel (x,y,z) in $\mathbb{E}^n$ with $y\neq x\neq z$, the *Euclidean angle* between $[x,y]$ and $[x,z]$ is the unique number $\angle_x(y,z)\in[0,\pi]$ such that

$$\cos\angle_x(y,z) = \frac{\langle y-x, z-x\rangle}{\|y-x\|\cdot\|z-x\|}.$$

(iii) A *hyperplane in* $\mathbb{E}^n$ is an $(n-1)$-dimensional affine subspace in $\mathbb{E}^n$.

(iv) Let $H\in\mathbb{E}^n$ be a hyperplane in $\mathbb{E}^n$. Then there is a unique isometry r_H of $\mathbb{E}^n$ fixing H and interchanging the two connected components of $\mathbb{E}^n\setminus H$, called the *reflection through* H. For any $p\in H$ and $u\in H^\perp$ a unit vector orthogonal to H, we can write r_H as

$$\begin{aligned}r_H\colon \mathbb{E}^n &\longrightarrow \mathbb{E}^n\\ x &\longmapsto x - 2\cdot\langle x-p,u\rangle\cdot u.\end{aligned}$$

(v) For any two distinct points $x,y\in\mathbb{E}^n$, the set

$${}_xH_y := \big\{z\in\mathbb{E}^n \;\big|\; d^n_{\mathbb{E}}(x,z) = d^n_{\mathbb{E}}(y,z)\big\}$$

is a hyperplane, called *the hyperplane bisector between x and y*. Clearly, $r_{{}_xH_y}(x) = y$. Conversely, if $H\subset\mathbb{E}^n$ is a hyperplane, and $x\in\mathbb{E}^n\setminus H$, then H is the hyperplane bisector of x and $r_H(x)$.

Proposition 1.2.3 (Elementary properties of Euclidean spaces)**.** *Fix* $n \in \mathbb{N}_{>0}$.

(i) The Euclidean Law of Cosines *holds: Let* (x, y, z) *be a tripel in* $\mathbb{E}^n$ *with* $y \neq x \neq z$ *and set* $c := d^n_{\mathbb{E}}(y, z)$, $b := d^n_{\mathbb{E}}(x, z)$ *and* $a := d^n_{\mathbb{E}}(x, y)$. *Let* $\gamma := \angle_x(y, z)$ *be the Euclidean angle between* $[x, y]$ *and* $[x, z]$. *Then*

$$c^2 = a^2 + b^2 - 2 \cdot a \cdot b \cdot \cos\gamma.$$

(ii) The Euclidean n*-space* $\mathbb{E}^n$ *is a uniquely geodesic metric space.*

(iii) The geodesic segments in $\mathbb{E}^n$ *are exactly the straight line segments.*

(iv) The Euclidean angle coincides with the angle defined in Definition 1.1.14.

(v) Every (closed or open) ball in $\mathbb{E}^n$ *is convex.*

Proof. Check if you still remember how to prove these elementary results! □

1.2.2 The Sphere

Definition 1.2.4 (The n-sphere)**.** Fix $n \in \mathbb{N}$. Recall that the (standard) n-sphere is the set

$$S^n := \{x \in \mathbb{R}^{n+1} \mid \langle x, x\rangle = 1\}.$$

We define a metric $d^n_{\mathbb{S}}$ on S^n by setting

$$\begin{aligned} d^n_{\mathbb{S}} \colon S^n \times S^n &\longrightarrow \mathbb{R}_{\geq 0} \\ (x, y) &\longmapsto \angle_0(x, y), \end{aligned}$$

and write $\mathbb{S}^n := (S^n, d^n_{\mathbb{S}})$ for the sphere endowed with this metric.

Remark 1.2.5. That this is indeed a metric follows for instance from Proposition 1.1.18. The metric $d^n_{\mathbb{S}}$ coincides with the length metric induced by the Euclidean metric $d^{n+1}_{\mathbb{E}}$. Exercise!

Definition 1.2.6 (Basic concepts in spherical geometry)**.** Fix $n \in \mathbb{N}_{>0}$.

(i) For $x, u \in \mathbb{S}^n$ with $\langle x, u\rangle = 0$ and $a \in (0, \pi]$, we call the set

$$\{\cos t \cdot x + \sin t \cdot u \mid t \in [0, a]\} \subset \mathbb{S}^n$$

the *minimal great arc from* x *to* $(\cos a \cdot x + \sin a \cdot u)$ with initial vector u. For $x, y \in \mathbb{S}^n$ with $d^n_{\mathbb{S}}(x, y) \in (0, \pi)$, there is a unique minimal great arc from x to y and the initial vector of this arc is given by

$$\frac{y - \langle x, y\rangle \cdot x}{\|y - \langle x, y\rangle \cdot x\|}.$$

If $x, y \in \mathbb{S}^n$ are antipodal, i.e. if $d_{\mathbb{S}^n}(x, y) = \pi$, then there are many minimal great arcs from x to y, one for any given initial vector.

(ii) Let c, c' be two minimal great arcs in $\mathbb{S}^n$ issuing from a common point in $\mathbb{S}^n$ with initial vectors u and v respectively. The *spherical angle between c and c'* is defined as $\angle_0(u, v)$.

(iii) A *hyperplane in* $\mathbb{S}^n$ is the intersection of an n-dimensional subspace of $\mathbb{R}^{n+1}$ with $\mathbb{S}^n$. Any hyperplane in $\mathbb{S}^n$ is isometrically isomorphic to $\mathbb{S}^{n-1}$

(iv) Let $H \in \mathbb{S}^n$ be a hyperplane in $\mathbb{S}^n$. The restriction r_H to $\mathbb{S}^n$ of the Euclidean reflection through the Euclidean hyperplane spanned by H is called the *reflection through* H. It is the unique isometry of $\mathbb{S}^n$ fixing H and interchanging the two connected components of $\mathbb{S}^n \setminus H$.

(v) For any two distinct points $x, y \in \mathbb{S}^n$, the set

$$ {}_xH_y := \left\{ z \in \mathbb{S}^n \;\middle|\; d^n_{\mathbb{S}}(x,z) = d^n_{\mathbb{S}}(y,z) \right\} $$

is a hyperplane, called *the hyperplane bisector between x and y*. Clearly, $r_{{}_xH_y}(x) = y$. Conversely, if $H \subset \mathbb{S}^n$ is a hyperplane, and $x \in \mathbb{S}^n \setminus H$, then H is the hyperplane bisector of x and $r_H(x)$.

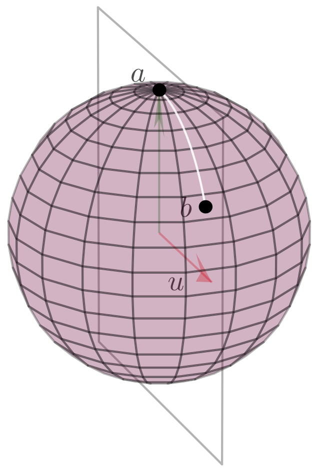

Figure 1.1: Minimal great arc between two points a and b in S^2

Proposition 1.2.7 (Elementary properties of spheres)**.** *Fix $n \in \mathbb{N}_{>0}$.*

(i) The Spherical Law of Cosines *holds: Let (x, y, z) be a triple in $\mathbb{S}^n$ with $y \neq x \neq z$ and set $c := d^n_{\mathbb{S}}(y,z)$, $b := d^n_{\mathbb{S}}(x,z)$ and $a := d^n_{\mathbb{S}}(x,y)$. Let c_y and c_z be two minimal great arcs from x to y and z respectively and let γ be the spherical angle between c_y and c_z. Then*

$$ \cos c = \cos a \cdot \cos b + \sin a \cdot \sin b \cdot \cos \gamma. $$

(ii) The n-sphere $\mathbb{S}^n$ is a geodesic metric space.

(iii) The geodesic segments in $\mathbb{S}^n$ are exactly the minimal great arcs.

(iv) The spherical angle coincides with the angle defined in Definition 1.1.14.

(v) For $x, y \in \mathbb{S}^n$ with $d^n_{\mathbb{S}}(x, y) < \pi$, there is a unique geodesic between x and y.

(vi) Every open ball in $\mathbb{S}^n$ of radius smaller or equal $\pi/2$ is convex.

Proof. This is similar to the proof of Proposition 1.2.11. Exercise! □

1.2.3 The Hyperbolic Space

As the standard sphere can be seen as the set of unit vectors with respect to the Euclidean scalar product, the hyperbolic space admits a model that can be described analogously via an appropriate bilinear form:

Definition 1.2.8 (Minkowski bilinear form)**.** For $n \in \mathbb{N}$, we call the symmetric bilinear form

$$\begin{aligned} \langle \cdot, \cdot \rangle_M \colon \mathbb{R}^{n+1} \times \mathbb{R}^{n+1} &\longrightarrow \mathbb{R} \\ (x, y) &\longmapsto -x_n \cdot y_n + \sum_{i=0}^{n-1} x_i \cdot y_i, \end{aligned}$$

the *Minkowski bilinear form on* $\mathbb{R}^{n+1}$, where $x = (x_0, \dots, x_n)$ and $y = (y_0, \dots, y_n)$.

The set of points $x \in \mathbb{R}^{n+1}$ satisfying $\langle x, x \rangle_M = -1$ is a *two-sheeted hyperboloid* (Figure 1.2). As a model for the hyperbolic n-space, we consider just the upper sheet of this hyperboloid:

Definition 1.2.9. For $n \in \mathbb{N}$, we set

$$H^n := \{x \in \mathbb{R}^{n+1} \mid \langle x, x \rangle_M = -1,\ x_n > 0\}.$$

For all $x, y \in H^n$, let $d^n_{\mathbb{H}}(x, y)$ be the unique number in $\mathbb{R}_{\geq 0}$ such that

$$\cosh d^n_{\mathbb{H}}(x, y) = -\langle x, y \rangle_M.$$

We call the pair $\mathbb{H}^n := (H^n, d^n_{\mathbb{H}})$ the *hyperbolic n-space*.

First, we define some fundamental notions hyperbolic geometry, before showing in Proposition 1.2.11, that $d^n_{\mathbb{H}}$ is indeed a metric.

Definition 1.2.10 (Basic concepts in hyperbolic geometry)**.** Fix $n \in \mathbb{N}_{>0}$.

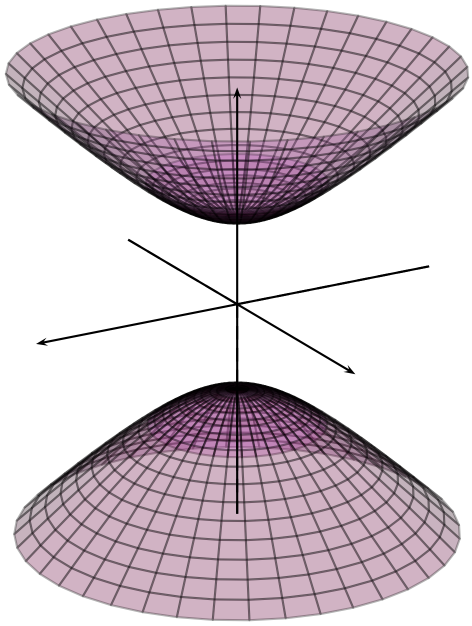

Figure 1.2: The two-sheeted hyperboloid H^2

(i) For $x, y \in \mathbb{H}^n$ with $x \neq y$, let $u \in \mathbb{R}^{n+1}$ be the unit vector (i.e. $\langle u, u\rangle_M = 1$) in the direction of $(y + \langle x, y\rangle_M \cdot x)$. We call the set

$$[x, y] := \big\{ \cosh t \cdot x + \sinh t \cdot u \mid t \in [0, d^n_{\mathbb{H}}(x, y)]\big\}$$

the *hyperbolic segment* from x to y and u the *initial vector* of $[x, y]$.

(ii) Let c, c' be two hyperbolic segments in $\mathbb{H}^n$ issuing from a common point in $\mathbb{H}^n$ with initial vectors u and v respectively. The *hyperbolic angle between c and c'* is defined as the unique number $\alpha \in [0, \pi]$, such that

$$\cos\alpha = \langle u, v\rangle_M.$$

(iii) A *hyperplane in* $\mathbb{H}^n$ is a non-empty intersection of an n-dimensional subspace of $\mathbb{R}^{n+1}$ with $\mathbb{H}^n$. Any hyperplane in $\mathbb{H}^n$ is isometrically isomorphic to $\mathbb{H}^{n-1}$.

(iv) Let $H \subset \mathbb{H}^n$ be a hyperplane in $\mathbb{H}^n$ and E be the subspace of $\mathbb{R}^{n+1}$ generated by H. Then there is a unique isometry r_H of $\mathbb{H}^n$ fixing H and interchanging the two connected components of $\mathbb{H}^n \setminus H$, called the *reflection through* H. For any unit vector $u \in \mathbb{R}^{n+1}$ (i.e. $\langle u, u\rangle_M = 1$)

orthogonal to E (i.e. $\langle u, v\rangle_M = 0$ for all $v \in E$), we can write r_H as

$$\begin{aligned} r_H \colon \mathbb{H}^n &\longrightarrow \mathbb{H}^n \\ x &\longmapsto x - 2 \cdot \langle x, u\rangle_M \cdot u. \end{aligned}$$

(v) For any two distinct points $x, y \in \mathbb{H}^n$, the set

$$_xH_y := \{z \in \mathbb{H}^n \mid d^n_{\mathbb{H}}(x, z) = d^n_{\mathbb{H}}(y, z)\}$$

is a hyperplane, called *the hyperplane bisector between x and y*. Clearly, $r_{_xH_y}(x) = y$. Conversely, if $H \subset \mathbb{H}^n$ is a hyperplane, and $x \in \mathbb{H}^n \setminus H$, then H is the hyperplane bisector of x and $r_H(x)$.

Proposition 1.2.11 (Elementary properties of hyperbolic n-spaces)**.** *Fix $n \in \mathbb{N}_{>0}$.*

(i) The Hyperbolic Law of Cosines *holds: Let (x, y, z) be a triple in $\mathbb{H}^n$ with $y \neq x \neq z$ and set $c := d^n_{\mathbb{H}}(y, z)$, $b := d^n_{\mathbb{H}}(x, z)$ and $a := d^n_{\mathbb{H}}(x, y)$. Let γ be the hyperbolic angle between $[x, y]$ and $[x, z]$. Then*

$$\cosh c = \cosh a \cdot \cosh b - \sinh a \cdot \sinh b \cdot \cos\gamma.$$

(ii) The hyperbolic n-space $\mathbb{H}^n$ is a uniquely geodesic metric space.

(iii) The geodesic segments in $\mathbb{H}^n$ are exactly the hyperbolic segments.

(iv) The hyperbolic angle coincides with the angle defined in Definition 1.1.14.

(v) If the intersection of $\mathbb{H}^n$ with a 2-dimensional subspace in $\mathbb{R}^{n+1}$ is non-empty, then it is a geodesic line and all geodesic lines in $\mathbb{H}^n$ are of this form.

(vi) Every (open or closed) ball in $\mathbb{H}^n$ is convex.

Proof.

(i) This is a short calculation: Let u be the initial vector of $[x, y]$ und v be the initial vector of $[x, z]$. Then

$$\begin{aligned} \cosh c &= -\langle y, z\rangle_M \\ &= -\langle \cosh a \cdot x + \sinh a \cdot u, \cosh b \cdot x + \sinh b \cdot v\rangle_M \\ &= -\langle \cosh a \cdot x, \cosh b \cdot x\rangle_M - \langle \sinh a \cdot u, \cosh b \cdot x\rangle_M \\ &\quad - \langle \cosh a \cdot x, \sinh b \cdot v\rangle_M - \langle \sinh a \cdot u, \sinh b \cdot v\rangle_M \\ &= \cosh a \cdot \cosh b - \sinh a \cdot \sinh b \cdot \langle u, v\rangle_M \\ &= \cosh a \cdot \cosh b - \sinh a \cdot \sinh b \cdot \cos\gamma. \end{aligned}$$

(ii) Consider $x, y, z \in \mathbb{H}^n$ with $y \neq x \neq z$ and let $a, b, c \in \mathbb{R}$ be as in (i). Define a function

$$\begin{aligned} \varphi\colon [0,\pi] &\longrightarrow \mathbb{R} \\ \gamma &\longmapsto \cosh a \cdot \cosh b - \sinh a \cdot \sinh b \cdot \cos\gamma. \end{aligned}$$

A short calculation, using the additivity theorem for hyperbolic trigonometric functions, shows that φ is a strictly increasing function with image $[\cosh(b-a), \cosh(b+a)]$. In particular, $\cosh c \leq \cosh(b+a)$ and thus $c \leq b + a$. Therefore, $\mathbb{H}^n$ is a metric space. Since φ is strictly monotonic, equality holds if and only if the hyperbolic angle between $[x, y]$ and $[x, z]$ is π, thus if and only if $x \in [y, z]$. Therefore, $\mathbb{H}^n$ is uniquely geodesic.

(iii) Follows from (ii) after verifying that the hyperbolic segments are indeed geodesic segments.

(iv) This is a somewhat lengthy calculation. We refer to [3, Proposition I.2.9].

(v) Follows directly by comparing geodesic segments with the intersections of $\mathbb{H}^n$ with planes in $\mathbb{R}^{n+1}$.

(vi) Consider a triple $x, y, z \in \mathbb{H}^n$ and $r \in \mathbb{R}_{>0}$ and assume that

$$d^n_{\mathbb{H}}(x, y), d^n_{\mathbb{H}}(x, z) < r.$$

By (v), any point $w \in [y, z]$ can be written as a sum $\lambda \cdot y + \mu \cdot z$ with $\lambda, \mu \in \mathbb{R}_{\geq 0}$ and $\lambda + \mu \leq 1$ (using that hyperbolae are convex functions). Thus, we have

$$\begin{aligned} \cosh d^n_{\mathbb{H}}(x, w) &= -\langle x, \lambda \cdot y + \mu \cdot z\rangle_M \\ &= -\lambda \cdot \langle x, y\rangle_M - \mu \cdot \langle x, z\rangle_M \\ &< \cosh r. \end{aligned}$$ □

1.2.4 Model Spaces of Constant Curvature

Now we have available our standard models $\mathbb{S}^n$, $\mathbb{E}^n$ and $\mathbb{H}^n$ of constant curvature 1, 0 and -1 respectively (in fact, this statement is tautological from our point of view, since we will define curvature in terms of these spaces in the next chapter). By scaling the metrics on $\mathbb{S}^n$ and $\mathbb{H}^n$ by a constant, we can define model spaces of arbitrary constant curvature:

Definition 1.2.12. For $n \in \mathbb{N}$ and $\kappa \in \mathbb{R}$, we define a metric space $\mathbb{M}^n_\kappa$, called the *model n-space of constant curvature κ*, by setting

$$\mathbb{M}^n_\kappa := \begin{cases} (S^n, \frac{1}{\sqrt{\kappa}} \cdot d^n_{\mathbb{S}}) & \text{if } \kappa > 0 \\ \mathbb{E}^n & \text{if } \kappa = 0 \\ (H^n, \frac{1}{\sqrt{-\kappa}} \cdot d^n_{\mathbb{H}}) & \text{if } \kappa < 0. \end{cases}$$

We sometimes write D_κ for the diameter of $\mathbb{M}^n_k$, so $D_\kappa = \pi/\sqrt{\kappa}$ for $\kappa \in \mathbb{R}_{>0}$ and $D_\kappa = \infty$ for $\kappa \in \mathbb{R}_{\leq 0}$.

Notions and properties similar to the ones discussed in the last sections can easily be deduced for the general model spaces of constant curvature.

Remark 1.2.13. In Riemannian geometry, there are several related notions of curvature, the most elementary of which is *sectional curvature*. One can show that, up to isometry, for any $n \in \mathbb{N}_{\geq 2}$ (the cases $n = 0, 1$ are not very interesting) and $\kappa \in \mathbb{R}$ there is a unique complete, simply connected, Riemannian n-manifold of constant sectional curvature κ. The spaces we have constructed above are exactly the underlying metric spaces of (a model of) these manifolds.

Definition 1.2.14. Let (X, d) be a metric space, (a, b, c) a triple of points in X and $\kappa \in \mathbb{R}$. A *comparison triangle for* (a, b, c) *in* $\mathbb{M}^2_\kappa$ is a geodesic triangle $\Delta(\bar{a}, \bar{b}, \bar{c})$ in $\mathbb{M}^2_\kappa$ satisfying

$$\begin{aligned} d_{\mathbb{M}^2_\kappa}(\bar{a}, \bar{b}) &= d(a, b) \\ d_{\mathbb{M}^2_\kappa}(\bar{a}, \bar{c}) &= d(a, c) \\ d_{\mathbb{M}^2_\kappa}(\bar{b}, \bar{c}) &= d(b, c). \end{aligned}$$

Similar to the Euclidean case, one can compare geodesic triangles in a metric space with triangles in one of the model spaces. This will be very important in the next chapter, where we will define curvature for metric spaces by comparing the shape of triangles with the triangles in our model spaces.

Lemma 1.2.15. Let (X, d) be a metric space, (a, b, c) a triple of points in X and $\kappa \in \mathbb{R}$, such that $d(a, b) + d(a, c) + d(b, c) < 2 \cdot D_\kappa$. Then there exists a comparison triangle for (a, b, c) in $\mathbb{M}^2_\kappa$ and this comparison triangle is unique up to isometry in $\mathbb{M}^2_\kappa$.

Proof. Exercise! □

One simple feature of model spaces that we will need later is that points that behave approximately like midpoints are close to midpoints:

Lemma 1.2.16 (Approximate midpoints)**.** For every $\kappa \in \mathbb{R}$, $l \in [0, D_\kappa)$ and $\varepsilon \in \mathbb{R}_{>0}$, there exists a $\delta \in \mathbb{R}_{>0}$, such that: Let $x, y \in \mathbb{M}^2_\kappa$ be points in $\mathbb{M}^2_\kappa$ with $d_{\mathbb{M}^2_\kappa}(x, y) < l$ and $m \in [x, y]$ the midpoint of $[x, y]$. Then for all $m' \in \mathbb{M}^2_\kappa$ with

$$\max\{d_{\mathbb{M}^2_\kappa}(x, m'), d_{\mathbb{M}^2_\kappa}(y, m')\} \leq \frac{1}{2} \cdot d(x, y) + \delta,$$

we have $d_{\mathbb{M}^2_\kappa}(m, m') < \varepsilon$.

Proof. Exercise using Lemma 1.2.17 below. □

1.2.5 Isometry Groups of Model Spaces

The next lemma is fundamental for the study of the geometry of our model spaces. It tells us that the action of the isometry group is highly transitive and hence the model spaces are highly homogeneous. Even more, we can achieve transitivity by using only reflections, a result that will imply that the isometry groups of the model spaces are generated by reflections.

Lemma 1.2.17 (Transitivity and reflections)**.** Let $k \in \mathbb{N}_{>0}$ be a number and $A_1, \ldots, A_k$ and $B_1, \ldots, B_k$ be points in $\mathbb{M}^n_\kappa$, such that for all $i, j \in \{1, \ldots, k\}$, we have $d(A_i, A_j) = d(B_i, B_j)$. Then there is a product of at most k reflections $\varphi \in \operatorname{Isom}(\mathbb{M}^n_\kappa)$, such that $\varphi(A_i) = B_i$ for all $i \in \{1, \ldots, k\}$.

Proof. We show the result by induction on $k \in \mathbb{N}_{>0}$. If $k = 1$, then $r_{{}_{A_1}H_{B_1}}(A_1) = B_1$. For $k \in \mathbb{N}_{>1}$ arbitrary, by induction there is a product of $k-1$ reflections $\psi \in \operatorname{Isom}(\mathbb{M}^n_\kappa)$, such that $\psi(A_i) = B_i$ for all $i \in \{1, \ldots, k-1\}$. For any $i \in \{1, \ldots, k-1\}$, we have

$$d(\psi(A_k), B_i) = d(\psi(A_k), \psi(A_i)) = d(A_k, A_i) = d(B_k, B_i).$$

Hence, for all $i \in \{1, \ldots, k-1\}$, the point B_i is contained in the hyperplane bisector of $\psi(A_k)$ and B_k. Therefore, $r_{{}_{\psi(A_k)}H_{B_k}} \circ \psi$ maps $\{A_1, \ldots, A_k\}$ to $\{B_1, \ldots, B_k\}$ as desired. ☐

Proposition 1.2.18 (Isometries and fixed points)**.** *Fix $n \in \mathbb{N}$ and $\kappa \in \mathbb{R}$. Let $\varphi \in \operatorname{Isom}(\mathbb{M}^n_\kappa)$ be an isometry of $\mathbb{M}^n_\kappa$.*

(i) Either the fixed point set $\operatorname{Fix}(\varphi)$ is contained in a hyperplane of $\mathbb{M}^n_\kappa$, or φ is the identity.

(ii) If $\operatorname{Fix}(\varphi)$ is a hyperplane, then φ is the reflection through this hyperplane.

(iii) The map φ can be written as a product of $(n+1)$ or less reflections in $\mathbb{M}^n_\kappa$.

Proof.

(i) If φ is not the identity, then there exists an $a \in \mathbb{M}^n_\kappa$ such that $\varphi(a) \neq a$. For any $b \in \operatorname{Fix}(\varphi)$, we have

$$d(\varphi(a), b) = d(\varphi(a), \varphi(b)) = d(a, b).$$

Hence $\operatorname{Fix}(b)$ is contained in the hyperplane bisector of $\varphi(a)$ and a.

(ii) As we have seen, for any $a \in \mathbb{M}^n_\kappa \setminus \operatorname{Fix}(\varphi)$, we have that $\operatorname{Fix}(\varphi)$ is contained in the hyperplane bisector of a and $\varphi(a)$, and since $\operatorname{Fix}(\varphi)$ is a hyperplane, it is equal to this bisector. In particular, we have $r_{\operatorname{Fix}(\varphi)}(a) = \varphi(a)$.

(iii) Take $(n+1)$ points $a_0, \ldots, a_n \in \mathbb{M}^n_\kappa$ that are not contained in a hyperplane of $\mathbb{M}^n_\kappa$. Let $r \in \mathrm{Isom}(\mathbb{M}^n_\kappa)$ be a product of at most $(n+1)$ reflections, mapping $\{a_0, \ldots, a_n\}$ to $\{\varphi(a_0), \ldots, \varphi(a_n)\}$ as in Lemma 1.2.17. Then $\{a_0, \ldots, a_n\} \subset \mathrm{Fix}(r^{-1} \circ \varphi)$ and therefore by (i), we have $r = \varphi$.

□

Definition 1.2.19. Let $n \in \mathbb{N}$ be a number.

(i) We write $O(n)$ for the group of orthogonal n-matrices, i.e.,

$$O(n) = \{A \in \mathrm{GL}(n, \mathbb{R}) \mid A^t \cdot A = E_n\}.$$

By definition, $O(n)$ is the group of matrices preserving the Euclidean scalar product on $\mathbb{R}^n$.

(ii) We write $O(n, 1)$ for the group

$$O(n, 1) = \{A \in \mathrm{GL}(n+1, \mathbb{R}) \mid A^t \cdot J_n \cdot A = J_n\};$$

where J_n is the diagonal matrix with entries $(1, \ldots, 1, -1)$. By definition, $O(n, 1)$ is the group of matrices preserving the Minkowski bilinear form on $\mathbb{R}^{n+1}$.

(iii) We denote by $O(n, 1)_+$ the subgroup of maps preserving the upper half-sheet of the hyperboloid H^n, i.e., we set

$$O(n, 1)_+ := \{A \in O(n, 1) \mid a_{n,n} > 0\}.$$

Knowing that any isometry in a model space is a product of reflections, we can now easily calculate the isometry groups of these spaces. Since scaling a metric by a constant clearly does not change the isometry group, it suffices to consider $\mathbb{H}^n, \mathbb{E}^n$ and $\mathbb{S}^n$:

Theorem 1.2.20 (Isometry groups of model spaces)**.** *For $n \in \mathbb{N}_{>1}$, we have*

$$\begin{aligned} \mathrm{Isom}(\mathbb{H}^n) &\cong O(n, 1)_+ \\ \mathrm{Isom}(\mathbb{E}^n) &\cong \mathbb{R}^n \rtimes O(n) \\ \mathrm{Isom}(\mathbb{S}^n) &\cong O(n+1). \end{aligned}$$

In any of these cases, the stabilizer of any point (i.e., the subgroup of isometries fixing this point) is isomorphic to $O(n)$.

Proof. In the light of Proposition 1.2.18, it suffices to show that each of these groups corresponds to a subgroup of the isometry group of the model space that contain all reflections:

(i) By definition, the group $O(n,1)_+$ preserves both the upper sheet of the hyperboloid and the Minkowski bilinear form. Hence, restricting to H^n defines a map $O(n,1)_+ \longrightarrow \mathrm{Isom}(\mathbb{H}^n)$. This map is an injection since H^n spans $\mathbb{R}^{n+1}$ (and hence a map in $\mathrm{GL}(n+1,\mathbb{R})$ is uniquely defined by its restriction to H^n). From our concrete description of hyperbolic reflections, it is easy to see that $O(n,1)_+$ contains all reflections.

(ii) The group $\mathbb{R}^n \rtimes O(n)$ is isomorphic to the subgroup of $\mathrm{Isom}(\mathbb{R}^n)$ generated by translations and linear isometries. Since any Euclidean reflection is the conjugate of a linear reflection by a translation, $\mathbb{R}^n \rtimes O(n)$ contains all reflections.

(iii) This is completely analogous to the hyperbolic case.

That the point stabilizers equal $O(n)$ is clear in the Euclidean and spherical case and follows by an easy calculation in the hyperbolic one. □

1.2.6 More Models for the Hyperbolic Space

If you look at say a Renaissance portrait or a photography, you will see a 2-dimensional effigy of objects in 3-dimensional space. Mathematically, this way of representation can be expressed by the following simple map:

Definition 1.2.21 (Perspectivic projection)**.** For $n \in \mathbb{N}$ and $r, s \in \mathbb{R}$ with $r < s$, we call the continuous map

$$\begin{aligned} \mathrm{Pers}^n_{r,s} \colon \mathbb{R}^n \times \mathbb{R}_{\geq s} &\longrightarrow \mathbb{R}^n \\ (x,z) &\longmapsto \frac{s-r}{z-r} \cdot x \end{aligned}$$

the *perspectivic projection with drawing parameters r and s.*

The picture one gets drawing perspectivicly depends essentially on the point of view of the artist and the place of the canvas. We will now consider an interesting exemplification of this fact given by two rather different "perspectivic drawings" of the hyperboloid:

Proposition 1.2.22. *Fix $n \in \mathbb{N}$.*

(i) The perspectivic projection $\mathrm{Pers}^n_{0,1}$ restricts to a homeomorphism

$$p^n_K \colon H^n \longrightarrow B^n.$$

Its inverse map is given by

$$\begin{aligned} B^n &\longrightarrow H^n \\ x &\longmapsto \frac{1}{\sqrt{1-\|x\|_2^2}} \cdot (x,1). \end{aligned}$$

Here, B^n denotes the open n-ball $B(0,1)$ in $\mathbb{R}^n$.

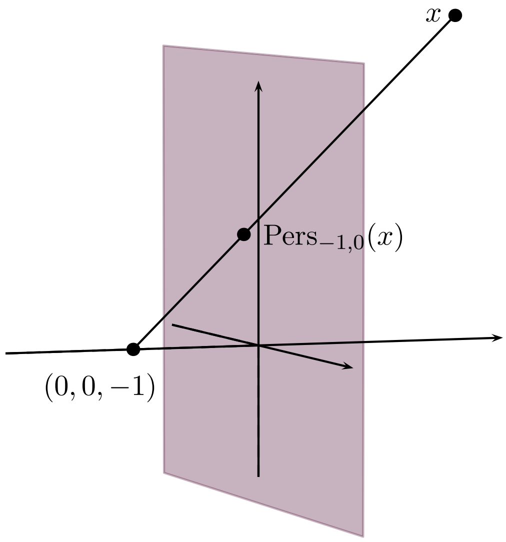

Figure 1.3: The perspectivic projection

(ii) *The perspectivic projection* $\mathrm{Pers}^n_{-1,0}$ *restricts to a homeomorphism*

$$p^n_P \colon H^n \longrightarrow B^n.$$

Its inverse map is given by

$$\begin{aligned} B^n &\longrightarrow H^n \\ x &\longmapsto \frac{1}{1-\|x\|_2^2} \cdot (2 \cdot x, 1 + \|x\|_2^2). \end{aligned}$$

Proof. All the maps are clearly continuous and a short calculation shows that each pair is mutually inverse. □

Definition 1.2.23. Let (X, d) be a metric space, Y a set and $\varphi \colon X \longrightarrow Y$ be a bijection. We call the metric $\varphi_* d$ on Y defined by the relation

$$\forall_{x,y \in X} \quad d(x, y) = \varphi_* d(\varphi(x), \varphi(y))$$

the *pushforward of d by φ*. Clearly, φ defines an isometry between (X, d) and $(Y, \varphi_* d)$.

Definition 1.2.24. Fix $n \in \mathbb{N}$.

(i) We call B^n together with the pushforward metric $d^n_{\mathrm{KB}} := (p^n_K)_* d^n_{\mathbb{H}}$ the *Klein model for* $\mathbb{H}^n$.

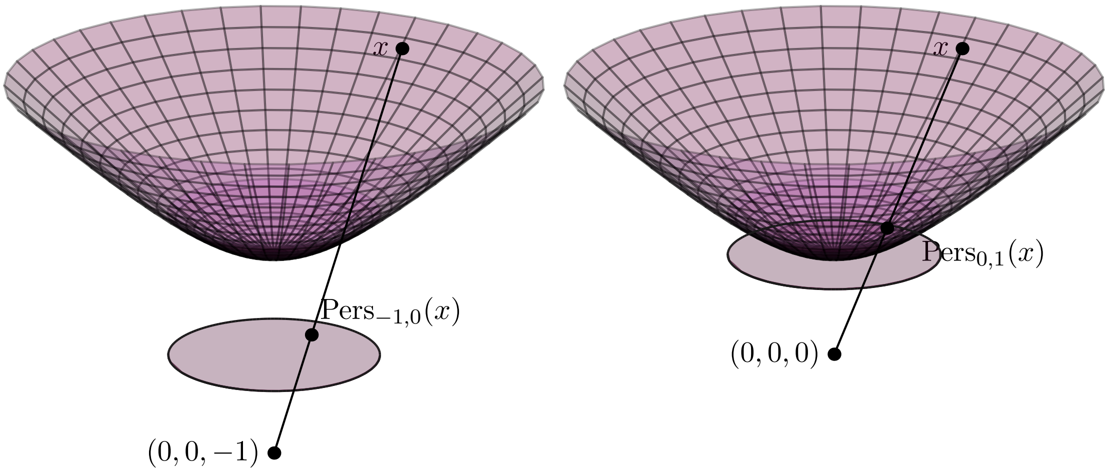

Figure 1.4: The perspectivic projection for the Poincaré ball and the Klein model

(ii) We call B^n together with the pushforward metric $d^n_{\mathrm{PB}} := (p^n_P)_* d^n_{\mathbb{H}}$ the *Poincaré ball model for $\mathbb{H}^n$*.

By a straightforward calculation, we can express these two metrics a bit more concretely:

Proposition 1.2.25. *Fix $n \in \mathbb{N}$. Let $x, y \in B^n$ be two points.*

(i) The distance between x and y in the Klein metric is given by the equation

$$\cosh d^n_{\mathrm{KB}}(x, y) = \frac{1 - \langle x, y\rangle}{\sqrt{1 - \|x\|_2^2} \cdot \sqrt{1 - \|y\|_2^2}}.$$

(ii) The distance between x and y in the Poincaré metric is given by the equation

$$\cosh d^n_{\mathrm{PB}}(x, y) = 1 + \frac{2 \cdot \|x - y\|_2^2}{(1 - \|x\|_2^2) \cdot (1 - \|y\|_2^2)}.$$

Definition 1.2.26. Fix $n \in \mathbb{N}$.

(i) We write $\widehat{\mathbb{R}^n} = \mathbb{R}^n \cup \{\infty\}$ for the one point compactification of $\mathbb{E}^n$ (the topology is such that the stereographic projection $\mathbb{S}^n \longrightarrow \widehat{\mathbb{R}^n}$ is a homeomorphism).

(ii) For any $k \in \{0, \ldots, n-1\}$, a *(generalised) k-sphere in $\widehat{\mathbb{R}^n}$* is either a k-sphere in $\mathbb{E}^n$ or the union of an affine k-subspace of $\mathbb{E}^n$ with $\{\infty\}$.

(iii) Let $S \subset \mathbb{E}^n$ be an $(n-1)$-sphere with center $a \in \mathbb{E}^n$ and radius $r \in \mathbb{R}_{>0}$. The *inversion with respect to S* is the involution

$$I\colon \widehat{\mathbb{R}^n} \longrightarrow \widehat{\mathbb{R}^n}$$
$$x \longmapsto \begin{cases} a & \text{if } x = \infty \\ \infty & \text{if } x = a \\ \frac{r^2}{\|x-a\|_2^2} \cdot (x-a) + a. & \text{else.} \end{cases}$$

(iv) Let $H \in \mathbb{E}^n$ be a hyperplane. The *inversion with respect to* $S := H \cup \{\infty\}$ is the involution

$$I_S\colon \widehat{\mathbb{R}^n} \longrightarrow \widehat{\mathbb{R}^n}$$
$$x \longmapsto \begin{cases} \infty & \text{if } x = \infty \\ r_H(x) & \text{else.} \end{cases}$$

Proposition 1.2.27. *Let $S \subset \widehat{\mathbb{R}^n}$ be a generalised sphere.*

(i) The inversion I_S maps generalised spheres to generalised spheres.

(ii) The inversion I_S preserves Euclidean angles between generalised spheres.

Proof. Short calculation. Exercise! □

For $n \in \mathbb{N}_{>0}$, set $\mathrm{HS}^n := \mathbb{R}^{n-1} \times \mathbb{R}_{>0}$. The homeomorphism C^n defined in the following proposition is also called *Cayley transformation*:

Proposition 1.2.28. *Fix $n \in \mathbb{N}$ and let $S \subset \mathbb{R}^{n+1}$ be the sphere of radius $\sqrt{2}$ and centre $(0, \dots, 0, -1)$. Then the inversion I_S maps $\partial B^{n+1} = S^n$ to $\partial\,\mathrm{HS}^n = (\mathbb{R}^n \times \{0\}) \cup \{\infty\}$ and restricts to a homeomorphism*

$$C^n\colon B^n \longrightarrow \mathrm{HS}^n .$$

Proof. The inversion I_S maps the sphere S^n to a generalised n-sphere S'. Since, $(0, \dots, 0, -1) \in S^n$ is mapped to ∞ by I_S, we have that S' is the union of ∞ with a hyperplane in $\mathbb{R}^{n+1}$. Since $S^n \cap \mathrm{Fix}(I_S) = S^n \cap S$ spans the hyperplane $\mathbb{R}^n \times \{0\}$, we have $I_S(S^n) = (\mathbb{R}^n \times \{0\}) \cup \{\infty\}$. Both $\widehat{\mathbb{R}^{n+1}} \setminus S^n$ and $\widehat{\mathbb{R}^{n+1}} \setminus \mathbb{R}^n \times \{0\} \cup \{\infty\}$ have exactly two connected components, one of which is HS^n and B^n respectively, hence I_S restricts to the homeomorphism C^n. □

Definition 1.2.29. We call HS^n together with the push-forward metric $d^n_{\mathrm{HS}} := C^n_* d^n_{\mathrm{PB}}$ the *Poincaré half space model for the hyperbolic n-space.*

We can use our description of the geodesic lines in the hyperboloid model and the concrete isometries between the different models for the hyperbolic space to calculate the shape of the geodesic lines in all the models:

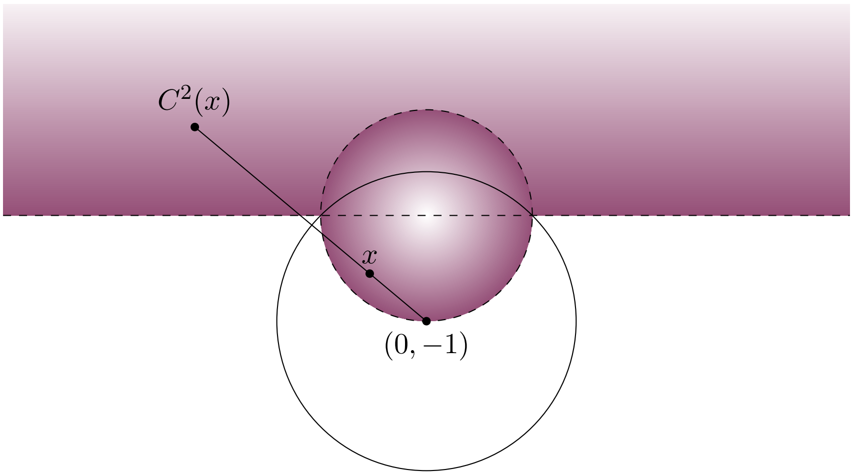

Figure 1.5: The isometric identification via the Cayley transformation C^2.

Theorem 1.2.30 (Geodesics in different models)**.** *For all $n \in \mathbb{N}_{>0}$, we have the following:*

(i) The geodesic lines in the Klein model are exactly the intersection of B^n with straight lines.

(ii) The geodesic lines in the Poincaré ball model are exactly the intersections of B^n with generalised circles that are orthogonal to the boundary ∂B^n.

(iii) The geodesic lines in the half space model are exactly the intersections of HS^n with generalised circles that are orthogonal to the boundary $\partial \, \mathrm{HS}^n$.

Proof.

(i) Since p_K^n is the restriction of the perspectivic projection $\mathrm{Pers}^n_{0,1}$, it maps the intersection of a 2-plane P in $\mathbb{R}^{n+1}$ with B_n to the intersection of P with H^n.

(ii) We show that the composition $l := p_K \circ p_P^{-1} \colon B^n \longrightarrow B^n$ maps the intersections of B^n with generalised circles intersecting ∂B^n orthogonally to the intersections of affine lines with B^n. We consider B^n embedded as $B^n \times \{0\}$ into $\mathbb{R}^{n+1}$. Note that with this identification, l is given by

$$\begin{aligned} l \colon B^n \times \{0\} &\longrightarrow B^n \times \{0\} \\ (x, 0) &\longmapsto \left(\frac{2 \cdot x}{1 + \|x\|^2}, 0 \right). \end{aligned}$$

Hence, l is given by the restriction to $B^n \times \{0\}$ of the composition $p \circ I_S$, where I_S is the inversion through the sphere of radius $\sqrt{2}$ and centre $(0, \ldots, 0, -1)$ and p is the orthogonal projection $\mathbb{R}^{n+1} \longrightarrow \mathbb{R}^n \times \{0\}$. Since I_S is an inversion and $\partial B^n \subset S$, it maps generalised circles lying in $B^n \times \{0\}$ orthogonal to ∂B^n to generalised circles lying in S^n orthogonal to ∂B^n and the orthogonal projection p maps such circles to affine lines in B^n.

(iii) Since C^n is the restriction of an inversion, it maps generalised circles intersecting ∂B^n orthogonally to generalised circles intersecting $\partial\, \mathrm{HS}^n$ orthogonally. □

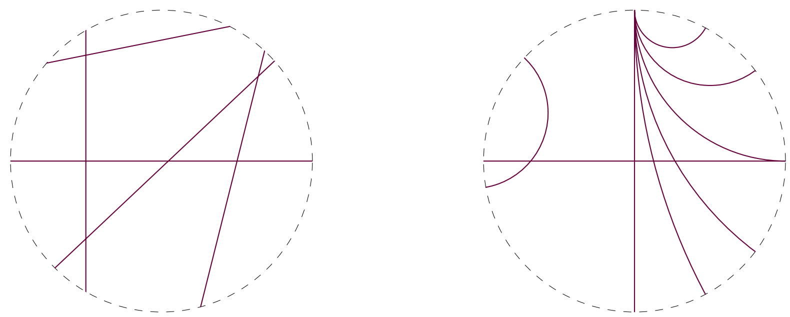

Figure 1.6: The Klein model (left-hand side) and the Poincaré ball model (right-hand side) and some of their geodesics.

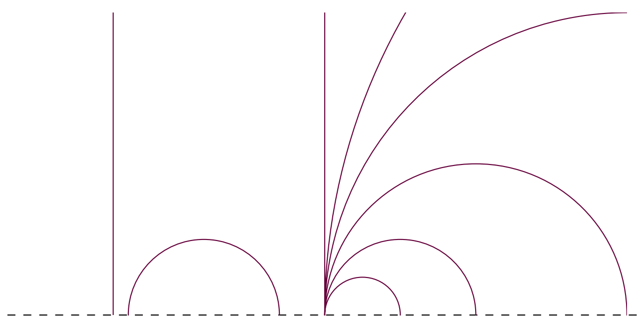

Figure 1.7: The half-plane model and some of its geodesics.

Remark 1.2.31. One advantage of the Poincaré models is that they are conformal in the sense that the depicted Euclidean angles correspond to the hyperbolic angles in the model.

We finish this section by mentioning some of the (partially psychological) advantages and disadvantages of the four models of the hyperbolic n-space we have discussed so far:

Remark 1.2.32. We took the hyperboloid model as the starting point of our observations since the metric can be defined rather elegantly and in a way that supports the idea that the hyperbolic space is a sort of "anti-sphere". Furthermore, the hyperbolic reflections can be handled via linear algebra quite similarly to the well-known Euclidean case. A disadvantage of this model is that it does not give a good representation of the boundary (which we will discuss later) and is not as visually accessible as the other models.

The Klein model easily depicts the boundary and the geodesics are straight line segments, but the angles between geodesics do not correspond to the Euclidean angles in the representation.

More often, one will use the Poincaré ball model, where the geodesics are still very easy (generalised circles) and the depicted angles in the Euclidean representation correspond to the actual angle between the geodesics.

The half space model is a rather minor modification of the Poincaré ball model sharing many of its advantages while sometimes allowing a simpler description of isometries.

1.3 Spaces of curvature bounded from above

1.3.1 Basic definitions and examples

In this section, we will introduce the basic objects we want to work with in the remainder of this chapter, namely spaces with curvature bounded from above by some constant $\kappa \in \mathbb{R}$. To define what this means, we will compare geodesic triangles in such a space with triangles in the model space $\mathbb{M}^2_\kappa$ and demand that the former are not bigger than the latter in a certain sense.

To make this more precise, we recall the notion of comparison triangles (Definition 1.2.14 and Lemma 1.2.15): Let x, y, z be three points in a metric space X which are joined by geodesic segments in X. Using the law of cosines for the model spaces, we always find points $\bar{x}, \bar{y}, \bar{z} \in \mathbb{M}^2_\kappa$ such that

$$d(x,y) = d(\bar{x},\bar{y}) \quad \text{and} \quad d(x,z) = d(\bar{x},\bar{z}) \quad \text{and} \quad d(y,z) = d(\bar{y},\bar{z})$$

which are unique up to isometry and these points are joined by unique geodesics in the model spaces. In the case $\kappa > 0$ we have to demand the necessary and sufficient condition

$$\operatorname{perim}(x,y,z) = d(x,y) + d(x,z) + d(y,z) < 2D_\kappa$$

to achieve this. Recall that we have defined $D_\kappa = \infty$ if $\kappa \leq 0$, so we could have succinctly demanded this for any $\kappa \in \mathbb{R}$. We call such points $\bar{x}, \bar{y}, \bar{z}$ comparison points and the resulting geodesic triangle

$$\bar{\Delta}(x,y,z) = \Delta(\bar{x},\bar{y},\bar{z})$$

a comparison triangle. If a is a point on the chosen geodesic segment between the points x, y then the unique point $\bar{a}$ on the geodesic segment between $\bar{x}, \bar{y}$ such that $d(x, a) = d(\bar{x}, \bar{a})$ is called the comparison point of a.

Definition 1.3.1. Let $\kappa \in \mathbb{R}$ and let X be a metric space. We say that X is a *CAT(κ) space* if it is D_κ-geodesic and if all geodesic triangles in X satisfy the *CAT(κ) inequality*, i.e., if for all points a, b on such a triangle and corresponding comparison points $\bar{a}, \bar{b}$ on a comparison triangle in $\mathbb{M}^2_\kappa$ we have

$$d(a, b) \leq d(\bar{a}, \bar{b}).$$

We say that X has *curvature* $\leq \kappa$ if it is locally CAT(κ), i.e. if for each $x \in X$ there is a ball around x on which the induced metric from X is CAT(κ).

Remark 1.3.2. We will mainly be interested in the case $\kappa \leq 0$. For simplicity, we will state most of the results in this section for this case only. They also hold in the case $\kappa > 0$ with minor modifications, namely, one has to keep in mind that the diameter of $\mathbb{M}^2_\kappa$ for $\kappa > 0$ is bounded by $D_\kappa < \infty$ and geodesics are only unique when their endpoints x, y satisfy $d(x, y) < D_\kappa$. See e.g. [3] for the correct statements in the case $\kappa > 0$.

We want to give some equivalent characterizations of the CAT(κ) curvature notion. Before we state them, we have to make an observation which we will need also for later considerations:

Observation 1.3.3 (Alexandrov's Lemma)**.** In $\mathbb{M}^2_\kappa$ for $\kappa \leq 0$, consider two geodesic triangles which share a common side as in the following picture:

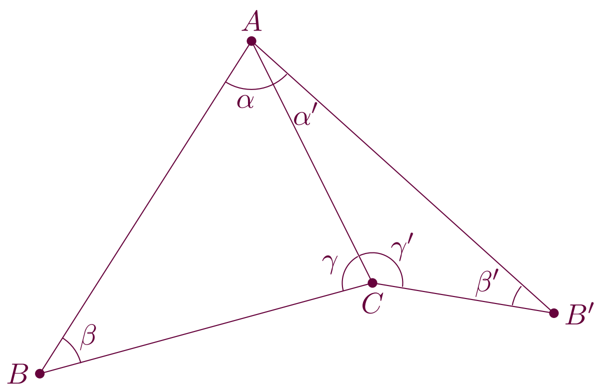

The geodesic line spanned by the two vertices A, C separates $\mathbb{M}^2_\kappa$ into two connected components. Assume that B and B' lie in different connected components. Assume further that $\gamma + \gamma' \geq \pi$. Let B'' be the unique point on the geodesic ray spanned by the vertices B, C such that $d(C, B'') = d(C, B')$ and such that B'' lies on the same side relative to the geodesic line spanned by the points A, C as B' does:

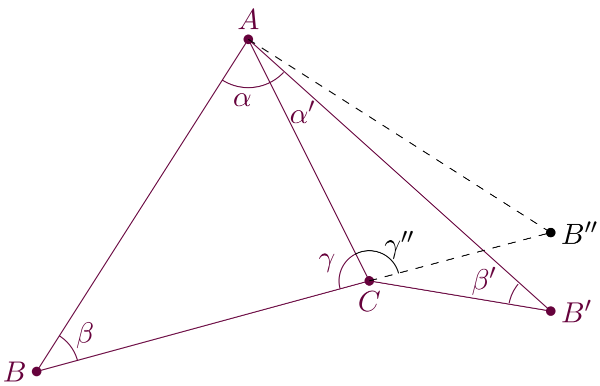

Since $\gamma+\gamma' \geq \pi$ we must have $\gamma'' \leq \gamma'$. By the law of cosines, it follows that

$$d(B,A)+d(A,B') \geq d(B,A)+d(A,B'') \geq d(B,B'') = d(B,C)+d(C,B'')$$

and thus:

(i) $d(B,A)+d(A,B') \geq d(B,C)+d(C,B')$.

Now assume we want to move the vertex C a little bit, so that the angle $\gamma+\gamma'$ becomes equal to π, without changing the lengths of the sides of the quadrilateral A,B,C,B':

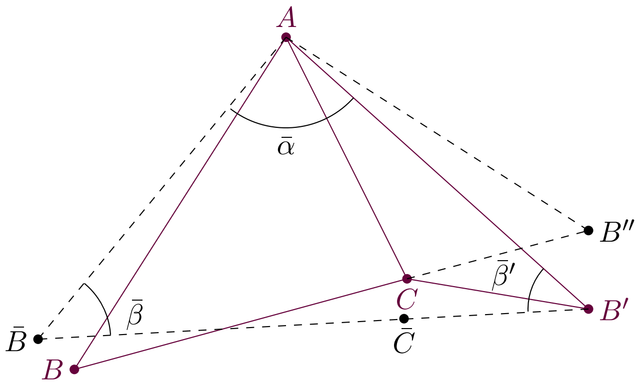

More generally, we could observe any geodesic triangle with side lengths $d(A,B)$, $d(A,B')$ and $d(B,C)+d(C,B')$. We have already seen above that $d(A,B'') \leq d(A,B')$. So by the law of cosines we get $\bar{\beta} \geq \beta$. Analogously, weg obtain $\bar{\beta}' \geq \beta'$. Furthermore, by the triangle inequality, $d(\bar{B},B') \geq d(B,B')$. So by the law of cosines again, we get $\bar{\alpha} \geq \alpha+\alpha'$. Last but not least, we obtain $d(A,\bar{C}) \geq d(A,C)$ from $\bar{\beta} \geq \beta$ and the law of cosines once more. We summarize

(ii) $\bar{\alpha} \geq \alpha+\alpha' \quad \bar{\beta} \geq \beta \quad \bar{\beta}' \geq \beta' \quad d(A,\bar{C}) \geq d(A,C)$

It is clear that if $\gamma + \gamma' = \pi$, then we have equalities everywhere. Conversely, if we have equality somewhere, then $\gamma + \gamma' = \pi$ (and thus equality everywhere).

Proposition 1.3.4. *Let $\kappa \in \mathbb{R}_{\leq 0}$ and X be a geodesic space. Then the following are equivalent:*

(i) X is CAT(κ).

For each geodesic triangle Δ in X with vertices p, q, r and comparison triangle $\bar{\Delta}$ in $\mathbb{M}^2_\kappa$ with vertices $\bar{p}, \bar{q}, \bar{r}$ we have:

(ii) If $x \neq p \neq y$ are points on the sides $[p, q]$ and $[p, r]$ of Δ, then the angle at $\bar{p}$ in a comparison triangle $\bar{\Delta}(p, x, y)$ is smaller or equal to the angle at $\bar{p}$ in $\bar{\Delta}$.

(iii) All the angles in Δ are smaller or equal to the corresponding angles in the comparison triangle $\bar{\Delta}$.

(iv) If α is the angle at p in Δ and if $\tilde{p}, \tilde{q}, \tilde{r}$ are points in $\mathbb{M}^2_\kappa$ such that $d(p, q) = d(\tilde{p}, \tilde{q})$, $d(p, r) = d(\tilde{p}, \tilde{r})$ and the angle at $\tilde{p}$ between $\tilde{q}$ and $\tilde{r}$ is α, then $d(\tilde{q}, \tilde{r}) \leq d(q, r)$.

(v) If x is a point on $[q, r]$, then $d(p, x) \leq d(\bar{p}, \bar{x})$.

(vi) If x is the midpoint of $[q, r]$, then $d(p, x) \leq d(\bar{p}, \bar{x})$.

Proof. Item (v) obviously implies (vi). The converse implication is an exercise. So (v) and (vi) are equivalent. It is also not hard to see, using the law of cosines in $\mathbb{M}^2_\kappa$, that items (i) and (ii) as well as items (iii) and (iv) are equivalent. Below we will show the implications (ii)⇒(iii)⇒(v)⇒(ii) which is then a full cycle among the statements.

(ii)⇒(iii): We have noted in Proposition 1.2.11(iv) that the hyperbolic angle equals the general angle as defined in Definition 1.1.14. It follows from this observation that in the lim sup of Definition 1.1.14(ii), we can use the hyperbolic angle in comparison triangles in $\mathbb{M}^2_\kappa$ instead of the Euclidean angles in flat comparison triangles. Item (ii) says that all these hyperbolic comparison triangles appearing in the limiting process computing the angle at p in Δ are smaller or equal to the angle at $\bar{p}$ in $\bar{\Delta}$. Consequently, also the angle at p in Δ is smaller or equal to the angle at $\bar{p}$ in $\bar{\Delta}$.

(iii)⇒(v): Let $x \in [q, r]$ and choose a geodesic from x to p. In $\mathbb{M}^2_\kappa$ choose points $\tilde{p}, \tilde{q}, \tilde{r}, \tilde{x}$ such that the triangles with vertices $\tilde{p}, \tilde{q}, \tilde{x}$ and $\tilde{p}, \tilde{r}, \tilde{x}$ are comparison triangles for the triangles in X with corresponding vertices. Arrange them such that $\tilde{q}$ and $\tilde{r}$ lie on different sides of the line spanned by $\tilde{p}$ and $\tilde{x}$. Let γ, γ' be the angles at x between q and p and between r and p respectively (measured using the chosen geodesic between x and p). Denote by $\tilde{\gamma}, \tilde{\gamma}'$ the corresponding angles in the comparison triangles. Then we have

$\gamma + \gamma' \geq \pi$ and, by (iii), $\tilde{\gamma} \geq \gamma$ and $\tilde{\gamma'} \geq \gamma'$. Consequently, the hypothesis of Alexandrov's Lemma above is satisfied and we obtain

$$d(p, x) = d(\tilde{p}, \tilde{x}) \leq d(\bar{p}, \bar{x})$$

since the geodesic triangle with vertices $\bar{p}, \bar{r}, \bar{q}$ and comparison point $\bar{x}$ is the straightened version of the quadrilateral with vertices $\tilde{p}, \tilde{q}, \tilde{r}, \tilde{x}$ as described in Alexandrov's Lemma.

(v)⇒(ii): Let $x \neq p \neq y$ as in (ii). In $\mathbb{M}^2_\kappa$ we choose two comparison triangles $\bar{\Delta}' = \Delta(\bar{p}', \bar{x}', \bar{y}')$ and $\bar{\Delta}'' = \Delta(\bar{p}'', \bar{x}'', \bar{r}'')$ of the points p, x, y and p, x, r respectively. Denote by $\bar{\alpha}, \bar{\alpha}', \bar{\alpha}''$ the angles at $\bar{p}, \bar{p}', \bar{p}''$ in the triangles $\bar{\Delta}, \bar{\Delta}', \bar{\Delta}''$. By (v) we have $d(\bar{x}', \bar{y}') = d(x, y) \leq d(\bar{x}'', \bar{y}'')$ and so $\bar{\alpha}' \leq \bar{\alpha}''$ by the law of cosines. Similarly, we have $d(\bar{x}'', \bar{r}'') = d(x, r) \leq d(\bar{x}, \bar{r})$ and thus $\bar{\alpha}'' \leq \bar{\alpha}$. This proves (ii). □

Some authors prefer the following concise definition of CAT(0) spaces:

Corollary 1.3.5 (CN inequality of Bruhat and Tits)**.** Let X be a geodesic space. Then X is CAT(0) if and only if for all points $q, r \in X$ there is $m \in X$ such that for all $p \in X$, we have the inequality

$$d(p, m)^2 \leq \frac{1}{2}\Big(d(p, q)^2 + d(p, r)^2\Big) - \frac{1}{4} d(q, r)^2$$

Note that any $m \in X$ as above necessarily is a midpoint of q, r, i.e.

$$d(q, m) = d(m, r) = d(q, r)/2$$

Furthermore, if m, m' are two midpoints of q, r, then the inequality applied to $p = m'$ implies $d(m, m') = 0$ and so $m = m'$.

Example 1.3.6.

(i) *(Hilbert spaces)* A normed real vector space is a pre-Hilbert space (i.e. the norm comes from a scalar product or, equivalently, the parallelogram law holds) if and only if it is CAT(0).

(ii) *(Subsets in Euclidean space)* A subset of $\mathbb{E}^n$ (together with the induced metric) is CAT(0) if and only if it is convex.

Observe the subset

$$X_n := \big\{(x_1, \ldots, x_n) \in \mathbb{E}^n \mid \exists_{i \in \{1,\ldots,n\}}\ x_i \geq 0\big\}$$

together with the induced *length* metric from $\mathbb{E}^n$. It is CAT(0) if $n = 2$ but is *not* CAT(0) if $n = 3$ (not even CAT(κ) for any κ).

Let L be a geodesic segment (straight line) in $\mathbb{E}^2$. Consider the metric completion of the space $\mathbb{E}^2 \setminus L$ with the induced length metric. This space has non-positive curvature, i.e. curvature ≤ 0, but is *not* CAT(0).

(iii) *(Trees)* A metric simplicial tree, i.e. the geometric (metric) realisation of a (combinatorial) tree, is CAT(0). It is also CAT(κ) for any κ. Conversely, a geodesic space which is CAT(κ) for every κ need not be a simplicial tree, but it is close to such a space: A metric space X is CAT(κ) for every κ if and only if it is a so-called $\mathbb{R}$-tree which is characterized by the following two properties:

- X is uniquely geodesic.
- The union of two geodesic segments which meet at exactly one point is again a geodesic segment.

1.3.2 Comparison of different curvature notions

We state some facts, without proof, concerning the relationship between the CAT(κ) curvature condition for different κ, the δ-hyperbolicity condition of Gromov and the sectional curvature of Riemannian manifolds. We first have the intuitive fact:

- If $\kappa_1 \leq \kappa_2$, then a CAT(κ_1) space is also a CAT(κ_2) space.

This is also the reason why we will concentrate mainly on CAT(0) spaces in the rest of this chapter. As an example, hyperbolic n-space $\mathbb{H}^n = \mathbb{M}^n_{-1}$ is a CAT(-1) space, so it is also a CAT(0) space. We also have some kind of upper semi-continuity for the CAT(κ) curvature condition:

- If X is CAT(κ) for all $\kappa > \kappa_0$, then it is also CAT(κ_0).

The study of curvature plays also an important role in Riemannian Geometry. The curvature of a smooth curve in the plane at a point p is $1/r$ where r is the radius of a circle which is tangent to the curve at p and approaches the curve most tightly, the so-called osculating circle. Imagine driving a car and at some point p lock the steering wheel. The circle you drive then is the osculating circle at p of the curve you were driving. The curvature of a smooth surface embedded into $\mathbb{R}^3$ at the point p is the product of the two principal curvatures of that surface at the point p. The principal curvatures are obtained by looking at all smooth curves through p which are contained in the surface and also in a hyperplane of $\mathbb{R}^3$. The maximum and minimum of the curvatures of these curves at p are the principal curvatures, where a negative sign is given to one the curvatures if the two oscullating circles of the two corresponding curves lie on different sides of the surface. Now, Gauss observed in his Theorema Egregium that this curvature notion is an intrinsic notion of the surface, i.e. the curvature can be described solely in terms of the Riemannian metric induced from $\mathbb{R}^3$ via the embedding. In particular, this curvature notion can be generalized to abstract smooth 2-dimensional manifolds with a Riemannian metric. Finally, the sectional curvature of a smooth n-dimensional Riemannian manifold M is defined as follows: Pick a point $p \in M$ and a 2-dimensional subspace S of the tangent space T_pM. Consider a 2-dimensional totally geodesic submanifold N of M tangent to S at p. The curvature of M at p in direction S is the curvature of N at p. Now we have the following result:

- Let X be a smooth Riemannian manifold. Then the length metric induced by the Riemannian metric is of curvature $\leq \kappa$, i.e. locally CAT(κ), if and only if the sectional curvature of X is everywhere $\leq \kappa$.

Recall Gromov's δ-hyperbolicity condition ($\delta > 0$): A geodesic space is δ-hyperbolic if all geodesic triangles are δ-thin in the sense that each side of such a triangle is contained in the δ-neighborhood of the union of the other two sides. One can show that the hyperbolic plane $\mathbb{H}^2$ is δ-hyperbolic with $\delta = \ln(1+\sqrt{2})$. By the CAT(κ) inequality we obtain:

- Every CAT(κ) space with $\kappa < 0$ is $\delta(\kappa)$-hyperbolic.

In general, δ-hyperbolicity is not related to the CAT(0) curvature condition: Each bounded geodesic space is δ-hyperbolic for some δ, but not CAT(0) if it is not contractible (see Proposition 1.3.8 below). Conversely, the flat plane $\mathbb{E}^2 = \mathbb{M}^2_0$ is clearly CAT(0) but not δ-hyperbolic for any δ. However, if the space in question is already CAT(0) as well as proper and cocompact (i.e. there is a group of isometries acting cocompactly), then flat planes are the only obstruction to being δ-hyperbolic:

- (Flat Plane Theorem) Let X be a proper and cocompact CAT(0) space. Then X is δ-hyperbolic for some δ if and only if X does not contain an isometrically embedded flat plane.

1.3.3 Basic properties

In the following, we will frequently have to deal with geodesics which are linearly reparameterized. These are exactly the paths $c\colon [a,b] \to X$ such that $d(c(t), c(t')) = \lambda|t-t'|$ for some $\lambda > 0$. These paths are called *constant speed geodesics* or *cs-geodesics* for short. There is also the analogous notion of a *cs-local-geodesic*: There is $\lambda > 0$ such that for all $t_0 \in [a,b]$ there is $\varepsilon > 0$ such that $d(c(t), c(t')) = \lambda|t-t'|$ for all $t, t' \in [t_0-\varepsilon, t_0+\varepsilon]$ (we could define the λ to be dependent on t_0, but then one proves that in fact it is not).

A great deal (but not everything) of the rich structure of CAT(0) spaces comes from the fact that the metric of CAT(0) spaces is convex:

Lemma 1.3.7. The metric of a CAT(0) space X is *convex*, i.e. if c_1 and c_2 are two cs-geodesics $[0,1] \to X$, then for all $t \in [0,1]$ we have

$$d\big(c_1(t), c_2(t)\big) \leq (1-t)\cdot d\big(c_1(0), c_2(0)\big) + t\cdot d\big(c_1(1), c_2(1)\big)$$

Proof. Exercise! □

Proposition 1.3.8. *Let X be a CAT(0) space. Then we have:*

(i) X is uniquely geodesic.

(ii) Geodesics vary continuously with their endpoints.

(iii) Local geodesics are geodesics.

(iv) Every ball in X is convex.

(v) X and every ball in X is contractible.

(vi) Approximate midpoints are close to midpoints.

A few explanations before we turn to the proof: Let c be a geodesic from x to y in a uniquely geodesic space X and let c_n be a sequence of geodesics from x_n to y_n such that $x_n \to x$ and $y_n \to y$. Assume that, after linear reparameterization, that the paths c, c_n are cs-geodesics defined on a fixed intervall, say $[0, 1]$. By definition, geodesics in X *vary continuously with their endpoints* if $c_n \to c$ uniformly on $[0, 1]$ in every such situation.

Next, recall that a topological space X is called *contractible* if there is a continous map $H\colon X \times [0, 1] \to X$ and a point $x_0 \in X$ such that $H(x, 0) = x$ and $H(x, 1) = x_0$ for all $x \in X$.

Item (vi) means exactly the same thing as Lemma 1.2.16 with $\mathbb{M}^2_\kappa$ replaced with the CAT(0) space X.

Proof.

(i) By definition, X is geodesic. Let $x, y \in X$ and let $c, c'\colon [0, l] \to X$ be two geodesics from x to y. Let $t \in (0, l)$ and $z = c(t)$. Observe the geodesic triangle Δ with vertices x, z, y formed by the geodesics $c|_{[0,t]}, c|_{[t,l]}, c'|_{[0,l]}$ and a comparison triangle $\bar{\Delta}$ in $\mathbb{E}^2$ with vertices $\bar{x}, \bar{y}, \bar{z}$. The angle at z in Δ is π, so also the angle at $\bar{z}$ in $\bar{\Delta}$ has to be π (Proposition 1.3.4(iii)). So $\bar{\Delta}$ is degenerate in the sense that it is contained in a geodesic segment. So if $a := c'(t)$, it follows that $\bar{z} = \bar{a}$. By the CAT(0) inequality, we obtain $d(a, z) = 0$ and thus $c(t) = c'(t)$.

(ii) Let c, c_n, x, y, x_n, y_n be as in the paragraph right after the proposition. For each n, let $c'_n\colon [0, 1] \to X$ be the cs-geodesic from x_n to y. By the triangle inequality and the convexity of the metric (Proposition 1.3.7), we obtain:

$$\begin{aligned} d(c(t), c_n(t)) &\leq d(c(t), c'_n(t)) + d(c'_n(t), c_n(t)) \\ &\leq (1 - t) \cdot d(c(0), c'_n(0)) + t \cdot d(c'_n(1), c_n(1)) \\ &\leq d(x, x_n) + d(y, y_n) \end{aligned}$$

and so the claim follows.

(iii) Let $c\colon [0,l] \to X$ be a local geodesic and consider the set S of points $t \in [0,l]$ such that $c|_{[0,t]}$ is a geodesic. Since c is a local geodesic, S contains a neighborhood of 0 in $[0,l]$. It is easy to see that S is closed in $[0,l]$. We have to show that it is open: Let $t \in S$ with $t > 0$ and $z = c(t)$. Then since c is a local geodesic, there is an $\varepsilon > 0$ such that $c|_{[t-\varepsilon,t+\varepsilon]}$ is a geodesic. Consider the geodesic triangle Δ formed by the geodesics $c|_{[0,t]}, c|_{[t,t+\varepsilon]}$ and another geodesic from $c(0) = x$ to $c(t+\varepsilon) = y$. Consider a comparison triangle $\bar{\Delta}$ in $\mathbb{E}^2$ with vertices $\bar{x}, \bar{z}, \bar{y}$. The angle at z in Δ is π and so, similarly as in (i), we have that $\bar{\Delta}$ is degenerate. By the CAT(0) inequality, also Δ has to be degenerate. Thus, we can compute

$$\begin{aligned} l\big(c|_{[0,t+\varepsilon]}\big) &= l\big(c|_{[0,t]}\big) + l\big(c|_{[t,t+\varepsilon]}\big) \\ &= d(x,z) + d(z,y) \\ &= d(\bar{x},\bar{z}) + d(\bar{z},\bar{y}) \\ &= d(\bar{x},\bar{y}) \\ &= d(x,y) \end{aligned}$$

Consequently, $c|_{[0,t+\varepsilon]}$ is a geodesic.

(iv) Let $x \in X$ and $r > 0$. Let $y, z \in B_r(x)$. Consider a geodesic triangle with vertices x, y, z and a comparison triangle in $\mathbb{E}^2$ with vertices $\bar{x}, \bar{y}, \bar{z}$. Consider any point a on the geodesic segment between y and z. By 1.3.4(v) we have

$$d(x,a) \le d(\bar{x},\bar{a}) \le \max\{d(\bar{x},\bar{y}), d(\bar{x},\bar{z})\} = \max\{d(x,y), d(x,z)\} < r$$

(v) Choose a point $x_0 \in X$. For each $x \in X$ let $c_x\colon [0,1] \to X$ be the unique cs-geodesic from x to x_0. It follows from (ii) that the map $H\colon X \times [0,1] \to X$ given by $(x,t) \mapsto c_x(t)$ is continuous. Thus, X is contractible. The contractibility of balls follows similarly using that balls are convex.

(vi) Let $\varepsilon \in \mathbb{R}_{>0}$ and $l \in \mathbb{R}_{\ge 0}$. Let $\delta > 0$ be as in Lemma 1.2.16 for $\mathbb{E}^2$. Let $x, y \in X$ with $d(x,y) < l$ and m be the midpoint on the geodesic segment between x and y. Let m' be a δ-approximate midpoint for x, y, i.e.

$$\max\{d(x,m'), d(y,m')\} \le \frac{1}{2} \cdot d(x,y) + \delta$$

Consider the geodesic triangle with vertices x, y, m' and a comparison triangle in $\mathbb{E}^2$ with vertices $\bar{x}, \bar{y}, \bar{m}'$. Then $\bar{m}'$ is a δ-approximate midpoint of $\bar{x}, \bar{y}$ and by the Lemma we obtain $d(\bar{m}', \bar{m}) < \varepsilon$. By Proposition 1.3.4(vi) we have $d(m', m) \le d(\bar{m}', \bar{m}) < \varepsilon$. □

Remark 1.3.9. It also follows from the proof of item (iii) that the union of two geodesic segments forming an angle of π at a common initial point is again a geodesic segment. We used the fact that this is true in $\mathbb{E}^2$ (and more generally in $\mathbb{M}^n_\kappa$ which follows by direct calculation with spherical resp. Euclidean resp. hyperbolic segments). It is not true in general metric spaces.

For Hilbert spaces one can prove the existence of unique orthogonal projections: For $V \subset X$ a non-trivial closed subspace of a Hilbert space, there is a unique self-adjoint projection $P\colon X \to X$ onto V, i.e. P is linear and bounded, $P(X) = V$, $P^2 = P$ and $P^* = P$. We generalize this construction to CAT(0) spaces in the following proposition. The unique metric map π constructed there will simply be called the *projection* onto C. Recall that a map $d\colon X \to Y$ of metric spaces is called *metric* if $d(f(x), f(x')) \leq d(x, x')$ for all $x, x' \in X$ (these are the most natural morphisms between metric spaces if one views a metric space simply as a category enriched over the reals with the order relation as morphisms).

Proposition 1.3.10. *Let X be a CAT(0) space and let C be a convex subset which is complete in the induced metric.*

(i) For each $x \in X$ there is a unique $\pi(x) \in C$ such that

$$d(x, \pi(x)) = d(x, C) = \inf_{c \in C} d(x, c)$$

(ii) If $y \in [x, \pi(x)]$, then $\pi(x) = \pi(y)$.

(iii) If $x \notin C$ and $\pi(x) \neq y \in C$, then the angle at $\pi(x)$ between the unique geodesics from $\pi(x)$ to x and y respectively is greater or equal to $\pi/2$.

(iv) The map $\pi\colon X \to X$ is metric.

Proof.

(i) Let y_n be a sequence in C such that $d(y_n, x) \to d(x, C) =: D$. If we show that the sequence y_n is Cauchy, then we can set $\pi(x) := \lim_{n\to\infty} y_n \in C$. So let $\varepsilon > 0$ be small. Let $\delta > 0$ be the positive solution of the quadratic equation $\delta^2 + 2D\delta - \varepsilon^2/4 = 0$. Let N be sufficiently big so that for all $n \geq N$ we have $d(y_n, x) < D + \delta$. For $n, m \geq N$ consider the geodesic triangle with vertices y_n, y_m, x and a comparison triangle in $\mathbb{E}^2$ with vertices $\bar{y}_n, \bar{y}_m, \bar{x}$. For each $z \in [y_n, y_m]$ we have $z \in C$ since C is convex. Furthermore, we have $d(\bar{x}, \bar{z}) \geq D$, else, by Proposition 1.3.4(v), we would have $d(z, x) \leq d(\bar{x}, \bar{z}) < D$ which contradicts the definition of D. So the line segment $[\bar{y}_n, \bar{y}_m]$ entirely lies in the annulus around $\bar{x}$ with inner radius D and outer radius $D+\delta$. An elementary calculation shows that the length of such a line is bounded by $2\sqrt{2\delta D+\delta^2}$. Thus, we obtain

$$d(y_n, y_m) = d(\bar{y}_n, \bar{y}_m) \leq 2\sqrt{2\delta D + \delta^2} = \varepsilon$$

(ii) We have to show that $\pi(x) \in C$ minimizes the distance to y. Assume there is a $c \in C$ with $d(c, y) < d(\pi(x), y)$. Then we have

$$\begin{aligned} d(c, x) &\leq d(c, y) + d(y, x) \\ &< d(\pi(x), y) + \big(d(x, \pi(x)) - d(\pi(x), y)\big) \\ &= d(x, \pi(x)) \end{aligned}$$

which contradicts the fact that $\pi(x)$ minimizes the distance to x.

(iii) Assume that the angle in question is $< \pi/2$. Let $z = \pi(x)$. By the definition of the angle, we find points $x' \in [z, x]$ and $y' \in [z, y]$ such that the (Euclidean angle) in a comparison triangle $\bar{\Delta}(z, x', y') = \Delta(\bar{z}, \bar{x}', \bar{y}')$ also is $< \pi/2$. Then we find a point $\bar{p} \in [\bar{z}, \bar{y}']$ such that $d(\bar{p}, \bar{x}') < d(\bar{z}, \bar{x}')$. By Proposition 1.3.4(v) we obtain

$$d(p, x') \leq d(\bar{p}, \bar{x}') < d(\bar{z}, \bar{x}') = d(z, x') = d(\pi(x'), x') = d(x', C)$$

a contradiction to $p \in [z, y'] \subset [z, y] \subset C$.

(iv) Let x_1, x_2 be two points not in C with $z_1 := \pi(x_1) \neq \pi(x_2) =: z_2$. Let $\bar{\Delta}(x_1, z_1, z_2) = \Delta(\bar{x}_1, \bar{z}_1, \bar{z}_2)$ and $\bar{\Delta}(x_1, x_2, z_2) = \Delta(\bar{x}_1, \bar{x}_2, \bar{z}_2)$ be two comparison triangles with $\bar{z}_1$ and $\bar{x}_2$ on different sides of the common diagonal. We already know that the angles at z_1 and z_2 are $\geq \pi/2$. So by the triangle inequality for angles and Proposition 1.3.4(iii) the same is true for the Euclidean angles at $\bar{z}_1$ and $\bar{z}_2$. By elementary considerations in the Euclidean plane we deduce

$$d(x_1, x_2) = d(\bar{x}_1, \bar{x}_2) \geq d(\bar{z}_1, \bar{z}_2) = d(z_1, z_2) = d(\pi(x_1), \pi(x_2))$$

and so the claim follows. □

We finish this section by proving a fixed point theorem for CAT(0) spaces which generalizes a famous theorem of Cartan in the Riemannian setting. For this we need the notion of circumcenters: Let X be a metric space and $\emptyset \neq A \subset X$ bounded. For $x \in X$ arbitrary define

$$r(x, A) := \sup_{a \in A} d(x, a)$$

The *circumradius* of A is defined as

$$r(A) := \inf_{x \in X} r(x, A)$$

Each $x \in X$ with $r(x, A) = r(A)$ is called a *circumcenter* of A.

Proposition 1.3.11. *Let X be a complete CAT(0) space and $\emptyset \neq A \subset X$ bounded. Then there exists a unique circumcenter c_A of A.*

Proof. Let $q, r \in X$ and $m \in X$ be the (unique) midpoint. By the CN inequality of Bruhat and Tits (Corollary 1.3.5), we have

$$d(p,m)^2 \leq \frac{1}{2}\Big(d(p,q)^2 + d(p,r)^2\Big) - \frac{1}{4}d(q,r)^2$$

for each $p \in X$. Taking the sup over $p \in A$, we get

$$r(A)^2 \leq r(m,A)^2 \leq \frac{1}{2}\Big(r(q,A)^2 + r(r,A)^2\Big) - \frac{1}{4}d(q,r)^2$$

and thus

$$d(q,r) \leq \sqrt{2\Big(r(q,A)^2 + r(r,A)^2 - 2r(A)^2\Big)}$$

If q, r are both circumcenters of A, we immediately get $d(q,r) = 0$ and thus $q = r$. On the other hand, if x_n is a sequence of points with $r(x_n, A) \to r(A)$, then by the inequality above with $q = x_n$ and $r = x_m$, we see that x_n is a Cauchy sequence and so $x_n \to x$ with a unique $x \in X$. This x is the unique circumcenter of A. □

Corollary 1.3.12 (Bruhat-Tits Fixed Point Theorem)**.** Let G be a group acting on a complete CAT(0) space X by isometries such that there is a non-empty bounded $A \subset X$ with $g \cdot A = A$ for all $g \in G$. Then G fixes the circumcenter c_A of A. In particular, if there is a finite orbit or if G is finite, then there is a global fixed point.

1.3.4 Flatness Theorems

Theorem 1.3.13 (Flat Triangle Theorem)**.** *Let X be CAT(0). If at least one of the angles in a geodesic triangle Δ in X is equal to the corresponding angle in a comparison triangle $\bar{\Delta}$ in $\mathbb{E}^2$, then Δ is flat in the sense that the convex hull $C(\Delta)$ of Δ (i.e. the intersection of all convex subsets containing Δ) is isometric to the convex hull $C(\bar{\Delta})$ of $\bar{\Delta}$.*

Proof. Step 1: Let p, q, s be the vertices of Δ and $\bar{p}, \bar{q}, \bar{s}$ be the corresponding vertices in $\bar{\Delta}$ and assume that the angles at p resp. $\bar{p}$ coincide. Let $r \in [q, s]$. By the CAT(0) property we have $d(p,r) \leq d(\bar{p}, \bar{r})$ but we want to show that they are equal. To this end, consider comparison triangles $\tilde{\Delta}(p,s,r) = \Delta(\tilde{p}, \tilde{s}, \tilde{r})$ and $\tilde{\Delta}(p,r,q) = \Delta(\tilde{p}, \tilde{r}, \tilde{q})$ such that $\tilde{s}$ and $\tilde{q}$ lie on different sides of the common edge. By the CAT(0) property, the sum of the two angles at the common vertex $\tilde{r}$ is at least π. By the triangle equality for angles, the CAT(0) property and Alexandrov's Lemma, we have

$$\angle_p(q,s) \leq \angle_p(q,r) + \angle_p(r,s) \leq \angle_{\tilde{p}}(\tilde{q},\tilde{r}) + \angle_{\tilde{p}}(\tilde{r},\tilde{s}) \leq \angle_{\bar{p}}(\bar{q},\bar{s}) = \angle_p(q,s)$$

So we have equality everywhere and the quadrilateral $\tilde{p}, \tilde{q}, \tilde{s}, \tilde{r}$ is already a triangle isometric to $\bar{\Delta}$. So we have $d(p, r) = d(\tilde{p}, \tilde{r}) = d(\bar{p}, \bar{r})$ as we wanted to show.

Step 2: Define $j \colon C(\bar{\Delta}) \to X$ by requiring that the geodesic segments $[\bar{p}, \bar{r}]$ are isometrically mapped to the geodesic segments $[p, r]$ (this is possible since we have shown in the first step that their lengths coincide). We first want to show that j is an isometric embedding. Let $\bar{r}_1 \neq \bar{r}_2$ such that $\bar{s}, \bar{r}_1, \bar{r}_2, \bar{q}$ are aligned on a straight line. Consider two points $\bar{x}_i$ on the segments $[\bar{p}, \bar{r}_i]$ and define $x_i = j(\bar{x}_i)$. The segments $[p, r_i]$ divide Δ into three triangles and the same is true for $\bar{\Delta}$. In step 1 we have shown that $d(p, r_i) = d(\bar{p}, \bar{r}_i)$ so the three triangles in $\bar{\Delta}$ are indeed comparison triangles for the three triangles in Δ. Let δ_j resp. $\bar{\delta}_j$ for $j = 1, 2, 3$ be the angles at p resp. $\bar{p}$ of the three triangles. Then by the CAT(0) property we have

$$\angle_p(q, s) \leq \sum_j \delta_j \leq \sum_j \bar{\delta}_j = \angle_{\bar{p}}(\bar{q}, \bar{s}) = \angle_p(q, s)$$

So we have equality everywhere and consequently $\delta_j = \bar{\delta}_j$, in particular, $\delta_2 = \bar{\delta}_2$. It follows from Proposition 1.3.4(iv) that $d(x_1, x_2) \geq d(\bar{x}_1, \bar{x}_2)$. On the other hand, by the CAT(0) inequality (Definition 1.3.1), we have $d(x_1, x_2) \leq d(\bar{x}_1, \bar{x}_2)$. So j is an isometric embedding.

Step 3: Finally, want to prove $j(C(\bar{\Delta})) = C(\Delta)$. It is clear that $j(\bar{\Delta}) = \Delta$, so if C is a convex subset containing Δ, it follows from the definition of j that $j(C(\bar{\Delta})) \subset C$. Thus, $j(C(\bar{\Delta})) \subset C(\Delta)$. Conversely, if we show that $j(C(\bar{\Delta}))$ is convex, we have equality. Let x_1, x_2 be as in step 2 and let γ be a geodesic from x_1 to x_2. The geodesic from $\bar{x}_1$ to $\bar{x}_2$ is mapped to a geodesic from x to x' since j is an isometric embedding. By uniqueness, this geodesic has to be γ, so γ is contained in the image of j. □

Theorem 1.3.14 (Flat Quadrilateral Theorem)**.** *Let X be CAT(0). Consider four points in X together with four geodesic segments joining these points to form a quadrilateral in X. Then the sum of the angles in such a quadrilateral is at most 2π and if it is equal to 2π then its convex hull is isometric to the convex hull of a convex quadrilateral in $\mathbb{E}^2$.*

Proof. Introduce a diagonal in the quadrilateral in X. Then use the Flat Triangle Theorem to deduce that the two resulting triangles are flat via isometries j and j'. Then show that these two isometries combine to give an isometry onto the convex hull of the given four points in X. The details are left to the reader as an exercise. □

We say that two geodesic lines $c, c' \colon \mathbb{R} \to X$ are *asymptotic* if we have $\sup_{t \in \mathbb{R}} d(c(t), c'(t)) < \infty$. The proof of the following theorem shows in particular that asymptotic lines in CAT(0) spaces are even *parallel* in the sense that there is a constant D such that $d(c(t), c'(t)) = D$ for all $t \in \mathbb{R}$. Furthermore, c, c' can be parameterized in such a way that $D = d(c(\mathbb{R}), c'(\mathbb{R}))$.

Theorem 1.3.15 (Flat Strip Theorem)**.** *Let X be CAT(0). If two geodesic lines c, c' in X are asymptotic then the convex hull of $c(\mathbb{R}) \cup c'(\mathbb{R})$ is isometric to a flat strip, i.e. to a convex hull of two parallel lines in $\mathbb{E}^2$.*

Proof. Consider the projection $\pi\colon X \to X$ onto $c(\mathbb{R})$. We can assume that $\pi(c'(0)) = c(0)$. Below we will show that $\pi(c'(t)) = c(t)$ for all $t \in \mathbb{R}$. Similarly then, the projection π' onto $c'(\mathbb{R})$ satisfies $\pi'(c(t)) = c'(t)$. For any two $t_1 < t_2 \in \mathbb{R}$, consider now the quadrilateral given by the geodesic segments $[c(t_1), c'(t_1)]$, $[c(t_2), c'(t_2)]$, $c([t_1, t_2])$ and $c'([t_1, t_2])$. By Proposition 1.3.10(iii) all the angles in this quadrilateral are at least $\pi/2$. Thus by the Flat Quadrilateral Theorem, it is flat. By this observation it is easy to manufacture an isometry as claimed.

We now prove the remaining claim: For any $a, b \in \mathbb{R}$ the function $t \mapsto d\big(c(t + a), c'(t + b)\big)$ is convex and bounded, hence constant. In particular, since $\pi(c'(0)) = c(0)$ is the closest point to $c'(0)$ on $c(\mathbb{R})$, we have

$$d(c(t + a), c'(t)) = d(c(a), c'(0)) \geq d(c(0), c'(0)) = d(c(t), c'(t))$$

This says that $c(t)$ is the closest point to $c'(t)$ on $c(\mathbb{R})$, hence $\pi(c'(t)) = c(t)$. □

There is an important corollary to the Flat Strip Theorem which concerns product decompositions of CAT(0) spaces. Recall that the *product* of two metric spaces X_1, X_2 is the set $X_1 \times X_2$ endowed with the metric

$$d\big((x_1, x_2), (y_1, y_2)\big) := \sqrt{d(x_1, y_1)^2 + d(x_2, y_2)^2}$$

Note that $\mathbb{E}^n \times \mathbb{E}^m$ is then isometric to $\mathbb{E}^{n+m}$. The conclusion of the Flat Strip Theorem can be slightly restated saying that the convex hull of two asymptotic geodesic rays being isometric to $\mathbb{R} \times [0, D]$ where D is the distance between the two geodesic rays.

Corollary 1.3.16. Let X be a CAT(0) space and $c\colon \mathbb{R} \to X$ a geodesic line. Let X_c be the union of the images of all geodesic lines asymptotic (and thus parallel) to c. Then X_c is convex in X. Furthermore, there is a convex subset $Y \subset X_c$ and an isometry $\psi\colon Y \times \mathbb{R} \to X_c$.

Sketch of proof. Let $\pi\colon X_c \to X_c$ be the projection onto $c(\mathbb{R})$. Set $Y := \pi^{-1}(c(0))$. Define a bijection $\psi\colon Y \times \mathbb{R} \to X_c$ as follows: Let $x \in Y$ and $t \in \mathbb{R}$. There is a unique geodesic line c_x asymptotic to c and containing x and by the Flat Strip Theorem, we can assume that c_x is parameterized such that $d(c(t), c_x(t))$ is constant equal to the distance between $c(\mathbb{R})$ and $c_x(\mathbb{R})$. This implies $\pi(c_x(0)) = c(0)$ and thus $x = c_x(0)$. Now define $\psi(x, t) = c_x(t)$.

Let $x \neq x' \in Y$ and $t, t' \in \mathbb{R}$. Then by the Flat Strip Theorem, the two geodesics $c_x, c_{x'}$ span a flat strip isometric to $\mathbb{R} \times [0, D]$ where (cf. [3, Lemma II.2.15])

$$D = d(c_x(\mathbb{R}), c_{x'}(\mathbb{R})) = d(c_x(0), c_{x'}(0)) = d(x, x')$$

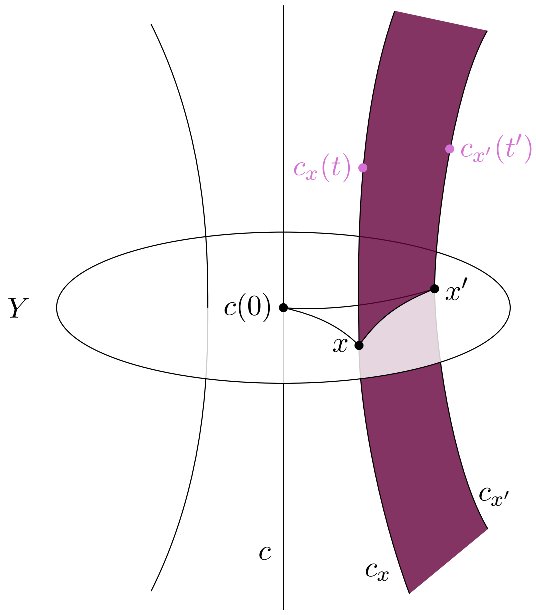

Figure 1.8: Product Decomposition for CAT(0) spaces

Thus the distance between $\psi(x, t)$ and $\psi(x', t')$ is given by the product distance between (x, t) and (x', t'). Consequently, ψ is an isometry. □

1.4 The Cartan-Hadamard Theorem

1.4.1 Motivation

Recall that a geodesic space is called CAT(0) if all geodesic triangles are smaller than in flat space, whereas we say that it is of curvature ≤ 0 if this is the case only locally. So the first condition is a global one, whereas the second condition is of local nature. One can ask, how big is the difference between the global and local curvature conditions. In proposition 1.3.8 we have seen that CAT(0) spaces are always contractible. This does not have to be the case for non-positively curved spaces as the following example shows (we choose this example because it is instructive to think of it in the proof of the Cartan-Hadamard Theorem).

Let Ω be the closed annulus

$$\Omega := \{x \in \mathbb{R}^2 \mid 1 \leq \|x\| \leq 2\}$$

Endow Ω with the metric

$$d(x,y) = \sqrt{\left(\|x\| - \|y\|\right)^2 + \angle_0(x,y)^2}$$

where $\angle_0(x,y)$ is simply the Euclidean angle between x and y. Note that (Ω, d) is just a thickened version of the circle with the angular metric, or a cylinder $S^1 \times [1,2]$ over the circle with the flat product metric. For each $v \in S^1 \subset \mathbb{R}^2$, the open subset $\{x \in \Omega \mid x/\|x\| \neq v\}$ is isometric to a flat strip with side lengths 1 and 2π, for example to the subset

$$\left\{(x,y) \in \mathbb{R}^2 \mid 0 \leq x \leq 1, 0 < y < 2\pi\right\}$$

of the flat plane $\mathbb{E}^2$. Thus, Ω is non-positively curved. However, it is not CAT(0): Observe the geodesic triangle with vertices $1, e^{i\frac{2\pi}{3}}, e^{i\frac{4\pi}{3}}$, where we identify $\mathbb{R}^2$ with $\mathbb{C}$. The unique geodesic segments joining these points are circle segments of radius 1. Each angle in this triangle is π, so it cannot be smaller than a triangle in $\mathbb{E}^2$. Moreover Ω is not contractible. The obstruction to being contractible here, and also to triangles being smaller than in $\mathbb{E}^2$, is the big hole in the middle, which is detected by a non-trivial fundamental group (see below). Basically, the Cartan-Hadamard Theorem states that if the space in question is a complete, connected length space of non-positive curvature, then this is the only obstruction to being CAT(0).

Using that CAT(0) spaces are not only simply connected (i.e. the fundamental group is trivial or there are no “holes”) but even contractible, the Cartan-Hadamard is also an instance of a local-to-global phenomenon: A local condition, namely non-positive curvature, together with a (relatively) weak global condition, namely simply connectedness, implies the strong global condition contractibility.

The theorem which we are presenting here for metric spaces has its origins in Riemannian geometry due to Cartan and Hadamard: The universal covering of a complete connected Riemannian manifold of non-positive sectional curvature is diffeomorphic to $\mathbb{R}^n$. The generalization to metric spaces is due to Gromov and Ballmann.

1.4.2 Fundamental groups and coverings

In this subsection we recall the theory of fundamental groups, covering spaces and their metric counterparts.

Let X, Y be topological spaces and $f, g\colon X \to Y$ two continuous maps. These two maps are called *homotopic* if there is a continuous map

$$H\colon X \times [0,1] \to Y$$

such that $H(x,0) = f(x)$ and $H(x,1) = g(x)$ for all $x \in X$. If $A \subset X$, we say that f, g are homotopic relative to A if $H(a,t) = f(a) = g(a)$ for all $a \in A$ and $t \in [0,1]$. This is an equivalence relation.

Let X be a space and $x_0 \in X$. We set

$$\pi_1(X, x_0) = \{\gamma\colon [0,1] \to X \mid \gamma \text{ continuous with } \gamma(0) = x_0 = \gamma(1)\}/\!\sim$$

where $\sim$ is homotopy relative to $\{0,1\} \subset [0,1]$. We can make this set into a group by concatenation of paths, the unit element is then represented by the constant path at x_0. This is called the *fundamental group* of (X, x_0).

We obtain, in the obvious way, a functor

$$\pi_1\colon \texttt{TOP}_* \to \texttt{GROUPS}$$

where $\texttt{TOP}_*$ is the category of pointed spaces (X, x_0) with basepoint preserving maps. This functor is homotopy invariant in the sense that if f, g are homotopic relative to the basepoint, then $\pi_1(f) = \pi_1(g)$.

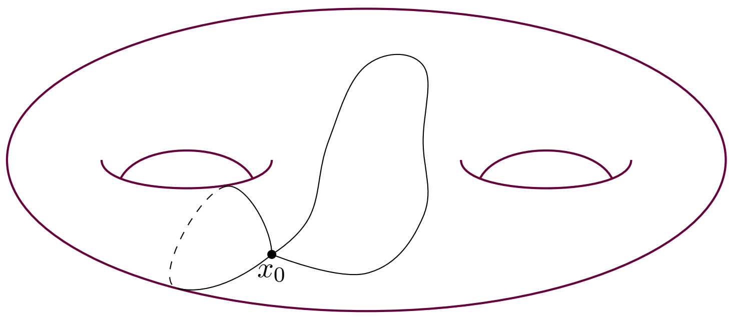

Figure 1.9: A trivial and a non-trivial loop on the surface of genus 2.

We call a space X *simply connected* if it is path connected and if $\pi_1(X, x_0)$ is trivial for some and therefore for any $x_0 \in X$.

Definition 1.4.1. Let $f\colon X \to Y$ be a map of spaces. We say that f is

i) a *local homeomorphism* if each point $x \in X$ has an open neighborhood U such that f sends U homeomorphically onto an open subset of Y.

ii) a *covering* if it is surjective and if for each point $y \in Y$ there is an open neighborhood V such that $f^{-1}(V)$ is a disjoint union of open sets in X which are sent homeomorphically onto V by f.

Theorem 1.4.2 (Lifts)**.** *Let $p\colon (Y, y_0) \to (X, x_0)$ be a covering and let Z be a path connected and locally path connected (i.e. every neighborhood of a point contains a path connected neighborhood). Let $f\colon (Z, z_0) \to (X, x_0)$ be a map. Then the lifting problem*

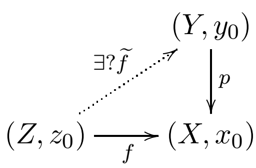

is solvable if and only if $\operatorname{im}\pi_1(f) \subset \operatorname{im}\pi_1(p)$*. In this case, the lift is unique.*

In particular, if $p\colon Y \to X$ is a covering and $\gamma\colon [0,1] \to X$ a path in X, then there is a unique lift $\widetilde{\gamma}$ of γ to Y provided the lift of the starting point $\gamma(0)$ is given. It also follows from the lifting proposition that a covering $p\colon (Y, y_0) \to (X, x_0)$ induces an injection on π_1, i.e. $\pi_1(p)$ is injective. Thus, $\pi_1(Y, y_0)$ can be seen as a subgroup of $\pi_1(X, x_0)$.

A space X is called *semilocally simply connected* if every point has an open neighborhood U such that the map $\pi_1(U \hookrightarrow X)\colon \pi_1(U, y) \to \pi_1(X, y)$ is trivial for each $y \in U$.

Theorem 1.4.3 (Classification of coverings)**.** *Let* (X, x_0) *be path connected, locally path connected and semilocally simply connected. Then for every subgroup* $G < \pi_1(X, x_0)$ *there is exactly one connected covering*

$$p^G\colon (Y^G, y_0^G) \to (X, x_0)$$

such that $\operatorname{im}\pi_1(p) = G$*. If* $H < G < \pi_1(X, x_0)$ *then there is exactly one covering map* $p\colon (Y^H, y_0^H) \to (Y^G, y_0^G)$ *such that* $p^G \circ p = p^H$*.*

The uniqueness statement is true if Y is path connected and locally path connected, without any assumptions on X.

The coverings associated to normal subgroups (i.e. $\operatorname{im}\pi_1(p)$ is normal in $\pi_1(X, x_0)$) are called *normal coverings*. The covering associated to the trivial subgroup (i.e. $\operatorname{im}\pi_1(p)$ is trivial or, in other words, Y is simply connected) is called the *universal covering* since it covers every other covering. The universal covering $(\widetilde{X}, \widetilde{x}_0)$ of an (X, x_0) as in the theorem usually is constructed as the set of homotopy classes (with homotopies fixing the start and end point) of paths $\gamma\colon [0,1] \to X$ with $\gamma(0) = x_0$ together with a suitable topology. The base point $\widetilde{x}_0$ is represented by the constant path at x_0. The covering map then sends a point of $\widetilde{X}$ to $\gamma(1)$ where γ is a representative of that point.

Remark 1.4.4. In the situation of Theorem 1.4.3, the index of G in $\pi_1(X, x_0)$ equals the cardinality of any fiber $(p^G)^{-1}(x)$.

Definition 1.4.5. Let $p\colon Y \to X$ be a covering. A homeomorphism $\varphi\colon Y \to Y$ is called a *deck transformation* if $p \circ \varphi = \varphi$. Denote the group of deck transformations by $\mathcal{D}(p)$.

Definition 1.4.6. An action of a group G on a space X is called

(i) *proper* if for each compact subset $K \subset X$ the set $\{g \in G \mid g{\cdot}K \cap K \neq \emptyset\}$ is finite.

(ii) *properly discontinuous* if for each $x \in X$ there is an open neighborhood U such that $g \cdot U \cap U = \emptyset$ for each $g \neq 1$.

Proposition 1.4.7. *Let $p\colon (Y, y_0) \to (X, x_0)$ be a covering with Y path connected and locally path connected. Then $\mathcal{D}(p)$ acts properly discontinuously on Y. It acts transitively on each fiber $p^{-1}(x)$ if and only if p is a normal covering. In this case, we have*

$$\mathcal{D}(p) \cong \pi_1(X, x_0)\big/ \operatorname{im} \pi_1(p)$$

und p induces a homeomorphism $\mathcal{D}(p)\backslash Y \to X$.

Conversely to this proposition, normal coverings can be obtained from group actions:

Proposition 1.4.8. *Let G be a group acting properly discontinuously on a path connected and locally path connected space Y. Then the quotient map $p\colon Y \to G\backslash Y$ is a normal covering.*

Remark 1.4.9.

(i) In the situation of Proposition 1.4.8, if $X := G\backslash Y$, $y_0 \in Y$ and $x_0 = p(y_0)$, we obtain $G = \mathcal{D}(p) \cong \pi_1(X, x_0)\big/ \operatorname{im} \pi_1(p)$. In particular, $G \cong \pi_1(X, x_0)$ if Y is simply connected.

(ii) In the situation of Proposition 1.4.7, if p is not normal, then we still have
$$\mathcal{D}(p) \cong N(\operatorname{im} \pi_1(p))\big/ \operatorname{im} \pi_1(p)$$
where $N(\operatorname{im} \pi_1(p))$ is the normalizer in $\pi_1(X, x_0)$. The isomorphism sends a loop γ representing an element in $N(\operatorname{im} \pi_1(p))$ to the unique deck transformation sending y_0 to the endpoint of the lift $\widetilde{\gamma}$ starting at y_0. Furthermore, p induces a covering $\mathcal{D}(p)\backslash Y \to X$ and the pair of coverings $Y \to \mathcal{D}(p)\backslash Y \to X$ corresponds exactly to the sequence
$$\operatorname{im} \pi_1(p) \triangleleft N(\operatorname{im} \pi_1(p)) < \pi_1(X, x_0)$$

We now want to treat metric versions of the notions above.

Definition 1.4.10. Let X, Y be metric spaces and $f\colon X \to Y$. We say that f is

(i) a *local isometry* if for every x there is an $r > 0$ such that f maps $B(x, r)$ isometrically onto $B(f(x), r)$.

(ii) a *metric covering* if it is surjective and if for every point $y \in Y$ there is an $r > 0$ such that for each $x, z \in f^{-1}(y)$ we have $B(x, r) \cap B(z, r) = \emptyset$ and f maps $B(x, r)$ isometrically onto $B(y, r)$.

Lemma 1.4.11. Let X be a length space. Let Y be a Hausdorff space and $p\colon Y \to X$ a local homeomorphism. Then there is unique length metric d on Y such that p is a local isometry. It is called the *induced length metric*. The original topology on Y is induced by d.

In general, there are many more non-length metrics which make p into a local isometry.

Proof. For $y_1, y_2 \in Y$ we define

$$d(y_1, y_2) = \inf \left\{ l(p \circ \gamma) \mid \gamma\colon [0,1] \to Y \text{ is a path from } y_1 \text{ to } y_2 \right\}$$

This is a metric since we assumed that Y is Hausdorff. If we know that p is a local isometry, then it follows that p preserves the length of curves (since measuring the length of curves is a local procedure), and that Y is a length space.

We show that p is a local isometry: First note that p is metric by construction. Let $y \in Y$ and V, U open subsets such that $y \in V$ and $p|_V : V \to U$ is a homeomorphism with local inverse $s\colon U \to V$. Let $r > 0$ be so small so that $B(p(y), 2r) \subset U$. We claim that

$$s' := s|_{B(p(y),r)} \colon B(p(y), r) \to B(y, r)$$

is an isometry. First let $a \in B(y, r)$, then $d(p(y), p(a)) \leq d(y, a) < r$ and $s(p(a)) = a$, so s' is surjective. Let $\epsilon > 0$ and $a, b \in B(p(y), r)$. Since X is a length space, there exists a path c in $B(p(y), 2r)$ from a to b with $l(c) < d(a, b) + \epsilon$. By the definition of the metric on Y we obtain $d(s(a), s(b)) \leq l(c) < d(a, b) + \epsilon$. In the limit we obtain $d(s(a), s(b)) \leq d(a, b)$. Since p is metric, we indeed have equality here.

If d' is any other length metric on Y such that $p\colon (Y, d') \to (X, d)$ is a local isometry, then also $\mathrm{id}\colon (Y, d') \to (Y, d)$ is a local isometry und thus preserves lengths of curves. It follows $d = d'$ because both metrics are length metrics.

Let $\mathcal{O}$ be the original topology on Y and $\mathcal{O}_d$ the one induced by d. Then $\mathrm{id}\colon (Y, \mathcal{O}) \to (Y, \mathcal{O}_d)$ is a bijective local homeomorphism and every such map is a homeomorphism. □

Now let X be a length space and $p\colon Y \to X$ a covering. Since then also Y is Hausdorff, we can apply the above lemma to deduce that there is a unique length metric on Y making p into a local isometry. It follows from the proof of the lemma that $p\colon Y \to X$ is even a metric covering.

We also note that the group of deck transformations then acts on Y by isometries because of the following facts: Let $\varphi \in \mathcal{D}(p)$. Then $\varphi\colon Y \to Y$ is a local isometry and a homeomorphism. Since any bijective local isometry of length spaces is already an isometry, φ is an isometry.

Definition 1.4.12. Let G be a group acting on a metric space X by isometries. This action is called *metrically proper* if for each $x \in X$ there is $r > 0$ such that the set $\{g \in G \mid g \cdot B(x,r) \cap B(x,r) \neq \emptyset\}$ is finite.

Every metrically proper action is proper. The converse implication is true if X is proper. However, in general, being metrically proper is a stronger condition.

An action $G \curvearrowright X$ is metrically proper and free if and only if it is properly discontinuous. To see that it is properly discontinuous, choose a point $x \in X$ and $r > 0$ such that $g_1, \ldots, g_k$ are the only non-trivial group elements with $g_i \cdot B(x,r) \cap B(x,r) \neq \emptyset$. Then observe the ball around x with radius $r' := \frac{1}{2} \min\{d(x, g_i x) \mid i = 1, \ldots, k\} > 0$. Thus, the projection $X \to G\backslash X$ is a (normal) covering if X is path connected and locally path connected. To clarify the metric counterpart of this situation, we have to introduce the metric counterpart of the quotient topology.

Definition 1.4.13. Let Y be a metric space, X a set and $p \colon Y \to X$ a surjective function. Then there is a unique pseudo-metric on X such that p is metric and the following universal property is fulfilled: In every commutative diagram of pseudo-metric spaces and functions

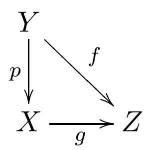

the function f is metric if and only g is metric. This pseudo-metric is given by the formula

$$d(x, x') = \inf \sum_{i=1}^{n} d(a_i, b_i)$$

where the infimum runs over all tuples $(a_1, b_1, a_2, b_2, \ldots, a_n, b_n)$ of points in Y such that $p(a_1) = x$, $p(b_n) = x'$ and $p(b_i) = p(a_{i+1})$ for all $i = 1, \ldots, n-1$.

Remark 1.4.14. Provided that the quotient pseudo-metric is a metric, the induced topology is the same as the quotient topology. Furthermore, in this case, the quotient X is a length space if Y is one.

Applying this definition to the projection $X \to G\backslash X$ of an action of G on X by isometries, we obtain a pseudo-metric d on the quotient $G\backslash X$. This metric is given by the formula

$$d(Gx, Gx') = \inf\{d(z, z') \mid z \in Gx, z' \in Gx'\}$$

which is just the distance between the orbits Gx, Gx' as subsets of X. To see this, note that each tuple $(a_1, b_1, \ldots, b_n)$ with $n > 1$ as in the formula above can be replaced by the shorter tuple $(a_1, gb_2, ga_3, \ldots, ga_n, gb_n)$ with g such that $b_1 = ga_2$ because we have $d(ga_i, gb_i) = d(a_i, b_i)$ as well as $d(a_1, gb_2) \leq d(a_1, b_1) + d(a_2, b_2)$.

Now assume that the action $G \curvearrowright X$ is metrically proper and X is a path connected and locally path connected length space. Then $G\backslash X$ is a metric space (not only a pseudo-metric space) and thus also a length space. Assume further that the action is free, then the projection $X \to G\backslash X$ is a local isometry. We have already noted that $X \to G\backslash X$ is a covering and we know that it is even a metric covering with the induced length metric on X. But from the uniqueness statement in Lemma 1.4.11 it follows that the induced length metric on X has to be the original one. Summarizing, we obtain the metric counterpart of Proposition 1.4.8: *The quotient of a metrically proper and free action by isometries on a path connected and locally path connected length space is a metric covering and the base space is also a path connected and locally path connected length space.*

There is also a metric counterpart to Proposition 1.4.7: Let $p \colon Y \to X$ be a metric covering with path connected and locally path connected length spaces Y, X. We have already noted that deck transformations are isometries. Proposition 1.4.7 says that the action by deck transformations is metrically proper and free. If p is normal, then p induces a homeomorphism $\mathcal{D}(p)\backslash Y \to X$. This is also a local isometry and consequently an isometry. Summarizing, we have: *The group of deck transformations of a metric covering of path connected and locally path connected length spaces acts metrically properly and freely by isometries on the total space and, if the covering is normal, the quotient is isometric to the base space.*

Example 1.4.15. The standard example is that of the circle $\mathbb{S}^1 = \mathbb{M}^1_1$. Its fundamental group is isomorphic to $\mathbb{Z}$ (counting the number how often a loop turns around the circle). The subgroups of $\mathbb{Z}$ are the trivial subgroup and the subgroups $n\mathbb{Z}$ for $n \in \mathbb{N}_{\geq 1}$. The cover of $\mathbb{S}^1$ corresponding to the subgroup $n\mathbb{Z}$ is

$$\mathbb{M}^1_{1/n^2} \to \mathbb{S}^1 \qquad t \mapsto t^n$$

where we think of S^1 embedded into $\mathbb{C}$. The universal cover is given by

$$\mathbb{E}^1 = \mathbb{M}^1_0 \to \mathbb{S}^1 \qquad t \mapsto \exp(it)$$

In the following, it is often instructive to think of the annulus (Ω, d) from the introduction. The universal cover is just a thickened version of the universal cover of $\mathbb{S}^1$:

$$\mathbb{E}^1 \times [1, 2] \to \Omega \qquad (t, r) \mapsto r\exp(it)$$

1.4.3 The main theorem and outline of the proof

Theorem 1.4.16 (Cartan-Hadamard-Gromov). *Let X be a complete connected length space. If X is of curvature ≤ 0, then the universal covering $\widetilde{X}$*

of X is CAT(0).

Remark 1.4.17.

(i) Since each CAT(0) space is contractible and X is locally CAT(0), X is locally path connected and semilocally simply connected. Since every connected length space is path connected, it follows that the universal covering $\widetilde{X}$ exists.

(ii) As explained in the preceding subsection, the universal covering is also a metric covering by implicitly endowing $\widetilde{X}$ with the induced length metric.

(iii) More generally, if $\kappa \leq 0$, curvature $\leq \kappa$ implies that the universal cover is CAT(κ).

Corollary 1.4.18. A complete simply connected length space of non-positive curvature is a CAT(0) space.

Let X be as in the theorem. We are going to reconstruct the universal covering directly: Choose a base point $x_0 \in X$. Define

$$\begin{aligned}\widetilde{X}_{x_0} = &\big\{c\colon [0,1] \to X \mid c \text{ cs-local geodesic with } c(0) = x_0\big\}\cup \\ &\big\{c\colon [0,1] \to X \mid \forall_t\ c(t) = x_0\big\}\end{aligned}$$

together with the metric

$$d(c, c') = \sup\big\{d(c(t), c'(t)) \mid t \in [0,1]\big\}$$

Define furthermore

$$\exp\colon \widetilde{X}_{x_0} \to X \qquad c \mapsto c(1)$$

If we knew that each homotopy class of paths would contain exactly one cs-local-geodesic, we would already know that exp is the universal covering of (X, x_0) by the standard construction of the universal covering (see preceding subsection). However, we will show directly that exp is the universal covering and deduce from this that each homotopy class contains exactly one cs-local-geodesic. For this we need two technical lemmas which we will prove in subsection 1.4.5:

Lemma 1.4.19. Let X be a metric space such that the metric is locally CAT(0) and locally complete. Let $c\colon [0,1] \to X$ be a cs-local-geodesic from x to y. Then there is an $\epsilon > 0$ small enough so that whenever $d(x, \bar{x}) < \epsilon$ and $d(y, \bar{y}) < \epsilon$, then there is exactly one cs-local-geodesic $\bar{c}\colon [0,1] \to X$ from $\bar{x}$ to $\bar{y}$ such that

$$\sup_{t\in[0,1]} d(c(t), \bar{c}(t)) < \epsilon$$

This cs-local-geodesic satisfies $l(\bar{c}) \le l(c) + d(x, \bar{x}) + d(y, \bar{y})$. Moreover, for any two cs-local-geodesics $c', c'' \colon [0,1] \to X$ with

$$\sup_{t\in[0,1]} d(c(t), c'(t)) < \epsilon \quad \text{and} \quad \sup_{t\in[0,1]} d(c(t), c''(t)) < \epsilon$$

the function $t \mapsto d(c'(t), c''(t))$ is convex.

Lemma 1.4.20. Let Y, X be metric spaces such that Y is complete, X is path connected and let $p \colon Y \to X$ be a local isometry. Assume furthermore that X is locally uniquely geodesic and these geodesics vary continuously with their endpoints. Then p is a covering.

We return to the proof of the main theorem.

Proposition 1.4.21.

(i) $\widetilde{X}_{x_0}$ *is contractible and complete.*

(ii) $\exp \colon \widetilde{X}_{x_0} \to X$ *is a local isometry and the universal covering of* X

Proof.

(i) To show that $\widetilde{X}_{x_0}$ is contractible, consider the homotopy

$$\widetilde{X}_{x_0} \times [0,1] \to \widetilde{X}_{x_0} \qquad (c, s) \mapsto r_s(c) := \big[[0,1] \ni t \mapsto c(st)\big]$$

Continuity follows from the estimate

$$\begin{aligned} d(r_s(c), r_{s'}(c')) &= \sup_t d(c(st), c'(s't)) \\ &\le \sup_t d(c(st), c(s't)) + \sup_t d(c(s't), c'(s't)) \\ &\le \sup_t \lambda|st - s't| + \sup_t d(c(t), c'(t)) \\ &\le \lambda|s - s'| + d(c, c') \end{aligned}$$

For completeness, let c_n be a Cauchy sequence in $\widetilde{X}_{x_0}$. Then for each $t \in [0,1]$, $c_n(t)$ is a Cauchy sequence in X and hence converges to a point $c(t) \in X$ which defines a function $c \colon [0,1] \to X$ with $c(0) = x_0$. We have to show that c is a cs-local-geodesic. Fix some $t \in [0,1]$ and choose $r > 0$ such that $B(c(t), r)$ is CAT(0). For N sufficiently large and all $n \ge N$, the points $c_n(t)$ lie in $B(c(t), r/2)$. It follows from Proposition 1.3.8(iii) that the parts of the cs-local-geodesics c_n for $n \ge N$ lying in $B(c(t), r)$ are cs-geodesics. We find $\epsilon > 0$ small enough such that for $I := [t - \epsilon, t + \epsilon]$ and $n \ge N$ we have $c_n(I) \subset B(c(t), r)$ as well as $c|_I \subset B(c(t), r)$. By Proposition 1.3.8(ii) the cs-geodesics $c_n|_I$ vary continuously with their endpoints and therefore converge uniformly to the unique cs-geodesic $\gamma \colon I \to B(c(t), r)$ from $c(t - \epsilon)$ to $c(t + \epsilon)$. Consequently, $\gamma = c|_I$ and $c|_I$ is a cs-geodesic. Hence c is a cs-local-geodesic.

(ii) First we want to show that exp is a local isometry: Let $c \in \widetilde{X}_{x_0}$. Let $\epsilon > 0$ as in Lemma 1.4.19 applied to c. Then the lemma implies that the restriction of exp to the ball $B(c, \epsilon)$ is a bijection onto the ball $B(c(1), \epsilon)$. Moreover, the convexity statement in the lemma implies that if $c', c'' \in B(c, \epsilon)$ then $d(c', c'') = d(c'(1), c''(1))$. So the restriction of exp is an isometry.

From Proposition 1.3.8 we know that CAT(0) spaces are uniquely geodesic and geodesics vary continuously with their endpoints. Let $x \in X$. Since X is locally CAT(0), we can choose $r > 0$ such that $B(x, r)$ is CAT(0). Let $a, b \in B(x, r/2)$. Then the unique geodesic in $B(x, r)$ from a to b must already lie in $B(x, r/2)$ since balls in CAT(0) spaces are convex. Any other path in X from a to b which is not contained in $B(x, r)$ must have length at least r and thus cannot be a geodesic because $d(a, b) < r$. It follows that X is locally uniquely geodesic and these geodesics vary continuously with their endpoints. We see that all the assumptions in Lemma 1.4.20 are satisfied, so exp is a covering. From (i) it follows that exp is the universal covering. □

Proposition 1.4.22. *Each two points $p, q \in \widetilde{X}_{x_0}$ are joined by a unique cs-local-geodesic $[0, 1] \to \widetilde{X}_{x_0}$ and these cs-local-geodesics vary continuously with their endpoints.*

Proof. We first prove the claim in the case $p = \widetilde{x}_0$ the constant path at x_0: Let $\widetilde{q}\colon [0, 1] \to \widetilde{X}_{x_0}$ be a cs-local-geodesic from $\widetilde{x}_0$ to q, then $\exp \circ \widetilde{q}$ is a cs-local-geodesic from x_0 to $q(1)$ (since exp is a local isometry). It follows from Lemma 1.4.19 that there is exactly one such cs-local-geodesic, namely q itself. So we must have $\widetilde{q}(s)(1) = q(s)$ for each $s \in [0, 1]$. Since each $\widetilde{q}(s)$ is itself a cs-local-geodesic, by uniqueness again, we must have $\widetilde{q}(s)(t) = q(st)$ for all $s, t \in [0, 1]$. Conversely, with $\widetilde{q}$ so defined, we have that $\exp \circ \widetilde{q} = q$ is a cs-local-geodesic. Hence also $\widetilde{q}$ is a cs-local-geodesic.

Now we claim that every homotopy class (fixing the endpoints) of paths with starting point x_0 contains a unique cs-local-geodesic: Let $c\colon [0, 1] \to X$ be such a path. Let $\widetilde{c}$ be the lift of c along the universal covering map exp with $\widetilde{c}(0) = \widetilde{x}_0$. By covering theory, the paths from x_0 to $c(1)$ homotopic to c are in bijective correspondence to the paths in $\widetilde{X}_{x_0}$ from $\widetilde{x}_0$ to $\widetilde{c}(1)$. Furthermore, since exp is a local isometry, being a cs-local-geodesic is preserved by this correspondence. The claim follows now from the first part of the proof.

Since x_0 is arbitrary, the same conclusion holds for all paths, i.e. each homotopy class (fixing the endpoints) of paths (with arbitrary start point) contains a unique cs-local-geodesic. The same argument as above (in the reverse direction) now shows that each two points in $\widetilde{X}_{x_0}$ are joined by a unique cs-local-geodesic.

By now, we know that $\widetilde{X}_{x_0}$ is complete and locally CAT(0), so Lemma

1.4.19 applied to $\widetilde{X}_{x_0}$ yields that the unique cs-local-geodesics vary continuously with their endpoints. □

At this point we emphasize that $\widetilde{X}_{x_0}$ is *not* a length space in general. However, since $\widetilde{X}_{x_0}$ is locally CAT(0), it is a length space locally. In particular, the metric d on $\widetilde{X}_{x_0}$ and the associated length metric $\bar{d}$ are locally the same. Consequently, the topology on $(\widetilde{X}_{x_0}, \bar{d})$ (henceforth written as $\overline{X}_{x_0}$) is the same as on $\widetilde{X}_{x_0}$. Moreover,

$$\mathrm{Exp}\colon \overline{X}_{x_0} \to X$$

is still a local isometry and the universal covering (topologically). But since $\overline{X}_{x_0}$ is now a length space, it is exactly the universal metric covering from the main theorem (see Remark 1.4.17(ii) and Lemma 1.4.11). So we have to show that $\overline{X}_{x_0}$ is CAT(0).

At first, we mention that the local geodesics in $\overline{X}_{x_0}$ and in $\widetilde{X}_{x_0}$ are the same. So the conclusions of Proposition 1.4.22 still hold. Furthermore, $\overline{X}_{x_0}$ is still complete because of the following facts: The identity $\overline{X}_{x_0} \to \widetilde{X}_{x_0}$ is metric and a homeomorphism and every such map $Y \to Z$ transfers completeness from Z to Y. Last but not least, $\overline{X}_{x_0}$ is locally CAT(0). These observations are used in the following proposition:

Proposition 1.4.23. *Each cs-local-geodesic in $\overline{X}_{x_0}$ is already a cs-geodesic.*

Proof. For $p, q \in \overline{X}_{x_0}$ denote by c_{pq} the unique cs-local-geodesic from p to q. We want to show that $l(c_{pq}) \leq l(\gamma)$ for each curve γ from p to q. Since $\overline{X}_{x_0}$ is a length space it follows $d(p,q) = l(c_{pq})$ and hence that c_{pq} is a cs-geodesic.

So let $\gamma\colon [0,1] \to \overline{X}_{x_0}$ be any rectifiable curve. We want to show $l(c_{\gamma(0)\gamma(t)}) \leq l(\gamma|_{[0,t]})$ for each $t \in [0,1]$ with a connectedness argument: Define

$$S := \left\{ t_0 \in [0,1] \mid \forall_{t \leq t_0}\ l(c_{\gamma(0)\gamma(t)}) \leq l(\gamma|_{[0,t]}) \right\}$$

Choose a ball around $\gamma(0)$ such that the induced metric on the ball is CAT(0). This shows that S contains an open neighborhood around 0. Furthermore, $S \subset [0,1]$ is closed. If we show that it is also open, then $S = [0,1]$ and we are done. Let $t_0 \in S$ with $t_0 > 0$. By Lemma 1.4.19 we know that there is $\epsilon > 0$ small enough such that when $t' \in (t_0 - \epsilon, t_0 + \epsilon)$ we have

$$\begin{aligned} l\big(c_{\gamma(0)\gamma(t')}\big) &\leq l\big(c_{\gamma(0)\gamma(t_0)}\big) + d\big(\gamma(t_0), \gamma(t')\big) \\ &\leq l\big(\gamma|_{[0,t_0]}\big) + l\big(\gamma|_{[t_0,t']}\big) \\ &= l\big(\gamma|_{[0,t']}\big) \end{aligned}$$

Consequently, $S \subset [0,1]$ is open. □

We know now that $\overline{X}_{x_0}$ is even uniquely geodesic and these geodesics vary continuously with their endpoints. Proposition 1.4.24 below applied to $Y = \overline{X}_{x_0}$ will complete the proof.

1.4.4 Alexandrov's Patchwork

Proposition 1.4.24. *Let Y be a metric space of curvature ≤ 0 and assume that points are joined by unique geodesics which vary continuously with their endpoints. Then Y is CAT(0).*

Let p, q_0, q_1 be three distinct points in Y and consider the unique geodesic triangle Δ formed by these three points. We want to show that the angles in this geodesic triangle are smaller or equal to the corresponding angles in a comparison triangle in $\mathbb{E}^2$ (henceforth called the angle condition). This proves the proposition because of Proposition 1.3.4(iii). Let $q\colon [0,1] \to Y$ be the unique cs-geodesic from q_0 to q_1 and for each $s \in [0,1]$ let $c_s\colon [0,1] \to Y$ be the unique cs-geodesic from p to $q(s)$. Since geodesics in Y vary continuously with their endpoints, the function $c\colon [0,1] \times [0,1] \to Y$ given by $c(s,t) = c_s(t)$ is continuous. Since the image of c is compact and Y is locally CAT(0), we can cover $\operatorname{im} c$ by finitely many balls on which the induced metric is CAT(0). Hence we find

$$0 = s_0 < s_1 < \ldots < s_k = 1$$
$$0 = t_0 < t_1 < \ldots < t_l = 1$$

such that each $c\big([s_i, s_{i+1}] \times [t_j, t_{j+1}]\big)$ is contained in such a CAT(0) ball.

Case 1: Assume first that $k = 2$ and $l = 1$. Denote by Δ_1 the triangle with vertices $p, q_0, c_{s_1}(1)$ and by Δ_2 the triangle with vertices $p, c_{s_1}(1), q_1$. Then Δ_1 and Δ_2 satisfy the angle condition (since they both lie in a CAT(0) ball by construction). By observing comparison triangles for Δ_1 and Δ_2 in $\mathbb{E}^2$ which are joined at a common edge and by applying Alexandrov's Lemma, we see that also the triangle Δ satisfies the angle condition (compare with the proof of (iii)$\Rightarrow$(v) in Proposition 1.3.4). This observation is applied several times in the following.

Case 2: Now assume the case $k = 1$ and $l > 1$ and join the points $c(s_i, t_j)$ with geodesics as depicted in the following picure:

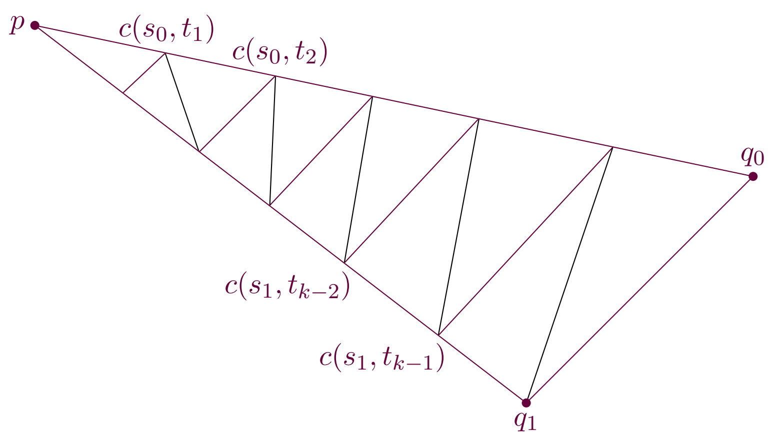

Each triangle in this picture lies in a CAT(0) ball and thus satisfies the angle condition. Now we apply Alexandrov's Lemma (just as in the preceding paragraph) successively to the following pairs of triangles to show that the union of the two triangles again satisfies the angle condition:

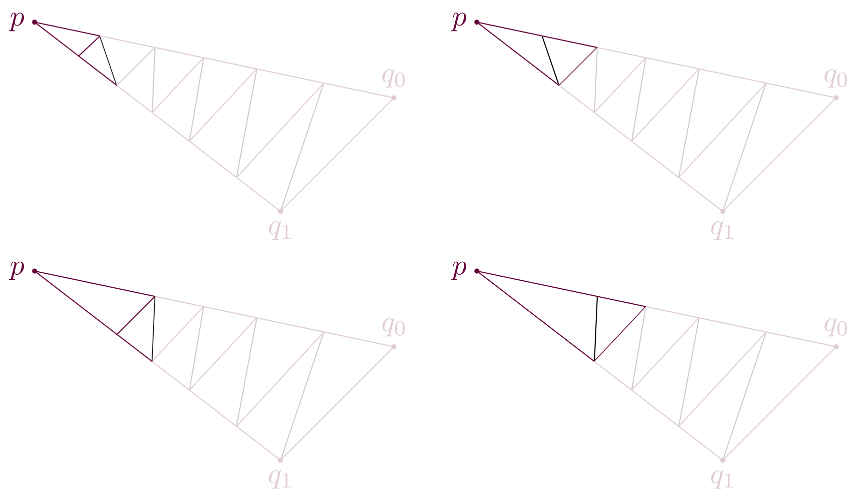

We can repeat this until we have shown that the big triangle Δ satisfies the angle condition.

Case 3: Now we assume the general case $k > 1$ and $l > 1$. Each triangle with vertices $q, c_{s_i}(1), c_{s_{i+1}}(1)$ satisfies the angle condition by the argument in the preceding paragraph. Let Δ_i be the triangle with vertices $q, q_0, c_{s_i}(1)$ for $i = 1, \ldots, k$. By applying Alexandrov's Lemma, we deduce that Δ_2 satisfies the angle condition. Then, by applying it again, we deduce that Δ_3 satisfies the angle condition. Repeating this, we end up with Δ satisfying the angle condition and we are done.

1.4.5 Proofs of the lemmata

We sketch the proofs of the remaining lemmata.

Proof of Lemma 1.4.19. Since $c([0,1])$ is compact and X is locally CAT(0) and locally complete, we find $\epsilon > 0$ small enough so that $\bar{B}(c(t), 2\epsilon)$ is CAT(0) and complete. This will be the ϵ as required in the lemma.

Let c', c'' be as in the lemma. Then by the choice of ϵ, the function $t \mapsto d(c'(t), c''(t))$ is locally convex (see Lemma 1.3.7) und therefore convex. If additionally $c'(0) = c''(0)$ and $c'(1) = c''(1)$ then it follows $c' = c''$. This is the uniqueness statement in the lemma.

Now only assume $c'(0) = c''(0)$. For small $t > 0$, the paths $c'|_{[0,t]}$ and

$c''|_{[0,t]}$ are geodesics. Hence, we obtain

$$\begin{aligned}
tl(c'') &= l(c''|_{[0,t]}) \\
&= d(c''(0), c''(t)) \\
&= d(c'(0), c''(t)) \\
&\leq d(c'(0), c'(t)) + d(c'(t), c''(t)) \\
&\leq l(c'|_{[0,t]}) + td(c'(1), c''(1)) \\
&= tl(c') + td(c'(1), c''(1))
\end{aligned}$$

where we have also used the convexity of $t \mapsto d(c'(t), c''(t))$. Consequently, $l(c'') \leq l(c') + d(c'(1), c''(1))$.

Assume for the moment that we have already shown the existence statement in the lemma (to be done below). Then there is a unique cs-local-geodesic $\gamma\colon [0,1] \to X$ which is ϵ-near to c and joining $\bar{x}$ to y. Then by the preceding paragraph:

$$\begin{aligned}
l(\bar{c}) &\leq l(\gamma) + d(\bar{y}, y) \\
&\leq l(c) + d(\bar{x}, x) + d(\bar{y}, y)
\end{aligned}$$

Remains to construct the cs-local-geodesic $\bar{c}\colon [0,1] \to X$ from $\bar{x}$ to $\bar{y}$ and ϵ-near to c: Observe the following property named $P(A)$ for $A > 0$:

> For all $a, b \in [0,1]$ with $0 < b - a \leq A$ and all $p \in B(c(a), \epsilon)$, $q \in B(c(b), \epsilon)$ there is a cs-local-geodesic $\gamma\colon [a,b] \to X$ from p to q and ϵ-near to c, i.e. $d(\gamma(t), c(t)) < \epsilon$ for all $t \in [a,b]$.

We know $P(A)$ is true for $A = \epsilon/l(c)$, because then $d(c(a), c(b)) \leq \epsilon$ and the points p, q lie in $B(c(a), 2\epsilon)$ which is CAT(0) and therefore uniquely geodesic. The unique geodesic from p to q is ϵ-near to c since the metric on the ball is also convex. We now sketch how we can prove $P(3A/2)$ if we assume that $P(A)$ holds. This then completes the proof.

So let $a, b \in [0,1]$ such that $0 < b - a \leq 3A/2$. Divide $[a,b]$ into three equal parts with endpoints $a < a' < b' < b$. Set $p_0 = c(a')$ and $q_0 = c(b')$. Using $P(A)$ we construct cs-local-geodesics $\alpha_0\colon [a, b'] \to X$ from p to q_0 and $\beta_0\colon [a', b] \to X$ from p_0 to q. Define $p_1 = \alpha_0(a')$ and $q_1 = \beta_0(b')$. Continuing inductively, we construct cs-local-geodesics α_n and β_n which are point-wise Cauchy. Thus they converge pointwise in the complete balls $\bar{B}(c(t), \epsilon)$. It turns out that they converge uniformly to two cs-local-geodesics which coincide on $[a', b']$. The union of these two cs-local-geodesic is the cs-local-geodesic we were looking for (see Figure 1.10). $\square$

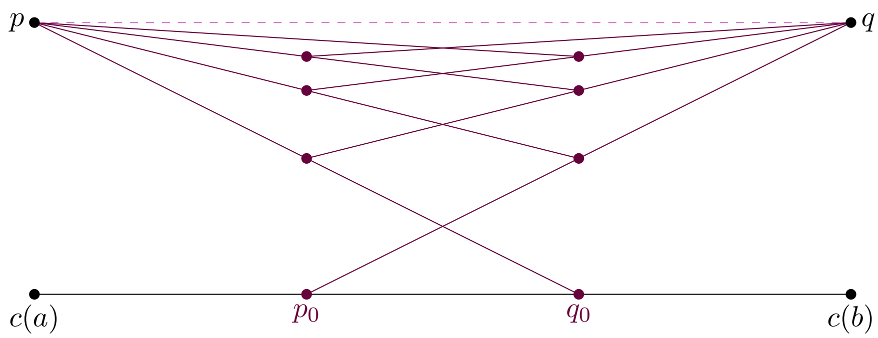

Figure 1.10: Building bridges for Cartan-Hadamard

Proof of Lemma 1.4.20. Step 1: First we show that p satisfies the unique path lifting property, i.e. if $\gamma\colon [0,1] \to X$ is path and $y \in Y$ with $p(y) = \gamma(0)$ is given, then there is a unique path $\widetilde{\gamma}\colon [0,1] \to Y$ with $p \circ \widetilde{\gamma} = \gamma$ and $\widetilde{\gamma}(0) = y$. Uniqueness follows from p being a local homeomorphism, because then the set $S \subset [0,1]$ of points where two lifts of γ coincide is non-empty, open and closed. Let $a \in (0,1]$ such that the lift exists on $[0,a)$. Such an a always exists, since p is a local homeomorphism. Let x_n be a sequence in $[0,a)$ converging to a. Then we have (since p is a local isometry):

$$d\big(\widetilde{\gamma}(x_n), \widetilde{\gamma}(x_m)\big) \le l\big(\widetilde{\gamma}|_{[x_n,x_m]}\big) = l\big(\gamma|_{[x_n,x_m]}\big)$$

Hence $\widetilde{\gamma}(x_n)$ is a Cauchy sequence which converges in Y and $\widetilde{\gamma}$ can be extended to $[0,a]$. Since p is a local homeomorphism, it can be extended beyond a. Thus, a lift has to exist on the whole interval $[0,1]$. Since X is path connected, it also follows from the (unique) path lifting property that p is surjective.

Step 2: Let $x \in X$ and $r > 0$ so small such that between any two points in $B(x,r)$ there is a unique geodesic in X which already lies in $B(x,r)$ and geodesics in $B(x,r)$ vary continuously with their endpoints. We claim that $B(x,r)$ is homeomorphically covered by disjoint open subsets of Y. Let $y \in p^{-1}(x)$. Define $s_y\colon B(x,r) \to Y$ by $s_y(z) = \widetilde{\gamma}_z(1)$ where γ_z is the unique geodesic from x to z and $\widetilde{\gamma}_z$ is the unique lift with start point y. Note that, by uniqueness of path lifts, $\operatorname{im} s_y \cap \operatorname{im} s_{y'} = \emptyset$ if $y \neq y'$. In step 3 we will sketch a proof that s_y is continuous for each y. It then follows that the maps s_y are homeomorphisms onto their open images and the claim follows.

Step 3: So let $y \in p^{-1}(x)$ and $z \in B(x,r)$. The image of $\widetilde{\gamma}_z$ can be covered by finitely many balls $B_1, \dots, B_k$ such that p maps each B_i isometrically onto balls in X. Then the image of γ_z is covered by the balls pB_i. If z' is near to z, then the geodesic $\gamma_{z'}$ is uniformly near to γ_z since geodesics in $B(x,r)$ vary continuously with their endpoints. In particular, we can assume that $\gamma_{z'}$ lies in the union of the balls pB_i. We then can use the local (isometric) inverses $(p|_{B_i})^{-1}$ to describe the lift $\widetilde{\gamma}_{z'}$ and conclude that $\widetilde{\gamma}_{z'}(1) = s_y(z')$ has to be near to $\widetilde{\gamma}_z(1) = s_y(z)$ which means that s_y is continuous at z. □

1.5 Isometries of CAT(0) spaces

1.5.1 Types of Isometries

In this section, we will first study individual isometries of CAT(0) spaces and show that the can be naturally divided into three classes, *elliptic, hyperbolic* and *parabolic* isometries,and study the different behaviour of these classes. We also prove several results that were already mentioned last semester [14, Section 7 and 8].

Definition 1.5.1. Let X be a metric space.

(i) Let $\gamma \in \mathrm{Isom}(X)$ be an isometry. We call the function

$$\begin{aligned} d_\gamma \colon X &\longrightarrow \mathbb{R} \\ x &\longmapsto d(x, \gamma(x)) \end{aligned}$$

the *displacement function of* γ.

(ii) We call the function

$$\begin{aligned} |\cdot| \colon \mathrm{Isom}(X) &\longrightarrow \mathbb{R}_{\geq 0} \\ \gamma &\longmapsto \inf_{x \in X} d_\gamma(x) \end{aligned}$$

the *translation length*.

(iii) Let $\gamma \in \mathrm{Isom}(X)$ be an isometry. We define

$$\mathrm{Min}(\gamma) := \{x \in X \mid d_\gamma(x) = |\gamma|\};$$

and call γ *semi-simple* if $\mathrm{Min}(\gamma)$ is not empty.

Definition 1.5.2. Let X be a metric space and $\gamma \in \mathrm{Isom}(X)$ an isometry.

(i) We call γ *elliptic*, if it fixes a point, i.e., if γ is semi-simple and $|\gamma| = 0$.

(ii) We call γ *hyperbolic* (or *axial*) if γ is semi-simple and $|\gamma| > 0$.

(iii) We call γ *parabolic* if $\mathrm{Min}(\gamma)$ is empty, i.e., if γ is not semi-simple.

Clearly any element in $\mathrm{Isom}(X)$ falls into exactly one of these classes.

Proposition 1.5.3. *Let X be a metric space and let $\gamma \in \mathrm{Isom}(X)$ be an isometry.*

(i) The set $\mathrm{Min}(\gamma)$ *is* γ*-invariant.*

(ii) Let $\alpha \in \mathrm{Isom}(X)$ *be an isometry. Then* $|\alpha\gamma\alpha^{-1}| = |\gamma|$ *and* $\mathrm{Min}(\alpha\gamma\alpha^{-1}) = \alpha \cdot \mathrm{Min}(\gamma)$*. In particular, if two isometries are conjugated in* $\mathrm{Isom}(X)$*, then they belong to the same class in Definition 1.5.2 and their translation length is equal, i.e., translation length is a class function.*

(iii) If the metric on X *is convex, e.g., if* X *is a CAT(0) space, then* d_γ *is also convex and* $\mathrm{Min}(\gamma)$ *is a closed convex subspace of* X*.*

(iv) Let X *be CAT(0). If* $C \subset X$ *is a non-empty complete convex,* γ*-invariant subset of* X*, then* $|\gamma|_C| = |\gamma|$ *and* $\mathrm{Min}(\gamma)$ *is non-empty if and only if* $\mathrm{Min}(\gamma|_C) = \mathrm{Min}(\gamma) \cap C$ *is non-empty.*

Proof. (i), (ii) and (iii) are obvious. Regarding (iv), consider the projection $p\colon X \longrightarrow C$ onto C. Then for any $x \in X$, we have

$$d_\gamma(x) = d(x, \gamma x) \geq d(p(x), p(\gamma x)) = d(p(x), \gamma p(x)) = d_\gamma(p(x)). \qquad \square$$

Proposition 1.5.4. *Let* X *be a complete CAT(0) space and* $\gamma \in \mathrm{Isom}(X)$ *an isometry. Then* γ *is elliptic if and only if it has a bounded orbit. Furthermore, if there exists an* $n \in \mathbb{N}_{>0}$ *such that* γ^n *is elliptic, then also* γ *is elliptic.*

Proof. The first statement is the Bruhat-Tits fixed point theorem, Corollary 1.3.12. For the second part, notice that if the γ^n-orbit of x is bounded, then so is the γ-orbit of x. $\square$

Proposition 1.5.5. *Let* X *be a CAT(0) space and* $\gamma \in \mathrm{Isom}(X)$ *an isometry. Then* γ *is hyperbolic if and only if there exists a geodesic line* $c\colon \mathbb{R} \longrightarrow X$ *that is translated non-trivially by* γ*, i.e. there exists an* $a_c \in \mathbb{R}_{>0}$ *such that* $\gamma \cdot c(t) = c(t + a_c)$ *for all* $t \in \mathbb{R}$*. For any such line* c*, we have that* $a_c = |\gamma|$*.*

Proof. If there exists such a c, then by Proposition 1.5.3(iv), γ is hyperbolic and $a_c = |\gamma|$. Suppose conversely that γ is hyperbolic and pick $x \in \mathrm{Min}(\gamma)$. Consider the family $([\gamma^n x, \gamma^{n+1} x])_{n\in\mathbb{Z}}$ of geodesic segments in X. We will show that the concatenation of these segments is a geodesic line, since then it is clearly translated by γ. Since X is CAT(0), local geodesics are global geodesics, so it suffices to show that we can glue two segments together. Let $m \in X$ be the midpoint of $[x, \gamma x]$. Then γm is the midpoint of $[\gamma x, \gamma^2 x]$. The concatenation of these two intervals is a geodesic segment if and only if $d(m, \gamma m) = d(m, \gamma x) + d(\gamma x, \gamma m)$. Because $\mathrm{Min}(\gamma)$ is convex and contains x and γx, we have that m is in $\mathrm{Min}(\gamma)$ and thus

$$d(m, \gamma m) = |\gamma| = d(x, \gamma x) = d(x, m) + d(m, \gamma x) = d(m, \gamma x) + d(\gamma x, \gamma m). \ \square$$

Definition 1.5.6. Let X be a CAT(0) space, $\gamma \in \mathrm{Isom}(X)$ a hyperbolic isometry and $c\colon \mathbb{R} \longrightarrow X$ a geodesic line that is translated by γ. We call $c(\mathbb{R}) \subset X$ an *axis of* γ *in* X.

Proposition 1.5.7. *Let X be a CAT(0) space and $\gamma \in \mathrm{Isom}(X)$ a hyperbolic isometry. Then:*

(i) The axes of γ are parallel to each other, i.e. if c and c' are two geodesic lines translated by γ, then $t \longmapsto d(c(t), c'(t))$ is constant.

(ii) The union of the axes of γ is $\mathrm{Min}(\gamma)$.

(iii) There exists a metric space Y such that $\mathrm{Min}(\gamma)$ is isometric to $Y \times \mathbb{R}$, and under this identification, γ acts on $\mathrm{Min}(\gamma)$ via $\gamma(y,t) = (y, t+|\gamma|)$ for all $y \in Y$ and $t \in \mathbb{R}$.

(iv) Let $\alpha \in \mathrm{Isom}(X)$ be an isometry that commutes with γ. Then the restriction of α to $\mathrm{Min}(\gamma) \cong Y \times \mathbb{R}$ splits as $\alpha = (\alpha', \alpha'')$, where α' is an isometry of Y and α'' is a translation on $\mathbb{R}$.

Proof.

(i) Let c and c' be two geodesic lines that are translated by γ. Then for all $t \in \mathbb{R}$, we have

$$d(c(t+|\gamma|), c'(t+|\gamma|)) = d(\gamma \cdot c(t), \gamma \cdot c'(t)) = d(c(t), c'(t)).$$

Therefore, the convex function $t \longmapsto d(c(t), c'(t))$ is bounded and thus constant, so c and c' are parallel.

(ii) Every $x \in \mathrm{Min}(\gamma)$ is contained in the axis corresponding to the family $(\gamma^n \cdot x)_{n\in\mathbb{Z}}$.

(iii) Since $\mathrm{Min}(\gamma)$ is the union of all lines parallel to an axis, we can apply the product decomposition theorem, Corollary 1.3.16, to see that it is isometric to a product $Y \times \mathbb{R}$ and each $\{y\} \times \mathbb{R}$ corresponds to an axis.

(iv) Since α commutes with γ, it leaves $\mathrm{Min}(\gamma)$ invariant and maps axes to axes. Hence under the identification $\mathrm{Min}(\gamma) \cong Y \times \mathbb{R}$, it splits as $\alpha = (\alpha', \alpha'')$. Since α'' is an isometry of $\mathbb{R}$ commuting with the translation induced by γ, it must be a translation. □

Proposition 1.5.8. *Let X be a complete CAT(0) space and $\gamma \in \mathrm{Isom}(X)$. If γ^n is hyperbolic for some $n \in \mathbb{N}$, so is γ.*

Proof. By Proposition 1.5.7, $\mathrm{Min}(\gamma^n)$ splits as a product $Y \times \mathbb{R}$. Since γ commutes with γ^n, it restricts under this identification to an isometry (γ', γ'') of $Y \times \mathbb{R}$. Since Y is a closed convex subset of a complete CAT(0) space, it is also complete and CAT(0). Because $\gamma^n|_Y = \mathrm{id}_Y$, we have that γ' is periodic, hence by Bruhat-Tits has a fixed point $y \in Y$. Then $\{y\} \times \mathbb{R}$ is an axis of γ and γ is hyperbolic. □

The next proposition shows, that in groups acting properly and cocompactly on proper spaces, parabolic elements actually do not occur and the set of translation lengths is discrete:

Proposition 1.5.9. *Let X be a proper metric space and let G be a group acting properly and cocompactly by isometries on X.*

(i) The action of G on X is semi-simple, i.e., there are no parabolic elements in G.

(ii) The set $\{|\gamma| \mid \gamma \in G\}$ of translation lengths of G is discrete in $\mathbb{R}$.

Proof. Let $K \subset X$ be a compact subset with $X = \bigcup_{g \in G} g \cdot K$.

(i) Let $(x_n)_{n\in\mathbb{N}}$ be a sequence in X such that $(d_\gamma(x_n))_{n\in\mathbb{N}}$ converges to $|\gamma|$ and let $(\gamma_n)_{n\in\mathbb{N}}$ be a sequence in G such that $y_n := \gamma_n \cdot x_n$ is in K for every $n \in \mathbb{N}$. Then for every $n \in \mathbb{N}$, we have

$$d(\gamma_n \cdot \gamma \cdot \gamma_n^{-1} \cdot y_n, y_n) = d(\gamma \cdot x_n, x_n)$$

hence $(d(\gamma_n \cdot \gamma \cdot \gamma_n^{-1} \cdot y_n, y_n))_{n\in\mathbb{N}}$ converges to $|\gamma|$. In particular, for every $x \in K$, the sequence $(d(\gamma_n \cdot \gamma \cdot \gamma_n^{-1} \cdot x, x))_{n\in\mathbb{N}}$ is bounded. Therefore, by the properness of X and of the action, there exists an $\alpha \in G$ such that for infinitely many $n \in \mathbb{N}$, we have $\gamma_n \cdot \gamma \cdot \gamma_n^{-1} = \alpha$. After passing to a subsequence, we can thus assume that $\gamma_n \cdot \gamma \cdot \gamma_n^{-1} = \alpha$ for all $n \in \mathbb{N}$. Since K is compact, we can also assume that $(y_n)_{n\in\mathbb{N}}$ converges to some $y \in K$. Now d_γ assumes its minimum in $\gamma_0^{-1} \cdot y$, because

$$\begin{aligned} d(\gamma \cdot \gamma_0^{-1} \cdot y, \gamma_0^{-1} \cdot y) &= d(\gamma_0 \cdot \gamma \cdot \gamma_0^{-1} \cdot y, y) \\ &= d(\alpha \cdot y, y) \\ &= \lim_{n\to\infty} d(\alpha \cdot y_n, y_n) \\ &= \lim_{n\to\infty} d(\gamma_n \cdot \gamma \cdot \gamma_n^{-1} \cdot y_n, y_n) \\ &= \lim_{n\to\infty} d(\gamma \cdot x_n, x_n) \\ &= |\gamma|. \end{aligned}$$

(ii) Let $(\gamma_n)_{n\in\mathbb{N}}$ be a sequence in G such that $(|\gamma_n|)_{n\in\mathbb{N}}$ converges. By the first part, we can choose for each $n \in \mathbb{N}$ an $x_n \in \mathrm{Min}(\gamma_n)$. After replacing γ_n with suitable conjugates, we can also assume that $x_n \in K$ for all $n \in \mathbb{N}$. But then we have that for any $x \in K$, the sequence $d(\gamma_n \cdot x_n, x)$ is bounded, so by the properness of the action and of X, we have that $\{\gamma_n \mid n \in \mathbb{N}\}$ is finite and therefore $(|\gamma_n|)_{n\in\mathbb{N}_{\geq m}}$ is constant for $m \in \mathbb{N}$ sufficiently big. So $\{|\gamma| \mid \gamma \in G\}$ is discrete in $\mathbb{R}$. □

In some very important examples, one thus only has to consider hyperbolic elements:

Example 1.5.10.

(i) Let M be a compact Riemannian manifold. Then the fundamental group $\pi_1(M)$ acts via deck transformations on the Riemannian universal cover $\widetilde{M}$. With respect to this action, any element in $\pi_1(M) \setminus \{1\}$ is hyperbolic.

(ii) More generally, let $p\colon Y \longrightarrow X$ be a normal metric covering between path connected and locally path connected length spaces, and assume that X is compact. Then the action of $\pi_1(X)$ on Y via deck transformations is free (i.e. there are no non-trivial elliptic elements) and satisfies the conditions of Proposition 1.5.9.

As a first application we get a result relating extensions with Abelian centres and actions on CAT(0) spaces:

Theorem 1.5.11. *Let X be a CAT(0) space and let G be a finitely generated group acting isometrically on X. Let $A \leq G$ be a central subgroup of G such that the induced action on X is hyperbolic and faithful and assume that $A \cong \mathbb{Z}^n$ for some $n \in \mathbb{N}$. Then G contains a finite index subgroup that splits as $A \times B$ for some subgroup $B \leq G$.*

Proof. We prove the result by induction on $n \in \mathbb{N}$. For $n = 0$, there is nothing to show. So assume that $n \in \mathbb{N}_{>0}$. Let $b \in A$ be a non-trivial element. Since A is central, G preserves $\mathrm{Min}(b) \cong Y \times \mathbb{R}$ and each element in G restricted to $\mathrm{Min}(b)$ splits as (g', g''). This induces a map

$$\begin{aligned} G &\longrightarrow \mathrm{Isom}(\mathbb{R}) \\ g &\longmapsto g''. \end{aligned}$$

The image of this map is a subgroup of translations that is finitely generated, since G is finitely generated, hence it induces a surjection $G \longrightarrow \mathbb{Z}^m$ for some $m \in \mathbb{N}_{>0}$ and the image of b and hence of A under this map is non-trivial. Hence we can compose it with a projection $\mathbb{Z}^m \longrightarrow \mathbb{Z}$ to get a surjection $\varphi\colon G \longrightarrow \mathbb{Z}$, such that $\varphi(A)$ is non-trivial. Let $a \in A$ be an element such that $\varphi(A) = \varphi(a) \cdot \mathbb{Z}$. Let H' be the kernel of the map $G \longrightarrow \mathbb{Z}/\varphi(a)\cdot\mathbb{Z}$ induced by φ. Then $H' \leq G$ is a finite index subgroup containing A and we get a split exact sequence $1 \longrightarrow \ker\varphi \longrightarrow H' \longrightarrow \varphi(a)\cdot\mathbb{Z} \longrightarrow 1$, where a splitting is given by mapping the generator $\varphi(a)$ to a. Hence H' splits as a semi-direct product of $\ker\varphi$ and $\langle a \rangle$ and since a is central, we have $H' = \ker\varphi \times \langle a \rangle$. Set $A' := \ker\varphi \cap A$. Since H' is finitely generated as a finite index subgroup of G, so is $\ker\varphi$. By induction, there exists a finite index subgroup $H'' \leq \ker\varphi$, such that $H'' = A' \times B$ for some subgroup $B \leq G$. Then $H := H'' \times \langle a \rangle$ is a finite index subgroup of G and $H = A \times B$.

□

Corollary 1.5.12. Let $1 \longrightarrow \mathbb{Z}^n \longrightarrow G \longrightarrow Q \longrightarrow 1$ be a central extension and assume that G acts faithfully and hyperbolically on a CAT(0) space X. Then G is virtually a product $\mathbb{Z}^n \times Q'$ for some group Q'.

1.5.2 Examples

In this section, we will study the isometries of $\mathbb{H}^n$ and $\mathbb{E}^n$ under the point of view developed in the previous section.

Recall that $\mathrm{Isom}(\mathbb{E}^n) \cong \mathbb{R}^n \rtimes O(n)$ and that any isometry of $\mathbb{R}^n$ can be written as $x \longmapsto A \cdot x + b$ with $A \in O(n)$ and $b \in \mathbb{R}^n$. Clearly rotations and affine reflections are examples of elliptic isometries and translations are examples of hyperbolic isometries.

Example 1.5.13. Consider the isometry $\gamma \in \mathrm{Isom}(\mathbb{E}^2)$ given by $\gamma(x, y) = (x + 1, -y)$, i.e., γ translates by 1 the first coordinate and then reflects through the first axis. Then γ is a hyperbolic isometry and $\mathrm{Min}(\gamma) = \mathbb{R} \times \{0\}$.

The next proposition shows that this example is typical, every Euclidean isometry splits into a translation and an elliptic part.

Proposition 1.5.14 (Isometries of $\mathbb{E}^n$)**.** *For any $n \in \mathbb{N}$, every isometry in $\mathrm{Isom}(\mathbb{E}^n)$ is semi-simple. If $\gamma \in \mathrm{Isom}(\mathbb{E}^n)$ is hyperbolic, then $\mathrm{Min}(\gamma)$ is an affine subspace of $\mathbb{E}^n$ and γ splits on $\mathbb{E}^n = \mathrm{Min}(\gamma) \times \mathrm{Min}(\gamma)^\perp$ (where $\mathrm{Min}(\gamma)^\perp$ is a choice of an orthogonal affine complement of $\mathrm{Min}(\gamma)$) as $\gamma = (\gamma', \gamma'')$ where γ' is a non-trival translation and γ'' an elliptic isometry of $\mathrm{Min}(\gamma)^\perp$.*

Sketch of Proof. If γ is not elliptic, then A has eigenvalue 1. If v is a 1-eigenvector of A, then γ splits on $\langle v \rangle \times \langle v \rangle^\perp$ as (γ_0, γ_1) where γ_0 is a translation and γ_1 an isometry of $\langle v \rangle^\perp$. The claim follows now by induction on n. ☐

Example 1.5.15. We present some typical examples of isometries of $\mathbb{H}^2$:

(i) For any $\lambda \in (0, 1) \cup (1, \infty)$, the map

$$\begin{aligned} \gamma_\lambda \colon \mathrm{HS}^n &\longrightarrow \mathrm{HS}^n \\ x &\longmapsto \lambda \cdot x. \end{aligned}$$

is a hyperbolic isometry of the half plane model of $\mathbb{H}^2$. Furthermore, we have $|\gamma_\lambda| = |\log \lambda|$ and $\mathrm{Min}(\gamma_\lambda)$ is the vertical straight line segment $\{0\} \times \mathbb{R}_{>0}$.

(ii) The map

$$\begin{aligned} \mathrm{HS}^2 &\longrightarrow \mathrm{HS}^2 \\ x &\longmapsto x + (1, 0). \end{aligned}$$

is a parabolic isometry of the half plane model for $\mathbb{H}^2$ (since it clearly does not fix a point and $\inf_{n \in \mathbb{N}}(d^2_{\mathbb{H}}(0, n), (1, n)) = 0$).

(iii) For $\lambda \in [0, 2 \cdot \pi)$, let φ_λ be the restriction to B^2 of the rotation in $\mathbb{R}^2$ around 0 of angle λ. Then φ_λ is an elliptic isometry of the Poincaré ball model for $\mathbb{H}^2$.

Last semester, we have introduced the Gromov boundary for hyperbolic spaces [14, Section 8.3] and have mentioned a classification of isometries via fixed points at the boundary. Before we can prove this result, we show that for the Poincaré models, the Gromov boundary canonically coincides with the topological boundary:

Proposition 1.5.16. *Fix $n \in \mathbb{N}_{>0}$. Then the Gromov boundary $\partial\mathbb{H}^n$ is homeomorphic to S^{n-1}.*

Proof. For $s \in S^{n-1}$, let $r_s \colon \mathbb{R}_{\geq 0} \longrightarrow B^n$ be the geodesic ray in the Poincaré ball model given by the intersection of the straight line through s and 0 with B^n. Then the map

$$\begin{aligned} S^{n-1} &\longrightarrow \partial\mathbb{H}^n \\ s &\longmapsto [r_s] \end{aligned}$$

is a homeomorphism, since: It is surjective since any geodesic ray issuing from 0 is of the form r_s for some $s \in S^{n-1}$. It is injective since for $s_0 \neq s_1 \in S^{n-1}$, we have $\lim_{t\to\infty}(d(r_{s_0}(t), r_{s_1}(t)) = \infty$. It is a homeomorphism since a sequence of points $(s_n)_{n\in\mathbb{N}}$ converges in S^{n-1} if and only if the sequence $(r_{s_n})_{n\in\mathbb{N}}$ converges uniformly on compact subsets of $\mathbb{R}_{\geq 0}$. □

Remark 1.5.17. We set now $\overline{\mathbb{H}}^n := B^n \cup S^{n-1}$. We have seen in Proposition 1.2.28 that the Cayley transformation extends to a homeomorphism from $B^n \cup S^{n-1}$ to $\mathrm{HS}^n \cup \mathbb{R}^n \times \{0\} \cup \{\infty\}$ preserving the boundary, hence we can change between these models if suitable. Any isometry $\gamma \in \mathrm{Isom}(\mathbb{H}^n)$ is in the Poincaré model given as a product of inversions through spheres orthogonal to S^{n-1}, so there is a natural extension of γ to a homeomorphism $\overline{\gamma}$ of $\overline{\mathbb{H}}^n$. We say that γ fixes a point in $\partial\mathbb{H}^n$ if $\overline{\gamma}$ does.

As we have mentioned last semester, via the Gromov boundary one can intrinsically define a natural compactification $\overline{\mathbb{H}}^n$ of $\mathbb{H}^n$ using only the metric on $\mathbb{H}^n$ without any reference to an embedding. We will stick to the concrete point of view here and just remark that the intrinsic and extrinsic definitions coincide for the Poincaré models.

Proposition 1.5.18. *Fix $n \in \mathbb{N}$ and let $\gamma \in \mathrm{Isom}(\mathbb{H}^n)$ be an isometry. Then:*

(i) The isometry γ is hyperbolic if and only if it fixes exactly two points in $\partial\mathbb{H}^n$ and none in $\mathbb{H}^n$.

(ii) The isometry γ is parabolic if and only if it fixes exactly one point in $\partial\mathbb{H}^n$ and none in $\mathbb{H}^n$.

Proof. Since $\overline{\gamma}\colon \overline{\mathbb{H}}^n \longrightarrow \overline{\mathbb{H}}^n$ is a homeomorphism of the n-ball, by the Brouwer-Fixed-Point Theorem [12, Corollary 2.15], it fixes at least one point in $\overline{\mathbb{H}}^n$, so a non-elliptic isometry fixes at least one point in $\partial\mathbb{H}^n$. If it fixes at least two different points $a, b \in \partial\mathbb{H}^n$, then it leaves invariant the geodesic line corresponding to the circle through a and b orthogonal to $\partial\mathbb{H}^n$ and is thus hyperbolic. Conversely, if γ is hyperbolic, it fixes the two endpoints in $\partial\mathbb{H}^n$ of the orthogonal circle corresponding to its axis. Assume now that γ fixes more than two points in $\partial\mathbb{H}^n$, in particular it is hyperbolic. We use the half space model so that one of the fixed points is ∞. Let $a \neq b \in \partial\mathbb{H}^n$ be two other fixed points. Then γ leaves invariant the vertical lines through a and b, so these lines are axis, but they are not parallel, contradicting Proposition 1.5.7. □

Now we study products of two reflections in $\mathrm{Isom}(\mathbb{H}^2)$ (those are exactly the orientation preserving isometries of $\mathbb{H}^2$ in the sense of topology or differential geometry).

Proposition 1.5.19. *Let $\gamma \in \mathrm{Isom}(\mathbb{H}^2) \setminus \{\mathrm{id}_{\mathbb{H}^2}\}$ be a non-trivial product of two reflections and let C and D be the corresponding generalised circles. Then (Figure 1.11):*

- *(i) The isometry γ is elliptic if and only if C and D intersect in $\mathbb{H}^2$.*
- *(ii) The isometry γ is parabolic if and only if C and D intersect in $\partial\mathbb{H}^2$.*
- *(iii) The isometry γ is hyperbolic if and only if C and D do not intersect.*

Proof. We use the half space model for $\mathbb{H}^2$ so that D corresponds to a vertical line $\{s\} \times \mathbb{R}_{>0}$ and C to the circle of radius 1 with centre 0 (this we can assume after conjugating with an isometry). Clearly 0 and ∞ are not fixed by $\overline{\gamma}$. We have for all $(a, 0) \in \partial\mathbb{H}^2 \setminus \{0, \infty\}$

$$\overline{\gamma}(a, 0) = I_D(I_C(a, 0)) = I_D(1/a, 0) = (2 \cdot s - 1/a, 0).$$

Hence the points in $\partial\mathbb{H}^2$ fixed by $\overline{\gamma}$ are exactly the solutions of the equation

$$(a - s)^2 = s^2 - 1.$$

But this equation has no solution if and only if $|s| < 1$, i.e., if and only if C and D intersect in $\mathbb{H}^2$, has exactly one solution if and only if $|s| = 1$, i.e., if and only if C and D intersect in $\partial\mathbb{H}^2$ and two solutions else. But if γ fixes a point in $\partial\mathbb{H}^2$, it cannot fix a point in $\mathbb{H}^2$, since else it would fix the corresponding geodesic line and thus be a reflection, Proposition 1.2.18. Hence the claim follows from Proposition 1.5.18. □

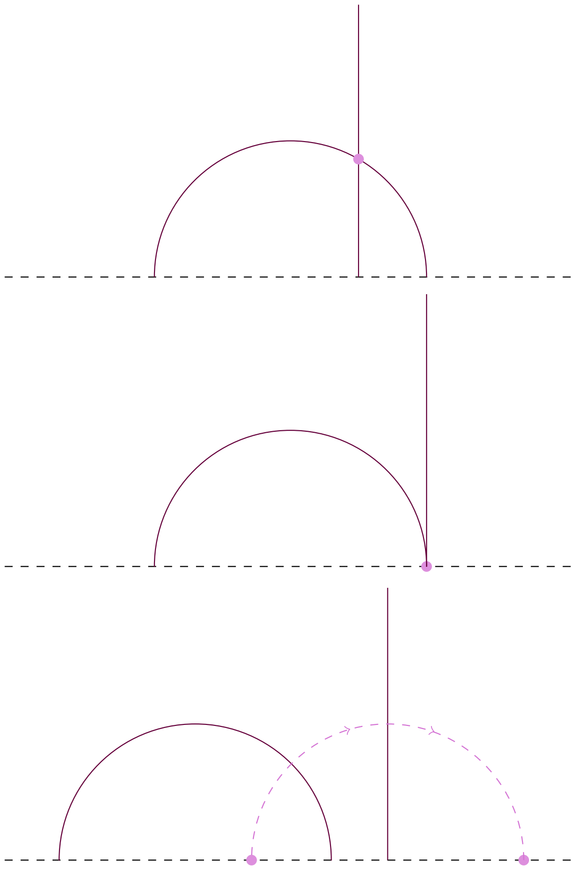

Figure 1.11: Trichotomy of orientation preserving isometries of $\mathbb{H}^2$: Elliptic, Parabolic and Hyperbolic isometries (drawn with fixed points in $\overline{\mathbb{H}}^2$ and axis)

Remark 1.5.20. We can get a better picture of how the three classes of isometries in $\mathbb{H}^2$ act on points by considering appropriate decompositions $\mathbb{H}^2 = \amalg_{\lambda \in \Lambda} S_\lambda$ where each S_λ is invariant under the chosen isometry. For the rotation around 0 in the Poincaré ball model, we can choose the S_λ to be the family of concentric circles around 0 of radius $\lambda \in [0, 1]$. Clearly any rotation around 0 leaves each of this circles invariant. Other decompositions of this type are illustrated in Figure 1.12. (To see how this decompositions are achieved, consider the obvious decompositions for the hyperbolic and parabolic isometries in Examples 1.5.15. Then use Cayley transformations to change ∞ to another point of the boundary and recall that Cayley transformations preserve generalised circles). In the pictures one can for instance see directly that parabolic and elliptic isometries do not preserve a geodesic line and that there is exactly one axis for the hyperbolic isometries of $\mathbb{H}^2$.

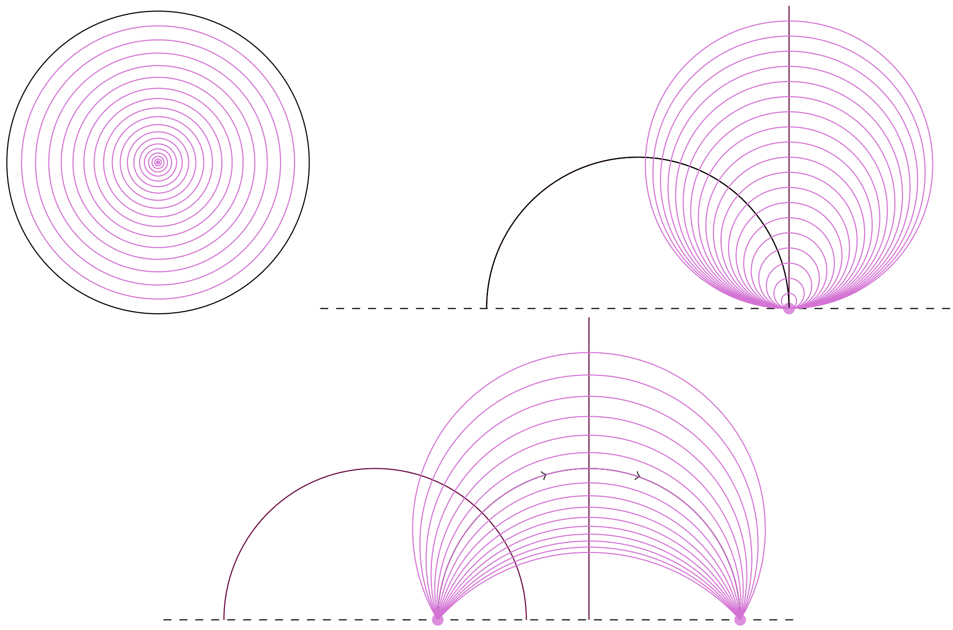

Figure 1.12: Decomposition of $\mathbb{H}^2$ into invariant sets for an elliptic, a parabolic and a hyperbolic Isometry.

Remark 1.5.21. In this section we have concentrated on a geometric perspective on the isometries of $\mathbb{H}^2$. Alternatively, one can also use a linear algebraic approach to study $\mathrm{Isom}(\mathbb{H}^2)$. Exercises!

1.5.3 Flat Torus Theorem and Applications

In this section we will see how the existence of a proper action of a free Abelian group via semi-simple isometries on a CAT(0) space X enforces the existence of flat subspaces in X (the Flat Torus Theorem) and derive some algebraic and topological consequences.

Theorem 1.5.22 (Flat Torus Theorem)**.** *Let A be a free Abelian group of rank $n \in \mathbb{N}$ that acts properly via semi-simple isometries on a CAT(0) space X. Then we have:*

(i) The set $\mathrm{Min}(A) := \bigcap_{\alpha \in A} \mathrm{Min}(\alpha)$ splits as a metric product $Y \times \mathbb{E}^n$ for some non-empty metric space Y.

(ii) Every element of A respects this splitting and splits into the identity on Y and a non-trivial translation on $\mathbb{E}^n$.

(iii) For each $y \in Y$, the quotient of $\{y\} \times \mathbb{E}^n$ by A is an n-torus.

(iv) *Assume that* $\gamma \in \operatorname{Isom}(X)$ *normalises* A*, i.e., that* $\gamma \cdot A \cdot \gamma^{-1} = A$. *Then* γ *leaves* $\operatorname{Min}(A)$ *invariant and preserves the splitting* $\operatorname{Min}(A) = Y \times \mathbb{E}^n$.

(v) *Let* $G \leq \operatorname{Isom}(X)$ *be a subgroup that contains* A *as a normal subgroup. Then there exists a finite index subgroup* $G' \leq G$ *that contains* A *as a central subgroup (i.e., elements of* A *commute with elements of* G'*). Moreover, if* G *is finitely generated, then it is virtually a product* $A \times B$ *for some group* B.

Proof. We first show parts (i), (ii) and (iii) by induction on $n \in \mathbb{N}$. For $n = 0$, the result is trivial with $Y = X$. Fix $n \in \mathbb{N}_{>0}$ and let $\alpha_1, \dots, \alpha_n \in A$ be a primitive generating set of A. Since the action is proper, any elliptic element is torsion and therefore A acts hyperbolically on X. Hence we have that $\operatorname{Min}(\alpha_1) \cong Y_1 \times \mathbb{E}^1$ and since $\alpha \in A$ commutes with α_1, we have that α preserves $Y_1 \times \mathbb{E}^1$ and splits as (α', α''). Let N be the kernel of the map

$$\begin{aligned} \varphi \colon A &\longrightarrow \operatorname{Isom}(Y_1) \\ \alpha &\longmapsto \alpha'. \end{aligned}$$

Then N acts properly by translations on $\{y\} \times \mathbb{E}^1$ and is thus cyclic. Since it contains the primitive generator α_1, we have $N = \langle \alpha_1 \rangle$. So A/N is a free Abelian group of rank $n-1$ that acts properly and hyperbolically via φ on Y_1. Since Y_1 is a convex subset of a CAT(0) space, it is also CAT(0) and therefore, by the induction hypothesis, $\operatorname{Min}(A/N)$ splits as $Y \times \mathbb{E}^{n-1}$, each element of A/N preserves this splitting and the quotient of each $\{y\} \times \mathbb{E}^{n-1}$ by A/N is a torus. Furthermore, we have

$$\begin{aligned} Y \times \mathbb{E}^n &\cong \operatorname{Min}(A/N) \times \mathbb{E}^1 \\ &\cong \bigcap_{\alpha \in A} \operatorname{Min}(\alpha') \times \mathbb{E}^1 \\ &\cong \bigcap_{\alpha \in A} \operatorname{Min}(\alpha') \times \operatorname{Min}(\alpha'') \\ &\cong \bigcap_{\alpha \in A} \operatorname{Min}(\alpha|_{\operatorname{Min}(\alpha_1)}) && \text{(Exercise 9.3)} \\ &\cong \bigcap_{\alpha \in A} \operatorname{Min}(\alpha) \cap \operatorname{Min}(\alpha_1) && \text{(Proposition 1.5.3)} \\ &\cong \bigcap_{\alpha \in A} \operatorname{Min}(\alpha). \end{aligned}$$

Each $\alpha \in A$ splits on $(Y \times \mathbb{E}^{n-1}) \times \mathbb{E}^1$ as (α', α'') and then α' splits by induction on $Y \times \mathbb{E}^{n-1}$.

Now (iv): Since for all $\alpha \in A$, we have $\gamma \cdot \operatorname{Min}(\alpha) = \operatorname{Min}(\gamma \cdot \alpha \cdot \gamma^{-1})$ and $\gamma \cdot \alpha \cdot \gamma^{-1} \in A$, we have that γ leaves $\operatorname{Min}(A)$ invariant. Since it normalises A, it maps A-orbits to A-orbits and thus also convex hulls of A-orbits to the convex hulls of A-orbits. By (ii), the convex hulls of A-orbits of points in $\mathrm{Min}(A)$ are all of the form $\{y\} \times \mathbb{E}^n$, hence γ preserves also the splitting.

Regarding (v): Since A is normal in G, there is a group homomorphism

$$\begin{aligned} \varphi\colon G &\longrightarrow \mathrm{Aut}(A) \\ g &\longmapsto (a \longmapsto g \cdot a \cdot g^{-1}); \end{aligned}$$

and $G' := \ker \varphi$ is the centraliser of A in G. Since translation length is a class function, for each $g \in G$, we see that $\varphi(g)$ preserves translation length. However, by (iii), for any $r \in \mathbb{R}_{\geq 0}$, there exist only finitely many $a \in A$ with $|a| = r$. Hence the image of φ must be finite since A is finitely generated free Abelian and thus $G' \leq G$ is a finite index subgroup. The second statement follows from Theorem 1.5.11. □

Definition 1.5.23. Let G be a group. A subgroup $H \leq G$ is called *characteristic in* G, if for each automorphism $\varphi \in \mathrm{Aut}(G)$, we have $\varphi(H) = H$.

Since H is left invariant by inner automorphisms, it is normal in G. The reason for considering characteristic subgroups is that while being a normal subgroup is not a transitive relation, characteristic subgroups are better behaved and we can often pass to characteristic subgroups:

Proposition 1.5.24.

(i) Let G be a finitely generated group and $n \in \mathbb{N}$. Then there are only finitely many subgroups $H \leq G$ of index n.

(ii) Let G be a finitely generated group and $H \leq G$ a finite index subgroup. Then there exists a finite index subgroup $H' \leq H$ that is characteristic in G.

(iii) Let $A \leq H \leq G$ be groups with A characteristic in H and H normal in G. Then A is normal in G.

Proof.

(i) For any subgroup $H \leq G$ of index n, fix a bijection $\sigma_H\colon G/H \longrightarrow \{1, \dots, n\}$ with $\sigma_H(H) = 1$ which by conjugation induces a bijection $\tau_H\colon \mathrm{Sym}(G/H) \longrightarrow \mathrm{Sym}_n$. The action of G by left-translations on the set of left-H-cosets induces a group morphism $t_H\colon G \longrightarrow \mathrm{Sym}(G/H)$. Hence we get a map

$$\begin{aligned} \Phi\colon \{H \leq G \mid [G : H] = n\} &\longrightarrow \mathrm{Map}_{\mathsf{Grp}}(G, \mathrm{Sym}_n) \\ H &\longmapsto \tau_H \circ t_H. \end{aligned}$$

Clearly $\Phi(H)^{-1}(\{\sigma \in \mathrm{Sym}_n \mid \sigma(1) = 1\}) = H$, so Φ is injective. Since G is finitely generated and Sym_n finite, there are only finitely many maps $\varphi\colon G \longrightarrow \mathrm{Sym}_n$. So $\{H \leq G \mid [G : H] = n\}$ is finite.

(ii) Set $H' := \bigcap\{A \leq G \mid [G : A] = [G : H]\}$. By (i), this is a finite intersection of finite index subgroups and hence of finite index in G and H. Since any automorphism of G maps subgroups of index $[G : H]$ to subgroups of the same index, H' is characteristic.

(iii) Since H is normal in G, conjugation by an element in G defines an automorphism of H. Since $A \leq H$ is characteristic, this automorphism leaves invariant A, so A is normal in G. □

Theorem 1.5.25 (Flat Torus Theorem for virtually free Abelian groups)**.** *Let G a group that is virtually $\mathbb{Z}^n$ and acts properly via semi-simple isometries on a complete CAT(0) space X. Then*

(i) There exists a closed convex subspace in $Z \subset X$ that splits as a product $Y \times \mathbb{E}^n$ for some non-empty metric space Y. Every element in G leaves Z invariant and preserves the splitting, acting as the identity on Y and cocompactly on $\mathbb{E}^n$.

(ii) Every $\gamma \in \mathrm{Isom}(X)$ that normalises G preserves Z and the splitting.

Proof.

(i) Let $B \subset G$ be a finite index subgroup isomorphic to $\mathbb{Z}^n$. By Proposition 1.5.24, there exists a subgroup $A \subset B$ of finite index in B and characteristic in G. In particular, A is also isomorphic to $\mathbb{Z}^n$. By Theorem 1.5.22, $\mathrm{Min}(A)$ splits non-trivially as $Y' \times \mathbb{E}^n$ and G leaves $\mathrm{Min}(A)$ invariant and preserves the splitting. Since A acts trivially on Y', this induces an action of G/A on Y'. Let Y be the fixed point set of this action. Since G/A is finite, by Bruhat-Tits, Corollary 1.3.12, since Y' is closed and thus complete, Y is a non-empty convex subset of Y'. Hence G preserves $Y \times \mathbb{E}^n$ and the splitting, acting by the identity on Y and cocompactly (since A acts already cocompactly) on $\mathbb{E}^n$.

(ii) If γ normalises G, it also normalises A since A is characteristic, so γ leaves $\mathrm{Min}(A)$ invariant and preserves the splitting $Y' \times \mathbb{E}^n$. Since γ normalises G, the action of γ on Y' leaves also invariant Y, the set of fixed points of the action of G on Y'. □

Remark 1.5.26. We summarize the key features of Theorems 1.5.22 and 1.5.25 in a single concise statement: Let G be a (virtually) free abelian group of rank n acting properly and semi-simply on a (complete) CAT(0) space X. Then we find a closed convex G-invariant subset $C \subset X$ isometric to $\mathbb{E}^n$ such that $G\backslash C$ is compact. Any subgroup A of G isomorphic to $\mathbb{Z}^n$ acts by translations on C and $A\backslash C$ is homeomorphic to the n-torus T^n.

Definition 1.5.27.

(i) Let X be a metric space. We say that a group G *acts geometrically on* X if it acts properly and cocompactly by isometries on X.

(ii) We call a group G a *CAT(0) group (a CAT(-1) group)*, if it acts geometrically on a proper (non-empty) CAT(0) space (a CAT(-1) space).

Remark 1.5.28. As we have seen and will see, groups that act nicely on CAT(0) spaces have many interesting properties and such actions arise in many contexts, making the study of such actions an important part of geometric group theory. There are however several factors that limit the importance of the *term* CAT(0) group:

(i) In contrast to say hyperbolicity, being CAT(0) is not an intrinsic feature of a finitely generated group, in the sense that the Cayley graph of such a group will in general not be a CAT(0) space. (More concretely, it can only be a CAT(0) space if the group is free or $\mathbb{Z}/2\mathbb{Z}$, since these are the only groups admitting a simply connected Cayley graph.)

(ii) While by the Svarč-Milnor lemma [14, Corollary 5.3.7], every CAT(0) group is finitely generated and quasi-isometric to a CAT(0) space, the converse is not true. Being a CAT(0) group is also *not* a quasi-isometry invariant.

(iii) As we have already seen in Theorem 1.5.11, assuming that the action on the CAT(0) space is geometric is often an unnecessarily strong assumption.

Sketch of a proof for (ii). Consider the groups G that can be expressed as a central extension of the form

$$1 \longrightarrow \mathbb{Z} \longrightarrow G \longrightarrow \Sigma_2 \longrightarrow 1,$$

where Σ_2 denotes the fundamental group of a closed oriented surface of genus 2. Using bounded cohomology, Gersten [9, Corollary 3.8] showed that any such G is quasi-isometric to $\mathbb{Z} \times \Sigma_2$. In the exercise session, it was shown that Σ_2 is the fundamental group of a compact locally CAT(0) space, hence it is CAT(0) and so is $\mathbb{Z} \times \Sigma_2$ (which acts geometrically on the CAT(0) space $\mathbb{E}^1 \times \mathbb{H}^2$). Using the classification of central extensions via group cohomology [5], one can show however that there are group extensions G that do not split as a product, and which do not even have finite index subgroups that do. Such a group G cannot be CAT(0) by Theorem 1.5.11. □

Definition 1.5.29. Let X be a CAT(0) space. The *rank of* X is defined as the supremum of the dimensions of flats in X, i.e.,

$$\operatorname{rk}(X) := \sup\{n \in \mathbb{N} \mid \text{there exists an isometric embedding } \mathbb{E}^n \longrightarrow X\}.$$

Lemma 1.5.30. Let X be a proper cocompact CAT(0) space, i.e., assume that Isom(X) acts cocompactly on X. Then rk(X) is finite.

Proof. Let $B(x, 2 \cdot r)$ be a ball in X such that the translates of $B(x, 2 \cdot r)$ cover X. Since $\overline{B}(x, 2 \cdot r)$ is compact, there exists a number $p \in \mathbb{N}$, such that $\overline{B}(x, 2 \cdot r)$ is covered by p balls of radius $r/2$. We claim that $p/2$ is an upper bound for rk(X). If there is an n-flat in X, since translates of $B(x, 2 \cdot r)$ cover X, we can assume that there exists a convex subspace $D \subset B(x, 2 \cdot r)$ that is isometric to a Euclidean n-ball of radius r. Considering the points in D that correspond to the r-unit vectors in Euclidean space, we find $2 \cdot n$ points in D, each pair of which is at least $\sqrt{(2)} \cdot r > r$ far apart. So each of these points is contained in a different ball of radius $r/2$, and hence $2 \cdot n < p$. □

Proposition 1.5.31 (Elementary Subgroups)**.** *Let G be a CAT(0) group.*

(i) There is a uniform bound on the rank of free Abelian subgroups in G, that is, $\sup\{\operatorname{rk} A \mid A \leq G,\ A \text{ free Abelian group}\} < \infty$.

(ii) There are only finitely many conjugacy classes of finite subgroups in G.

(iii) In particular, there is a uniform bound on the order of finite subgroups in G, that is, $\sup\{|B| \mid B \leq G,\ |B| < \infty\} < \infty$.

Proof.

(i) Follows from Lemma 1.5.30 and the Flat Torus Theorem.

(ii) Assume that G acts geometrically on the proper CAT(0) space X. Let $K \subset X$ be a compact subset with $G \cdot K = X$. Since the action is proper, for any $x \in K$, there exists an $r_x \in \mathbb{R}_{>0}$, such that the set $\{\gamma \in G \mid \gamma \cdot B(x, r_x) \cap B(x, r_x) \neq \emptyset\}$ is finite. Cover K with finitely many such balls $B(x_1, r_{x_1}), \ldots, B(x_n, r_{x_n})$ and set

$$\Gamma = \bigcup_{i=1}^{n} \{\gamma \in G \mid \gamma \cdot B(x_i, r_{x_i}) \cap B(x_i, r_{x_i}) \neq \emptyset\}.$$

Let $H \leq G$ be a finite subgroup. By Bruhat-Tits, H has a global fixed point x. Pick $\gamma \in G$ such that $\gamma \cdot x \in K$. Then $\gamma \cdot x$ is a fixed point of $\gamma \cdot H \cdot \gamma^{-1}$, hence $\gamma \cdot H \cdot \gamma^{-1}$ is contained in the finite set Γ.

(iii) Follows directly from (ii). □

Theorem 1.5.32. *Let G be a CAT(0) group and* $H_0 \leq H_1 \leq \cdots \leq G$ *an ascending sequence of virtually Abelian subgroups. Then this sequence stabilises, i.e., there exists an* $n \in \mathbb{N}$ *such that for all* $m \in \mathbb{N}_{\geq n}$*, we have that* $H_n = H_m$.

Corollary 1.5.33. Every Abelian subgroup of a CAT(0) group is finitely generated.

Proof. Follows directly from Theorem 1.5.32. □

Proof of Theorem 1.5.32. From any strictly ascending sequence of virtually Abelian groups we could construct a strictly ascending sequence of finitely generated virtually Abelian groups, so it suffices to consider the case that the $(H_i)_{i\in\mathbb{N}}$ are finitely generated. By Proposition 1.5.31 (i), there is an upper bound on the rank of free Abelian subgroups of G, hence we can assume (after dropping finitely many H_i), that there exists an $n \in \mathbb{N}$, such that all H_i contain a finite index free Abelian subgroup A_i of rank n. If $n = 0$, the claim follows from Proposition 1.5.31 (iii), so assume $n \in \mathbb{N}_{>0}$. For $j \in \mathbb{N}$ set $B_j := A_0 \cap A_j$. Then B_j is of finite index in A_0, since B_j is the isotropy group of A_j of the action of A_0 via left-translations on the finite set H_j/A_j. Thus B_j is of rank n and so also of finite index in A_j and hence in H_j. So A_0 is of finite index in H_j for all $j \in \mathbb{N}$. Thus it only remains to show that the index $[H_j : A_0]$ is uniformly bounded. Let G act geometrically on the proper CAT(0) space X. It is straightforward to see (Exercise!) that since the action of G on X is proper and cocompact, for any $r \in \mathbb{R}_{>0}$, there exists an $n_r \in \mathbb{N}$, such that for any $x \in X$, we have

$$|\{\gamma \in G \mid \gamma \cdot x \in B(x, r)\}| \leq n_r.$$

Fix $j \in \mathbb{N}$. By the Flat Torus Theorem for virtually Abelian groups applied to H_j, there exists a subspace of X that splits non-trivially as $Y \times \mathbb{E}^n$, such that H_j leaves this space invariant, preserves the splitting and A_0 acts on $\mathbb{E}^n$ via translations. Pick an $R \in \mathbb{R}_{>0}$, such that the translates via A_0 of any ball of radius R around a point in $Y \times \{0\}$ cover $\mathbb{E}^n$. For any $h \in H_j$ and $x \in Y \times \{0\}$, there is thus an $a \in A_0$, such that $d(x, a \cdot h \cdot x) < R$. Thus $[H_j : A_0] = |H_j/A| \leq n_R$ and is therefore bounded independent of $j \in \mathbb{N}$. □

Lemma 1.5.34. Let G be a finitely generated group with finite commutator subgroup. Then G is virtually Abelian.

Proof. Let $H \leq G$ be the centraliser of $[G, G]$ in G, i.e., the kernel of the map $G \longrightarrow \mathrm{Aut}([G, G])$ given by the action of G on $[G, G]$ by conjugation. Since $[G, G]$ is finite, so is $\mathrm{Aut}([G, G])$ and hence H has finite index in G. Since $[G, G] \cap H$ is central in G, we have for all $g, h \in H$ and $n = |[G, G]|$

$$g^n \cdot h = h \cdot (g \cdot [g, h])^n = h \cdot g^n \cdot [g, h]^n = h \cdot g^n.$$

If $Z(H)$ is the centre of H, this implies that $H/Z(H)$ is torsion, Abelian, and it is finitely generated as the quotient of a finitely generated group. So $H/Z(H)$ is finite and therefore the Abelian group $Z(H)$ has finite index in H and thus in G. □

Theorem 1.5.35 (Solvable Subgroup Theorem, Gromoll-Wolf, Lawson-Yau)**.** *Let G be a CAT(0) group. Then every virtually solvable subgroup of G is finitely generated and virtually Abelian.*

Proof. Let $S \leq G$ be a solvable subgroup. First assume that S is finitely generated. By induction on the length of S (i.e., the minimal length of a derived series for S), we can assume that the commutator subgroup S' of S is finitely generated and virtually Abelian. Hence by Proposition 1.5.24, S' contains a free Abelian characteristic subgroup A of finite index. Since A is characteristic in S', it is normal in S. Thus by Theorem 1.5.22, there exists a finite index subgroup $T \leq S$ that splits as a product $A \times B$. Seeing that for any $(a, b), (a', b') \in A \times B$, we have $[(a, b), (a', b')] = (0, [b, b'])$, we get that $[T, T] \cap A = \{e\}$. Because $[T, T] \leq S'$, we get an injective map $[T, T] \longrightarrow S'/A$. Because A is of finite index in S', this implies that $[T, T]$ is finite. Hence by Lemma 1.5.34, T is virtually Abelian, and hence so is S.

For general S, we argue as follows: By Svarč-Milnor, G is finitely generated and hence countable and thus S is also countable. Since S is countable, it can be written as an ascending union of finitely generated subgroups and these subgroups are solvable. Hence S is the ascending union of finitely generated virtually Abelian groups and thus by Theorem 1.5.32 itself finitely generated and virtually Abelian. □

We will see now that one can drop the cocompactness assumption in Theorem 1.5.35 if one restricts attention to polycyclic groups instead of solvable groups:

Definition 1.5.36 (Polycyclic groups)**.** A *polycyclic group* is a solvable group in which all subgroups are finitely generated.

Corollary 1.5.37. A polycyclic group acts properly and semi-simply on a complete CAT(0) space if and only if it is virtually Abelian.

Proof. The *if*-part is Exercise 11.4. For the other implication: In the proof of Theorem 1.5.35, we needed cocompactness of the action only to show that subgroups of the virtually solvable group are finitely generated, but this holds by definition for polycyclic groups. □

We have seen how nice actions of free Abelian groups on a CAT(0) space X imply the existence of flat subspaces in X. We discuss now a property of geodesic spaces that implies that in such a situation the whole of X is actually flat.

Definition 1.5.38. We say that a geodesic space X has *the geodesic extension property*, if every non-trivial local geodesic can be extended a bit, i.e., if for each local geodesic $c\colon [a, b] \longrightarrow X$ with $a \neq b$, there exists an $\varepsilon \in \mathbb{R}_{>0}$, and an extension of c to a local geodesic $c'\colon [a, b + \varepsilon] \longrightarrow X$.

Remark 1.5.39. A complete CAT(0) space X has the geodesic extension property if and only if X is geodesically complete, i.e., if for each geodesic $c\colon [a, b] \longrightarrow X$ there exists an extension to a geodesic line $c'\colon \mathbb{R} \longrightarrow X$.

Proof. Easy exercise. □

Example 1.5.40. Let M be a complete connected Riemannian manifold. Then M has the geodesic extension property.

Proposition 1.5.41. *Let X be a complete CAT(0) space with the geodesic extension property. Let G be a group that acts cocompactly via isometries on X. Then the only G-invariant closed convex subspaces of X are X and $\emptyset$.*

Proof. Let $C \subset X$ be a non-empty closed convex G-invariant subset of X. Fix $c \in C$ and let $K \subset X$ be a compact subset with $G \cdot K = X$ containing c. Then any point in X is contained in a diam(K)-neighborhood of C $(*)$ since any point in X is contained in $g \cdot K$ for some $g \in G$ and $g \cdot c \in (g \cdot K) \cap C$. If there existed a $y \in X \setminus C$, then one could extend the non-trivial geodesic segment $[\pi(x), x]$, where π denote the projection to the closed convex subspace C, to a geodesic ray $c\colon [0, \infty) \longrightarrow X$. By Proposition 1.3.10, we have $d(x, \pi(y)) = d(x, C)$ for all $x \in [\pi(x), x]$, thus c satisfies $d(c(t), C) = t$ for all $t \in [0, d(\pi(y), y)]$. But then, since $l_C\colon t \longmapsto d(c(t), C)$ is a convex function, we have $d(c(t), C) = t$ for all $t \in \mathbb{R}_{\geq 0}$, contradicting $(*)$. In more detail: Pick any $s \in \mathbb{R}_{>0}$ with $d(c(s), C) = s$. Then for all $t \in \mathbb{R}_{\geq s}$, we have by the convexity of l_C that

$$s = l_C(s) \leq \frac{t-s}{t} \cdot l_C(0) + \frac{s}{t} \cdot l_C(t) = \frac{s}{t} \cdot l_C(t).$$

Hence $t \leq l_C(t) \leq t$. □

Corollary 1.5.42. Let X be a compact geodesic space of non-positive curvature with the geodesic extension property. Let $\pi_1(X)$ be virtually solvable. Then $\pi_1(X)$ is finitely generated and virtually Abelian and X is a flat manifold, i.e., the metric universal cover of X is isometric to $\mathbb{E}^n$, where n is the rank of $\pi_1(X)$

Proof. By Cartan-Hadamard, $\widetilde{X}$ is a CAT(0) space. The fundamental group $\pi_1(X)$ acts properly by isometries on $\widetilde{X}$ and since X is compact, also cocompactly. Hence by the Solvable Subgroup Theorem, $\pi_1(X)$ is finitely generated and virtually Abelian. By the Flat Torus Theorem for virtually Abelian groups, there is a $\pi_1(X)$ invariant flat $\{y\} \times \mathbb{E}^n$ in $\widetilde{X}$. Since X has the geodesic extension property, and this is a local property, also $\widetilde{X}$ has the geodesic extension property. Hence by Proposition 1.5.41, $\widetilde{X}$ is isometric to $\mathbb{E}^n$. □

CHAPTER 2

Exotic groups

2.1 Introduction

After Gromov's slogan, the class of all finitely generated groups is so large that every statement holding for all finitely generated groups should either be trivial or false. In other words, collecting some properties of groups which are not obviously incompatible, there should always be a finitely generated group satisfying these properties. Here are some examples of questions in this spirit: Is there a finitely generated infinite group G

Q1 which is torsionfree, ∞-dimensional and of type F_∞?

Q2 such that there is an integer $n > 1$ with $g^n = 1$ for every $g \in G$?

Q3 which is non-amenable but contains no non-abelian free subgroup?

Q4 which is simple?

Q5 which has intermediate growth, i.e. the growth function is subexponential but not polynomial?

Q6 non-Hopfian, i.e. there is a non-trivial normal subgroup $N \triangleleft G$ such that $G \cong G/N$?

The answer to all these questions is *yes*:

A1 Being of type F_∞ means that there is a classifying space with finitely many cells in each dimension (see below for explanations). Being ∞-dimensional means that there is no classifying space which is of finite dimension. The Thompson group F was the first known infinite torsionfree group with these two properties (Brown and Geoghegan 1984).

A2 This is known as Burnside's problem. The Burnside groups $B(m, n)$ were introduced to treat this and related problems. Adian and Novikov solved the Burnside problem in 1968. In 1979, Ol'shanskii proved that there exists a so-called Tarski monster, an infinite finitely generated group such that each non-trivial proper subgroup is cyclic of order a fixed prime p. In 1991, Ol'shanskii showed that for each non-elementary torsionfree hyperbolic group there is an $n \in \mathbb{N}$ large enough so that G/G^n is infinite. Later, Ivanov and Ol'shanskii could drop the torsionfreeness assumption from this theorem (1996). Nowadays, this result can be obtained as a corollary to the group theoretic Dehn filling theorem for hyperbolically embedded subgroups.

A3 This is known as the von Neumann problem. Ol'shanskii's monster is the first example for such a group. Monod gave another beautiful and easy example in 2013. In 2014 Lodha proved that a certain subgroup of this group, which is still an example for the von Neumann problem, is even of type F_∞. It is known that the Thompson group F contains no non-abelian free subgroups and it was conjectured that it is also non-amenable before Ol'shanskii contructed his monster group. This problem is still open at the time of writing.

A4 Higman introduced the first example of such a group 1951. Another elegant example is a variant of F, the Thompson group V which is even of type F_∞ and contains every finite group as a subgroup.

A5 Grigorchuk proved in 1984 that his finitely generated infinite group (nowadays called the Grigorchuk group) has intermediate growth.

A6 Baumslag and Solitar provided examples in 1962, the Baumslag Solitar groups $BS(m,n)$.

2.2 Thompson's group F

2.2.1 Basic definitions and properties

Definition 2.2.1. *Thompson's group F* is the group of all piecewise linear homeomorphisms $f\colon [0,1] \to [0,1]$ of the unit interval such that:

- There are finitely many breakpoints and all are dyadic rationals, i.e. of the form $\frac{k}{2^n}$ for $k,n \in \mathbb{N}$.
- The slope of each linear part is a power of 2, i.e. of the form 2^k for $k \in \mathbb{Z}$.

A closed interval in $[0,1]$ with dyadic rationals as endpoints is called a *dyadic interval*. The intervals of the form

$$\left[\frac{k}{2^n}, \frac{k+1}{2^n}\right] \qquad k,n \in \mathbb{N}$$

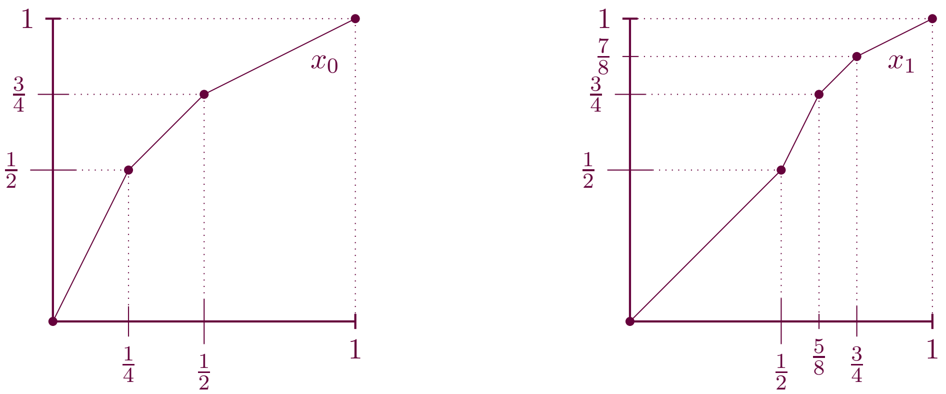

Figure 2.1: Elements of Thompson's group F

which are contained in $[0, 1]$ are called the *standard dyadic intervals*. A *dyadic subdivision* of $[0, 1]$ is a partition of $[0, 1]$ into standard dyadic intervals. Such a subdivision can be obtained by successively cutting the unit interval and resulting intervals in half (see the axes in Figure 2.1). Two dyadic subdivisions D_1, D_2 with the same number of subintervals already determine an element of F: We map the i'th interval of the subdivision D_1 linearly onto the i'th interval of the subdivision D_2. Conversely, let $f \in F$. Choose N large enough such that:

- f is linear on the standard dyadic intervals of the form
$$\left[\frac{k}{2^N}, \frac{k+1}{2^N}\right]$$
- On each such standard dyadic interval, f has slope 2^m for some $m \in \mathbb{Z}$.

Then f maps the dyadic subdivision consisting of the standard dyadic intervals of length $1/2^N$ linearly to another dyadic subdivision. Note that the dyadic subdivisions are not uniquely determined by the group element f. However, it is not hard to see that there are unique minimal dyadic subdivisions representing a given element $f \in F$.

Dyadic subdivisions are purely combinatorial objects. For example, they can be represented as subtrees of the binary tree. These trees describe how one obtains the dyadic subdivision by successively cutting intervals in half. This allows for a purely combinatorial description of F.

Proposition 2.2.2. *F is infinite and torsionfree.*

Proof. Let $1 \neq f \in F$ and set
$$t_0 := \inf\{t \in [0, 1] \mid f(t) \neq t\}$$

Then $f(t_0) = t_0$, f is the identity on the left of t_0 (if $t_0 > 0$) and has slope 2^m on the right of t_0 with $m \neq 0$. Then f^n is the identity on the left of t_0 and has slope 2^{mn} on the right of t_0. Hence all the f^n are distinct. □

Proposition 2.2.3. *F contains a free abelian subgroup of infinite rank.*

Proof. For each $n \geq 1$ choose an element $\gamma_n \in F$ which is not the identity on the standard dyadic interval

$$J_n := \left[\frac{1}{2^n}, \frac{1}{2^{n-1}}\right]$$

and the identity on the rest of $[0, 1]$. Then by the previous proposition, all the elements γ_n have infinite order. But they also commute since each pair of the intervals J_n is either disjoint or only meets at one point. Hence, the elements γ_n generate a free abelian subgroup of infinite rank. □

Proposition 2.2.4. *F contains no non-abelian free subgroup.*

Proof. It suffices to show that F contains no free subgroup of rank 2. So let $f, g \in F$ and we want to show that they do not generate a free subgroup of rank 2.

Step 1: We will first have a look at the commutator $[f, g] = f^{-1}g^{-1}fg$. Let 2^k be the slope of g near 0 and 2^l be the slope of f near zero. Then the slope of $[f, g]$ near 0 is $2^{-l-k+l+k} = 1$. The same goes for the slope near 1. So $[f, g]$ is the identity on neighborhoods of 0 and 1 or, in other words, the *support* of $[f, g]$ is contained in an open dyadic interval (a, b) with $0 < a < b < 1$.

Step 2: Assume for the moment that 0 and 1 are the only common fixed points x of f, g, i.e. $f(x) = x = g(x)$. Let $t \in (0, 1)$ and $t_0 := \inf\big(\langle f, g\rangle t\big)$. We claim that $t_0 = 0$. Assume $t_0 > 0$. Then either $f(t_0) \neq t_0$ or $g(t_0) \neq t_0$. Assume without loss of generality that $f(t_0) < t_0$. Then we find an open neighborhood U of t_0 such that $f(\tau) < t_0$ for all $\tau \in U$. In particular, we find an element $\tau \in U$ which also lies in the orbit $\langle f, g\rangle t$. Thus we have $f(\tau) < t_0$ and since $f(\tau)$ itself is in this orbit, this contradicts the definition of t_0.

Step 3: By the previous step, we find $h \in \langle f, g\rangle$ with $b' := h(b) < a$ and consequently also $a' := h(a) < b'$. The support of the conjugate $h[f, g]h^{-1}$ is contained in (a', b') which is disjoint from the support of $[f, g]$. So $h[f, g]h^{-1} \in \langle f, g\rangle$ commutes with $[f, g] \in \langle f, g\rangle$ and $\langle f, g\rangle$ cannot be a free group of rank 2.

Step 4: Now assume that there is a common fixed point of f and g other than 0 or 1. There are unique dyadic intervals $J_1, \ldots, J_l$ such that the endpoints of each J_i are common fixpoints, each point not in one of the J_i is a common fixpoint and each point in the interior of one of the J_i is not a common fixpoint. For simplicity we assume $l = 2$ (the other cases are similar). Each element of $\langle f, g\rangle$ restricted to J_i yields and element of F by conjugating with a dyadic map $J_i \to [0,1]$. This way, we obtain a monomorphism

$$\mu\colon \langle f, g\rangle \to F \times F$$

By the previous steps, the images of $\mathrm{pr}_1 \circ \mu$ and $\mathrm{pr}_2 \circ \mu$ are not free or, in other words, the kernels N_i of the compositions

$$F_2 \twoheadrightarrow \langle f, g\rangle \xrightarrow{\mathrm{pr}_i \circ \mu} F$$

are non-trivial. Then the kernel of $F_2 \twoheadrightarrow \langle f, g\rangle$ is $N = N_1 \cap N_2$ and is non-trivial because if $N_1 \cap N_2 = \{1\}$ then $[n_1, n_2] \in N_1 \cap N_2$ for $n_i \in N_i$ would always be trivial and hence we find two non-trivial elements n_1, n_2 in F_2 which commute. □

Figure 2.1 pictures two elements x_0 and x_1 of F which together generate F. Our goal in the following is to show that F is even of type F_∞. For this, we first have to recall some basic definitions of homotopy theory.

2.2.2 Higher homotopy groups and cell complexes

We first review a natural generalization of the fundamental group to higher dimensions, the higher *homotopy groups* (see [12, Chapter 4] for details): Let (X, x_0) be a pointed space. For $n \geq 0$ set

$$\pi_n(X, x_0) = \big\{f\colon S^n \to X \;\big|\; f \text{ continuous with } f(s_0) = x_0\big\}/\sim$$

where $s_0 = (1, 0, \dots, 0) \in \mathbb{R}^{n+1}$ and $\sim$ is homotopy relative to x_0. For $n = 0$ we only get a set (the set of path connected components of X), and for $n = 1$ we get the usual fundamental group as reviewed in Subsection 1.4.2. For $n \geq 2$ we even get abelian groups with the following group structure: Let $f, g\colon S^n \to X$ be representatives of two elements in $\pi_n(X, x_0)$. In S^n smash the equator $S^{n-1} \subset S^n$ to a point to obtain two S^n joined at a common point. Let $h\colon S^n \to X$ be f on the upper S^n and g on the lower S^n after smashing the equator and then define $[f] \cdot [g] = [h]$. The unit element is represented by the constant map at x_0.

As in the case $n = 1$, we get homotopy invariant functors

$$\pi_n\colon \mathtt{TOP}_* \to \mathtt{GROUPS}$$

for all $n \in \mathbb{N}$ (where GROUPS has to be replaced by SETS in the case $n = 0$).

We now generalize this even more to the concept of *relative homotopy groups*. If $A \subset B \subset X$ and $U \subset V \subset Y$ we write $f\colon (X, B, A) \to (Y, V, U)$ to mean a map $X \to Y$ satisfying $f(A) \subset U$ and $f(B) \subset V$. This is in accordance with earlier notations such as maps of pointed spaces $(X, x_0) \to$

(Y, y_0) where $B = \{x_0\}$, $A = \emptyset = U$, $V = \{y_0\}$. Now for $n \geq 1$ and $x_0 \in A \subset X$ we define

$$\pi_n(X, A, x_0) = \big\{ f \colon (D^n, S^{n-1}, s_0) \to (X, A, x_0) \bigm| f \text{ continuous } \big\}/\!\sim$$

where $\sim$ is homotopy of maps of the type $(D^n, S^{n-1}, s_0) \to (X, A, x_0)$ (i.e. for each time $t \in [0,1]$ the map $H(\_, t)$ maps s_0 to x_0 and S^{n-1} into A). This time, multiplication is done by smashing $D^{n-1} \subset D^n$ to a point. This way, we get a group structure on $\pi_n(X, A, x_0)$ for $n \geq 2$ which is abelian in the case $n \geq 3$. We obtain homotopy invariant functors

$$\pi_n \colon \mathtt{TOP}^2_* \to \mathtt{GROUPS}$$

on the category of pointed pairs of spaces (again replacing GROUPS by SETS in the case $n = 1$).

Proposition 2.2.5. *Let $p \colon (Y, y_0) \to (X, x_0)$ a covering map. Then p induces isomorphisms $\pi_n(Y, y_0) \cong \pi_n(X, x_0)$ for $n \geq 2$.*

Sketch of proof. This follows from the Lifting Theorem 1.4.2 because the spaces S^n and $S^n \times [0,1]$ for $n \geq 2$ are simply connected. □

Theorem 2.2.6 (Long exact homotopy sequence)**.** *Let (X, A, x_0) be a pointed pair of spaces. Let*

$$\partial \colon \pi_n(X, A, x_0) \to \pi_{n-1}(A, x_0)$$

be the map induced by restricting maps $(D^n, S^{n-1}, x_0) \to (X, A, x_0)$ to S^{n-1}. It is a homomorphism of groups if $n \geq 2$ and a map of pointed sets else and the sequence

$$\cdots \to \pi_n(A, x_0) \to \pi_n(X, x_0) \to \pi_n(X, A, x_0) \xrightarrow{\partial} \pi_{n-1}(A, x_0) \to \cdots \to \pi_0(X, x_0)$$

where all the unlabeled maps are induced by inclusion, is exact. More generally, we have an exact sequence

$$\cdots \to \pi_n(A, B, x_0) \to \pi_n(X, B, x_0) \to \pi_n(X, A, x_0) \to \pi_{n-1}(A, B, x_0) \to \cdots$$

ending at $\pi_1(X, A, x_0)$ for a triple $x_0 \in B \subset A \subset X$ of spaces.

A sequence

$$A \xrightarrow{f} B \xrightarrow{g} C$$

of groups is called exact at B if $\ker g = \operatorname{im} f$. If this is just a sequence of pointed sets A, B, C with base points a, b, c, this means that $g^{-1}(c) = f(A)$, and in this sense the sequences in the theorem are exact near the end.

Definition 2.2.7. A space X is called (-1)-*connected* if it is not empty. It is called n-*connected* for $n \geq 0$ if it is not empty and if $|\pi_k(X, x_0)| = 1$ for each $k \leq n$ and some (and therefore every) $x_0 \in X$. In particular, a non-empty space X is 0-connected if and only if it is path connected and 1-connected if and only if it is simply connected.

More generally, a pair (X, A) of spaces is called (-1)-connected if A is not empty and n-connected for $n \geq 0$ if in addition $|\pi_k(X, A, x_0)| = 1$ for $k \leq n$.

A space X is called *aspherical* if it is 0-connected and if $|\pi_k(X, x_0)| = 1$ for each $k \geq 2$ and some (and therefore every) $x_0 \in X$.

Next we want to recall the notion of *cell complexes* (*CW-complexes*). Let X, Y be spaces, $A \subset Y$ closed and $f\colon A \to X$ a continuous map. Then we say that the space Z is obtained by attaching Y to X along f if

$$Z := (X \sqcup Y)/{\sim}$$

is given the quotient topology and $\sim$ is the equivalence relation generated by $f(a){\sim}a$ for all $a \in A$. We say that a space Z is obtained from X by attaching an n-cell, if $Y = D^n$ and $A = S^{n-1}$. More generally, we say that Z is obtained from X by attaching n-cells if Y is a disjoint union of disks D^n and A the disjoint union of the boundaries S^{n-1}.

Definition 2.2.8. A 0-dimensional cell complex is a set with the discrete topology. Inductively, we define an n-dimensional cell complex to be a space Z which is obtained from a $(n-1)$-dimensional cell complex by attaching n-cells. A cell complex Z is the colimit of a sequence of such constructions, i.e. if $X_0 \subset X_1 \subset \dots$ where each X_n is obtained from X_{n-1} by attaching n-cells, then $Z = \bigcup_n X_n$ together with the weak topology, i.e. a subset $U \subset Z$ is open if and only if $U \cap X_n$ is open in X_n for each n. The space X_n is called the n-skeleton of Z.

Cell complexes are nicely behaved in terms of homotopy. For example, a connected cell complex $Z \neq \emptyset$ is contractible if and only if $|\pi_n(Z)| = 1$ for all $n \geq 0$. This is a consequence of *Whitehead's Theorem*:

Theorem 2.2.9 (Whitehead). *If $f\colon X \to Y$ is a continuous map of connected cell complexes which induces isomorphisms $\pi_n(f)\colon \pi_n(X) \to \pi_n(Y)$ for all n, then f is a homotopy equivalence.*

A *homotopy equivalence* of spaces X, Y is a continuous map $f\colon X \to Y$ which has an inverse up to homotopy, i.e. there is a continuous $g\colon Y \to X$ such that $g \circ f$ is homotopic to id_X and $f \circ g$ is homotopic to id_Y. Note that a space is contractible if and only if it is homotopy equivalent to a point.

We obtain the notion of a Δ-*complex* by restricting the possible gluing maps: As n-cells, we take n-simplices, the convex hull of $n+1$ points in

$\mathbb{R}^n$ in general position. The boundary $\partial\Delta^n$ of an n-simplex Δ^n is the union of $n+1$ $(n-1)$-simplices, called the faces of the n-simplex. We demand that the gluing map $\partial\Delta^n \to X_{n-1}$ of an n-simplex Δ^n maps the faces of Δ^n linearly and homeomorphically onto $(n-1)$-simplices in X_{n-1}. (*Note:* We have taken the name Δ-complex from [12]. However, our Δ-complexes do not remember the orientation on the n-simplices.)

By additionally imposing the following rules, we arrive at the notion of a *simplicial complex*:

- No two faces of an n-simplex are glued onto the same $(n-1)$-simplex in X_{n-1}.
- Two n-simplices in X_n are either disjoint or meet at exactly one m-simplex with $m < n$.

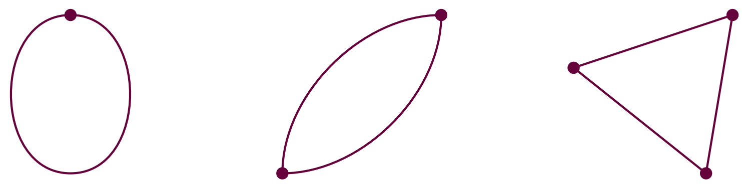

Figure 2.2: Two Δ-complexes which are not simplicial (left) and one simplicial complex (right)

Simplicial complexes allow for a purely combinatorial description:

Definition 2.2.10. A (abstract) simplicial complex is a set V of vertices together with a set S of non-empty finite subsets of V (the simplices) such that

- $\{v\} \in S$ for each $v \in V$
- If $s \in S$ and $\emptyset \neq s' \subset s$, then also $s' \in S$.

If $s \in S$ and $s' \subset s$ of cardinality one less, then s' is a face of s and this describes how to glue n-simplices to obtain a concrete simplicial complex as above, called the geometric realisation of the abstract simplicial complex (V, S).

Different abstract simplicial complexes can give the same space after geometric realisation. For example, we can apply *barycentric subdivision* by introducing a midpoint for each simplex (see Figure 2.4).

So every simplicial complex is a cell complex. Conversely, every cell complex can be made simplicial by homotoping the gluing maps on the boundaries S^{n-1} of cells D^n and triangulating the cells D^n fine enough, for example by viewing D^n as an n-simplex and applying barycentric subdivision often enough.

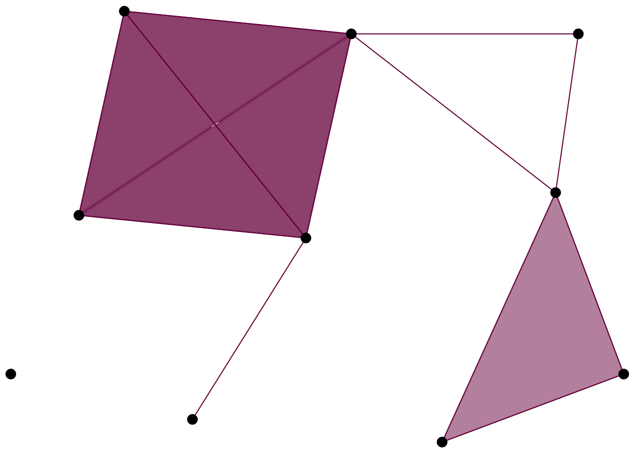

Figure 2.3: A simplicial complex

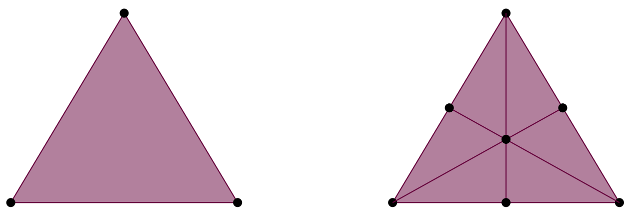

Figure 2.4: Barycentric subdivision

Definition 2.2.11. Let v be a vertex in a simplicial complex. The *star* $st(v)$ of v is the collection of all simplices containing v together with all their subsimplices. The *link* $lk(v)$ is the collection of all simplices in $st(v)$ which do not contain v.

Definition 2.2.12. A *flag complex* is a simplicial complex such that the following is satisfied: Whenever n vertices are pairwise joined by edges, then the simplicial complex already contains the $(n-1)$-simplex spanned by these vertices.

Note that a flag complex is already determined by its vertices and edges.

A map of abstract simplicial complexes $(V, S) \to (V', S')$ is a map $f\colon V \to V'$ such that $f(s) \in S'$ for each $s \in S$. Since geometric realisation of abstract simplicial complexes is functorial, we obtain the corresponding notion of a simplicial map for (concrete) simplicial complexes. Up to homotopy and subdivision of simplices, every continuous map of simplicial complexes is simplicial:

Theorem 2.2.13 (Simplicial approximation)**.** *Let X, Y be simplicial complexes and $f\colon X \to Y$ continuous. Then, after applying barycentric subdivi-*

sion a finite number of times to X, the map f is homotopic to a simplicial map. A similar statement holds for maps $f\colon (X,A) \to (Y,B)$ of pairs of simplicial complexes.

There is an analogous theorem for cell complexes. A map $X \to Y$ of cell complexes is called cellular if $f(X_n) \subset Y_n$ for each n.

Theorem 2.2.14 (Cellular approximation)**.** *Let X, Y be cell complexes and $f\colon X \to Y$ continuous. Then f is homotopic to a cellular map. A similar statement holds for maps $f\colon (X,A) \to (Y,B)$ of pairs of cell complexes.*

2.2.3 Finiteness properties of groups

Definition 2.2.15. Let G be a group. A classifying space for G is an aspherical cell complex BG with $\pi_1(BG) \cong G$.

In general, there are a lot of different models for the classifying space of a group, but one can show that the classifying space is at least unique up to homotopy equivalence.

One can think of the classifying space of a group G to be a natural "space-like" version of the group G. This is supported by the following standard construction of the classifying space which yields rather large Δ-complexes in general: Take a single 0-simplex v and attach a loop (1-simplex) to v for each element $1 \neq g \in G$. Then attach a 2-simplex for each pair $(g_1, g_2) \in G \times G$ with $g_i \neq 1$ to the 1-simplices given by g_1, g_2 and g_1g_2. More generally, attach an n-simplex for each sequence $(g_1, \ldots, g_n)$ of non-trivial elements in G to the $(n-1)$-simplices given by

$$(g_2, g_3, \ldots, g_n)$$
$$(g_1, g_2, \ldots, g_{n-1})$$
$$(g_1, \ldots, g_{i-1}, g_ig_{i+1}, g_{i+2}, \ldots, g_n)$$

for each $i = 1, \ldots, n-1$.

A smaller model can often be obtained by choosing a presentation $\langle S \mid R \rangle$ of the group G and perform the following construction: Again take a single 0-simplex v and for each generator $s \in S$ attach a loop to v and remember the orientation on each such loop. This is the 1-skeleton X_1. Then for each relation in $r \in R$, which is a word of length k in the generators in S and their inverses, attach a 2-cell D^2 by gluing the boundary S^1 to the loops at v in the following manner: Subdivide S^1 into k equal parts and map the i'th part linearly to the loop corresponding to the generator s where s is the i'th letter in r, in reverse direction if the i'th letter in r is decorated with a -1 exponent. This yields the 2-skeleton X_2. By construction and the Cellular Approximation Theorem, we have $\pi_1(X_2, v) \cong G$. We can now make this space into an aspherical one by further attaching cells in higher dimensions:

For each $1 \neq [f] \in \pi_2(X_2, v)$, attach a 3-cell to X_2 by gluing the boundary S^2 via the map f and obtain the space X_3. By cellular approximation, $\pi_1(X_2, v) \cong \pi_1(X_3, v)$ and $\pi_2(X_3, v) = 0$. We can iterate this in higher dimensions to obtain an aspherical cell complex X as desired.

Definition 2.2.16. Let G be a group and $n \in \mathbb{N}_{\geq 1}$. We say that G is of *type F_n* if there exists a classifying space BG of G with n-skeleton having only finitely many cells. We say that G is of *type F_∞* if it is of type F_n for all $n \in \mathbb{N}_{\geq 1}$.

One can show that a group is of type F_∞ if and only if there is a classifying space BG of G with finitely many cells in each dimension (see [8, Proposition 7.2.2]).

Proposition 2.2.17. *A group G is of type F_1 if and only if it is finitely generated and of type F_2 if and only if it is finitely presented.*

Sketch of proof. Let X be a classifying space with finitely many 1-cells. Then the 1-skeleton X_1 is a finite graph. Choose a spanning tree of X_1 and collapse it to a point v. The remaining loops $\gamma_1, \ldots, \gamma_k$ at v generate $\pi_1(X_1, v)$. Since each loop in X based at v can be homotoped into X_1 by cellular approximation, these loops also generate $\pi_1(X, v)$.

Assume now that also X_2 has only finitely many cells. Then, since homotopies of loops in X_1 can themselves be homotoped into X_2 (cellular approximation), we can read off a full set of relations from the 2-cells: Let $f\colon S^1 \to X_1$ be the gluing map of a 2-cell. Going around S^1 counter-clockwise, we record how often S^1 wraps around the loops γ_i via f. This gives a word in the generators γ_i and their inverses which is a relation in $\pi_1(X, v)$.

Conversely, the second construction after Definition 2.2.15 yields a type F_1 or F_2 classifying space whenever G is finitely generated or presented. □

So we see that being of type F_∞ is a very strong finiteness property for groups.

Remark 2.2.18. One can show that being of type F_n is invariant under quasi-isometries of finitely generated groups (see [10, $1.C_2'$]).

Definition 2.2.19. The *(geometric) dimension* of a group G is the dimension of a classifying space BG of minimal dimension.

It follows from Proposition 2.2.3 that F is infinite dimensional. To see this, first note that a classifying space for $\mathbb{Z}^n$ is given by the n-torus

$$T^n = S^1 \times \ldots \times S^1 = \prod_{i=1}^{n} S^1$$

which is universally covered by $\mathbb{R}^n$ via

$$\mathbb{R}^n \to T^n \qquad (t_1, \ldots, t_n) \mapsto \big(\exp(it_1), \ldots, \exp(it_n)\big)$$

It can be shown using homological methods that there can't be any classifying space of lower dimension, so the dimension of $\mathbb{Z}^n$ is n. Furthermore, $\dim(H) \leq \dim(G)$ if $H < G$. Hence the dimension of F cannot be finite.

2.2.4 Methods of discrete Morse Theory

Lemma 2.2.20 (Fundamental Morse Lemma)**.** Let X be a simplicial complex and v_0 a vertex. Let X^{-v_0} be the complex obtained from X after deleting all simplices containing v_0. Assume that $lk(v_0) \subset X^{-v_0}$ is $(n-1)$-connected. Then the pair (X, X^{-v_0}) is n-connected.

Proof. Let x_0 be some other vertex than v_0 and

$$f\colon (D^n, S^{n-1}, s_0) \to (X, X^{-v_0}, x_0)$$

a continuous map. We want to show that f can be homotoped to a map with image lying in X^{-v_0}. First, we triangulate (D^n, S^{n-1}) fine enough so that we can homotope f to a simplicial map by the Simplicial Approximation Theorem.

Step 1: First we want to see that we can homotope f such that if $f(v) = v_0 = f(v')$ for two vertices $v, v' \in D^n$, then v and v' are not joined by an edge in D^n. If there are two such vertices, then there is also a vertex v with $f(v) = v_0$ which is both joined to a vertex v' with $f(v') = v_0$ and a vertex w with $f(w) \neq v_0$. We then can define another simplicial map $f'\colon D^n \to X$ with $f'(v) = f(w)$ and $f'(x) = f(x)$ for any other vertex x. The map f is homotopic to f' by sliding the vertex v from v_0 to $f(w)$ along the edge joining these two vertices.

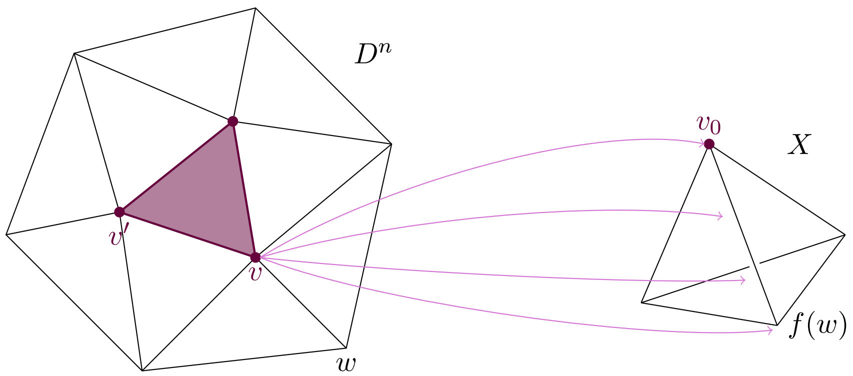

Iterating this often enough, we arrive at a situation as claimed above.

Step 2: Let v be a vertex in D^n such that $f(v) = v_0$. After step 2, no vertex in D^n which is joined to v by an edge is mapped to v_0. Hence we have

$$f(lk(v)) \subset lk(v_0) \subset X^{-v_0}$$

Since $lk(v)$ is homeomorphic to S^{n-1} and $lk(v_0)$ is assumed to be $(n-1)$-connected, we can extend the map $f|_{lk(v)}\colon lk(v) \to lk(v_0)$ to D^n and obtain a continuous map $f'\colon D^n \to lk(v_0)$. Then $f|_{st(v)}$ can be homotoped in $st(v_0)$ to f' while fixing the boundary $lk(v)$. If we repeat this with all vertices v with $f(v) = v_0$, we are done. □

Now assume we have a simplicial complex X with vertices V and a full subcomplex A spanned by a subset of vertices $V_A \subset V$, i.e. A is the simplicial complex with vertices V_A and with every simplex in X which has vertices lying in V_A. Let $V' = V \setminus V_A$ and $f\colon V' \to \mathbb{N}$ a function such that $f(v_1) \neq f(v_2)$ whenever v_1 and v_2 are joined by an edge in X. Such a function is called a *Morse function* on the pair (X, A).

The reason why we introduce such a function is that we can build up X from A by successively adding vertices in order of increasing Morse height f. Start with a vertex $v_1 \in V'$ with minimal Morse height $f(v)$ and observe the full subcomplex X_1 of X spanned by the vertices $V_A \cup \{v_1\}$. The pair (X_1, A) has a link as in Lemma 2.2.20, called the *descending link* with respect to f. Then we add another vertex v_2 of minimal Morse height (if there are any). Note that the descending links of adding v_1 and v_2 does not depend on the order because of the defining property of Morse functions. Then we add vertices of the next bigger Morse height and so on, until we obtain all of X.

The upshot of this procedure is: If we have enough control of the connectivities of the descending links and we know the connectivity of either A or X, then we can make deductions on the connectivity of X or A using Lemma 2.2.20. The success of this procedure depends crucially on a clever choice of the Morse function f.

This method will be applied several times in the following.

2.2.5 The main theorem and outline of the proof

Theorem 2.2.21. *Thompson's group F is of type F_∞. In particular, it is finitely presented.*

The first step of the proof is to construct the universal cover X_F of a certain model for BF directly as a simplicial complex.

Definition 2.2.22.

(i) Let $n \in \mathbb{N}_{\geq 1}$. Denote by I the unit interval and by nI the (ordered) disjoint union of n unit intervals. A *dyadic subdivision* on nI is a dyadic subdivision on each interval of nI. We order the subintervals primarily by the order of the corresponding I in nI and secondarily by the natural order on I.

(ii) Let $n, m \in \mathbb{N}_{\geq 1}$. A *dyadic map* $f\colon nI \to mI$ is a continuous and surjective map of the following form: There are dyadic subdivisions on nI and on mI such that f maps the i'th subinterval of the subdivision on nI homeomorphically, linearly and order preservingly onto the i'th subinterval of the subdivision on mI. We denote by $F(n, m)$ the set of dyadic maps $nI \to mI$.

(iii) We call such a dyadic map a *merge map* if the dyadic subdivision on the domain nI can be chosen to be trivial (i.e. no interval I is subdivided).

(iv) A merge map is called *simple* if the dyadic subdivisions on the intervals I of the codomain mI are either trivial (i.e. do not subdivide the interval) or are just simple halvings of the interval.

Construction of X_F

We now give the definition of the simplicial complex X_F: We have one vertex for each dyadic map in $\bigcup_{n=1}^{\infty} F(n, 1)$. If $f_n \in F(n, 1)$ and $f_m \in F(m, 1)$ with $n > m$, then we join f_n and f_m by an edge if there is a merge map $g \in F(n, m)$ such that $f_m \circ g = f_n$. This g is unique if it exists. Now we define X_F to be the full simplicial flag complex generated by this graph.

Contractibility of X_F

We want to show that X_F is contractible.

Lemma 2.2.23. Let $f_1, \ldots, f_l$ be vertices in X_F. Then there is a vertex g which is joined by an edge to f_i for each i.

Proof. Let D_i be the dyadic subdivision of the codomain interval of the dyadic map $f_i\colon n_i I \to I$. Let k_i be the largest natural number such that $1/2^{k_i}$ is the length of a subinterval of D_i. Set

$$k := \max\{k_1, \ldots, k_l\}$$

Then let D be the dyadic subdivision of I such that each subinterval has length $1/2^k$. Then D refines all subdivision D_i and the unique merge map $g\colon 2^k I \to I$ with D as dyadic subdivision is easily seen to satisfy the criterion in the lemma. □

Now assume that we have a continuous map $f\colon S^n \to X_F$ which we can assume to be simplicial. Then, since S^n is compact, f meets only finitely many vertices $f_1, \ldots, f_l$ in X_F. Let g be the vertex as in the lemma. Then, since X_F is a flag complex, the full subcomplex spanned by $\operatorname{im} f$ and g is a cone over $\operatorname{im} f$. So f can be homotoped into the point g. It follows that all homotopy groups of X_F vanish and hence that it is contractible.

The action of F on X_F

We define a simplicical action (i.e. by simplicial automorphisms) of F on X_F by

$$\gamma \cdot f := \gamma \circ f \in F(n,1)$$

if $\gamma \in F$ and $f \in F(n,1)$. Since each vertex f is a surjective function, the action so defined is free. Since it is by simplicial automorphisms, it follows that the action is properly discontinuous. Hence,

$$p\colon X_F \to F\backslash X_F$$

is a normal covering and $\pi_1(F\backslash X_F) \cong F$ because X_F is contractible and thus simply connected. Note that $F\backslash X_F$ with the cell structure coming from X_F is not a simplicial complex anymore, but a Δ-complex.

The Morse function on X_F

Let V be the set of vertices of X_F. We define $\deg\colon V \to \mathbb{N}$ by $\deg(f) = n$ if $f \in F(n,1)$. By definition of X_F, this function is a Morse function on X_F. Let $(X_F^i)_i$ be the associated Morse filtration, i.e. X_F^i is the full subcomplex of X_F spanned by the vertices v with $\deg(v) \leq i$.

Lemma 2.2.24. For each k, The Δ-complex $F\backslash X_F^k$ has only finitely many cells.

Proof. We want to show that each two vertices of X_F with the same degree lie in the same F-orbit. The claim then follows since there are only finitely many merge maps $nI \to mI$ with $n \leq k$.

So let $f_1, f_2 \in F(n,1)$. Let f_i be defined by a dyadic subdivision D_i^d on the domain and D_i^c on the codomain. By further subdivision, we can assume that $D_1^d = D_2^d$. Let $\gamma \in F$ be defined by the dyadic subdivision D_1^c on the *domain* and D_2^c on the *codomain*. Then we have $\gamma \cdot f_1 = f_2$. □

In subsection 2.2.6 we will show that the connectivity of the descending link of each vertex v tends to infinity as $\deg(v) \to \infty$ and consequently, by the Fundamental Morse Lemma, that the connectivity of the pairs (X_F^{k+1}, X_F^k) tends to infinity as $k \to \infty$. More precisely, we have

$$\forall_m \exists_n \forall_{k \geq n}\ (X_F^{k+1}, X_F^k) \text{ is } m\text{-connected}$$

End of the proof

Note that for a triple of spaces $B \subset A \subset Z$, if (A,B) and (Z,A) are k-connected, then also (Z,B) is k-connected. This follows from the long exact homotopy sequence of that triple (Theorem 2.2.6). Furthermore, if $Z_1 \subset Z_2 \subset \ldots$ is a nested sequence of spaces with colimit Z and each

(Z_{i+1}, Z_i) is k-connected, then also (Z, Z_1) is k-connected. This follows from the previous remark and the fact that each map $K \to Z$ from a compact space K (e.g. $K = S^n$) factors through some Z_i. Applying this observation to our situation above, we obtain

$$\forall_m \exists_n \forall_{k \geq n}\ (X_F, X_F^k) \text{ is } m\text{-connected}$$

and this means

$$\forall_m \exists_n \forall_{k \geq n} \forall_{i \leq m}\ \pi_i(X_F, X_F^k) = 0$$

Using the long exact homotopy sequence of the pair (X_F, X_F^k), this is equivalent to

$$\forall_m \exists_n \forall_{k \geq n} \big(\forall_{i<m}\ \pi_i(X_F^k) \cong \pi_i(X_F) \text{ and } \pi_m(X_F^k) \twoheadrightarrow \pi_m(X_F)\big)$$

For some fixed $m \geq 1$, we now know that there is a k large enough such that

$$\pi_i(X_F^k) \cong \pi_i(X_F) = 0$$

for all $i \in \mathbb{N}$ with $0 \leq i \leq m$. In other words, X_F^k is m-connected. Define $Y^k := F \backslash X_F^k$. As above, $X_F^k \to Y^k$ is a covering and $\pi_1(Y^k) \cong F$. By Proposition 2.2.5, we have

$$\pi_i(Y^k) \cong \pi_i(X_F^k)$$

for $i \geq 2$ which is trivial for $2 \leq i \leq m$. Moreover, Y^k is a Δ-complex with only finitely many cells (Lemma 2.2.24). As in the second construction after Definition 2.2.15, we can now attach cells in dimensions $m+2$ and higher to kill all the homotopy groups above dimension m. Call this space Y_+^k. Then Y_+^k is a classifying space for F whose $(m+1)$-skeleton only has finitely many cells. Thus, Y_+^k is a witness that F is of type F_{m+1}. Since m was arbitrary, it follows that F is of type F_∞.

2.2.6 Connectivity of the links

Here we sketch the proof of the remaining claim that the connectivity of the descending link of each vertex v tends to infinity as $\deg(v) \to \infty$.

Fix a vertex v_0 in X_F and let $n_0 := \deg(v)$. We want to understand the descending link $lk_\downarrow(v_0)$ with respect to the Morse function deg above. Unraveling the definitions, one sees that $lk_\downarrow(v_0)$ is isomorphic to the flag complex spanned by the following graph: The vertices are given by the merge maps $n_0 I \to mI$ with $m < n_0$. We have an edge between two such merge maps if one is obtained from the other by a non-trivial merging, i.e. if $f_i \colon n_0 \to m_i I$ are the two merge maps with $m_1 < m_2 < n_0$, then there is a merge map $g \colon m_2 I \to m_1 I$ with $f_1 = g \circ f_2$. This g is unique if it exists.

In the first step we want to reduce $lk_\downarrow(v_0)$ to a homotopy equivalent subcomplex: Let $lk_\downarrow^*(v_0)$ be the full subcomplex of $lk_\downarrow(v_0)$ spanned by the vertices which are *simple* merge maps $n_0 I \to mI$ with $m < n_0$.

Proposition 2.2.25. *The inclusion* $i\colon lk_{\downarrow}^{*}(v_0) \to lk_{\downarrow}(v_0)$ *is a homotopy equivalence.*

Proof. The idea is to build up $lk_{\downarrow}(v_0)$ from $lk_{\downarrow}^{*}(v_0)$ using the Morse technique. If we introduce a suitable Morse function on the pair $(lk_{\downarrow}(v_0), lk_{\downarrow}^{*}(v_0))$ such that each descending link is contractible, then the inclusion i is a homotopy equivalence since $lk_{\downarrow}(v_0)$ is obtained from $lk_{\downarrow}^{*}(v_0)$ by successively coning off contractible subspaces. More precisely, the Fundamental Morse Lemma and the long exact homotopy sequence imply that i induces isomorphisms on all homotopy groups and thus is a homotopy equivalence by Whitehead's Theorem.

So we define a Morse function μ as follows: Let $f\colon n_0 I \to mI$ with $m < n_0$ a non-simple merge map, then set $\mu(f) = n_0 - m$. By definition, the descending link $lk_{\downarrow}^{\mu}(f)$ is the full subcomplex spanned by the merge maps $g\colon n_0 I \to lI$ with $m < l < n_0$ which are joined to f by an edge, or in other words, which are obtained from f by "splitting" some intervals in mI along some cut points in the dyadic subdivision. Let $g_0\colon n_0 I \to l_0 I$ be the unique merge map with $m < l_0 < n_0$ which is obtained from f by splitting all intervals I in mI with a non-trivial dyadic subdivision into two halves.

We claim that $lk_{\downarrow}^{\mu}(f)$ is almost a cone with tip g_0: The full subcomplex C spanned by the vertices which are joined to g_0 by an edge is certainly a cone with tip g_0 since $lk_{\downarrow}^{\mu}(f)$ is a flag complex. Let h be a vertex which is not joined to g_0. Then there is a unique merge map $H(h)\colon n_0 I \to kI$ in $lk_{\downarrow}^{\mu}(f)$ with maximal k such that both g_0 and h merge to $H(h)$, i.e. there are merge maps a, b with $a \circ g_0 = H(h) = b \circ h$. If h is a vertex in C, then we define $H(h) = h$. The maximality condition above implies that H induces a simplicial map $lk_{\downarrow}^{\mu}(f) \to lk_{\downarrow}^{\mu}(f)$. It turns out that H is a deformation retraction onto C. Roughly speaking, this means that when sliding each h not lying in C to $H(h)$ along the edge which joins both vertices, then we get a homotopy equivalence from $lk_{\downarrow}^{\mu}(f)$ to C. □

Figure 2.5 pictures the descending link $lk_{\downarrow}^{\mu}(f)$ in the proof for f being represented by the following dyadic subdivision:

Remark 2.2.26. For the inclined reader, we give a more precise explanation why H is a deformation retraction. Obviously, X_F is defined as a poset. Hence also $lk_{\downarrow}^{\mu}(f)$ is a poset. The maximality of k in the definition of $H(h)$ implies that H is a functor. Since we have an arrow $h \to H(h)$ for each h, we get a natural transformation from the identity to H. This is a homotopy on the level of spaces.

We now give a slightly different description of $lk_{\downarrow}^{*}(v_0)$: Let $l \geq 1$. We will define a simplicial flag complex $\mathcal{C}(l)$: The vertices are given by subsets

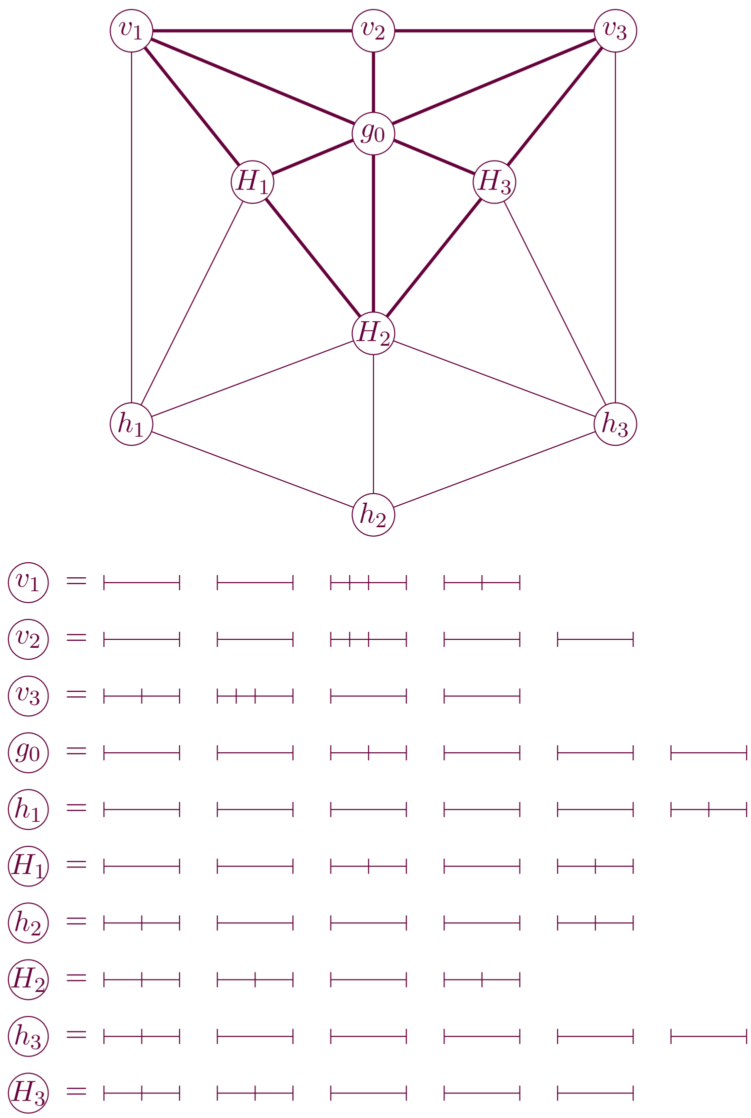

Figure 2.5: A descending link of a descending link

of $\{1, \dots, l\}$ of the form $\{i, i+1\}$ for $i \in \{1, \dots, l-1\}$. Two such vertices are joined by an edge if and only if they are disjoint as subsets.

Observe that $lk^*_\downarrow(v_0)$ is canonically isomorphic to the barycentric subdivision of $\mathcal{C}(l)$ where $l = n_0$: Let $f\colon n_0 I \to mI$ with $m < n_0$ be a simple merge map representing a vertex in $lk^*_\downarrow(v_0)$. Then there are $i_1, \dots, i_k$ such that the intervals of index i_j and $i_j + 1$ for each $j = 1, \dots, k$ are merged by f. The vertices $\{i_j, i_j + 1\}$ span a $(k-1)$-simplex in $\mathcal{C}(l)$ which is a vertex v_f in the barycentric subdivision of $\mathcal{C}(l)$. The map $f \mapsto v_f$ gives a simplicial isomorphism.

The following proposition concludes the proof of the main theorem.

Proposition 2.2.27. *The complex $\mathcal{C}(l)$ is $\nu(l)$-connected where*

$$\nu(l) = \left\lfloor \frac{l-2}{3} \right\rfloor - 1$$

Proof. We proceed by induction over l. If $l \geq 2$ then, obviously, $\mathcal{C}(l)$ is non-empty. For the induction step, consider the full subcomplex $\mathcal{A}$ of $\mathcal{C}(l)$ which is spanned by the vertices $\{i, i+1\}$ with $i > 2$. To build up $\mathcal{C}(l)$ from $\mathcal{A}$ we only have to add the vertices $v_1 = \{2, 3\}$ and $v_2 = \{1, 2\}$. The descending links do not depend on the order since v_1 and v_2 are not joined by an edge. The descending link of v_1 is canonically isomorphic to $\mathcal{C}(l-3)$ while the descending link of v_2 is canonically isomorphic to $\mathcal{C}(l-2)$. By induction, both descending links are at least $\nu(l-3)$-connected. By the Fundamental Morse Lemma it follows that the connectivity of the pair $(\mathcal{C}(l), \mathcal{A})$ is at least $\nu(l-3) + 1 = \nu(l)$.

We now want to show that the map $\pi_k(\mathcal{A}) \to \pi_k(\mathcal{C}(l))$ induced by the inclusion i is trivial. It then follows from $\pi_k(\mathcal{C}(l), \mathcal{A}) = 0$ for $k \leq \nu(l)$ and the long exact homotopy sequence that $\mathcal{C}(l)$ is $\nu(l)$-connected. So let $\varphi\colon S^k \to \mathcal{A}$ be a continuous map. We can assume that it is simplicial. Then also $\psi := i \circ \varphi\colon S^k \to \mathcal{C}(l)$ is simplicial. But every vertex of $\mathcal{C}(l)$ in the image of ψ is of the form $\{i, i+1\}$ with $i > 2$ and thus is joined to $\{1, 2\}$ by an edge. Hence we can homotope ψ within the star of $\{1, 2\}$ to a constant map. This proves the claim. □

2.3 Thompson's group V

2.3.1 Basic definitions and properties

Definition 2.3.1. *Thompson's group V* is the group of all right continuous bijections $f\colon [0, 1] \to [0, 1]$ of the unit interval such that:

- For all but finitely many $x \in [0, 1]$ there is an open neighborhood on which f is a linear function and the exceptional points are dyadic rationals.

- The slope of each linear part is a power of 2.

An element in V is already determined by the following data: A dyadic subdivision D_1 on the domain interval, a dyadic subdivision D_2 on the codomain interval with $|D_2| = |D_1| =: k$, and a permutation $\sigma \in \Sigma_k$ of k elements. Here, $|D|$ denotes the number of subintervals in D. This data defines an element in V by mapping the i'th interval of D_1 linearly onto the $\sigma(i)$'th interval of D_2 (modulo the endpoints which are determined by the

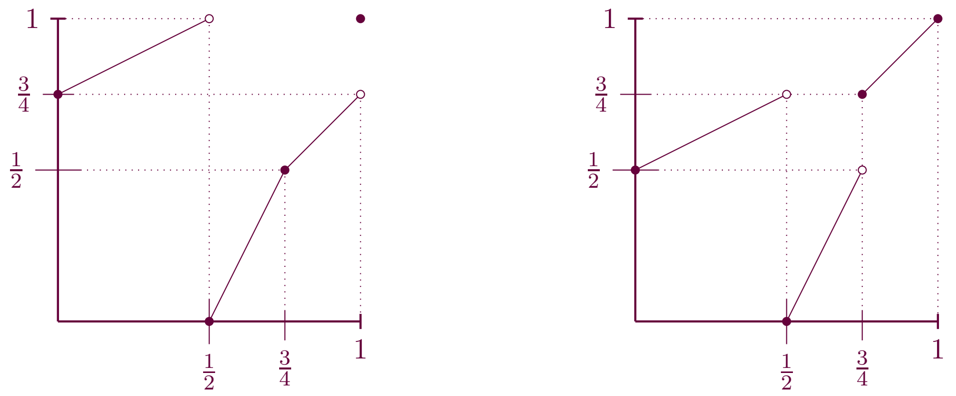

Figure 2.6: Elements of Thompson's group V

right continuity condition). Conversely, for every element $\gamma \in V$, we find dyadic subdivisions D_1, D_2 and a permutation σ which defines γ. The data (D_1, D_2, σ) is only uniquely determined by γ if we demand that D_1 and D_2 are minimal. In this case we say that the number $|D_1| = |D_2|$ is the *degree* of γ.

It is clear that F is a subgroup of V. Indeed, one can show that the two generators of F pictured in Figure 2.1 together with the two elements pictured in Figure 2.6 generate V. In particular, V is finitely generated. A similar but slightly more involved argument as for the group F shows that also V is of type F_∞.

Proposition 2.3.2. *Let G be a finite group. Then G can be embedded into V as a subgroup.*

Proof. Since

$$G \to \Sigma(G) \qquad g \mapsto \big[h \mapsto gh\big]$$

is an embedding into the symmetric group $\Sigma(G)$ on the set G, it suffices to embed the symmetric group $\Sigma_k \cong \Sigma(G)$ on $k = |G|$ elements into V. Let D be any dyadic subdivision of the unit interval with $|D| = k$. Then

$$\Sigma_k \to V \qquad \sigma \mapsto (D, D, \sigma)$$

is an embedding. ☐

Proposition 2.3.3. *V contains a non-abelian free subgroup and, consequently, is non-amenable.*

Proof. We want to play ping-pong in V. Let $X = [0,1]$ and

$$A_1 := \left(0, \frac{1}{2}\right) \subset X \qquad A_2 := \left(\frac{1}{2}, 1\right) \subset X$$
$$B_1 := \left(\frac{1}{2}, \frac{3}{4}\right) \subset X \qquad B_2 := \left(\frac{3}{4}, 1\right) \subset X$$

Let $\gamma_1 \in V$ be the element which interchanges A_1 with A_2 and $\gamma_2 \in V$ be the element which maps A_1 to B_2, B_1 to A_1 and B_2 to B_1. We have that γ_1 is of order 2 and γ_2 is of order 3. Furthermore, it is easy to see that $\gamma_1(A_2) \subset A_1$ as well as $\gamma_2(A_1) \subset A_2$ and $\gamma_2^2(A_1) \subset A_2$. By the ping-pong lemma, we have that

$$\langle \gamma_1, \gamma_2 \rangle \cong \langle \gamma_1 \rangle * \langle \gamma_2 \rangle \cong \mathbb{Z}_2 * \mathbb{Z}_3$$

It is known that the group $\mathbb{Z}_2 * \mathbb{Z}_3 \cong \mathrm{PSL}(2, \mathbb{Z})$ contains non-abelian free subgroups. For example, the commutator of this group is a free subgroup of rank 2 and index 6. □

2.3.2 Simplicity of V

Recall that a group G is simple if whenever $N \triangleleft G$ is a normal subgroup, then either $N = \{1\}$ or $N = G$.

Theorem 2.3.4. *Thompson's group V is simple.*

We start by picking a non-trivial normal subgroup $N \triangleleft V$ and want to show that $N = V$. Let γ be a non-trivial element in N, given by dyadic subdivisions and a permutation (D_1, D_2, σ) as above. Then there is a subinterval J_1 in D_1 which is mapped onto J_2 in D_2 and $J_1 \neq J_2$. If we replace J_1 with either the left or the right half of J_1 and J_2 with $\gamma(J_1)$ if necessary, we can assume that $J_1 \cap J_2 = \emptyset$ (modulo endpoints). Furthermore, we can assume that the length of the interval J_1 is $\leq \frac{1}{4}$. Last but not least, after shorten J_1 even more, we can assume that $J_0 := \gamma^{-1}(J_1)$ is a standard dyadic interval. So we are in the following situation:

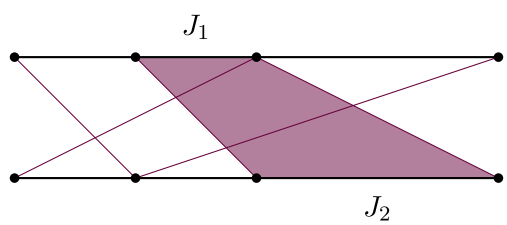

Let g be the element which transposes the two halves of J_1:

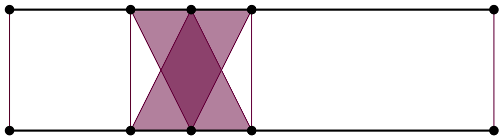

Then one can easily compute that the commutator $h := [\gamma, g] = \gamma^{-1}g^{-1}\gamma g \in N$ is given by transposing the two halves of J_1 and the two halves of J_0 (see Figure 2.7).

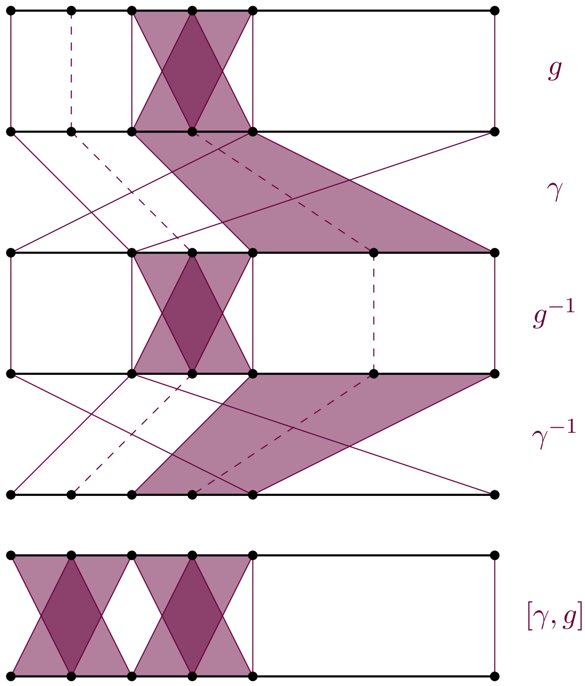

Figure 2.7: A commutator in V

Let k be the element which transposes the left half of J_0 with the left half of J_1:

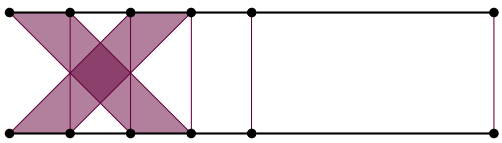

Then the commutator $[h, k] = h^{-1}k^{-1}hk \in N$ transposes the two intervals J_0 and J_1 (see Figure 2.8).

All in all, we have obtained a proper transposition in N (properness follows from the assumption that the length of the interval J_1 is $\leq \frac{1}{4}$). This means the following:

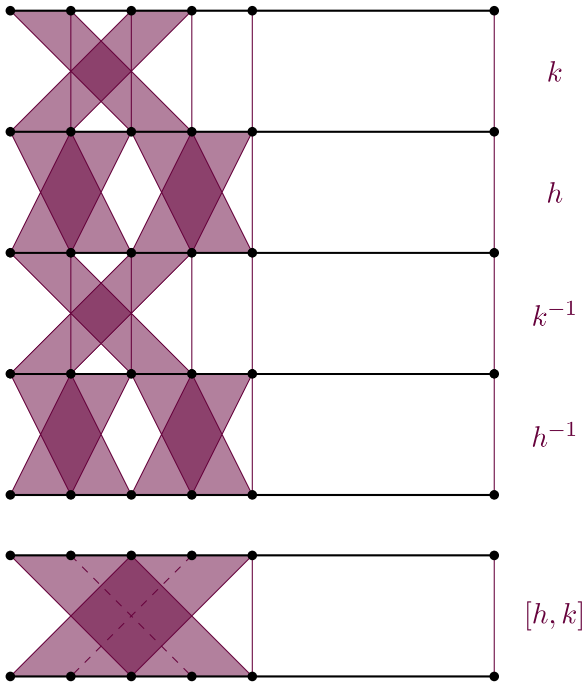

Figure 2.8: Another commutator in V

Definition 2.3.5. A *transposition* in V is an element that interchanges two standard dyadic subintervals of $[0, 1]$ and fixes the rest of $[0, 1]$. A transposition is called *proper* if it leaves at least one standard dyadic subinterval fixed, i.e. is not the transposition which interchanges the two halves of $[0, 1]$.

We proceed with the proof of Theorem 2.3.4.

Lemma 2.3.6. Let γ be a proper transposition. Then any other proper transposition τ is conjugate to γ.

Proof. Let I_1, I_2 be the two standard dyadic intervals which are swapped by γ and J_1, J_2 the two standard dyadic intervals which are swapped by τ. Then define $\nu \in V$ by sending J_1 to I_1 and J_2 to I_2 and complete it to an element of V in an arbitrary way. Note that this is only possible since both γ and τ are assumed to be proper. Then we have $\tau = \nu^{-1}\gamma\nu$, see Figure 2.9. □

By this lemma we now know that our normal subgroup N contains all proper transpositions. The following lemma says that then N has to contain any element of V. Hence $N = V$ and we are done.

Lemma 2.3.7. The proper transpositions generate V.

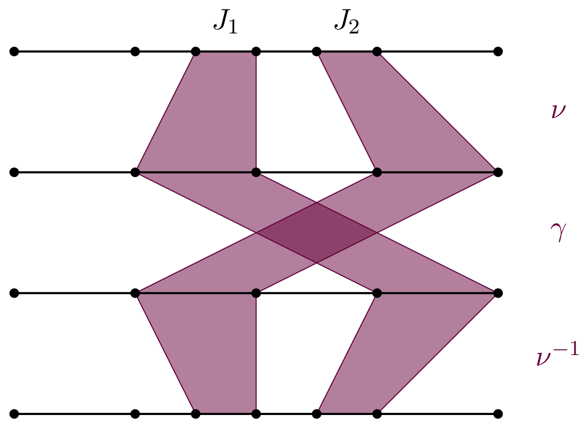

Figure 2.9: Proper transpositions in V

Proof. Step 1: Let $\gamma = (D_1, D_2, \sigma)$ be a non-trivial element in V with D_1 and D_2 minimal. In D_i, we can find two neighboring standard dyadic intervals J_i^l, J_i^r such that $J_i^l \cup J_i^r$ is another standard dyadic interval. Let $k = |D_1| = |D_2|$ and $j_i^l \in \{1, \ldots, k\}$ resp. $j_i^r \in \{1, \ldots, k\}$ be the index of J_i^l resp. J_i^r in the subdivision D_i. Now let $\sigma' \in \Sigma_k$ be a permutation with $\sigma'(j_1^l) = j_2^l$ and $\sigma'(j_1^r) = j_2^r$. Set $\tau = \sigma^{-1} \circ \sigma'$. Then we have

$$(D_1, D_2, \sigma) \cdot (D_1, D_1, \tau) = (D_1, D_2, \sigma \circ \tau) = (D_1, D_2, \sigma')$$

Note that in (D_1, D_2, σ') we can reduce each of the two subdivisions D_i by merging the subintervals J_i^l, J_i^r to one subinterval $J_i^l \cup J_i^r$. Hence we obtain $\gamma = \gamma' \cdot \alpha$ where α is of the form (D, D, ρ) and γ' is an element of V of lesser degree than γ. Repeating this with γ' we arrive at the trivial element after finitely many steps. So we have $\gamma = \alpha_k \cdots \alpha_1$ where each α_i is of the form (D_i, D_i, ρ_i).

Step 2: Now we want to see that an α of the form (D, D, ρ) can be written as a product of proper transpositions. By further subdividing D we can assume that D has at least 3 subintervals. Then write $\rho = \rho_k \circ \ldots \circ \rho_1$ where the ρ_i are transpositions. Then we have

$$(D, D, \rho) = (D, D, \rho_k) \cdots (D, D, \rho_1)$$

and the elements (D, D, ρ_i) are proper transpositions. $\square$

Index

Bibliography

[1] J. Belk. *Thompsons' group F.* ProQuest LLC, Ann Arbor, MI, 2004. Thesis (Ph.D.)–Cornell University.

[2] R. Benedetti and C. Petronio. *Lectures on hyperbolic geometry.* Universitext. Springer-Verlag, Berlin, 1992.

[3] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.

[4] M. Brin. Higher dimensional Thompson groups. *Geom. Dedicata*, 108:163–192, 2004.

[5] K. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1982.

[6] K. Brown. *Buildings*. Springer-Verlag, New York, 1989.

[7] D. Burago, Yu. Burago, and S. Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001.

[8] R. Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*. Springer, New York, 2008.

[9] S. M. Gersten. Bounded cocycles and combings of groups. *Internat. J. Algebra Comput.*, 2(3):307–326, 1992.

[10] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[11] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, english edition, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[12] A. Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[13] J. M. Lee. *Riemannian manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. An introduction to curvature.

[14] C. Löh. Geometric group theory, an introduction. Lecture notes, Regensburg, 2014.