Rigidity Theorems for Hyperbolic Groups

Part III Essay
submitted by

Tobias Hirsch

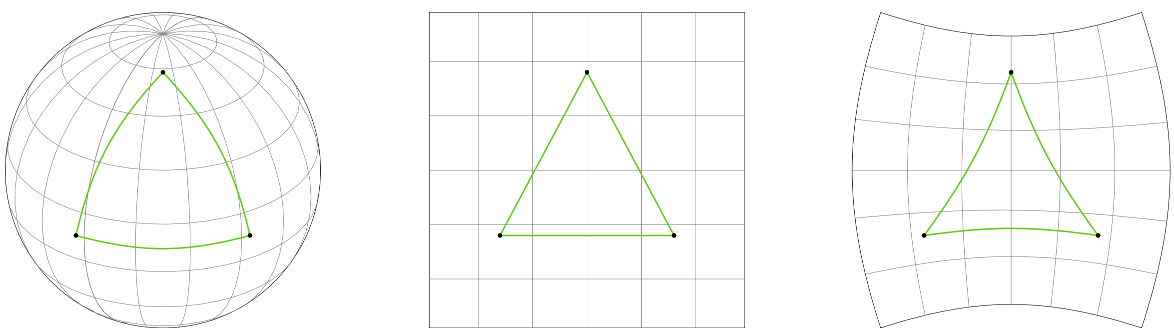

set by

Prof Henry Wilton

at the

University of Cambridge

1 Introduction

Rigidity phenomena appear in many places in geometric topology. They generally state that the type of some geometric structure possible on a manifold is heavily determined only by its topology. One such phenomenon is Mostow rigidity [Mos68: Theorem 12.1, Corollary 12.3] which implies in particular that a closed smooth manifold of dimension at least 3 admits at most one structure of a hyperbolic manifold and that this structure, if it exists, is already completely determined by the fundamental group:

Theorem (Mostow rigidity). *Let M and N be closed hyperbolic manifolds of dimension at least 3. If there exists an isomorphism $\pi_1(M) \to \pi_1(N)$, it is induced by a unique isometry $M \to N$.*

After observing that hyperbolic manifolds are Eilenberg-MacLane spaces of type $K(G, 1)$, it follows from Mostow rigidity that the isometry group of a hyperbolic manifold is isomorphic to its group of self-homotopy equivalences. Since the isometry group of a closed hyperbolic manifold is always finite, we obtain the following corollary:

Corollary 4.13. *A closed hyperbolic manifold of dimension at least 3 has finite group of self-homotopy equivalences.*

In this essay we will approach this statement from a different direction without using Mostow rigidity. The path to avoiding it is to notice that since a hyperbolic manifold is an Eilenberg-MacLane space, its group of self-homotopy equivalences is also isomorphic to the outer automorphism group of its fundamental group. Since we know that the fundamental groups of hyperbolic manifolds are precisely the discrete subgroups of $\mathrm{Isom}(\mathbb{H}^n)$ that act freely and properly discontinuously on $\mathbb{H}^n$ (see Theorem 4.7), we have now turned our original topological question into a question about outer automorphism groups of certain groups acting on $\mathbb{H}^n$. It turns out that it is fruitful to consider a somewhat larger class of groups encompassing any group acting cocompactly and properly discontinuously on a proper hyperbolic metric space – i.e. hyperbolic groups in the sense of Gromov [Gro87].
Their key property for us is that they also exhibit a rigidity phenomenon – namely, mostly having finite outer automorphism group:

Corollary 3.3. *A hyperbolic group has finite outer automorphism group unless it splits over a virtually cyclic subgroup.*

Proving this involves two steps: By work of Bestvina, Feighn [BF95] on the Rips machine, hyperbolic groups acting on $\mathbb{R}$-trees with virtually cyclic arc stabilizers and no global fixed points split over a virtually cyclic subgroup (see Theorem 3.2). We will not explore this part further. Instead, we focus on the other half of the required argument which is encapsulated by Paulin's Theorem [Pau91]:

Theorem 3.1. *Let Γ be a hyperbolic group. If the outer automorphism group $\mathrm{Out}(\Gamma)$ is infinite, Γ acts on an $\mathbb{R}$-tree with virtually cyclic arc stabilizers and no global fixed points.*

The central point in the proof we give for this is the idea of degenerating a sequence of actions on hyperbolic spaces to an action on an $\mathbb{R}$-tree: An infinite sequence of outer automorphisms of Γ induces a sequence of actions of Γ on its Cayley graph. In a suitable limit, these will converge to the desired action on an $\mathbb{R}$-tree.
The structure of this essay is now as follows: We develop the required background in an introductory chapter beginning with hyperbolic metric spaces. We define hyperbolicity for an arbitrary metric space, discuss that it is a quasi-isometry invariant on geodesic spaces and establish the connection between 0-hyperbolic spaces and $\mathbb{R}$-trees. We then turn to hyperbolic groups, where we are primarily interested in ideas pertinent to their subgroups. We study the notion of quasi-convex subgroups which allows us

to show that centralizers in hyperbolic groups are small leading to a proof that no hyperbolic group contains $\mathbb{Z}^2$.

The second chapter is dedicated to Paulin's Theorem: We give the statement and some immediate consequences and then dive into the proof as indicated above. This is divided into two parts: We first develop the necessary notion of convergence by associating to a Γ-space a pseudometric on Γ and prove a criterion for when sequences converge to $\mathbb{R}$-trees. After that we specialize to a sequence of Γ-actions on its Cayley graph and carefully analyse the arc stabilizers to complete the argument.

In the last chapter we bring these ideas together with hyperbolic manifolds. We first construct the pointed and unpointed group of self-homotopy equivalences of a topological space and study how they relate to each other and the fundamental group. We then calculate them for Eilenberg-MacLane spaces and use the argument indicated in the beginning to prove Corollary 4.13.

To give ourselves the space to properly explore the results of the later chapters, our path through the first chapter will be quite stringent and we will not take the time to discover these very interesting areas more than we need to for proving Paulin's Theorem.

Contents

2 General preparations

In this first chapter we set up some basic definitions needed later as stated in [Bes02: Section 2]. In this we generally follow the development of [BH99: Chapters III.H, III.Γ] with some adjustments.
Before we start properly, we establish some basic conventions and notations that do not warrant a dedicated introduction later:
Let X be a metric space. For a subset $A \subseteq X$ we denote by $\overline{\mathrm{N}}_\delta(A) \coloneqq \{x \in X \mid \exists a \in A : d(x,a) \leq \delta\}$ its *closed δ-neighbourhood*. The *diameter of* A is $\operatorname{diam}(A) \coloneqq \sup\{d(a,b) \mid a,b \in A\}$ where we take the diameter of the empty set to be 0.
For a group G and elements $g,h \in G$ we write $[g,h] \coloneqq ghg^{-1}h^{-1}$ for their *commutator*. We denote by $\mathrm{Z}(G) \coloneqq \{g \in G \mid \forall h \in G : [g,h] = e\}$ the *centre of* G and by $\mathrm{C}_G(g) \coloneqq \{h \in G \mid [g,h] = e\}$ the *centralizer of* g *in* G. For a generating set S of G we denote by $\mathrm{Cay}_S(G)$ the corresponding *Cayley graph*.
If a group acts on a metric space it is assumed that this action is by isometries, in particular, *G-space* always refers to a metric space with a G-action by isometries.

2.1 Hyperbolic metric spaces

We now properly begin by developing a notion of hyperbolicity for general metric spaces. We want to view groups acting on such hyperbolic metric spaces as a generalization of groups acting on the standard hyperbolic space $\mathbb{H}^n$. Since we consider discrete groups, we need to capture large scale geometric properties of $\mathbb{H}^n$ and translate them to the realm of metric spaces. To get an idea of how this can be achieved, we look to the 2-dimensional case:

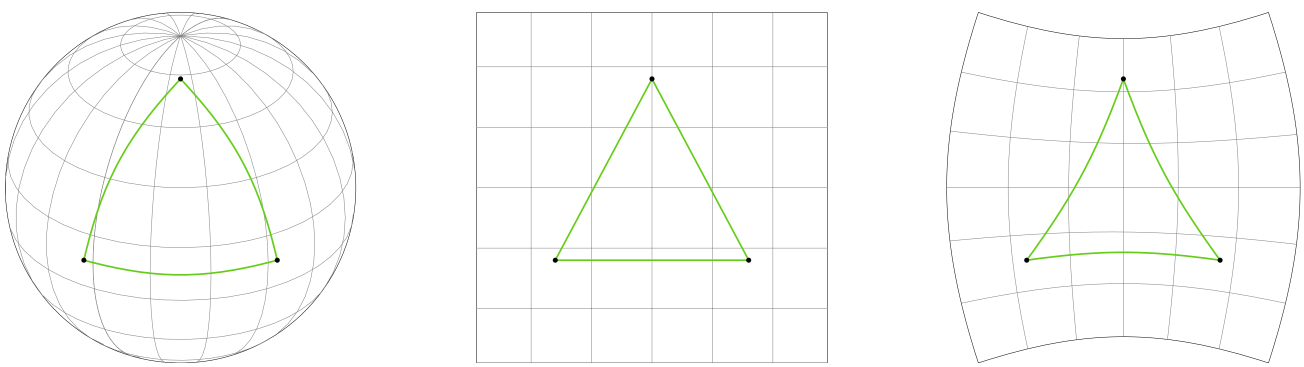

Here there are three possible geometries, distinguished by the appearance of their triangles (or really any polygon): In the usual uncurved Euclidean geometry the triangles look as we expect, in the positively curved elliptic geometry they look thicker and in the hyperbolic geometry they look thinner. Our goal is now to capture this intuitive idea precisely. We begin by defining which triangles we want to consider:

Definition. Let X be a metric space and $x,y \in X$. A *geodesic* from x to y is an isometric embedding $\gamma\colon [a,b] \to X$ with $\gamma(a) = x$ and $\gamma(b) = y$. A *geodesic segment* $[x,y]$ between x and y is the image of such a geodesic[1].
The metric space X is *geodesic* if there exists a geodesic between any two points in X.
A *geodesic triangle* spanned by $x,y,z \in X$ consists of geodesic segments $[x,y],[x,z],[y,z]$.

There are three different ways to capture the concept of thinness of triangles in a geodesic metric space. We establish them in the following construction. Eventually, it will conspire that they lead to equivalent notions.

[1] Some care is needed when using the notation $[x,y]$, as it might make the geodesic segment appear unique.

Construction 2.1. Let X be a geodesic metric space and Δ be a geodesic triangle in X spanned by $x, y, z \in X$

- slim triangles:
  The geodesic triangle Δ is δ-*slim* if $[a,b] \subseteq \overline{\mathrm{N}}_\delta([a,c] \cup [c,b])$ for all permutations (a,b,c) of (x,y,z).
- thin triangles:
  The *Gromov product* of y, z with respect to x is

$$(y \cdot z)_x \coloneqq \frac{1}{2}\big(d(x,y) + d(x,z) - d(y,z)\big)$$

  This is defined precisely such that

$$d(x,y) = (y \cdot z)_x + (x \cdot z)_y \qquad d(x,z) = (y \cdot z)_x + (x \cdot y)_z \qquad d(y,z) = (x \cdot z)_y + (x \cdot y)_z$$

  The *associated tripod* T_Δ of Δ is the metric space obtained by wedging intervals of length $(y \cdot z)_x, (x \cdot z)_y$ and $(x \cdot y)_z$. These equalities therefore imply that there is a map

$$\chi_\Delta \colon [x,y] \cup [x,z] \cup [y,z] \to T_\Delta$$

  that restricts to an isometry on $[x,y]$, $[x,z]$ and $[y,z]$. The geodesic triangle Δ is δ-*thin* if for all $t \in T_\Delta$ the diameter of $\chi_\Delta^{-1}(t)$ is at most δ.

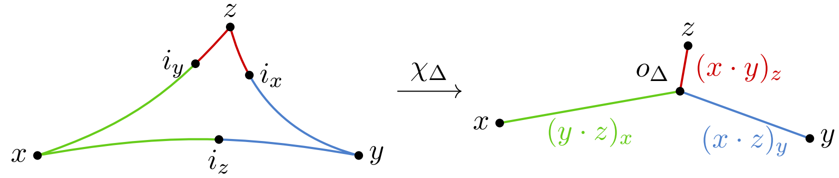

- insize of triangles:
  There is precisely one point $o_\Delta \in T_\Delta$ whose preimage $\chi_\Delta^{-1}(o_\Delta)$ contains a point on all three sides of the triangle, i.e. $\chi_\Delta^{-1}(o_\Delta) = \{i_x, i_y, i_z\}$ with $i_x \in [y,z], i_y \in [x,z], i_z \in [x,y]$. The i's are the *internal points* of Δ. The *insize* of Δ is the diameter of $\chi_\Delta^{-1}(o_\Delta)$.

With this terminology in place, we can now capture our intuitive idea of the sides of a triangle being thin in three ways by requiring that for a fixed $\delta \geq 0$ all geodesic triangles in X are δ-thin, or δ-slim or have insize at most δ.

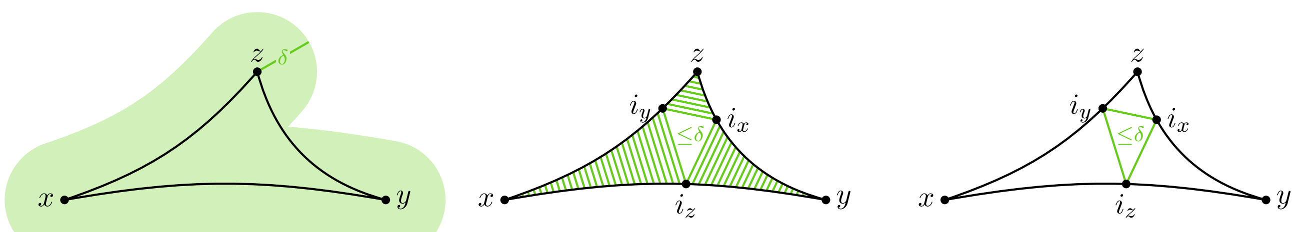

δ-slim $\qquad$ δ-thin $\qquad$ insize $\leq \delta$

There is a further way of capturing negative curvature in a metric space. It is a bit more abstract, but more general, since it applies even when the space is not geodesic. Hence, we will use it as our definition:

Definition. Let $\delta \geq 0$. A metric space X is δ-*hyperbolic* if for all $w, x, y, z \in X$

$$(x \cdot y)_w \geq \min\{(x \cdot z)_w, (y \cdot z)_w\} - \delta$$

We say X is *hyperbolic* if X is δ-hyperbolic for some $\delta \geq 0$.

The geometry behind this harder to intuit. If we assume that the minimum on the right is $(x \cdot z)_w$ it means that the signed distance $l = (x \cdot z)_w - (x \cdot y)_w$ in the figure below is at most δ.

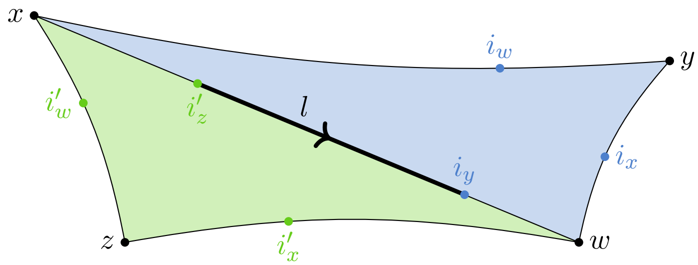

Thankfully, we do not have to work much with this abstract definition since we mostly care about geodesic spaces. In these the following proposition allows us to consider the much more concrete notions from Construction 2.1:

Proposition 2.2. *Let X be a metric space. The following are equivalent:*

(1) *There exists a $\delta \geq 0$ such that X is δ-hyperbolic.*
(2) *There exists a $\delta \geq 0$ such that X for all $w, x, y, z \in X$*

$$d(x, y) + d(w, z) \leq \max\{d(x, z) + d(w, y), d(y, z) + d(w, x)\} + 2\delta$$

If X is geodesic, the following are also equivalent to the above:

(3) *There exists a $\delta \geq 0$ such that all geodesic triangles in X are δ-slim.*
(4) *There exists a $\delta \geq 0$ such that all geodesic triangles in X are δ-thin.*
(5) *There exists a $\delta \geq 0$ such that all geodesic triangles in X have insize at most δ.*

These δ's depend on each other up to constant multiple and not on the metric space. In particular, if one of the δ's can be chosen to be 0, the others can be too.

Proof. The condition in (2) is obtained from the definition of hyperbolicity by multiplying both sides by -2 and adding $d(w, x) + d(w, y) + d(w, z)$. Hence, (1) and (2) are equivalent.
Since (1) and (2) can only easily be related to internal points, we show that they are equivalent to (5): Assume (5) holds and let $w, x, y, z \in X$. Wlog $\min\{(x \cdot z)_w, (y \cdot z)_w\} = (x \cdot z)_w$ and we need to show

$$(x \cdot y)_w \geq (x \cdot z)_w - \delta$$

This is clear if $d(w, i_y) = (x \cdot y)_w \geq (x \cdot z)_w = d(w, i'_z)$, so assume this does not hold. Consider geodesic triangles Δ spanned by x, w, z and Δ' spanned by x, w, y with interior points i_x, i_w, i_z and i'_x, i'_w, i'_y as in the figure above. Then

$$\begin{aligned} d(y, z) &= d(y, i_x) + d(i_x, i_y) + d(i_y, i'_z) + d(i'_z, i'_w) + d(i'_w, z) \overset{\text{triangle inequality}}{\leq} d(y, i_x) + d(i_y, i'_z) + d(i'_w, z) + 2\delta \\ &\underset{d(w, i_y) < d(w, i'_z)}{=} d(y, i_x) + \big(d(x, i_y) + d(w, i'_z) - d(x, w)\big) + d(i'_w, z) + 2\delta = d(x, y) - d(x, w) + d(w, z) + 2\delta \end{aligned}$$

Adding $d(x, w) + d(w, y) - d(y, z) - d(x, y)$ to both sides and dividing by -2 shows the claim.
Now assume (2) holds. Let Δ be a geodesic triangle spanned by $x, y, z \in X$. Let $i_x \in [y, z], i_y \in [x, z]$, $i_z \in [x, y]$ be its internal points and $P \coloneqq d(x, y) + d(y, z) + d(z, x)$ its perimeter. Observe that by construction of the internal points and the triangle inequality

$$d(z, i_x) + d(x, y) = d(y, i_x) + d(x, z) = d(x, i_z) + d(y, z) = \frac{P}{2}$$

$$d(x, i_x) + d(y, z) = \frac{1}{2}\Big(\underbrace{d(x, i_x) + d(i_x, z)}_{\geq d(x,z)} + \underbrace{d(x, i_x) + d(i_x, y)}_{\geq d(x,y)} + d(y, z)\Big) \geq \frac{P}{2}$$

Applying (2) to x, i_x, y, z therefore yields that

$$d(x, i_z) + d(y, z) \leq d(x, i_x) + d(y, z) \leq \frac{P}{2} + 2\delta = d(x, i_z) + d(y, z) + 2\delta$$

and thereby $0 \leq d(x, i_x) - d(x, i_z) \leq 2\delta$. Analogously, $0 \leq d(z, i_z) - d(z, i_x) \leq 2\delta$. Therefore by the triangle inequality

$$d(x, i_x) + d(z, i_z) \leq d(z, i_x) + d(x, i_z) + 4\delta = d(x, z) + 4\delta$$

Then by (2) applied to i_x, i_z, x, z

$$d(i_x, i_z) + d(x, z) \leq \max\{\overbrace{(i_x, z) + d(x, i_z)}^{=d(x,z)}, \overbrace{d(i_z, z) + d(x, i_x)}^{\leq d(x,z)+4\delta}\} + 2\delta \leq d(x, z) + 6\delta$$

implying $d(i_x, i_z) \leq 6\delta$. Analogous arguments show that the insize of Δ is at most 6δ.
The equivalence of (3)-(5) is a much nicer geometric argument:
Clearly, (4) implies (3). To show that (3) implies (5) consider again the geodesic triangle Δ. By assumption $i_x \in \overline{\mathrm{N}}_\delta([x, y])$ or $i_x \in \overline{\mathrm{N}}_\delta([x, z])$. Suppose $d(i_x, p) \leq \delta$ for some $p \in [x, y]$. Then

$$d(i_z, p) \overset{\substack{[x,y] \text{ geodesic} \\ \downarrow}}{=} |d(y, p) - d(y, i_z)| \overset{\substack{\text{construction of internal points} \\ \downarrow}}{=} |d(y, p) - d(y, i_x)| \overset{\substack{\text{triangle inequality} \\ \downarrow}}{\leq} d(p, i_x) \leq \delta$$

and therefore by the triangle inequality $d(i_x, i_z) \leq 2\delta$. Taking into account the case $i_x \in \overline{\mathrm{N}}_\delta([x, z])$, we have shown that $d(i_x, \{i_y, i_z\}) \leq 2\delta$. Since the same inequality holds for any permutation of x, y and z, it follows that the insize of Δ is at most 4δ.
To see that (5) implies (4) consider again Δ as above. Let $t \in T_\Delta$. We claim $\operatorname{diam} \chi_\Delta^{-1}(t) \leq \delta$: This is by assumption if $t = o_\Delta$, so assume wlog $t \neq o_\Delta$ is on the interval connecting o_Δ and x in T_Δ. Then $\chi_\Delta^{-1}(t) = \{p, q\}$ with $p \in [x, y]$ between x and i_z and $q \in [x, z]$ between x and i_y. Consider an isometry $f\colon [0, c] \to [x, y]$. For $h \in [0, C]$ let Δ_h be a geodesic triangle with two sides $[x, z]$ and $f([0, h])$. The internal point $i_h \in f([0, h])$ varies continuously along $[x, y]$ as it can be defined by the Gromov product. Since $i_0 = x$ and $i_c = i_z$, it follows from the Intermediate Value Theorem that there exists $h \in [0, c]$ such that $i_h = p$. Then q is also an internal point of Δ_h and the claim follows. ■

Having now developed the definition of hyperbolicity in metric spaces, we can give some (non-)examples:

Examples 2.3.
- The hyperbolic plane $\mathbb{H}^2$ is hyperbolic: A geodesic triangle Δ in $\mathbb{H}^2$ is δ-slim where δ is the radius of the largest semicircle within Δ. The area of such a semicircle is bounded above by the area of Δ which is less than π. Hence, δ is less than the radius of a semicircle with area π in $\mathbb{H}^2$.
  One can more generally show that the hyperbolic space $\mathbb{H}^n$ is indeed hyperbolic.
- Metric spaces of finite diameter δ – in particular, finite and compact spaces – are δ-hyperbolic.
- If a metric space X is a tree (i.e. it can be given the structure of a 1-dimensional CW-complex and two points in its 0-skeleton are connected by a unique arc), geodesic triangles are tripods. Hence, X is 0-hyperbolic.
- The last example in particular yields that $\mathbb{R}$ is hyperbolic. In contrast, $\mathbb{R}^n$ is not hyperbolic for $n \geq 2$ since an equilateral triangle of side length a has insize $a/2$.

Hyperbolicity in geodesic metric spaces has a further advantage, namely that it is a quasi-isometry invariant:

Theorem 2.4 ([BH99: Theorem 1.9])**.** *Let X, Y be geodesic metric spaces that are quasi-isometric. Then X is hyperbolic if and only if Y is.*

Sketch of a proof. We will not give a full proof since this would require developing the notions of quasi-geodesics and Hausdorff-distance between subsets of a metric space. But due to the importance of the theorem, we still want to sketch the required argument:
The key insight is the Morse Lemma which gives that in a hyperbolic metric space, quasi-geodesics are Hausdorff-close to actual geodesics. This implies that all geodesic triangles in a geodesic metric space X are thin if and only if all quasi-geodesic triangles in X are thin in a suitable sense. As a quasi-isometry maps quasi-geodesic triangles to quasi-geodesic triangles, the claim follows. ∎

As indicated hyperbolicity is not a quasi-isometry invariant when it comes to arbitrary metric spaces:

Example 2.5. The map $f\colon \mathbb{R} \to \mathbb{R}^2$ with $f(t) := (t, |t|)$ is a $(\sqrt{2}, 0)$-quasi-isometric embedding and $\mathbb{R}$ is 0-hyperbolic. But its image endowed with the subspace metric is not hyperbolic by a direct calculation.

The following proposition allows us to replace a metric space with a nicer – say, geodesic – metric space in some situations:

Proposition 2.6. *Let $f\colon X \to Y$ be an isometric embedding of metric spaces.*
(1) *If Y is δ-hyperbolic, then X also is δ-hyperbolic.*
(2) *If f is quasi-surjective and X is hyperbolic, then Y also is hyperbolic.*

Proof. (1) follows directly from the definition. For (2) let $K \geq 0$ such that for all $y \in Y$ there exists $x \in X$ with $d_Y(y, f(x)) \leq K$ and suppose X is δ-hyperbolic. Let $w, x, y, z \in Y$. There exist $w', x', y', z' \in Y$ such that $d_Y(w, f(w')), d_Y(x, f(x')), d_Y(y, f(y'))$ and $d_Y(z, f(z'))$ are it most K. Then by the triangle inequality

$$(y \cdot z)_x \geq \frac{1}{2}\big((d_X(x', y') - 2K) + (d_X(x, z) - 2K) - (d_X(y, z) + 2K)\big) = (y' \cdot z')_{x'} - 3K$$

Hence,

$$(x \cdot y)_w \geq (x' \cdot y')_{w'} - 3K \geq \min\{(x' \cdot z')_{w'}, (y' \cdot z')_{w'}\} - \delta - 3K \geq \min\{(x \cdot z)_w, (y \cdot z)_w\} - \delta - 6K$$

and Y is $(\delta + 6K)$-hyperbolic. ∎

In our later discussion of hyperbolic groups, we will need the following lemma on geodesics:

Lemma 2.7. *Let X be a metric space with a δ-slim geodesic triangle Δ spanned by $x, y, z \in X$. Let $\gamma\colon [0, l] \to X, \gamma'\colon [0, l'] \to X$ be the geodesics in Δ with $x = \gamma(0) = \gamma'(0)$. Then $d(\gamma(t), \gamma'(t)) \leq 2\delta + d(y, z)$ for $t \in [0, \max\{l, l'\}]$ where we set $\gamma(t) := \gamma(l)$ for $t \geq l$ and analogously for γ'.*

Proof. By δ-slimness of triangles we only need to consider the following cases:

Case 1: There exists $t' \in [0, l']$ such that $d(\gamma(t), \gamma'(t')) \leq \delta$.
Then

$$|t' - t| \overset{\gamma' \text{ geodesic}}{=} |d(\gamma'(0), \gamma'(t')) - d(\gamma'(0), \gamma'(t))| = |d(\gamma'(0), \gamma'(t')) - d(\gamma(0), \gamma(t))| \overset{\text{triangle inequality}}{\leq} d(\gamma(t'), \gamma(t)) \leq \delta$$

so $d(\gamma'(t'), \gamma'(t)) \leq \delta$. By the triangle inequality $d(\gamma(t), \gamma(t')) \leq 2\delta$.

Case 2: There exists $t' \in [0, l]$ such that $d(\gamma'(t), \gamma(t')) \leq \delta$.
Analogous to Case 1 with γ and γ' switched.

Case 3: There exist $p, q \in [\gamma(l), \gamma'(l')]$ such that $d(\gamma(t), p), d(\gamma'(t), q) \leq \delta$
Then clearly by the triangle inequality

$$d(\gamma(t), \gamma'(t)) \leq d(\gamma(t), p) + d(p, q) + d(\gamma'(t), q) \leq 2\delta + d(\gamma(l), \gamma(l'))$$ ∎

We now turn our attention to 0-hyperbolic spaces. We have seen that trees are examples for such spaces. It was somewhat arbitrary in that example to require the metric space to allow the structure of a CW-complex, and we therefore generalize this notion to arbitrary geodesic metric spaces:

Definition. Let X be a metric space and $x, y \in X$. An *arc* between x and y is the image of a topological embedding $\gamma\colon [a, b] \to X$ with $\gamma(a) = x$ and $\gamma(b) = y$. The arc is *degenerate* if it is a point.
An $\mathbb{R}$-*tree* is a geodesic metric space X such that for all $x, y \in X$ there is a unique arc between x and y.

The following example of an $\mathbb{R}$-tree shows that this notion is indeed more general than the graph-theoretic definition of a tree:

Example 2.8. Consider the set $X = \mathbb{R}^2$ with the metric

$$\begin{aligned} d\colon X \times X &\to \mathbb{R}^2 \\ \big((x_1, x_2), (y_1, y_2)\big) &\mapsto \begin{cases} |x_2| + |y_2| - |x_1 - y_1|, & \text{if } x_1 \neq y_1 \\ |x_2 - y_2|, & \text{if } x_1 = y_1 \end{cases} \end{aligned}$$

Intuitively, this may be understood as a generalization of the French railway metric with trains now operating along the x-axis and all vertical lines[2]. By direct calculation, this makes X an $\mathbb{R}$-tree. But X is not a graph-theoretic tree since in any such tree the set of all points whose complement has at least 3 components is a subset of the vertex set and in particular discrete – whereas in X this set is the x-axis.

Our interest in $\mathbb{R}$-trees stems from their close relation with 0-hyperbolic metric spaces. Essentially, the only obstruction for a 0-hyperbolic metric space to be an $\mathbb{R}$-tree is connectedness and one can turn any 0-hyperbolic metric space into an $\mathbb{R}$-tree in a unique way by connecting up the components:

Proposition 2.9 (Connecting-the-dots Proposition). *Let X be a 0-hyperbolic metric space. Then there exists an $\mathbb{R}$-tree T and an isometric embedding $i\colon X \to T$ such that*
- *no proper sub-$\mathbb{R}$-tree of T contains $i(X)$ and*
- *if $j\colon X \to T'$ is an isometric embedding of X into an $\mathbb{R}$-tree T', there is a unique isometric embedding $k\colon T \to T'$ such that $k \circ i = j$.*

In particular, T is unique up to isometry. If a group acts on X, the action extends to T.

Proof. The claim is clear if $X = \varnothing$, hence we may choose $* \in X$. If such a T exists, it must be the union of the geodesic from $*$ to the points of $i(X)$. Therefore,

$$Y \coloneqq \bigsqcup_{x \in X} \underbrace{\{(x, h) \mid h \in [0, d(*, x)]\}}_{=: I_x}$$

surjects onto T. This motivates the idea of defining T as a quotient of Y:
Consider the equivalence relation on Y generated by $(x, h) \sim (y, h)$ for $x, y \in X$ and $h \in [0, (x \cdot y)_*]$ and set $T \coloneqq Y/\sim$. To equip T with a metric, consider Y with the extended metric

$$\begin{aligned} \hat{d}\colon Y \times Y &\to \mathbb{R}_{\geq 0} \cup \{\infty\} \\ \big((x, h), (y, t)\big) &\mapsto \begin{cases} |h - t|, & \text{if } x = y \\ \infty, & \text{else} \end{cases} \end{aligned}$$

and consider the usual quotient metric d on T (see [BH99: Definitions 5.19, 5.23]).

[2]Note that the French railway metric itself also defines an $\mathbb{R}$-tree – but that one is in fact a graph-theoretic tree.

Since every point in Y has finite distance to only a compact set of points and every point in Y has finite distance to a point in the equivalence class $[*, 0]$, this is indeed a well-defined metric and the quotient topology on T aligns with the topology induced by d. More precisely, for $[x, h], [y, t] \in T$

$$d\big([x,h],[y,t]\big) = \inf_{m \in T}\Big\{\hat{d}((x,h),(z,s)) + \hat{d}((z',s'),(y,t)) \;\Big|\; [z,s],[z,s'] \in m\Big\} = \min_{s\in[0,(x\cdot y)_*]}\{|h-s|+|t-s|\}$$

and the quotient map $Y \to T$ restricts to an isometric embedding on each I_x. It now follows from the Intermediate Value Theorem that the unique arc between $[x, h], [y, t] \in T$ is the image under the quotient map of

$$\big\{(x,h') \;\big|\; h' \text{ between } h \text{ and } s\big\} \cup \big\{(y,t') \;\big|\; t' \text{ between } t \text{ and } s\big\}$$

where $s \in [0, (x \cdot y)_*]$ attains the minimum above. To see that this is a geodesic, observe that we can always take $s \in \{h, t, (x \cdot y_*)\}$. In the first two cases, it is already a geodesic in Y, in the last case, it is the union of two geodesics in Y that are glued together correctly by the definition of the Gromov product. Hence, T is an $\mathbb{R}$-tree. The two further claims of the proposition follow directly from the observation that T is the union of the geodesics in T between $*$ and $i(x)$ for $x \in X$.
The "in particular" statement follows in the usual manner when defining a space via a universal property. A group action on X induces a group action on Y that is compatible with $\sim$ since the Gromov product on X is invariant under the action, and thereby descends to a group action on T. It follows from the last equation that this action is compatible with d for the same reason. ■

This proposition gives us the following previously alluded – and surprisingly tricky – corollary:

Corollary 2.10. *A metric space X is an $\mathbb{R}$-tree if and only if it is path-connected and 0-hyperbolic. In particular, path-connected subspaces of $\mathbb{R}$-trees are $\mathbb{R}$-trees.*

Proof. Let X be a path-connected 0-hyperbolic metric space. By the Connecting-the-dots Proposition 2.9 there exists an $\mathbb{R}$-tree T and an isometric embedding $i\colon X \to T$ such that no proper sub-$\mathbb{R}$-tree of T contains $i(X)$. Suppose $i\colon X \to T$ is not surjective, i.e. there exists $p \in T \setminus i(X)$. Since $i(X)$ is path-connected, there exists a path-component T' of $T \setminus \{p\}$ such that $i(X) \subseteq T'$. By the Hahn–Mazurkiewicz Theorem (see [Wil70: Section 31]) T' is arc-connected, and therefore a proper sub-$\mathbb{R}$-tree of T. Contradiction! Hence, i is an isometry and X and $\mathbb{R}$-tree.
Now suppose X is an $\mathbb{R}$-tree. Then X is in particular path-connected. Let Δ be a geodesic triangle spanned by $x, y, z \in X$. By compactness, we can consider $p_x \in [x, y] \cap [x, z]$ with $d(x, p_x)$ maximal. We analogously define p_y, p_z. Considering only the part of Δ between the p's gives a geodesic triangle spanned by p_x, p_y, p_z whose sides do not intersect except in the vertices. Since a pair of points in X is connected by a unique arc, it follows that $p_x = p_y = p_z \in [x, y] \cap [x, z] \cap [y, z]$. Hence, the p's are the internal points of Δ which therefore has insize 0 and X is 0-hyperbolic by Proposition 2.2. ■

We will need one further technical construction:

Definition. Let X be a set. A *pseudometric* on X is a map $d\colon X \times X \to \mathbb{R}_{\geq 0}$ such that
- $d(x, x) = 0$ for all $x \in X$
- $d(x, y) = d(y, x)$ for all $x, y \in X$
- $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

A pseudometric is *trivial* if it is identically 0.
The pseudometric descends to a metric on $X/\!\sim$ where $\sim$ is the equivalence relation on X generated by $x \sim y$ if $d(x, y) = 0$. The metric space $X/\!\sim$ is the *associated metric space* of the pseudometric.
A pseudometric space is δ-hyperbolic for $\delta \geq 0$ if the associated metric space is δ-hyperbolic.

In principle, one could now generalize many of the previous results to pseudometric spaces, but we will not do so. Not least due to our definition of hyperbolicity by transferring it via the associated metric space, the few results that we will need, will be immediate.

2.2 Hyperbolic groups

We now turn our attention back to group theory. As motivated in the introduction, we want hyperbolic groups to generalize groups acting on $\mathbb{H}^n$, so it is natural to take the following definition:

Definition. A group Γ is *hyperbolic* if Γ acts properly discontinuously and cocompactly on a proper hyperbolic geodesic metric space.

The next lemma relates this definition closely with Cayley graphs and word metrics:

Proposition 2.11. *Let Γ be a group. The following are equivalent:*

(1) *Γ is hyperbolic.*
(2) *Γ has a finite generating set S such that Γ equipped with the word metric corresponding to S is a hyperbolic metric space.*
(2') *Γ is finitely generated and a hyperbolic metric space when equipped with the word metric corresponding to every finite generating set of Γ.*
(3) *Γ has a finite generating set S such that $\mathrm{Cay}_S(\Gamma)$ is hyperbolic.*
(3') *Γ is finitely generated and $\mathrm{Cay}_S(\Gamma)$ is hyperbolic for every finite generating set S of Γ.*

Proof. (2) and (3) are equivalent by Proposition 2.6, the same holds for (2') and (3'). (3), (3') and (1) are equivalent by the Schwarz-Milnor Lemma and Theorem 2.4. ■

Since we have already seen (non)-examples of hyperbolic metric spaces, it is now quite easy to give at least a few (non)-examples of hyperbolic groups:

Examples 2.12. The examples given here parallel the examples in Examples 2.3:

- Finite groups are hyperbolic because their Cayley graphs are finite.
- The non-abelian free groups are hyperbolic since their Cayley graphs are trees.
- The last example in particular yields that $\mathbb{Z}$ is hyperbolic. In contrast, $\mathbb{Z}^n$ is not hyperbolic for $n \geq 2$ since its Cayley graph with the usual generating set is quasi-isometric to $\mathbb{R}^n$.

Hyperbolic groups are quite a fruitful topic of study as they exhibit several quite interesting properties: For example, they satisfy the Tits alternative, are finitely presented with solvable word problem and play an important role in the study of hyperbolic manifolds. We will discuss the last relationship in Section 4.2, but otherwise only develop their theory in so far as we need to prove Paulin's Theorem. We therefore point to the seminal article [Gro87] or [BH99: Chapter III.Γ].
We will mainly be interested in subgroups of hyperbolic groups, in particular, in the question when are they also hyperbolic. Since the inclusion of a subgroup need not be an isometric embedding, this will not always be the case. For example, hyperbolic groups are finitely generated but already the free groups contain non-finitely generated subgroups. In fact, there are even examples of finitely generated and finitely presented non-hyperbolic subgroups of hyperbolic groups (see [Rip82] and [Bra99]).
We begin with the following obvious criterion for hyperbolic subgroups:

Lemma 2.13. *Let Γ be a hyperbolic group and $H \subseteq \Gamma$ a subgroup.*

(1) *If the inclusion $H \to \Gamma$ is a quasi-isometric embedding in the word metrics corresponding to some finite generating sets on H and Γ, then H also is hyperbolic.*
(2) *The inclusion of a finite index subgroup is a quasi-isometric embedding in all choices of word metrics corresponding to finite generating sets.*

Proof. (1) follows from Theorem 2.4, Proposition 2.6 and Proposition 2.11, (2) follows from the Schwarz-Milnor Lemma. ■

We want to look for more examples of quasi-isometrically embedded subgroups. We have already seen that it can be convenient to be able to work with geodesics. In view of this, we consider the following definition which will give such subgroups:

Definition. A subset A of a metric space X is *C-quasi-convex* for $C \geq 0$ if for $x, y \in A$ every geodesic segment between x and y in X is contained in $\overline{\mathrm{N}}_C(A)$. The subset A is *quasi-convex* if it is quasi-convex for some $C \geq 0$.
Let G be a finitely generated group. A subgroup $H \subseteq G$ is *quasi-convex* if $H \subseteq \mathrm{Cay}_S(G)$ is a quasi-convex for all finite generating sets S of G[3].

We note the following properties of quasi-convex subgroups:

Proposition 2.14. *Let G be a finitely generated group.*

(1) *If $H \subseteq G$ is a quasi-convex subgroup, then H is finitely generated and the inclusion $H \to G$ is a quasi-isometric embedding in all choices of word metrics on H and G corresponding to finite generating sets.*
(2) *If $H_1, H_2 \subseteq G$ are quasi-convex subgroups, then their intersection $H_1 \cap H_2 \subseteq G$ is also a quasi-convex subgroup.*

Proof.

(1) Let S be a generating set of G such that H is C-quasi-convex in $\mathrm{Cay}_S(G)$. We claim that the finite set T of elements of H with word length at most $2C + 1$ generates H:
Let $h \in H$. Then $[e, h]$ is an edge-path in $\mathrm{Cay}_S(G)$. Suppose the edges it traverses are labelled by $s_1, \ldots, s_n \in S$. Since H is C-quasi-convex there exist $u_0, \ldots, u_n \in G$ of word length at most C such that $h_i := u_{i-1} a_s u_n^{-1} \in H$ for $i \in \{1, \ldots, n\}$. By the triangle inequality $h_i \in T$. We can assume that $u_0 = u_n = e$, then $h = s_1 \cdots s_n = h_1 \cdots h_n$ is generated by T.
The inclusion $H \to G$ is quasi-isometric in the word metrics corresponding to T and A as for $h, h' \in H$ by the above
$$\frac{1}{2C+1} d_A(h, h') \leq d_T(h, h') \leq d_A(h, h')$$
The claim now follows since two word metrics corresponding to finite generating sets on the same group are quasi-isometric by the Schwarz-Milnor Lemma.
(2) Let S be a generating set of G such that H_1, H_2 are C-quasi-convex in $\mathrm{Cay}_S(G)$. We claim that $H_1 \cap H_2 \subseteq \mathrm{Cay}_S(G)$ is $(K + 1)$-quasi-convex for
$$K := |\underbrace{\{g \in G \mid d(e, g) \leq C\}}_{=:M}|^2$$
Let $h, h' \in H_1 \cap H_2$. Again $[h, h']$ is an edge-path in $\mathrm{Cay}_S(G)$ and we can suppose the vertices it traverses are labelled by $v_1, \ldots, v_n \in G$. For $i \in \{1, \ldots, n\}$ let $g \in G$ be the element of minimal length such that $v_i \cdot g \in H_1 \cap H_2$. Write $g = s_1 \cdots s_k$ for $s_i \in T$ with k minimal. For $1 \leq j \leq k$ there exist $t_j^{(1)}, t_j^{(2)} \in M$ such that
$$h_j^{(1)} := (v_i \cdot s_1 \cdots s_j) \cdot t_j^{(1)} \in H_1 \qquad \text{and} \qquad h_j^{(2)} := (v_i \cdot s_1 \cdots s_j) \cdot t_j^{(2)} \in H_2$$
since there exists a geodesic segment between elements of $H_1 \cap H_2$ containing $v_i \cdot s_1 \cdots s_j$.

[3]It is convenient for our purposes to require H to be quasi-convex with respect to all finite generating sets, but of course this condition is a priori hard to check. For a hyperbolic group, it suffices to assume H is quasi-convex with respect to one finite generating set, but we did not develop enough theory of hyperbolic spaces to show this (see [BH99: Corollary 3.6]).

Suppose $k > K$. Then by the pigeonhole principle there exist $1 \leq a < b \leq k$ such that $(t_a^{(1)}, t_a^{(2)}) = (t_b^{(1)}, t_b^{(2)})$ as there are only K possibilities for such pairs. Therefore

$$v_i \cdot (s_1 \cdots s_a \cdot s_{b+1} \cdot s_k) = h_a^{(1)} \cdot (t_a^{(1)})^{-1} \cdot (s_{b+1} \cdots s_k) = h_a^{(1)} \cdot (t_b^{(1)})^{-1} \cdot (s_{b+1} \cdots s_k)$$
$$= h_a^{(1)} \cdot (h_b^{(1)})^{-1} \cdot h' \in H_1$$

and analogously $v_i \cdot (s_1 \cdots s_a \cdot s_{b+1} \cdot s_k) \in H_2$. But then $s_1 \cdots s_a \cdot s_{b+1} \cdot s_k$ contradicts the minimality of g. Hence, $d(v_i, H_1 \cap H_2) = k \leq K$. Since by construction of the Cayley graph every point on γ is a distance of at most 1 from v_i for some i, the claim follows. ■

The next proposition gives a large class of examples of quasi-convex subgroups of hyperbolic groups:

Proposition 2.15. *Let Γ be a hyperbolic group and $g \in \Gamma$. The centralizer $\mathrm{C}_\Gamma(g) \subseteq \Gamma$ is quasi-convex.*

Proof. Let S be a finite generating set of G. By Proposition 2.11 and Proposition 2.2 there exists $\delta \geq 0$ such that all geodesic triangles in X are δ-slim.
We begin with the following claim:

Claim. There exists an increasing function $f\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that if $g, h \in \Gamma$ are conjugate and have word length at most M, there exists $v \in \Gamma$ of word length at most $f(M)$ such that $g = vhv^{-1}$.
Let $v \in \Gamma$ be of minimal length such that $g = vhv^{-1}$. Let $\gamma\colon [0, d(e,v)] \to \mathrm{Cay}_S(\Gamma)$ be a geodesic from e to v. It is an edge-path, suppose the vertices it transverses are labelled $v_1, \ldots, v_n$. Let $\gamma'\colon [0, d(e, vh)] \to \mathrm{Cay}_S(\Gamma)$ be a geodesic from e to vh. By Lemma 2.7 for $t \in [0, d(e,v)]$

$$d(\gamma(t), \gamma'(t)) \leq 2\delta + d(e,h)$$
$$d(g \cdot \gamma(t), \gamma'(d(e,vh) - (d(e,v) - t))) \leq 2\delta + d(e,g)$$

and by the triangle inequality

$$d(\gamma'(t), \gamma'(d(e,vh) - (d(e,v) - t))) = |d(e,vh) - d(e,v)| \leq d(e,h)$$

Taken together these inequalities show by the triangle inequality that

$$d(e, v_i^{-1} g v_i) = d(v_i, g v_i) \leq 4\delta + d(e,g) + 2d(e,h) \leq 4(\delta + M)$$

By minimality of v, the elements $v_i^{-1} g v_i \in \Gamma$ for $i \in \{1, \ldots, n\}$ must be pairwise different. The claim follows when defining $f(k)$ to be the number of elements with word length at most $4(\delta + k)$ in Γ since n is the word length of v. □

Let $h \in \mathrm{C}_\Gamma(g)$ and consider $p \in [1, h]$. By construction of the Cayley graph there exists $\overline{p} \in [1, h] \cap \Gamma$ such that $d(p, \overline{p}) \leq 1$. As in the proof of the Claim $d(e, \overline{p}^{-1} g \overline{p}) = d(\overline{p}, g\overline{p}) \leq 4(\delta + d(e,g))$. By the Claim there exists $v \in \Gamma$ such that $\overline{p}^{-1} g \overline{p} = v^{-1} g v$ and

$$d(e,v) \leq f\big(\max\{d(e, p^{-1}gp), d(e,g)\}\big) \leq f(\max\{4(\delta + d(e,g)), d(e,g)\}) =: C$$

Then $\overline{p}v^{-1} \in \mathrm{C}_\Gamma(g)$ and therefore

$$d(p, \mathrm{C}_\Gamma(g)) \leq d(p, \overline{p}v^{-1}) \leq d(\overline{p}, \overline{p}v^{-1}) + 1 = d(e,v) + 1 \leq C + 1$$

This proves the proposition since the definition of C does not depend on p. ■

With this knowledge on centralisers, we can now show that centralizers in hyperbolic groups are small:

Theorem 2.16. *Let Γ be a hyperbolic group and $g \in \Gamma$ of infinite order.*

1. *The inclusion $\langle g \rangle \to \Gamma$ is a quasi-isometric embedding in all choices of word metric corresponding to a finite generating set.*
2. *The subgroup $\langle g \rangle \subseteq \mathrm{C}_\Gamma(\gamma)$ has finite index.*

Proof.

(1) The inclusion $\mathrm{C}_\Gamma(g) \to \Gamma$ is a quasi-isometric embedding by Proposition 2.15 and Proposition 2.14. In particular, $\mathrm{C}_\Gamma(g)$ is hyperbolic by Lemma 2.13 and thereby has a finite generating set T. Since

$$\mathrm{Z}(\mathrm{C}_\Gamma(g)) = \bigcap_{t\in T} \mathrm{C}_\Gamma(t)$$

it follows from Proposition 2.14 that $\mathrm{Z}(\mathrm{C}_\Gamma(g))$ is also quasi-convex with the inclusion $\mathrm{Z}(\mathrm{C}_\Gamma(g)) \to G$ being a quasi-isometric embedding. Hence, again $\mathrm{Z}(\mathrm{C}_\Gamma(g))$ is hyperbolic. Since $\mathrm{Z}(\mathrm{C}_\Gamma(g))$ is abelian, it follows from the classification of finitely generated abelian groups and Lemma 2.13 that $\langle g\rangle \subseteq \mathrm{Z}(\mathrm{C}_\Gamma(g))$ is a finite index subgroup since g is not torsion. Going back up through the subgroups now proves (1) by Lemma 2.13.

(2) Let S be a generating set of Γ. By Proposition 2.11 there exists $\delta \geq 0$ such that all geodesic triangles in $\mathrm{Cay}_S(\Gamma)$ are δ-slim. Let $p, q \in \mathbb{Z}$ such that $t^{-1}g^p t = g^q$ for $t \in \Gamma$. Then $t^{-m}g^{p^m}t^m = g^{q^m}$ for all $m \geq 1$ implying that

$$\lambda|q|^m \cdot d(e, g) + C \leq d(e, g^{q^m}) \leq |p|^m \cdot d(e, g) + 2m \cdot d(1, t)$$

where the inclusion in (1) is a (λ, C)-quasi-isometric embedding with the generating set $\{g\}$ on $\langle g\rangle$. For large enough m this implies $|q| \leq |p|$. Together with a symmetrical argument, $|p| = |q|$. This shows that the powers of g fall into infinitely many distinct conjugacy classes. Hence, we can after replacing g with one of its powers assume that g is not conjugate to any element of distance at most $4\delta + 2$ from the identity as there are only finitely many (conjugacy classes of) such elements and $\mathrm{C}_\Gamma(g) \subseteq \mathrm{C}_\Gamma(g^n)$. Set $K \coloneqq 2d(e, g) + 4\delta$. We now claim $\mathrm{C}_\Gamma(g) \subseteq \overline{\mathrm{N}}_K(\langle g\rangle)$, proving (2):
Suppose there exists $h \in \mathrm{C}_\Gamma(g)$ with $d(h, \langle g\rangle) > K$. After replacing h with its multiplication by the inverse of its closest power of g, we may assume that $d(h, \langle g\rangle) = d(g, 1)$. Consider the quadrilateral with sides $[e, g], [e, h], g \cdot [e, h], h \cdot [e, g]$. This really is a quadrilateral since $hg = gh$.

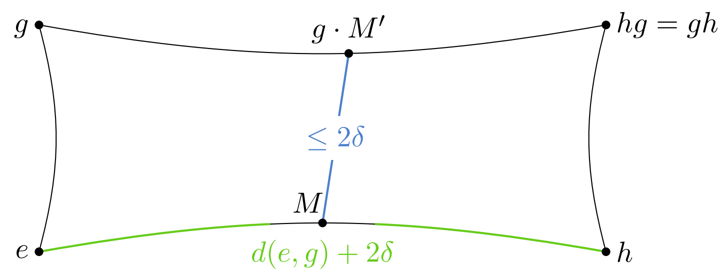

Let M be the midpoint of $[e, h]$. By applying δ-slimness of triangles twice, it follows that there is a point of distance at most 2δ from M on one of the other sides. By choice of K, M is at least a distance of $d(e, g) + 2\delta = d(h, gh) + 2\delta$ from e and h, hence by the triangle inequality this point must be on $g \cdot [e, h]$, i.e. it equals $g \cdot M'$ for some $M' \in [e, h]$. By the triangle inequality and since the sides are geodesics

$$d(M', M) = |d(\langle g\rangle, M) - d(\langle g\rangle, M')| = |d(\langle g\rangle, M) - d(\langle g\rangle, gM')| \leq 2\delta$$

and therefore $d(M, g \cdot M) \leq 4\delta$. By construction of the Cayley graph there exists $m \in \Gamma$ with $d(m, M) \leq 1$. By the triangle inequality

$$d(e, m^{-1}gm) = d(m, gm^{-1}) \leq d(M, g \cdot M) + 2 \leq 4\delta + 2$$

giving a contradiction. ∎

The last theorem restricts the size of centralizers in hyperbolic groups. This also restricts how large abelian subgroups can be:

Corollary 2.17. *No hyperbolic group contains* $\mathbb{Z}^2$ *as a subgroup.*

Proof. Let G be a group and $H \subseteq G$ a subgroup isomorphic to $\mathbb{Z}^2$. A non-trivial $h \in H$ has infinite order and index, and $H \subseteq \mathrm{C}_G(h)$. Therefore, $\langle h \rangle \subseteq \mathrm{C}_G(h)$ has infinite index in contrast to Theorem 2.16. ∎

More generally, along the same way one can argue that a hyperbolic group cannot contain any group with large (i.e. not virtually cyclic) centralizers – for example Baumslag-Solitar groups.

2.3 Miscellaneous group theory

This last preparatory section contains a somewhat eclectic collection of group-theoretic constructions that did not fit properly anywhere else but are nevertheless needed in further discussions:

Definition. Let G be a group. The *inner automorphisms group* of G is the normal subgroup of $\mathrm{Aut}(G)$ defined as

$$\mathrm{Inn}(G) := \left\{ \begin{array}{ccc} G & \to & G \\ x & \mapsto & gxg^{-1} \end{array} \;\middle|\; x \in G \right\}$$

The *outer automorphism group* of G is the quotient $\mathrm{Out}(G) := \mathrm{Aut}(G)/\,\mathrm{Inn}(G)$.

In general, it is a hard problem to determine outer automorphism groups. In the case of free groups, one can at least obtain some information by elementary means:

Example 2.18. Consider the non-abelian free group on n-generators $G = \langle a_1, \ldots, a_n \rangle$. We have a group homomorphism

$$\mathrm{Out}(G) \to \mathrm{Aut}(G_{ab}) \cong \mathrm{GL}_n(\mathbb{Z})$$
$$f \mapsto \left(\begin{array}{ccc} G_{ab} & \to & G_{ab} \\ [x] & \mapsto & [f(x)] \end{array} \right)$$

where $G_{ab} \cong \mathbb{Z}^n$ is the abelianization of G. It is surjective by the universal property of free groups. For $n = 2$ it even is an isomorphism. This is not the case for $n \geq 3$ since then its kernel will contain the non-inner automorphism fixing $a_1, \ldots, a_{n-1}$ and mapping $a_n \mapsto a_n \cdot [a_1, a_2]$.

A useful way of studying a group is to decompose it into smaller subgroups. Developing this further eventually leads to Bass-Serre Theory, see [Ser80]. We are only interested in the most basic way of splitting a group:

Definition. A group G *splits over a subgroup* H if it is isomorphic to

- a proper amalgamated product with amalgam H, i.e. there exist group monomorphisms $\varphi\colon H \to A$, $\psi\colon H \to B$ that are not isomorphisms such that

  $$G \cong A *_H B := A * B / \langle\langle \varphi(c) \cdot \psi(c)^{-1} \mid c \in H \rangle\rangle$$

  or
- an HNN-extension over H, i.e. there exist group monomorphisms $\varphi, \psi\colon H \to A$ such that

  $$G \cong A *_H := (A * \langle t \rangle)/\langle\langle t^{-1}\varphi(h)^{-1} t \psi(h) \mid h \in H \rangle\rangle$$

The splitting is *essential* if H is abelian and if $g^n \in H$ for some $g \in G,\ n \geq 1$ already implies $g \in H$.

The significance of the following definition stems from Geometric Group Theory being unable to distinguish between a group and its finite index subgroups as they have quasi-isometric Cayley graphs:

Definition. Let $\mathfrak{P}$ be a property of groups. A group G is *virtually* $\mathfrak{P}$ if there exists a finite index subgroup of G with property $\mathfrak{P}$.

We will mainly be interested in "virtually cyclic" or "virtually abelian". We consider these properties in the next three lemmas, beginning by noting that the two properties are equivalent for hyperbolic groups:

Lemma 2.19. *An abelian hyperbolic group is virtually cyclic. In particular, a virtually abelian hyperbolic group is virtually cyclic.*

Proof. Let Γ be an abelian hyperbolic group. By the classification of finitely generated abelian groups and Corollary 2.17, Γ is virtually cyclic. The "in particular" statement then follows from Lemma 2.13. ∎

The following criterion for detecting virtually abelian groups is [Pau91: Lemma 1.A]:

Lemma 2.20. *A finitely generated group G with $G' \coloneqq \{[a,b] \mid a,b \in G\}$ finite is virtually abelian.*

Proof. The group G acts on G' by conjugation since $g^{-1}[a,b]g = [g^{-1}ag, g^{-1}bg]$ for $a,b,g \in G$. The kernel H of this action is a finite index subgroup of G and hence finitely generated by the Schwarz-Milnor Lemma. Let $\{h_1,\dots,h_n\}$ be a generating set for H. For every $a \in G$ we have a finite subgroup $[a,H] \coloneqq \{[a,h] \mid h \in H\} \subseteq G$ as for $h,k \in H$ by definition of H

$$[a,h]\cdot[a,k] = aha^{-1}h^{-1}\cdot aka^{-1}k^{-1} = aha^{-1}\cdot aka^{-1}k^{-1}\cdot h^{-1} = ahka^{-1}k^{-1}h^{-1} = [a,hk]$$
$$[a,h^{-1}] = ah^{-1}a^{-1}h = h^{-1}\cdot hah^{-1}a^{-1}\cdot h = hah^{-1}a^{-1} = [a,h]^{-1}$$

This also shows that have homomorphisms

$$\begin{array}{rccc} \varphi_a\colon H & \to & [a,H] \\ h & \mapsto & [a,h] \end{array} \qquad \text{and} \qquad \begin{array}{rccc} \Phi\colon H & \to & [h_1,H]\times\cdots\times[h_n,H] \\ h & \mapsto & (\varphi_{h_1}(h),\dots,\varphi_{h_n}(h)) \end{array}$$

The kernel of Φ is the centre of H, hence the centre of H is a finite index subgroup. Therefore, H and by extension G are virtually abelian. ∎

Lastly, we note the following result for later usage:

Lemma 2.21. *A torsion-free virtually cyclic group is cyclic.*

Proof. Any virtually cyclic group G surjects $\{e\}$, $\mathbb{Z}$ or $(\mathbb{Z}/2\mathbb{Z})*(\mathbb{Z}/2\mathbb{Z})$ with finite kernel. This can be seen by algebraic calculation (see [Hem04: Lemma 11.4]) or geometrically by studying ends of groups (see [Sta71: Section 4.A.6]). If G is torsion-free this kernel must be trivial, i.e. G is isomorphic to one of the three given groups. This cannot be to $(\mathbb{Z}/2\mathbb{Z})*(\mathbb{Z}/2\mathbb{Z})$ since this group clearly has torsion. ∎

With this Lemma, we can also obtain another corollary to Theorem 2.16:

Corollary 2.22. *A non-cyclic torsion-free hyperbolic group has trivial centre.*

Proof. Let Γ be a torsion-free hyperbolic group and $\gamma \in \mathrm{Z}(\Gamma)$ non-trivial. Then $\mathrm{C}_\Gamma(\gamma) = \Gamma$ and it follows from Theorem 2.16 that Γ is virtually cyclic. Hence, Γ is cyclic by Lemma 2.21. ∎

Before finally coming to Paulin's Theorem, we again bring together geometry and group theory in the last definition of this section:

Definition. Let G be a group, X a G-space and $\gamma \subseteq X$ an arc. The *arc stabilizer* of γ is

$$\mathrm{Stab}(\gamma) \coloneqq \{g \in G \mid g\cdot\gamma = \gamma\}$$

If $\mathfrak{P}$ is a property of groups, X has $\mathfrak{P}$ *arc stabilizers* if the stabilizer of every non-degenerate arc in X has property $\mathfrak{P}$.

3 Paulin's theorem

3.1 Statement and Consequences

We have now developed enough of the general theory that we can come to the core of this essay. It is encapsulated in the following theorem originally proven in [Pau91]:

Theorem 3.1 (Paulin). *Let Γ be a hyperbolic group. If the outer automorphism group* $\mathrm{Out}(\Gamma)$ *is infinite,* Γ *acts on an* $\mathbb{R}$*-tree with virtually cyclic arc stabilizers and no global fixed points.*

We postpone the proof of this theorem to Section 3.3 and instead discuss some consequences. For this we will need the following theorem applying the Rips machine to actions on $\mathbb{R}$-trees. It picks up precisely where the last one left off, giving us an immediate corollary:

Theorem 3.2 ([BF95: Corollary 1.1])**.** *A hyperbolic group that has an action on an* $\mathbb{R}$*-tree with virtually cyclic arc stabilizers and no global fixed points splits over a virtually cyclic subgroup.*

Corollary 3.3. *Let Γ be a hyperbolic group. If the outer automorphism group* $\mathrm{Out}(\Gamma)$ *is infinite,* Γ *splits over a virtually cyclic subgroup.*

If we additionally assume the group to be torsion-free, the converse of this also holds:

Theorem 3.4. *Let Γ be a non-cyclic torsion-free hyperbolic group. The outer automorphism group* $\mathrm{Out}(\Gamma)$ *is infinite if and only if* Γ *splits essentially over a cyclic subgroup (i.e.* $\{e\}$ *or* $\mathbb{Z}$*).*

Proof. In the "only if"-direction Corollary 3.3 provides that Γ splits over a virtually cyclic subgroup H. By Lemma 2.21 H is trivial or infinite cyclic. That the splitting can be assumed to be essential requires a more careful study of the Rips machine, see [DG08: Theorem 5.6].
For the "if"-direction we consider the following possibilities separately. We implicitly use normal form theory for amalgamated products and HNN-extensions to do calculations (see [BH99: Lemma 6.4]):

Case 1: Γ splits as a amalgamated product over $\{e\}$
Then $\Gamma \cong A{*}B$ with A, B non-trivial. The inclusions of A and B into $A{*}B$ are isometric embeddings in the word metrics corresponding to choosing generating sets for A and B separately. Hence, A and B are hyperbolic by Proposition 2.11 and Theorem 2.4.
If A and B are abelian, we have $A \cong B \cong \mathbb{Z}$ by the classification of finitely generated abelian groups and Corollary 2.17. Therefore, by Example 2.18 $\mathrm{Out}(\Gamma) \cong \mathrm{GL}_2(\mathbb{Z})$ is infinite.
Hence, we may now wlog assume that A is non-abelian. Fix $x \in A$ non-trivial and consider the automorphism f of $A{*}B$ that restricts to the identity on B and conjugation by x on A. Suppose the subgroup generated by f in $\mathrm{Out}(A{*}B)$ is finite, i.e. there exists $n \in \mathbb{N}$ and $g \in G$ such that for all $a \in A, b \in B$

$$gag^{-1} = f^n(a) = x^n a x^{-n} \qquad \text{and} \qquad gbg^{-1} = b$$

It follows that $g \in A \cap B = \{e\}$ and therefore $x^n \in \mathrm{Z}(A)$. By Corollary 2.22 $x^n = e$. Contradiction!

Case 2: Γ splits as a amalgamated product over $\mathbb{Z}$
Then there are monomorphisms $\varphi\colon \mathbb{Z} \to A, \psi\colon \mathbb{Z} \to B$ that are not isomorphisms such that $\Gamma \cong A{*}_{\mathbb{Z}}B$. Again A and B are hyperbolic by Proposition 2.11 and Theorem 2.4 since their inclusions into $A{*}_{\mathbb{Z}}B$ are isometric embeddings in the word metrics corresponding to choosing generating sets for A and B separately where the generating set of A (respective B) contains $\varphi(1)$ (respective $\psi(1)$).
Suppose A or B is abelian, wlog A. As in Case 1 $A \cong \mathbb{Z}$. There exists $a \in A \setminus \varphi(\mathbb{Z})$, since φ is not an isomorphism. But since $A \cong \mathbb{Z}$, we have $a^n \in \varphi(\mathbb{Z})$ for some $n \in \mathbb{N}$, contradicting that the splitting is essential. Hence, A and B are non-abelian.

Consider the automorphism f of $A*_{\mathbb{Z}}B$ that restricts to conjugation by $c := \varphi(1)$ on A and the identity on B. Suppose the subgroup generated by f in $\mathrm{Out}(A*_{\mathbb{Z}}B)$ is finite, i.e. there exists $n \in \mathbb{N}$ and $g \in G$ such that for all $a \in A, b \in B$

$$gag^{-1} = f^n(a) = c^n a c^{-n} \qquad \text{and} \qquad gbg^{-1} = f^n(b) = b$$

Then $g \in \mathrm{Z}(B)$, i.e. $g = e$ by Corollary 2.22. Hence, analogously $c^n = e$. Contradiction!

Case 3: Γ splits as an HNN-extension over $\{e\}$
Then $\Gamma \cong A*_{\{e\}} \cong A*\mathbb{Z}$ for some group A. Since Γ is non-cyclic, A is non-trivial. Hence, we have already dealt with this in Case 1.

Case 4: Γ splits as an HNN-extension over $\mathbb{Z}$
Then there are monomorphisms $\varphi, \psi\colon \mathbb{Z} \to A$ such that $\Gamma \cong A*_{\mathbb{Z}}$. Again A is hyperbolic by Proposition 2.11 and Theorem 2.4 since its inclusion into $A*_{\mathbb{Z}}$ is an isometric embedding in the word metric corresponding to a generating set for A containing $\varphi(1), \psi(1)$. Since the splitting is essential, $A \ncong \mathbb{Z}$ as in Case 2. Set $c := \varphi(1)$. The automorphism of $A*\langle t\rangle$ restricting to the identity on A and mapping $t \mapsto c \cdot t$ descends to an automorphism f of $A*_{\mathbb{Z}}$. Suppose the subgroup generated by f in $\mathrm{Out}(A*_{\mathbb{Z}})$ is finite, i.e. there exists $n \in \mathbb{N}$ and $g \in G$ such that for all $a \in A$

$$gag^{-1} = f^n(a) = a \qquad \text{and} \qquad gtg^{-1} = f^n(t) = c^n \cdot t$$

Hence, $g \in \mathrm{Z}(A)$ is trivial by Corollary 2.22. Then $c^n = e$. Contradiction! ■

3.2 Compactness Theorem

As of now we have seen rather few examples of $\mathbb{R}$-trees. This has to change if we want to prove Paulin's Theorem 3.1. The idea behind the relevant construction is as follows: If a group G acts on a sequence of hyperbolic metric spaces stretching them more and more, the geodesic triangles in those spaces get drawn thinner and thinner by the action. Hence, one might hope that in the limit the action degenerates to an action on an 0-hyperbolic space – i.e. an $\mathbb{R}$-tree. This rough idea will be made precise in the Compactness Theorem 3.7. We follow [Bes02: Section 3] for this section.
To begin working towards our new goal, we need a notion of convergence of group actions. For pseudometrics there is already a natural sense of convergence:

Definition. Let G be a group. A pseudometric $d\colon G \times G \to [0, \infty)$ is *equivariant* if $d(ga, gb) = d(a, b)$ for all $g, a, b \in G$. Denote by $\mathcal{ED}_G$ the space of all non-trivial equivariant pseudometrics on G.
Scaling induces a free $\mathbb{R}^+$-action on $\mathcal{ED}_G$. The quotient space $\mathcal{PED}_G$ of this action is the space of *projectiviced equivariant pseudometrics* on G.
We topologize $\mathcal{ED}_G$ with the topology of pointwise convergence, i.e. by viewing it as a subspace of the $G \times G$-indexed product of $[0, \infty)$,[4] and subsequently $\mathcal{PED}_G$ with the quotient topology.

We note the following obvious lemma for later use:

Lemma 3.5. *Let $(d_n)_{n\in\mathbb{N}}$ be a convergent sequence in $\mathcal{ED}_G$. If d_n is δ_n-hyperbolic for $n \in \mathbb{N}$ and the sequence $(\delta_n)_{n\in\mathbb{N}}$ converges to $\delta \in [0, \infty)$, a limit of $(d_n)_{n\in\mathbb{N}}$ is δ-hyperbolic.*

Proof. Let d be a limit of $(d_n)_{n\in\mathbb{N}}$ in $\mathcal{ED}_G$. For any points in G the Gromov product with respect to d_n converges to the Gromov product with respect to d. Since the two sides of the "$\geq$"-inequality defining hyperbolicity are continuous in the involved Gromov products and δ's, the claim follows. ■

[4] This is of course the same as the compact-open topology on $\mathcal{ED}_G$ when viewing $G \times G$ as discrete as defined in [Bes02].

By associating a pseudometric to a G-space we can now transfer this notion of convergence to G-spaces:

Definition. Let G be a group. A *based G-space* is a pair (X, x) where x is a G-space and $x \in X$ is not fixed by every $g \in G$. A based G-space (X, x) *induces a non-trivial equivariant pseudometric* $d_{(X,x)}$ on G by setting

$$\begin{aligned} d_{(X,x)}\colon G \times G &\to [0, \infty) \\ (a, b) &\mapsto d_X(a \cdot x, b \cdot x) \end{aligned}$$

A sequence of based G-spaces $(X_n, x_n)_{n\in\mathbb{N}}$ *converges* to a based G-space (X, x) if their induced pseudometrics $(d_{(X_n,x_n)})_{n\in\mathbb{N}}$ converge to $d_{(X,x)}$ in $\mathcal{PED}_G$.

Here we only obtain pseudometrics since we do not assume that the basepoint has trivial stabilizer under the action in which case we could identify G with the orbit of x. But even in the more general setting we have an isometric map $G \to X$ leading to the following lemma:

Lemma 3.6. *Let G be a group and (X, x) a based G-space. If X is δ-hyperbolic, the induced pseudometric $d_{(X,x)}$ is also δ-hyperbolic.*

Proof. Acting on the basepoint x induces an isometric embedding from the metric space associated to $d_{(X,x)}$ into X, proving the claim by Proposition 2.6. ∎

The last two lemmas now set us on the right track for making the intuition from the beginning of this section precise: Suppose G acts on a sequence of δ-hyperbolic metric spaces $(X_n)_{n\in\mathbb{N}}$ for some fixed δ. We can scale down the metric on X_n by some λ_n without changing its equivalence class in $\mathcal{PED}_G$. If the λ_n tend to infinity, a resulting limit would have to be 0-hyperbolic. Unfortunately, there is a priori no reason why it should not be the degenerate $\mathbb{R}$-tree consisting of a single point. The first part of the Compactness Theorem 3.7 provides a sufficient condition to prevent this by ensuring that the basepoint is not fixed.
In Paulin's Theorem 3.1 we even want the action to not have any fixed point. To ensure this, we have to choose our basepoints carefully for the second part of the Compactness Theorem 3.7:

Definition. Let G be a group with a finite generating set S and X a G-space. A point $x \in X$ is *centrally located with respect to S* if the function

$$\begin{aligned} X &\to [0, \infty) \\ x &\mapsto \max_{s\in S}\{d(x, s \cdot x)\} \end{aligned}$$

attains a global minimum at x.

The reason why some delicacy is needed here, is that our definition of convergence only directly sees the basepoints and their translates under the action. The above definition hence makes the point moved the least amount by the action visible. If we additionally require the X_n to be geodesic, we can access sufficiently many points in X_n by going along geodesics between translates of the basepoints to allow us to prove that a limit of spaces with centrally located basepoints has again a centrally located basepoint.

Theorem 3.7 (Compactness Theorem). *Let G be a group with a finite generating set S and $\delta \geq 0$. Consider a sequence of δ-hyperbolic based G-spaces $(X_n, x_n)_{n\in\mathbb{N}}$. If*

$$\lambda_n := \max_{s\in S}\{d_{X_n}(x_n, s \cdot x_n)\}$$

is unbounded, a subsequence of $(X_n, x_n)_{n\in\mathbb{N}}$ converges to a based G-space (T, t) where T is an $\mathbb{R}$-tree. If X_n is geodesic and $x_n \in X_n$ is centrally located with respect to S for all $n \in \mathbb{N}$, then $t \in T$ is also centrally located. In particular, the G-action on T does not have a global fixed point.

Proof. After passing to a subsequence we may assume that $\lambda_n \neq 0$ for all $n \in \mathbb{N}$ and $(\lambda_n)_{n\in\mathbb{N}}$ tends to ∞. Define a sequence $(d_n)_{n\in\mathbb{N}}$ of pseudometrics on G by setting $d_n \coloneqq d_{(X_n,x_n)}/\lambda_n$.

Claim. A subsequence of $(d_n)_{n\in\mathbb{N}}$ *converges to some* $d \in \mathcal{ED}_G$.
Since S is finite, we may after passing to a subsequence assume that $\lambda_n = d_{X_n}(x_n, s_0 \cdot x_n)$ for all $n \in \mathbb{N}$ for some fixed $s_0 \in S$. By construction $d_n(e,s) = d_{X_n}(x_n, s\cdot x_n)/\lambda_n \leq 1$ for all $s \in S$. For $g \in G$ the triangle inequality then implies that $d_n(e,g)$ is at most the word length of g. Hence, we can iteratively construct subsequences of $(d_n)_{n\in\mathbb{N}}$, such that for the k-th one, $(d_n(e,g))_{n\in\mathbb{N}}$ converges for all $g \in G$ with word length at most k. By choosing the diagonal we can find a subsequence for which $(d_n(e,g))_{n\in\mathbb{N}}$ converges for all $g \in G$, then $(d_n(g,h))_{n\in\mathbb{N}}$ converges for all $g,h \in G$ by G-equivariance. Defining $d(g,h)$ as this limit, we obtain a pseudometric d on G which is the pointwise limit of a subsequence of $(d_n)_{n\in\mathbb{N}}$. By the choice at the beginning $d(e,s_0) = 1$, so d is non-trivial, i.e. $d \in \mathcal{ED}_G$. ☐

Since $d_n = d_{(X_n,x_n)}/\lambda_n$ is clearly δ/λ_n-hyperbolic, it follows from Lemma 3.6 and Lemma 3.5 that d is 0-hyperbolic. By applying the Connecting-the-dots Proposition 2.9 to the metric space associated to d we obtain an action of G on an $\mathbb{R}$-tree T and an isometric map $i\colon G \to T$ given by $i(g) = g\cdot t$ for $t \coloneqq i(e)$. In particular, the induced pseudometric $d_{(T,t)}$ equals d, i.e. $(X_n, x_n)_{n\in\mathbb{N}}$ converges to (T,t). Since d is non-trivial, (T,t) is a based G-space, proving the first part of the claim.
For the second part, consider the following construction: Let $a \in T, n \in \mathbb{N}$ and $F \subseteq G$ finite. We define $X_n(F,a) \subseteq X_n$ to be the set of all points $p \in X_n$ for which there exists $g,h \in F$ such that

- a lies on the geodesic segment $[g\cdot t, h\cdot t]$ in T.
- and p divides a geodesic segment between $g\cdot x_n$ and $h\cdot x_n$ in the same ratio as a divides $[g\cdot t, h\cdot t]$.

This construction has the following properties:

Claim.
1. *We have* $X_n(g\cdot F, g\cdot a) = g\cdot X_n(F,a)$ *for all* $F \subseteq G$ *finite,* $a \in T$, $n \in \mathbb{N}$ *and* $g \in G$.
2. *Let* $F, F' \subseteq G$ *finite. If* $F \subseteq F'$ *then* $X_n(F,a) \subseteq X_n(F',a)$ *for all* $a \in T$ *and* $n \in \mathbb{N}$.
3. *For every* $a \in T$ *there exists some* $F \subseteq G$ *finite with* $X_n(F,a) \neq \varnothing$ *for all* $n \in \mathbb{N}$.
4. *The sequence* $(\operatorname{diam} X_n(F,a)/\lambda_n)_{n\in\mathbb{N}}$ *converges to 0 for all* $F \subseteq G$ *finite and* $a \in T$.
5. *Let* $a,b \in T$ *and* $F \subseteq G$ *finite. The sequence* $(d_{X_n}(a_n,b_n)/\lambda_n)_{n\in\mathbb{N}}$ *converges to* $d_T(a,b)$ *for any choice of sequences* $(a_n)_{n\in\mathbb{N}}$, $(b_n)_{n\in\mathbb{N}}$ *with* $a_n \in X_n(F,a)$, $b_n \in X_n(F,b)$.

Statements (1) and (2) are clear by construction.
Statement (3) clearly holds if $a \in \Gamma\cdot t$, hence we may assume this is not the case. By the minimality condition in the Connecting-the-dots Proposition 2.9 $T\setminus\{a\}$ cannot be an $\mathbb{R}$-tree. Since it is 0-hyperbolic it must therefore have at least two path-components which are again $\mathbb{R}$-trees by Corollary 2.10. The minimality condition again implies that there exists $g \in \Gamma$ such that $g\cdot t$ and $e\cdot t$ are not in the same path-component. Hence, a lies on $[e\cdot t, g\cdot t]$ and $X_n(\{e,g\},a) \neq \varnothing$.
For (4) let $\epsilon > 0$. By Proposition 2.2 there exists $\delta' \geq 0$ such that for all $n \in \mathbb{N}$ all geodesic triangles in X_n are δ'-thin. Consider the finite set

$$P \coloneqq \{(g,h) \in F\times F \mid a \in [g\cdot t, h\cdot t]\}$$

For $(g,h) \in P$ there exists a unique $r_{g,h} \in [0,1]$ such that

$$d_T(g\cdot t, a) = r_{g,h}\cdot d_T(g\cdot t, h\cdot t)$$

Fix $p_n^{g,h} \in X_n(\{g,h\},a)$ for $(g,h) \in P, n \in \mathbb{N}$.

Claim. For all $(g,h),(g',h') \in P$ the sequence $(d_{X_n}(p_n^{g,h}, p_n^{g',h'})/\lambda_n)_{n\in\mathbb{N}}$ converges to 0.
Since T is an $\mathbb{R}$-tree, we have $(g,g') \in P$ or $(g,h') \in P$, wlog $(g,h') \in P$. Consider the geodesic triangle Δ indicated below where i_n, i'_n are internal points.

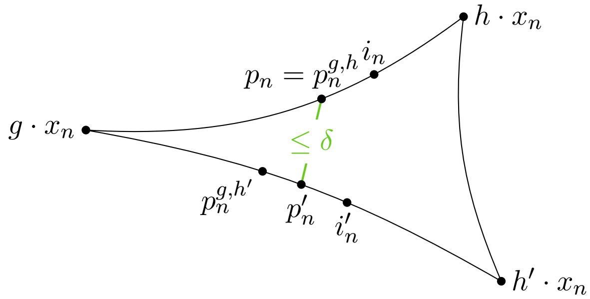

Observe that

$$\frac{d_{X_n}(g\cdot x_n, p_n^{g,h})}{\lambda_n} = r_{g,h} \cdot \frac{d_{X_n}(g\cdot x_n, h\cdot x_n)}{\lambda_n} \xrightarrow{n\to\infty} r_{g,h}\cdot d_T(g\cdot t, h\cdot t) = d_T(g\cdot t, a)$$

and by construction of the Gromov product

$$\frac{d_{X_n}(g\cdot x_n, i_1)}{\lambda_n} = \frac{1}{2}\left(\frac{d_{X_n}(g\cdot x_n, h\cdot x_n)}{\lambda_n} + \frac{d_{X_n}(g\cdot x_n, h'\cdot x_n)}{\lambda_n} - \frac{d_{X_n}(h\cdot x_n, h'\cdot x_n)}{\lambda_n}\right)$$
$$\xrightarrow{n\to\infty} \frac{1}{2}\left(d_T(g\cdot t, h\cdot t) + d_T(g\cdot t, h'\cdot t) - d_T(h\cdot t, h'\cdot t)\right) \geq d_T(g\cdot t, a)$$

Hence, there exists a sequence $(p_n)_{n\in\mathbb{N}}$ where $p_n \in [g\cdot x_n, i_n]$ such that $(d_{X_n}(p'_n, p_n^{g,h})/\lambda_n)_{n\in\mathbb{N}}$ converges to 0 (take $p_n \coloneqq p_n^{g,h}$ when possible and otherwise $p_n \coloneqq i_n$). By construction, the point p'_n corresponding to p_n under δ'-thinness of Δ lies on the geodesic between $g\cdot x_n$ and $h'\cdot x_n$. Then

$$\frac{d_{X_n}(p'_n, p_n^{g,h'})}{\lambda_n} = \frac{\left|d_{X_n}(g\cdot x_n, p'_n) - d_{X_n}(g\cdot x_n, p_n^{g,h'})\right|}{\lambda_n} = \left|\frac{d_{X_n}(g\cdot x_n, p_n)}{\lambda_n} - \frac{d_{X_n}(g\cdot x_n, p_n^{g,h'})}{\lambda_n}\right|$$
$$= \left|\underbrace{\frac{d_{X_n}(g\cdot x_n, p_n^{g,h})}{\lambda_n}}_{\to d(g\cdot t,a)} - \underbrace{\frac{d_{X_n}(p_n^{g,h}, p_n)}{\lambda_n}}_{\to 0} - \underbrace{\frac{d_{X_n}(g\cdot x_n, p_n^{g,h'})}{\lambda_n}}_{\to d(g\cdot t,a)}\right| \xrightarrow{n\to\infty} 0$$

Therefore, by the triangle inequality

$$\frac{d_{X_n}(p_n^{g,h}, p_n^{g,h'})}{\lambda_n} \leq \frac{d_{X_n}(p_n^{g,h}, p_n)}{\lambda_n} + \frac{d_{X_n}(p_n, p'_n)}{\lambda_n} + \frac{d_{X_n}(p'_n, p_n^{g,h'})}{\lambda_n} \xrightarrow{n\to\infty} 0$$

Analogously, the same holds with (g,h) replaced by (g',h') which taken together proves the claim. □

Since P is finite, the claim implies that we can find $N \in \mathbb{N}$ such that for all $n \geq N$

$$\forall_{(g,h),(g,h')\in P} \quad \frac{d_{X_n}(p_n^{g,h}, p_n^{g',h'})}{\lambda_n} < \frac{\epsilon}{2} \qquad \text{and} \qquad \frac{\delta'}{\lambda_n} < \frac{\epsilon}{4}$$

Then $\operatorname{diam} X_n(F,a)/\lambda_n < \epsilon$ for all $n \geq N$ proving (4):
Let $p, p' \in X_n(F,a)$. There exist $(g,h),(g',h') \in P$ such that p (respective p') originates from a geodesic from $g\cdot x_n$ to $h\cdot x_n$ (respective $g'\cdot x_n$ to $h'\cdot x_n$). Then $d_{X_n}(p, p_n^{g,h}) < \delta'$ (respective $d_{X_n}(p', p_n^{g',h'}) < \delta'$) by δ'-slimness. Hence, by the triangle inequality

$$\frac{d_{X_n}(p,p')}{\lambda_n} \leq \frac{d_{X_n}(p, p_n^{g,h})}{\lambda_n} + \frac{d_{X_n}(p_n^{g,h}, p_n^{g',h'})}{\lambda_n} + \frac{d_{X_n}(p_n^{g',h'}, p')}{\lambda_n} < \frac{\epsilon}{4} + \frac{\epsilon}{2} + \frac{\epsilon}{4} = \epsilon$$

For (5) take by (3) $g, g', h, h' \in G$ such that $a \in [g \cdot t, g' \cdot t]$, $b \in [h \cdot t, h' \cdot t]$. Since T is an $\mathbb{R}$-tree, $[a, b]$ intersects $[a, g \cdot t]$ or $[a, g' \cdot t]$ precisely in a, wlog $[a, g \cdot t]$. Analogously we can wlog assume that $[a, b] \cap [b, h \cdot t] = \{b\}$. Then $a, b \in [g \cdot t, h \cdot t]$. There exist $r_a, r_b \in [0, 1]$ such that

$$d_T(g \cdot t, a) = r_a \cdot d_T(g \cdot t, h \cdot t) \qquad \text{and} \qquad d_T(g \cdot t, b) = r_b \cdot d_T(g \cdot t, h \cdot t)$$

We can then define sequences $(a'_n)_{n\in\mathbb{N}}, (b'_n)_{n\in\mathbb{N}}$ with $a'_n \in X_n(\{g, h\}, a)$, $b'_n \in X_n(\{g, h\}, b)$ where a'_n and b'_n lie on the same geodesic connecting $g \cdot x_n$ and $h \cdot x_n$. Then

$$d_{X_n}(a'_n, b'_n) = |d_{X_n}(g \cdot x_n, a'_n) - d_{X_n}(g \cdot x_n, b'_n)| = |r_a - r_b| \cdot d_{X_n}(g \cdot x_n, h \cdot x_n)$$

and therefore

$$\frac{d_{X_n}(a'_n, b'_n)}{\lambda_n} = |r_a - r_b| \cdot \frac{d_{X_n}(g \cdot x_n, h \cdot x_n)}{\lambda_n} \xrightarrow{n\to\infty} |r_a - r_b| \cdot d_T(g \cdot t, h \cdot t) = d_T(a, b)$$

By the triangle inequality

$$\left| \frac{d_{X_n}(a_n, b_n)}{\lambda_n} - \frac{d_{X_n}(a'_n, b'_n)}{\lambda_n} \right| \leq \frac{d_{X_n}(a_n, a'_n)}{\lambda_n} + \frac{d_{X_n}(b'_n, b_n)}{\lambda_n}$$

By (4) applied to $F \cup \{g, h\}$ the right side go to 0 for n to infinity. Hence, the two sequences on the left have the same limit, proving (5). □

Let $x \in T$. By (3) there exists a finite subset $F \subseteq G$ with $X_n(F, x) \neq \varnothing$ for all $n \in \mathbb{N}$. Let $p_n \in X_n(F, x)$. Observe that for all $s \in S$ by (1) $X_n(s \cdot F, s \cdot x) = s \cdot X_n(F, x)$ and hence by (2) $p_n \in X_n(F \cup s \cdot F, x)$, $s \cdot p_n \in X_n(F \cup s \cdot F, s \cdot x)$. Then (5) and implies that

$$\frac{d_{X_n}(p_n, s \cdot p_n)}{\lambda_n} \xrightarrow{n\to\infty} d_T(x, s \cdot x)$$

Since $s \in S$ was arbitrary and S is finite

$$\max_{s\in S} \left\{ \frac{d_{X_n}(p_n, s \cdot p_n)}{\lambda_n} \right\} \xrightarrow{n\to\infty} \max_{s\in S} \{d_T(x, s \cdot x)\}$$

Specializing $x = t$ we deduce

$$\max_{s\in S} \left\{ \frac{d_{X_n}(x_n, s \cdot x_n)}{\lambda_n} \right\} \xrightarrow{n\to\infty} \max_{s\in S} \{d_T(t, s \cdot t)\}$$

Since $x_n \in X_n$ is centrally located, it follows by comparing these two limits that

$$\max_{s\in S} \{d_T(x, s \cdot x)\} \geq \max_{s\in S} \{d_T(t, s \cdot t)\}$$

Since $x \in T$ was arbitrary, this shows that $t \in T$ is centrally located. ■

3.3 Proof of Paulin's Theorem

The last section gave us the main theoretical tool to prove Paulin's Theorem 3.1, so we can now proceed to proving our main theorem. The remaining argument is now quite direct: Given a sequence of non-conjugate automorphisms of Γ, we can compose the usual action of Γ on its Cayley graph with them to obtain a sequence of Γ-spaces. We then show that the Compactness Theorem 3.7 can be applied to this sequence, giving us a Γ-action on an $\mathbb{R}$-tree without global fixed points. All that remains, is to prove that this action as virtually cyclic edge stabilizers which can be done by some hyperbolic geometry akin to the second half of proof of the Compactness Theorem 3.7 – but a bit more technically involved. In this section we again follow [Bes02: Theorem 7.3] using [BS94] to fill in some of the details.

Let us recall what we need to prove:

Theorem 3.1 (Paulin). *Let Γ be a hyperbolic group. If the outer automorphism group* $\mathrm{Out}(\Gamma)$ *is infinite, Γ acts on an $\mathbb{R}$-tree with virtually cyclic arc stabilizers and no global fixed points.*

Proof. Let S be a finite generating set for Γ and X the corresponding Cayley graph. Denote the usual metric on the Cayley graph by d. By Proposition 2.11 X is hyperbolic. By assumption there exists a sequence $(\varphi_n)_{n\in\mathbb{N}}$ of automorphisms of Γ that are pairwise non-conjugate. We get a sequence of Γ-spaces $(X_n)_{n\in\mathbb{N}}$ where the underlying metric space of X_n is still X but the action of $g \in \Gamma$ on X_n is given by $x \mapsto \varphi_n(g)\cdot x$. We always hence denote the metric on X_n by d.

Claim. For every $n \in \mathbb{N}$ there exists a centrally located $x_n \in X_n$.
Consider the function

$$\begin{aligned} f\colon X &\to [0,\infty) \\ x &\mapsto \max_{s\in S}\{d(x,\varphi_n(s)\cdot x)\} \end{aligned}$$

Let $e\colon [0,1] \to X$ be an isometric embedding onto an edge in X. For $h \in [0,1]$ and $s \in S$

$$d(e(h),\varphi_n(s)\cdot e(h)) = \min\begin{Bmatrix} d(e(0),\varphi_n(s)\cdot e(0)) + 2h, & d(e(0),\varphi_n(s)\cdot e(1)) + 1 \\ d(e(1),\varphi_n(s)\cdot e(0)) + 1, & d(e(1),\varphi_n(s)\cdot e(1)) + 2 - 2h \end{Bmatrix}$$

Since $d(e(a),\varphi_n(s)\cdot e(b)) \in \mathbb{Z}$ for $a,b \in \{0,1\}$ the case in which this minimum is attained can only change for $h \in \{0,\frac{1}{2},1\}$ and in between this function is linear in h with slope $0, 2$ or -2. Hence, $f\circ e$ is linear on $[\frac{i}{4},\frac{i+1}{4}]$ for $i \in \{0,\dots,3\}$, i.e. assumes its minimum on $\{0,\frac{1}{4},\frac{1}{2},\frac{3}{4},1\}$. Therefore, the minimum of $f\circ e$ lies in $(f\circ e)(\{0,\frac{1}{4},\frac{1}{2},\frac{3}{4},1\}) \subseteq \frac{1}{2}\mathbb{N}_0$. The claim follows as $\frac{1}{2}\mathbb{N}_0$ is well-ordered. ☐

Consider

$$\lambda_n \coloneqq \max_{s\in S}\{d(x_n,\varphi_n(s)\cdot x_n)\}$$

Claim. The sequence $(\lambda_n)_{n\in\mathbb{N}}$ is unbounded.
Suppose there is a bound $C \geq 0$ on $(\lambda_n)_{n\in\mathbb{N}}$. For every $n \in \mathbb{N}$ we can find some $g_n \in \Gamma$ with $d(x_n, g_n) < 1$ by the construction the Cayley graph. Then for $s \in S$ by the triangle inequality

$$d(e, g_n^{-1}\cdot\varphi_n(s)\cdot g_n) = d(g_n,\varphi_n(s)\cdot g_n) \leq d(x_n,\varphi_n(s)\cdot x_n) + 2 \leq C + 2$$

Hence, $g_n^{-1}\cdot\varphi(s)\cdot g_n \in \Gamma$ has word length at most $C+2$. This implies that $(g_n)_{n\in\mathbb{N}}$ has a subsequence on which $g_n^{-1}\cdot\varphi_n(s)\cdot g_n$ is constant. By applying the same procedure to all finitely many $s \in S$ we can iteratively find subsequences such that eventually $g_n^{-1}\cdot\varphi_n(s)\cdot g_n$ is constant over all $s \in S$ and $n \in \mathbb{N}$ in the last subsequence. In particular, there are distinct $n, m \in \mathbb{N}$ such that for all $s \in S$

$$g_n^{-1}\cdot\varphi_n(s)\cdot g_n = g_m^{-1}\cdot\varphi_m(s)\cdot g_m$$

Since $S \subseteq \Gamma$ is a generating set, this implies that φ_n and φ_m are conjugate. Contradiction! ☐

It now follows from the Compactness Theorem 3.7 that after possibly passing to a subsequence (X_n, x_n) converges to a based Γ-space (T,t) where T is an $\mathbb{R}$-tree and the Γ-action on T does not have a global fixed point.
It remains to prove that the arc stabilizers of this action are virtually cyclic: Let γ be a non-degenerate arc from a to b in T. We need to show that $\mathrm{Stab}(\gamma)$ is virtually cyclic. By definition $\mathrm{Stab}(\gamma)$ acts on γ, inducing a homomorphism $\rho\colon \mathrm{Stab}(\gamma) \to \mathrm{Isom}(\gamma)$. Since T is an $\mathbb{R}$-tree, γ is a geodesic segment and $\mathrm{Isom}(\gamma) \cong \mathbb{Z}/2\mathbb{Z}$. Therefore,

$$H \coloneqq \ker(\rho) = \{h \in \mathrm{Stab}(\gamma) \mid h\cdot x = x \text{ for all } x \in \gamma\}$$

has index at most 2 in $\mathrm{Stab}(\gamma)$ and it suffices to show that H is virtually cyclic.

Using Lemma 2.20 and Lemma 2.19 we only need to prove that $H' \coloneqq \{[h,k] \mid h,k \in H\}$ is finite. The intuition behind this is that on a geodesic segment in X_n approximating γ, H is close to acting by translations. Since translations along a geodesic segment commute, a commutator will be close to acting trivially giving us an upper bound on the number of commutators.
To make this precise, we need to recall the precise notion for approximating from the proof of the Theorem 3.7: Let $a \in T, n \in \mathbb{N}$ and $F \subseteq \Gamma$ finite. We define $X_n(F,a) \subseteq X_n$ to be the set of all points $p \in X_n$ for which there exists $g, h \in F$ such that

- a lies on the geodesic segment $[g \cdot t, h \cdot t]$ in T.
- and p divides a geodesic segment between $g \cdot x_n$ and $h \cdot x_n$ in the same ratio as a divides $[g \cdot t, h \cdot t]$.

We have already seen that this construction has the following properties (see p.19):

Claim.

(1) *We have* $X_n(g \cdot F, g \cdot a) = \varphi_n(g) \cdot X_n(F,a)$ *for all* $F \subseteq \Gamma$ *finite,* $a \in T$*,* $n \in \mathbb{N}$ *and* $g \in \Gamma$*.*
(2) *Let* $F, F' \subseteq \Gamma$ *finite. If* $F \subseteq F'$ *then* $X_n(F,a) \subseteq X_n(F',a)$ *for all* $a \in T$ *and* $n \in \mathbb{N}$*.*
(3) *For every* $a \in T$ *there exists some* $F \subseteq \Gamma$ *finite with* $X_n(F,a) \neq \varnothing$ *for all* $n \in \mathbb{N}$*.*
(4) *The sequence* $(\operatorname{diam} X_n(F,a)/\lambda_n)_{n\in\mathbb{N}}$ *converges to 0 for all* $F \subseteq \Gamma$ *finite and* $a \in T$*.*
(5) *Let* $a, b \in T$ *and* $F \subseteq \Gamma$ *finite. The sequence* $(d(a_n,b_n)/\lambda_n)_{n\in\mathbb{N}}$ *converges to* $d_T(a,b)$ *for any choice of sequences* $(a_n)_{n\in\mathbb{N}}$*,* $(b_n)_{n\in\mathbb{N}}$ *with* $a_n \in X_n(F,a)$*,* $b_n \in X_n(F,b)$*.*

By (2) and (3) there exists $F \subseteq \Gamma$ finite, such that $X_n(F,a)$ and $X_n(F,b)$ are non-empty for all $n \in \mathbb{N}$. Choose sequences $(a_n)_{n\in\mathbb{N}}, (b_n)_{n\in\mathbb{N}}$ with $a_n \in X_n(F,a)$, $b_n \in X_n(F,b)$. Consider a geodesic $\sigma_n \colon [0, D_n] \to X_n$ from a_n to b_n and let $c_n \coloneqq \sigma_n(D_n/2)$ be its midpoint.
Let $h \in H$ and let $\tau_{h,n} \coloneqq \max\{d(a_n, \varphi_n(h) \cdot a_n), d(n_n, \varphi_n(h) \cdot b_n)\}$ be the maximal amount one of the endpoints is moved by h. Then the map

$$\begin{aligned} \widehat{h}_n \colon [\tau_{h,n}, D_n - \tau_{h,n}] &\to \sigma_n([0,D_n]) \\ x &\mapsto \sigma_n(d(a_n, \varphi_n(h) \cdot b_n) - D_n + x) \end{aligned}$$

is well-defined (although possibly on an empty set) since by the triangle inequality for $x \in [\tau_{h,n}, D_n - \tau_{h,n}]$

$$d(a_n, \varphi_n(h) \cdot b_n) - d(a_n, b_n) + x \in [-d(b_n, \varphi_n(h) \cdot b_n), d(b_n, \varphi_n(h) \cdot b_n)] + x \subseteq [-\tau_{h,n}, \tau_{h,n}] + x \subseteq [0, D_n]$$

Note that $\widehat{h}_n$ is a translation by at most $\tau_{h,n}$ along the geodesic segment $\sigma_n([0, D_n])$.
We want to give a more geometric description of $\hat{h}_n$: Observe that by direct calculation

$$d(a_n, \hat{h}_n(x)) + d(\varphi_n(h) \cdot a_n, \varphi_n(h) \cdot b_n) - x = d(a_n, \varphi_n(h) \cdot b_n)$$

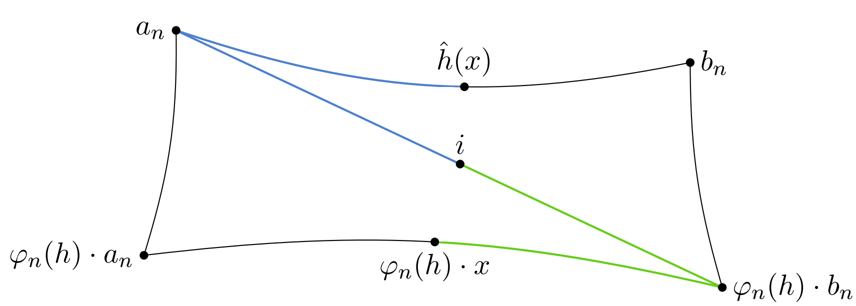

Hence geometrically $\hat{h}$ is given by taking the point i corresponding to $\varphi_n(h) \cdot x$ under δ-slimmness applied to the triangle spanned by $a_n, \varphi_n(h) \cdot b_n, \varphi_n(h) \cdot a_n$ and then letting $\hat{h}(x)$ be the point corresponding in turn to i under δ-slimmness applied to the triangle spanned by $a_n, b_n, \varphi_n(h) \cdot b_n$. The definition of $\tau_{h,n}$ ensures that i lies on the geodesic between a_n and $\varphi_n(h) \cdot b_n$.

This shows that the maps $\varphi_n(h)$ and $\widehat{h}_n$ are 2δ-close[5], where with Proposition 2.2 we have chosen $\delta \geq 0$ such that for all $n \in \mathbb{N}$ all geodesic triangles in X_n are δ-thin.
We next want to ensure that $\widehat{h}_n$ is defined on a sufficiently large set: By (1), (2) and (5) applied to F or $F \cup h \cdot F$

$$\frac{d(a_n, \varphi_n(h) \cdot a_n)}{\lambda_n} \xrightarrow{n\to\infty} d_T(a, h \cdot a) = 0 \qquad \frac{d(b_n, \varphi_n(h) \cdot b_n)}{\lambda_n} \xrightarrow{n\to\infty} d_T(b, h \cdot b) = 0$$
$$\frac{d(a_n, b_n)}{\lambda_n} \xrightarrow{n\to\infty} d_T(a, b) > 0$$

Hence, there exists $N_h \in \mathbb{N}$ such that for $n \geq N_h$

$$\frac{D_n}{\lambda_n} = \frac{d(a_n, b_n)}{\lambda_n} > 10 \cdot \max\left\{\frac{d(a_n, \varphi_n(h) \cdot a_n)}{\lambda_n}, \frac{d(b_n, \varphi_n(h) \cdot b_n)}{\lambda_n}\right\} = 10 \cdot \frac{\tau_{h,n}}{\lambda}$$

and therefore $D_n > 10\tau_{h_n}$ implying that $\widehat{h}_n$ is defined on a non-empty set.

Claim. There exists $C \geq 0$ such that for all $h, k \in H$ there exist $N \in \mathbb{N}$ such that $d(c_n, \varphi_n([h,k]) \cdot c_n) \leq C$ for all $n \geq N$.
Let $N := \max\{N_h, N_k\}$ and $\tau_n := \max\{\tau_{h,n}, \tau_{k,n}\}$. By the above construction $\widehat{h}_n, \widehat{k}_n$ (and therefore also $\widehat{h^{-1}}_n, \widehat{k^{-1}}_n$) are defined on the non-empty set $[\tau_n, D_n - \tau_n]$ and translate by at most τ_n. By the $10 = 2 \cdot (4 + 1)$ in the last inequality this implies that the composition of up to four of those maps is at least defined at the midpoint c_n. Since $\varphi_n(k^{-1})$ and $\widehat{k^{-1}}_n$ are 2δ-close, $\varphi_n(h^{-1}k^{-1})$ and $\varphi_n(h^{-1})\widehat{k^{-1}}_n$ are also 2δ-close. Hence, $\varphi_n(h^{-1}k^{-1})$ and $\widehat{h^{-1}}_n\widehat{k^{-1}}_n$ are 4δ-close. Iteratively, $\varphi_n([h,k])$ and $\widehat{h}_n\widehat{k}_n\widehat{h^{-1}}_n\widehat{k^{-1}}_n$ are 8δ-close. As the hat-maps are translations along the same geodesic segment and hence commute $\widehat{h}_n\widehat{k}_n\widehat{h^{-1}}_n\widehat{k^{-1}}_n = \widehat{h}_n\widehat{h^{-1}}_n\widehat{k}_n\widehat{k^{-1}}_n$ is also 8δ-close to $\varphi_n(hh^{-1}kk^{-1}) = \mathrm{id}$. All together this shows that $\varphi_n([h,k])$ is 16δ-close to the identity, proving the claim for $C = 16\delta$. □

Let $K \in \mathbb{N}$ be the number of elements in Γ with word length at most $C + 2$. Consider a finite subset $A \subseteq H$. By the Claim there exists $N \in \mathbb{N}$ such that $d(c_N, [h,k] \cdot c_N) < C$ for all $h, k \in A$. By construction of the Cayley graph, we can find $g \in \Gamma$ such that $d(g, c_N) < 1$. Then by the triangle inequality

$$d(e, g^{-1} \cdot \varphi_n([h,k]) \cdot g) = d(g, \varphi_n([h,k]) \cdot g) \leq d(c_n, \varphi_n([h,k]) \cdot c_n) + 2 \leq C + 2$$

Since $x \mapsto g^{-1} \cdot \varphi_n(x) \cdot g$ defines an automorphism of Γ, it follows that the set $A' := \{[h,k] \mid h, k \in A\}$ has cardinality at most K.
Since Γ is finitely generated, it is countable. Hence, H is also countable and there exists an ascending sequence of finite subsets $A_1 \subseteq A_2 \subseteq \cdots \subseteq H$ whose union is H. Then the ascending union of $A'_1 \subseteq A'_2 \subseteq \ldots$ is also H'. As each of the A'_n has at most K elements, H' also has at most K – in particular finitely many – elements. ■

[5] We say that two maps $f, g\colon I \to X$ are *η-close* if $d(f(x), g(x)) \leq \eta$ for $x \in I$ and let an element of Γ also refer to the map through which it acts.

4 Self-homotopy equivalences of manifolds

4.1 Two groups of self-homotopy equivalences

The goal of this section is to define the group of self-homotopy equivalences of a topological space with and without a basepoint and to study the relationship between these groups and the fundamental group of the space. The core issue will be to understand how precisely the various objects depend on the choice of basepoint. We begin by recalling this for the fundamental group (see [Hat02: Proposition 1.5]):

Construction 4.1. Let X be a topological space. Given paths $\alpha, \beta\colon [0,1] \to X$ with $\alpha(1) = \beta(0)$ we write $\alpha \cdot \beta$ for the path given by first traversing α followed by β and $\overline{\alpha}$ for the inverse path of α.
For every path $\gamma\colon [0,1] \to X$ there is a group isomorphism, the *change-of-basepoint along* γ,

$$\begin{aligned} \mathfrak{b}_\gamma\colon \pi_1(X, \gamma(1)) &\to \pi_1(X, \gamma(0)) \\ [\sigma] &\to [\gamma \cdot \sigma \cdot \overline{\gamma}] \end{aligned}$$

We note some basic properties of this construction:

- If $\gamma\colon [0,1] \to X$ is a loop, i.e. $x \coloneqq \gamma(0) = \gamma(1)$, it represents an element $[\gamma] \in \pi_1(X, x)$. Then $\mathfrak{b}_\gamma\colon \pi_1(X, x) \to \pi_1(X, x)$ is given by conjugation with $[\gamma]$.
- If $\alpha, \beta\colon [0,1] \to X$ are paths with $\alpha(1) = \beta(0)$, then $\mathfrak{b}_{\alpha\cdot\beta} = \mathfrak{b}_\alpha \cdot \mathfrak{b}_\beta$.
- The change-of-basepoint is natural in the sense that for a path $\gamma\colon [0,1] \to X$ and a continuous map $f\colon X \to Y$ between topological spaces the diagram

$$\begin{array}{ccc} \pi_1(X, \gamma(1)) & \xrightarrow{f_*} & \pi_1(Y, (f \circ \gamma)(1)) \\ \mathfrak{b}_\gamma \downarrow & & \downarrow \mathfrak{b}_{f\circ\gamma} \\ \pi_1(X, \gamma(0)) & \xrightarrow{f_*} & \pi_1(Y, (f \circ \gamma)(0)) \end{array}$$

  commutes.

Having refreshed our memory, we turn to the main object of interest for this section:

Definition. Let X be a topological space and $x \in X$ a basepoint. We consider two groups:
- $\mathcal{E}(X)$, the *group of self-homotopy equivalences of* X up to homotopy.
- $\mathcal{E}_x(X)$, the *group of basepoint preserving self-homotopy equivalences of* X up to pointed homotopy.

The group structure is in both cases given by composition.

At first glance it might seem as if $\mathcal{E}_x(X)$ were a subgroup of $\mathcal{E}(X)$. This is not the case since pointed homotopy is a more restrictive equivalence relation. Let us recall that in fact the opposite is the case (see [Hat02: Section 4.A] or [Rut97: Section 1]):

Construction 4.2. Let (X, x) be a well-pointed Hausdorff space, i.e. the pair (X, x) has homotopy lifting property. We associate to a loop $\gamma\colon [0,1] \to X$ based at x a homotopy equivalence $\varphi(\gamma)\colon X \to X$ in the following way: By the homotopy lifting property there exists a homotopy $F : X \times [0,1] \to X$ with $F_0 = \mathrm{id}$ and $\overline{\gamma}(t) = F(t, x)$ for $t \in [0,1]$. Set $\varphi(\gamma) \coloneqq F_1$. Since we chose $\overline{\gamma}(t) = F(t, x)$ (instead of $\gamma(t) = F(t, x)$) this defines a group homomorphism φ that fits into an exact sequence

$$\pi_1(X, x) \xrightarrow{\varphi} \mathcal{E}_x(X) \xrightarrow{p} \mathcal{E}(X) \longrightarrow 0$$

where p is the map given by forgetting the basepoint.

Such prepared, we return to exhibiting a relationship between the groups of self-homotopy equivalences of X and the fundamental group of X. A naive idea might be to associate to a homotopy equivalence $f\colon X \to X$ the map induced by it on the fundamental group – but this is only well-defined on a specific fundamental group if f respects a basepoint. For an arbitrary f one must correct this by a suitable change of basepoint which will only be well-defined up to conjugation:

Proposition 4.3. *Let X be a path-connected topological space and $x \in X$. We have homomorphisms*

$$\begin{aligned}\Psi_x\colon \mathcal{E}_x(X) &\to \mathrm{Aut}(\pi_1(X,x))\\ [f] &\mapsto \big(f_*\colon \pi_1(X,x) \to \pi_1(X,x)\big)\end{aligned}$$

$$\begin{aligned}\Psi\colon \mathcal{E}(X) &\to \mathrm{Out}(\pi_1(X,x))\\ f &\mapsto \big[\mathfrak{b}_\gamma \circ \big(f_*\colon \pi_1(X,x) \to \pi_1(X,f(x))\big)\big] \qquad \textit{for } \gamma\colon [0,1] \to X \textit{ a path from } f(x) \textit{ to } x\end{aligned}$$

If (X,x) is well-pointed Hausdorff, they fit into a commutative diagram with exact rows

$$\begin{array}{ccccccc}
\pi_1(X,x) & \xrightarrow{\varphi} & \mathcal{E}_x(X) & \xrightarrow{p} & \mathcal{E}(X) & \longrightarrow & 0\\
\downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \Psi_x} & & \downarrow{\scriptstyle \Psi} & & \\
\pi_1(X,x) & \xrightarrow{\mathcal{C}} & \mathrm{Aut}(\pi_1(X,x)) & \xrightarrow{q} & \mathrm{Out}(\pi_1(X,x)) & \longrightarrow & 0
\end{array}$$

where $\mathcal{C}$ is given by conjugation and q is the quotient map.

Proof. The well-definiteness of Ψ_x is a classical result, so let us turn to Ψ:
We begin by showing that the definition of Ψ is independent of the choice of path: Let $f \in \mathcal{E}(X)$ and $\alpha, \beta\colon [0,1] \to X$ be paths from $f(x)$ to x. Then $\alpha \cdot \overline{\beta}$ is a path from x to x and the change-of-basepoint along $\alpha \cdot \overline{\beta}$ corresponds to conjugation by $[\alpha \cdot \overline{\beta}]$ on $\pi_1(X,x)$. Hence, in $\mathrm{Out}(\pi_1(X,x))$

$$\Big[\mathfrak{b}_\beta \circ f_*\Big] = \Big[\mathfrak{b}_{\alpha\cdot\overline{\beta}} \circ \mathfrak{b}_\beta \circ f_*\Big] = \Big[\mathfrak{b}_{\alpha\cdot\overline{\beta}\cdot\beta} \cdot f_*\Big] = \Big[\mathfrak{b}_\alpha \circ f_*\Big]$$

showing that Ψ is well-defined.
To see that Ψ is a group homomorphism, let $f, g \in \mathcal{E}(X)$ and $\alpha\colon [0,1] \to X$ be a path from $f(x)$ to x, $\beta\colon [0,1] \to X$ a path from $g(x)$ to x. Then

$$\Phi(f) \circ \Phi(g) = (\mathfrak{b}_\alpha \circ f_*) \circ (\mathfrak{b}_\beta \circ g_*) = \mathfrak{b}_\alpha \circ \mathfrak{b}_{f\circ\beta} \circ f_* \circ g_* = \mathfrak{b}_{\alpha\cdot(f\circ\beta)} \circ (f \circ g)_* = \Phi(f \circ g)$$

Commutativity of the right square of the diagram follows by direct calculation: Given $[f] \in \mathcal{E}_x(X)$ we can choose a constant path $\gamma\colon [0,1] \to X$ at x to calculate

$$(\varphi \circ p)([f]) = [\mathfrak{b}_\gamma \circ f_*] = [f_*] = (q \circ \Psi_x)([f])$$

Commutativity of the left square is more involved: Let $[\sigma] \in \pi_1(X,x)$. As in Construction 4.2 there exists a homotopy $F\colon X \times [0,1] \to X$ with $F_0 = \mathrm{id}$, $F(t,x) = \overline{\sigma}(t)$ for $t \in [0,1]$ and $[F_1] = \varphi([\sigma])$. Then for $[\tau] \in \pi_1(X,x)$ by restricting the homotopy F

$$(\Psi_x \circ \varphi)([\sigma])([\tau]) = [F_1 \circ \tau] = [\sigma \cdot \tau \cdot \overline{\sigma}] = \mathcal{C}([\sigma])([\tau])$$

The top row in the given diagram is exact by Construction 4.2. The bottom row is exact by definition of the outer automorphism group ■

In general, there is of course no hope of these maps being an isomorphism since this would correspond to a correspondence between algebraic and topological properties. But of course, there is a class of spaces where such a correspondence does exist:

Definition. Let G be a group and $n \geq 1$. A path-connected topological space X is a *Eilenberg-MacLane space of type* $K(G,n)$ if X admits a CW-structure and

$$\pi_i(X) \cong \begin{cases} G, & \text{if } i = n \\ 0, & \text{otherwise} \end{cases}$$

The next proposition summarizes some properties of these spaces (see [Hat02: pp. 365-366, 390]):

Proposition 4.4. *Let G be a group and $n \geq 1$. If $n \geq 2$ we demand that G is abelian.*

(1) *An Eilenberg-MacLane space of type $K(G,n)$ exists and it is determined up to homotopy equivalence.*
(2) *We have a group isomorphism*

$$\begin{array}{rcl} \mathcal{E}_x(X) & \to & \mathrm{Aut}(\pi_n(K(G,n))) \\ {[f]} & \mapsto & f_* \end{array}$$

Combining this with Proposition 4.3 we obtain the following corollary:

Corollary 4.5. *Let G be a group and $n \geq 1$. If $n \geq 2$ we demand that G is abelian. Let X be an Eilenberg-MacLane space of type $K(G,n)$. Then $\mathcal{E}(X)$ and $\mathrm{Out}(G)$ are isomorphic.*

Proof. In the case $n = 1$, it follows from Proposition 4.4, the commutative diagram in Proposition 4.3 and the Five Lemma that $\Psi\colon \mathcal{E}(X) \to \mathrm{Out}(\pi_1(X))$ is an isomorphism.
For $n \geq 2$ we have $\pi_1(X) = 0$. This implies by the exact sequence in Construction 4.2 that $\mathcal{E}_x(X) \cong \mathcal{E}(X)$. The claim now follows directly from Proposition 4.4 since G is abelian. ∎

4.2 Hyperbolic manifolds

In this last section we want to return to the discussion from the introduction and give a finiteness condition for the group of self-homotopy equivalences of hyperbolic manifolds. We begin by giving some theory and examples for hyperbolic manifolds. We can however not fully explore the necessary background in geometry and instead refer to [Lee97].

Definition. A *hyperbolic* manifold is a complete Riemannian manifold of constant sectional curvature -1.

The class of hyperbolic manifolds is quite large. We give some examples in low dimensions:

Examples 4.6.

- Consider a closed (orientable or non-orientable) surface Σ of Euler characteristic χ. Such a surface can be defined by suitably gluing the edges of a $(4-2\chi)$-gon Δ. Under this gluing all vertices become one point, so to be able to perform this gluing geometrically (i.e. with a Riemannian metric on Δ descending to Σ) $(4-2\chi)$-times the interior angle of Δ needs to be 2π. Unless $\chi = 0$, this is not the case with the standard Euclidean metric on Δ.
  But a regular hyperbolic n-gon with all interior angles α exists if any only if its area $(n-2)\pi - n\alpha$ is positive, so there does exist a regular hyperbolic $(4-2\chi)$-gon Δ with interior angle $\alpha = 2\pi/(4-2\chi)$ when $\chi < 0$. Then one can perform the gluing geometrically on Δ giving Σ a hyperbolic structure. See [Thu97: p.27] for a more detailed account.
  If $\chi \geq 0$, such a hyperbolic $(4-2\chi)$-gon does not exist, which might lead us to expect that in this case there is no hyperbolic structure on the surface. We will return to this after developing some more theory.

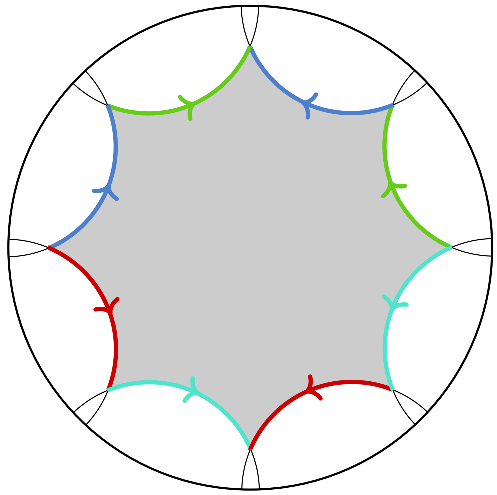

Gluing a hyperbolic octagon to build the orientable surface of genus 2

– In a similar manner, one can also obtain hyperbolic 3-manifolds by gluing polyhedrons: Consider a dodecahedron. Two opposing faces are misaligned by a twist of 1/10. Therefore, we can identify each face with its opposite with clockwise twist by 3/10 of a turn[6]:

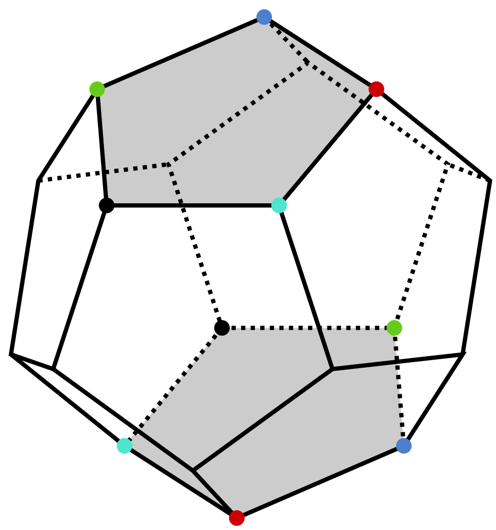

Notice that this operation is symmetrical, i.e. opposing faces get glued to each other in the same manner. On can check similarly to the surface case that the result is a closed 3-dimensional manifold – the Seifert–Weber dodecahedral space. All vertices become one point with small spherical neighbourhoods of them arranged as a subdivision of a icosahedron and the edges are identified in six groups of five.
The interior dihedral angle along an edge of the dodecahedron is $\arccos(-1/\sqrt{5}) \approx 2.03$ and therefore too large to perform this gluing geometrically on a Euclidean dodecahedron. But similar to the surface case one can argue that there exists a hyperbolic dodecahedron with interior angles $2\pi/5$. Gluing together this dodecahedron shows that the Seifert–Weber dodecahedral space can be given a hyperbolic structure. See again [Thu97: Example 1.4.5] for a more detailed account.

[6] Of course, we could also identify them with a 1/10 or 5/10 twist. The former leads to the Poincaré homology sphere (see [Thu97: Example 1.4.4]), the latter is the description of $\mathbb{RP}^3$ as a quotient of a ball.

– While the Seifert-Weber dodecahedral space is difficult to visualise, there are many 3-dimensional hyperbolic manifolds that can be depicted – if one accepts that they are no longer closed: In [Thu97: Examples 3.3.8-3.3.10, 3.4.16] there are instructions on how to glue hyperbolic polyhedra with some of their edges removed to obtain the complements in S^3 of the links indicated below.

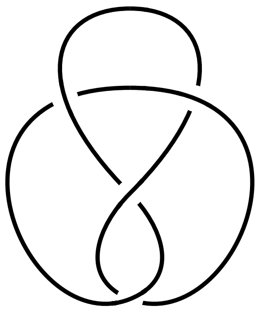

figure-8 knot

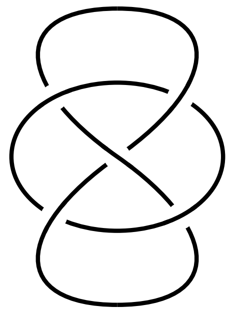

Whitehead link

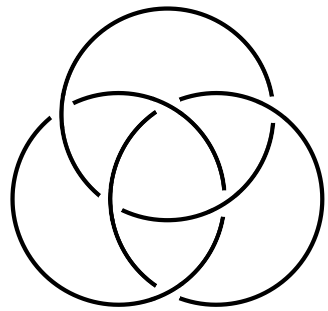

Borromean rings

In fact, the complements of many knots can be given a hyperbolic structure. This is a consequence of Thurston's Hyperbolization Theorem, see [Thu82: Corollary 2.5].

It is no accident that these examples were all given by suitably gluing sides of a hyperbolic polytopes. The classification of complete Riemannian manifolds of constant sectional curvature specialised to hyperbolic manifolds shows that any hyperbolic manifold is of this from:

Theorem 4.7 ([Lee97: Corollary 11.13])**.** *Let M be a connected hyperbolic manifold. Then M is isometric to $\mathbb{H}^n/\Gamma$ where $\Gamma \subseteq \mathrm{Isom}(\mathbb{H}^n)$ is a discrete subgroup isomorphic to $\pi_1(M)$ acting freely and properly discontinuously on $\mathbb{H}^n$.*

We can immediately deduce some properties of the fundamental group of a closed hyperbolic manifold:

Corollary 4.8. *Let M be a closed connected hyperbolic manifold.*
(1) *$\pi_1(M)$ is hyperbolic.*
(2) *$\pi_1(M)$ is torsion-free.*
(3) *M is an Eilenberg-MacLane space of type $K(\pi_1(M), 1)$.*

Proof.
(1) By Theorem 4.7 and the Schwarz-Milnor Lemma $\pi_1(M)$ is quasi-isometric to $\mathbb{H}^n$ which is a hyperbolic metric space by Examples 2.3. The claim follows from Theorem 2.4.
(2) By Theorem 4.7 $\pi_1(M)$ acts freely on $\mathbb{H}^n$, hence all elements of $\pi_1(M)$ act by parabolic or hyperbolic isometries which have infinite order.
(3) Since M is a smooth manifold, it admits a CW-structure. All higher homotopy groups of M vanish as its universal covering is contractible by Theorem 4.7. ■

Here the restriction to closed manifolds is essential to ensure the action of the fundamental group on $\mathbb{H}^n$ is cocompact so that we can apply the Schwarz-Milnor Lemma. In general, the fundamental group of a hyperbolic manifold confusingly need not be hyperbolic:

Example 4.9. It is a knot-theoretic consequence of the Loop Theorem that the inclusion of the torus bounding a tubular neighbourhood of a non-trivial knot into its complement induces an injection on the fundamental group (see [Lic97: Theorem 11.2]). This implies in particular that the fundamental group of the complement of a non-trivial knot contains $\mathbb{Z}^2$ as a subgroup and is therefore not hyperbolic by Corollary 2.17, in contrast to the existence of hyperbolic knot complements discussed in Examples 4.6.

With the next example we fulfil our promise to return to discussing hyperbolic structures on surfaces, confirming our suspicion from Examples 4.6 that surfaces of non-negative Euler characteristic do not allow hyperbolic structures:

Examples 4.10. A closed surface of Euler characteristic 0 is covered by a torus. Hence, its fundamental group contains $\mathbb{Z}^2$ as a subgroup and is therefore not hyperbolic by Corollary 2.17. By Corollary 4.8 the surface itself then also cannot be equipped with a hyperbolic structure.
The two surfaces of positive Euler characteristic have universal covering S^2. Hence, it follows from Theorem 4.7 that they also cannot carry a hyperbolic structure.

The next theorem brings together what we discussed. It shows that Paulin's Theorem 3.1 can be applied to closed hyperbolic manifolds:

Theorem 4.11. *A closed connected hyperbolic manifold M has infinite group of self-homotopy equivalences $\mathcal{E}(X)$ if and only if its fundamental group $\pi_1(M)$ essentially splits over a cyclic subgroup.*

Proof. Theorem 3.4 applies to $\pi_1(M)$ since
- $\pi_1(M)$ is hyperbolic and torsion-free by Corollary 4.8.
- $\pi_1(M)$ is not cyclic by Theorem 4.7 since cyclic groups do not act cocompactly on $\mathbb{H}^n$.

The claim follows since $\mathrm{Out}(\pi_1(M)) \cong \mathcal{E}(M)$ by Corollary 4.8 and Corollary 4.5. ∎

This gives us two corollaries. The first one deals with the surface case:

Corollary 4.12. *A closed hyperbolic surface has infinite group of self-homotopy equivalences.*

Proof. By Examples 4.6 and Examples 4.10 the closed hyperbolic surfaces are precisely those of negative Euler characteristic. Cutting along a non-separating curve shows that their fundamental group can be written as an HNN-extension over $\mathbb{Z}$, with the vertex group being the fundamental group of a surface of non-negative Euler characteristic minus two discs. Such a group is never $\mathbb{Z}$. Hence, this splitting is essential and the claim follows from Theorem 4.11. ∎

The second one deals with the situation in dimension at least 3. It is the corollary of Mostow rigidity we originally set out to prove and thus a fitting conclusion to this essay:

Corollary 4.13. *A closed hyperbolic manifold of dimension at least 3 has finite group of self-homotopy equivalences.*

Sketch of a proof. We tentatively only call this a "Sketch of a proof" since we do not have the space to establish the required background regarding the Gromov-boundary of a hyperbolic group (see [KB02]): Let M be a closed hyperbolic manifold of dimension $n \geq 3$. We need to show that $\pi_1(M)$ does not split essentially over a cyclic subgroup. We do this by analysing its Gromov-boundary $\partial\pi_1(M)$: It follows from Theorem 4.7, the Schwarz-Milnor Lemma and the quasi-isometry invariance of the Gromov-boundary that $\partial\pi_1(M)$ is homeomorphic to S^{n-1}. We can therefore derive the claim from the following two observations which can be deduced from the definition of the Gromov-boundary (see [KB02: Section 7]):
- If a hyperbolic group splits over $\{e\}$, its Gromov-boundary is disconnected.
- If a hyperbolic group splits essentially over $\mathbb{Z}$, its Gromov-boundary has a local cut point, i.e. a point whose removal makes all sufficiently small neighbourhoods of it disconnected. ∎

References

[Bes02] M. Bestvina. "$\mathbb{R}$-Trees in Topology, Geometry, and Group Theory". In: *Handbook of Geometric Topology*. North-Holland, Amsterdam, 2002, pp. 55–91.

[BF95] M. Bestvina, M. Feighn. "Stable actions of groups on real trees". In: *Invent. Math.* 121.2 (1995), pp. 287–321.

[BH99] M.R. Bridson, A. Haefliger. *Metric spaces of non-positive curvature*. Vol. 319. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 1999.

[Bra99] N. Brady. "Branched coverings of cubical complexes and subgroups of hyperbolic groups". In: *J. London Math. Soc. (2)* 60.2 (1999), pp. 461–480.

[BS94] M.R. Bridson, G.A. Swarup. "On Hausdorff-Gromov convergence and a theorem of Paulin". In: *Enseign. Math. (2)* 40.3-4 (1994), pp. 267–289.

[DG08] F. Dahmani, D. Groves. "The isomorphism problem for toral relatively hyperbolic groups". In: *Publ. Math. Inst. Hautes Études Sci.* 107 (2008), pp. 211–290.

[Gro87] M. Gromov. "Hyperbolic groups". In: *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. Springer, New York, 1987, pp. 75–263.

[Hat02] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002, pp. xii+544.

[Hem04] J. Hempel. *3-manifolds*. AMS Chelsea Publishing, Providence, RI, 2004, pp. xii+195.

[KB02] I. Kapovich, N. Benakli. "Boundaries of hyperbolic groups". In: *Combinatorial and geometric group theory*. Vol. 296. Contemp. Math. Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.

[Lee97] J.M. Lee. *Riemannian Manifolds*. Graduate texts in mathematics. Springer, New York, 1997, pp. xv+226.

[Lic97] W.B.R. Lickorish. *An introduction to knot theory*. Vol. 175. Graduate Texts in Mathematics. Springer, New York, 1997, pp. x+201.

[Mos68] G.D. Mostow. "Quasi-conformal mappings in n-space and the rigidity of hyperbolic space forms". In: *Inst. Hautes Études Sci. Publ. Math.* 34 (1968), pp. 53–104.

[Pau91] F. Paulin. "Outer automorphisms of hyperbolic groups and small actions on $\mathbb{R}$-trees". In: *Arboreal group theory*. Vol. 19. Math. Sci. Res. Inst. Publ. Springer, New York, 1991, pp. 331–343.

[Rip82] E. Rips. "Subgroups of small cancellation groups". In: *Bull. London Math. Soc.* 14.1 (1982), pp. 45–47.

[Rut97] J.W. Rutter. *Spaces of homotopy self-equivalences*. Vol. 1662. Lecture Notes in Mathematics. Springer, Berlin, 1997, pp. x+170.

[Ser80] J-P. Serre. *Trees*. Translated by J. Stillwell. Springer, Berlin-New York, 1980, pp. ix+142.

[Sta71] J. Stallings. *Group theory and three-dimensional manifolds*. Vol. 4. Yale Mathematical Monographs. Yale University Press, New Haven, Conn.-London, 1971, pp. v+65.

[Thu82] W.P. Thurston. "Three-dimensional manifolds, Kleinian groups and hyperbolic geometry". In: *Bull. Amer. Math. Soc. (N.S.)* 6.3 (1982), pp. 357–381.

[Thu97] W.P. Thurston. *Three-dimensional geometry and topology. Vol. 1*. Vol. 35. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 1997, pp. x+311.

[Wil70] S. Willard. *General topology*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1970, pp. xii+369.