Bachelor's thesis in mathematics

The Word Problem, Hyperbolic Groups and the Morse criterion

Submitted by
Marcel Wunderlich

Supervised by
Prof. Dr. Linus Kramer

Westfälische Wilhelms-Universität Münster

August 27, 2012

Contents

Preface

This thesis is divided in two parts, the first dealing with *hyperbolic groups*, as well as the *Word Problem* and the *Transformation Problem* in this type of groups. We will give an introduction to the geometry of hyperbolic spaces and geometric group theory. After the foundation for the definition of a hyperbolic group have been laid, we will state the Word and Conjugacy Problem and show how these can be solved in hyperbolic groups.

The second part introduces *Morse theory* on *cell complexes*, which is an useful tool to examine the topology of a given cell complex. The motivation for this is a criterion using Morse theory given by Noel Brady, which enables us to recognize free-by-cyclic groups by studying their 2-*presentation complex* with Morse theory. We will prove this criterion and show how one can apply it to a given group presentation by demonstrating it on some examples given by Brady.

1 Hyperbolic spaces and groups

Definition 1.1 (Hyperbolic groups)**.** A finitely generated group is called *hyperbolic* if its *Cayley graph* is a hyperbolic metric space.

This definition is the motivating theme for this section. Despite looking simple and innocuous, we have to take care that it is well-defined, since a group can have many different generating sets and thus many different Cayley graphs. We will show that the Cayley graphs w.r.t. a finite generating sets of a group are quasi-isometric to each other and that hyperbolicity is preserved under quasi-isometry. The latter will require some work, which we will use as opportunity to acquaint ourselves with the geometry of hyperbolic spaces and establish some useful results for the main theorems of this thesis.

First we introduce the Word metric of a group and definition of a Cayley graph, which allow us to study groups as geometric objects.

Definition 1.2 (Word metric)**.** Let G be a group with generating set $S \subset G$. Define for $g \in G, g \neq 1$ the *length of* g,

$$l_S(g) := \min\left\{n \mid g = t_1 \cdots t_n,\ t_i \in S \text{ or } t_i^{-1} \in S\right\}.$$

Furthermore let $l_S(1) = 0$. Then the following holds:

(i) $l_S(g) = 0 \Leftrightarrow g = 1$

(ii) $l_S(g) = l_S(g^{-1})$

(iii) $l_S(gh) \leq l_S(g) + l_S(h)$

If the generating set is clear from the context, we also write $l(g)$ to denote the length of a group element $g \in G$.

By means of this length l_S, we can define a metric d_S on G, the *word metric d_S on G w.r.t. S*:

$$d_S(g,h) := l_S(g^{-1}h).$$

This is clearly a metric, and moreover it is *left-invariant*, i.e. $d_S(g,h) = d_S(kg,kh)$. Another way to express is that G acts isometrically by left-multiplication on itself, i.e. $G \times G \to G, (k,g) \mapsto kg$ is an isometry.

Definition 1.3 (Cayley graph)**.** Let G be a group, $S \subset G$ a set of generators, i.e. $G = \langle S \rangle$. Then the *Cayley graph $\mathcal{C}_S(G)$ of G w.r.t. S* is the graph with vertices $V = G$ and edges $E = \{\{g,h\} \mid g,h \in G, g \neq h, g^{-1}h \in S^{\pm -1}\}$, where $S^{\pm -1} := S \cup S^{-1} := S \cup \{s^{-1} | s \in S\}$.

Thus we have a natural inclusion $G \hookrightarrow \mathcal{C}_S(G)$ and for convenience we do not differentiate between the group element $g \in G$ and the vertex $v \in \mathcal{C}_S(G)$ representing this element and simply write $g \in \mathcal{C}_S(G)$.

This allows a more accessible formulation: Two vertices g, h are joined by an edge if and only if there exists $s \in S$ such that $h = gs$ or $g = hs$ i.e. $h = gs^{-1}$.

Is $w = s_1 \cdots s_n$ a word given in generators $s_i \in S^{\pm -1}$, we can traverse the *path starting at a vertex g labelled with w* in $\mathcal{C}_S(G)$. Starting at g, we traverse the edge corresponding to s_1 connecting g to gs_1. Then we successively traverse the edges corresponding to s_i connecting $gs_1 \cdots s_{i-1}$ to $gs_1 \cdots s_i$.

So far this definition is a purely combinatorial one. To obtain a geometric object, we make the Cayley graph into a metric space by identifying edges with copies of the unit interval. Then the word metric on G naturally extends to a metric on $\mathcal{C}_S(G)$.

We will now define hyperbolic spaces. Remember that a *geodesic* in a metric space X is an isometric embedding $I \to X$ of an interval $I \subset \mathbb{R}$. A metric space is called *geodesic*, if for every pair of points $x, y \in X$ there exists a geodesic connecting x with y. We write $[x,y]$ to denote a geodesic segment joining x and y. Note that this geodesic is not uniquely determined in general, so unless specified otherwise, it refers to a choice of a geodesic segment connecting x and y. However this is not a problem, since we will not depend on the choice of this segment and only use the fact that it is geodesic. Later we will also see that hyperbolicity places an restriction on how far two geodesics connecting the same endpoints can be apart.

Definition 1.4 (δ-hyperbolic space)**.** Let (X,d) be a geodesic space. A geodesic triangle $\Delta = \Delta(p,q,r) \subset X$ with vertices $p,q,r \in X$ is called *δ-slim* for some $\delta \geq 0$, if each of its sides is contained in the δ-neighbourhood of the other two sides. This means that for each $x \in [p, q]$, where $[p, q] \subset \Delta$ denotes the geodesic segment joining p and q, there exists $y \in [q, r] \cup [r, p]$ such that $d(x, y) \leq \delta$ and analogous inequalities hold for the other two sides.

X is called *δ-hyperbolic*, if there exists a $\delta \geq 0$ such that every geodesic triangle is δ-slim. If we are not interested in the constant δ, we simply say that X is *hyperbolic*.

This definition of δ-hyperbolic is generally credited to Eliyahu Rips and therefore called *Rips definition*.

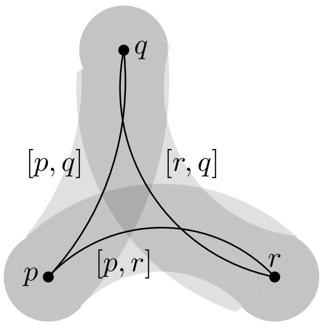

Figure 1: A slim triangle with the δ-neighbourhoods indicated in grey.

Although we still have to prove that the notion of hyperbolic groups is well-defined, the definitions given so far allow us to discuss some examples.

Example 1.5 (Free groups are hyperbolic)**.** Let F_m be the free group of rank m. Then F_m is 0-hyperbolic.

Proof. Let $S = \{x_1, \ldots, x_m\}$ be a free generating set of F_m. The Cayley graph $\mathcal{C}_S(F)$ is a tree. Geodesic triangles are (possibly degenerated) tripods, so each side lies in the 0-neighbourhood of the union of the other sides. □

Example 1.6 (Finite groups are hyperbolic)**.** Let G be a finite group. Then G is δ-hyperbolic, for some $\delta \geq 0$.

Proof. Because the vertices of $\mathcal{C}_S(G)$ are in one-to-one correspondence with group elements of G, $\mathcal{C}_S(G)$ is finite for any generating set $S \subset G$. Hence the diameter of $\mathcal{C}_S(G)$ gives an upper bound on δ. □

Definition 1.7 (Quasi-isometry)**.** Let $f : X \to Y$ be a map between metric spaces (X, d_X) and (Y, d_Y). Then the map f is a (λ, ε)*-quasi-isometric embedding*, if there are constants $\lambda \geq 1$ and $\varepsilon \geq 0$ such that for all $x, x' \in X$

$$\frac{1}{\lambda} d_X(x, x') - \varepsilon \leq d_Y(f(x), f(x')) \leq \lambda d_X(x, x') + \varepsilon.$$

Note that if this condition holds for some λ and ε, it also holds for λ' and ε' if $\lambda' \geq \lambda$ and $\varepsilon' \geq \varepsilon$. A map $f' : X \to Y$ is said to have *finite distance* from f, if there is a constant $c \geq 0$ such that for all $x \in X$

$$d_Y(f(x), f'(x)) \leq c.$$

Such an f is called a (λ, ε)-*quasi-isometry* if it is a (λ, ε)-quasi-isometric embedding for which there is a *quasi-inverse* quasi-isometric embedding, i.e. if there is a quasi-isometric embedding $g : Y \to X$ such that $g \circ f$ has finite distance from id_X and $f \circ g$ has finite distance from id_Y.

If the additive error ε is zero, f is also called a *bilipschitz embedding* if it is $(\lambda, 0)$-quasi-isometric embedding respectively a *bilipschitz equivalence* if it is a $(\lambda, 0)$-quasi-isometry.

This definition of quasi-isometry is good in the sense that it gives one an idea about the nature of quasi-isometries. However, checking that a map between two metric spaces is a quasi-isometry is cumbersome, since one has to show that there is a quasi-inverse quasi-isometric embedding and that the compositions of both maps have finite distance to the identity in the respective space. The next proposition, adapted from [Lö11], alleviates this by giving a criterion which ensures the existence of such quasi-inverses.

Proposition 1.8 (Alternative characterization of quasi-isometry)**.** Let $f : X \to Y$ be a map between metric spaces (X, d_X) and (Y, d_Y). Then f is a quasi-isometry if and only if it is a quasi-isometric embedding with *quasi-dense* image.

A map $f : X \to Y$ is said to have *quasi-dense* image if there is a constant $c \geq 0$ such that for all $y \in Y$ there exists a $x \in X$ satisfying $d_y(f(x), y) \leq c$.

Proof. Suppose that $f : X \to Y$ is a quasi-isometry. Then there exists a quasi-inverse quasi-isometric embedding $g : Y \to X$. Since $f \circ g$ has finite distance to id_Y, there exists $c \geq 0$ such that $d_Y(f \circ g(y), y) \leq c$ for all $y \in Y$. In particular we have that for all $y \in Y$ exists $x \in X$ (namely $g(y)$) such that $d_Y(f(x), y) \leq c$. Hence f has a quasi-dense image.

Now let $f : X \to Y$ be a (λ, ε)-quasi-isometric embedding with quasi-dense image. By increasing the constants λ, ε and c, if necessary, one can obtain a constant $k > 0$ (e.g. $k = \max\{\lambda, \varepsilon, c\}$) such that for all $x, x' \in X$

$$\frac{1}{k} d_X(x, x') - k \leq d_Y(f(x), f(x')) \leq k d_X(x, x') + k$$

and for all $y \in Y$ exists $x \in X$ satisfying $d_Y(f(x), y) \leq k$. One can construct a quasi-inverse quasi-isometric embedding in the following way: Define a map $g : Y \to X$, $y \mapsto x_y$ by invoking the axiom of choice to choose for every $y \in Y$ an element x_y with $d_Y(f(x_y), y) \leq k$.

Then g is a quasi-inverse to f, because by construction holds for all $y \in Y$

$$d_Y(f \circ g(y), y) = d_Y(f(x_y), y) \leq k.$$

Since f is a quasi-isometric embedding, one has for all $x, y \in X$

$$\frac{1}{k} d_X(x,y) - k \leq d_Y(f(x), f(y)) \Leftrightarrow d_X(x,y) \leq k d_Y(f(x), f(y)) + k^2.$$

This and the choice of x_y ensures that for all $x \in X$

$$d_X(g \circ f(x), x) = d_X(x_{f(x)}, x) \leq k d_Y(f(x_{f(x)}), f(x)) + k^2 \leq kk + k^2 = 2k^2.$$

So both $f \circ g$ and $g \circ f$ have finite distance from the respective identity maps. It remains to show that g is also a quasi-isometric embedding.

To this end, let $y, y' \in Y$. Then

$$d_X(g(y), g(y')) = d_X(x_y, x_{y'}) \leq k d_Y(f(x_y), f(x_{y'})) + k^2.$$

Using the triangle inequality, one can intersperse y and y'

$$= k\left(d_Y(f(x_y), y) + d_Y(y, y') + d_Y(f(x_{y'}), y')\right) + k^2.$$

Remembering that $x_y = g(y)$ and $x_{y'} = g(y')$ and that $f \circ g$ has finite distance k to id_Y, this becomes

$$\leq k\left(d_Y(y, y') + 2k\right) + k^2 = k d_Y(y, y') + 3k^2.$$

Similarly one can obtain a lower bound, however using the reverse triangle inequality:

$$\begin{aligned}
d_X(g(y), g(y')) = d_X(x_y, x_{y'}) &\geq \frac{1}{k} d_Y(f(x_y), f(x_{y'})) - 1 \\
&\geq \frac{1}{k}\left(d_Y(y, y') - d_Y(f(x_y), y) - d_Y(f(x_{y'}, y'))\right) - 1 \\
&\geq \frac{1}{k}\left(d_Y(y, y') - 2k\right) - 1 \\
&= \frac{1}{k} d_Y(y, y') - 3 \geq \frac{1}{k} d_Y(y, y') - 3k^2.
\end{aligned}$$

This shows that g is a quasi-isometric embedding. □

This allows us to show that Cayley graphs belonging to different finite generating sets of a finitely generated group are quasi-isometric to each other.

Lemma 1.9 (Quasi-isometry of Cayley graphs)**.** Let G be a finitely generated group, $S, S' \subset G$ two different finite generating sets of G and let $\mathcal{C}_S$ and $\mathcal{C}_{S'}$ be the Cayley graphs of G w.r.t. S and S'. Then the identity map id_G is a quasi-isometry between $\mathcal{C}_S$ and $\mathcal{C}_{S'}$.

Proof. First we show that (G, d_S) and $(G, d_{S'})$ are quasi-isometric to each other. In a second step, we show that $(G, d_S) \hookrightarrow (\mathcal{C}_S(G), d_S)$ is a quasi-isometry.

Since S is finite, the maximum

$$\lambda := \max_{a \in S} d_{S'}(1, a)$$

exists and is finite. Let $g, h \in G$ and let $n := d_S(g, h)$. Thus we can write $g^{-1}h = s_1 \cdots s_n$ for $s_1, \cdots, s_n \in S$. Using the triangle inequality and the left-invariance of the word metric, we can conclude that

$$\begin{aligned} d_{S'}(g,h) &= d_{S'}(g, gs_1 \cdots s_n) \\ &\leq d_{S'}(g, gs_1) + d_{S'}(gs_1, gs_1s_2) + \cdots + d_{S'}(gs_1 \cdots s_{n-1}, gs_1 \cdots s_n) \\ &= d_{S'}(1, s_1) + d_{S'}(1, s_2) + \cdots + d_{S'}(1, s_n) \\ &\leq n\lambda = \lambda d_S(g,h). \end{aligned}$$

Reversing the roles of S and S', we obtain that $\mathrm{id}_G : (G, d_S) \to (G, d_{S'})$ is a bilipschitz equivalence and therefore also a quasi-isometry.

Since quasi-isometry is an equivalence relation, it suffices to show that $(G, d_S) \hookrightarrow (\mathcal{C}_S(G), d_S)$ is a quasi-isometry: For every $y \in \mathcal{C}_S(G)$ exists a $x \in G$ such that $d_S(x, y) \leq 1$, thus the inclusion has a quasi-dense image. Since the metrics of G and $\mathcal{C}_S(G)$ coincide on the image of G, it is also a quasi-isometric embedding, thus, by Proposition 1.8, the inclusion is a quasi-isometry. □

To see that the definition of hyperbolic groups makes sense, it remains to show that hyperbolicity is preserved under quasi-isometry (the constant δ may change, though). As mentioned, this requires some more work.

The following observation and lemmata are adapted from [BH09], which covers many of the properties of hyperbolic spaces and groups. To build intuition for hyperbolic spaces, note the following.

Observation 1.10 (Geodesics in hyperbolic space stay close to each other)**.** Let (X, d) be a δ-hyperbolic space for some $\delta \geq 0$ and let $p, q \in X$. If c, c' are two geodesics with endpoints p and q, then $d(\mathrm{im}(c), \mathrm{im}(c')) \leq \delta$.

Proof. One can view c and c' as a degenerated geodesic triangle: Choose a point $x \in c$ and consider the geodesic triangle with sides $[p, x] \subset c$, $[x, q] \subset c$ and c'. Then c' is contained in the δ-neighbourhood of $[p, x] \cup [x, q] = c$. □

This observation generalizes to rectifiable curves and *quasi-geodesics* as well, where the latter are defined analogous to geodesics: A (λ, ε)-quasi-geodesic in a metric space is a (λ, ε)-quasi-isometric embedding of an interval of the real line. The following lemma shows that in hyperbolic space rectifiable curves behave well, i.e. they stay comparable close to geodesics connecting the same endpoints.

Lemma 1.11 (Rectifiable curves stay close to geodesics)**.** Let X be a δ-hyperbolic geodesic space and let c be a rectifiable path in X. If $[p, q]$ is a geodesic segment connecting the endpoints of c, then for every $x \in [p, q]$

$$d(x, \mathrm{im}(c)) \leq \delta |\log_2 l(c)| + 1.$$

Proof. The case that $l(c) \leq 1$ is trivial, thus suppose $l(c) > 1$: Without loss of generality let $c : [0, 1] \to X$ be parameterized proportional to arc length, $p = c(0)$ and $q = c(1)$. Let N be such that $l(c)/2^{N+1} < 1 \leq l(c)/2^N$. Let be $\Delta_1 = \Delta([c(0), c(1/2)], [c(1/2), c(1)], [c(0), c(1)])$ a geodesic triangle with the given geodesic $[c(0), c(1)]$.

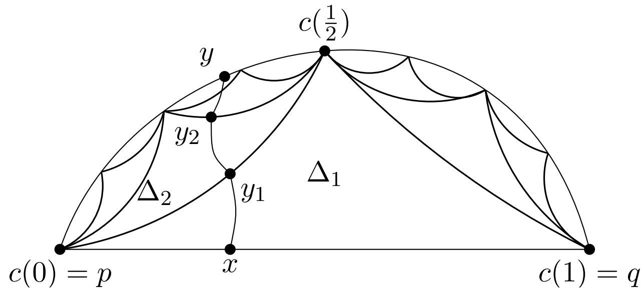

Figure 2: In hyperbolic space rectifiable paths stay close to geodesics. We see some of the geodesic triangles used in the proof of Lemma 1.11.

Application of the criterion for δ-slim triangles, we find y_1 in the union $[(c(0), c(1/2)] \cup [c(1/2), c(1)]$ such that $d(x, y_1) \leq \delta$. Repeat this process inductively with a geodesic triangle consisting of the geodesic containing y_{N-1} and two new sides chosen as above. We find in the N-th stage a y_N such that $d(x, y_N) \leq \delta N$ on an interval of length $l(c)/2^N$ with endpoints in $\mathrm{im}(c)$. Define y to be the closest endpoint of the interval. It follows that $l(c)/2^{N+1} < 1$, $2^N \leq l(c)$ and from this we can conclude that $d(x, y) \leq \delta |\log_2 l(c)| + 1$. ☐

Lemma 1.12 (Taming quasi-geodesics)**.** Let X be a geodesic space, $c : [a, b] \to X$ a (λ, ε)-quasi-geodesic. Then there exists a continuous (λ, ε')-quasi-geodesic $c' : [a, b] \to X$ with

(1) $c'(a) = c(a)$ and $c'(b) = c(b)$

(2) $\varepsilon' = 2(\lambda + \varepsilon)$

(3) $l(c'|_{[t,t']}) \leq k_1 d(c'(t), c'(t')) + k_2$ for all $t, t' \in [a, b]$ where $k_1 = \lambda(\lambda + \varepsilon)$ and $k_2 = (\lambda\varepsilon' + 3)(\lambda + \varepsilon)$

(4) $d_H(\mathrm{im}(c), \mathrm{im}(c')) \leq (\lambda + \varepsilon)$, where d_H is the Hausdorff distance.

Proof. Choose $c'(t) := c(t)$ for all $t \in \Sigma := \{a, b\} \cup (\mathbb{Z} \cap [a, b])$ and define c' as concatenation of linear reparameterizations of geodesic segments connecting these points. The length of each segment is at most $(\lambda + \varepsilon)$ and every point of $\mathrm{im}(c) \cup \mathrm{im}(c')$ lies in the $((\lambda + \varepsilon)/2)$-neighbourhood of $c(\Sigma)$, thus we have proven (4).

For $t \in [a, b]$ define $[t] \in \Sigma$ to be the point closest to t. For $t, t' \in [a, b]$ we have

$$\begin{aligned} d(c'(t), c'(t') &\leq d(c'([t]), c([t'])) + (\lambda + \varepsilon) \\ &\leq \lambda |[t] - [t']| + \varepsilon + (\lambda + \varepsilon) \\ &\leq \lambda(|t - t'| + 1) + (\lambda + 2\varepsilon). \end{aligned}$$

Because $\lambda \geq 1$, the following inequality holds:

$$\begin{aligned} \frac{1}{\lambda}|t - t'| - 2(\lambda + \varepsilon) &\geq \frac{1}{\lambda}|t - t'| - \lambda - \frac{1}{\lambda} - 2\varepsilon \\ &\leq \frac{1}{\lambda}(|t - t'| - 1) - (\lambda + 2\varepsilon) \\ &\leq \frac{1}{\lambda}|[t] - [t']| - (\lambda + 2\varepsilon) \\ &\leq d(c'([t]), c'([t'])) - (\lambda + \varepsilon) \\ &\leq d(c'(t), c'(t')) \end{aligned}$$

This shows that c' is a (λ, ε')-quasi-geodesic with $\varepsilon' = 2(\lambda + \varepsilon)$ as stated in (2).

For proving (3) let $n, m \in [a, b]$ be integers:

$$\begin{aligned} l(c'|_{[n,m]}) &= \sum_{i=n}^{m-1} d(c(i), c(i+1)) \geq \lambda + \varepsilon |m - n| \\ l(c'|_{[a,m]}) &\leq (\lambda + \varepsilon)(m - a + 1) \end{aligned}$$

It follows that for all $t, t' \in [a, b]$

$$\begin{aligned} l(c'|_{[t,t']}) &\leq (\lambda + \varepsilon)(|[t] - [t']| + 2) \\ d(c'(t), c'(t')) &\geq \frac{1}{\lambda}|t - t'| - \varepsilon' \geq \frac{1}{\lambda}(|[t] - [t']| - 1) - \varepsilon'. \end{aligned}$$

□

Lemma 1.13 (Stability of quasi-geodesics)**.** For all $\delta > 0, \lambda \geq 1, \varepsilon \geq 0$ there exists a constant $R = R(\delta, \lambda, \varepsilon)$ such that: If X is a δ-hyperbolic space, c a (λ, ε)-quasi-geodesic and $[p, q]$ a geodesic segment joining the endpoints of c, then $d_H([p, q], \mathrm{im}(c)) < R$, where d_H again is the Hausdorff distance.

Proof. First tame c by replacing it with c' as in Lemma 1.12. Let $[p, q]$ be a geodesic segment joining the endpoints of c'. Let $D := \sup\{d(x, \operatorname{im}(c')|x \in [p, q]\}$ and let $x_0 \in [p, q]$ be such that this supremum is attained. Then the intersection of the D-neighbourhood $B(x_0, D)$ of x_0 with $\operatorname{im}(c')$ is empty.

Choose $y \in [p, x_0]$ such that $d(y, x_0) = 2D$ (or if not possible let $y = p$) and $z \in [x_0, q]$ similarly (or $z = q$). Choose now $y', z' \in \operatorname{im}(c)$ such that $d(y, y') \leq D$ and $d(z, z') \leq D$ and geodesic segments $[y, y'], [z, z']$.

Consider the path γ from y to z traversing $[y, y']$, then following c' from y' to z' and finally traversing $[z', z]$. This path lies outside of $B(x_0, D)$ and it follows that

$$d(y', z') \leq d(y', y) + d(y, z) + d(z, z') \leq 6D.$$

By Lemma 1.12(3) we have that $l(\gamma) \leq 6Dk_1 + k_2 + 2D$.

From Lemma 1.11 follows that $d(x_0, \operatorname{im}(\gamma)) = D$. Therefore $D - 1 \leq \delta|\log_2(l(\gamma))| \leq \delta|\log_2(6Dk_1 + k_2 + 2D)|$, so we have an upper bound on D depending only on δ, λ and ε.

Let D_0 be such an upper bound. Then $\operatorname{im}(c')$ is contained in the $R' = D_0(k_1 + 1) + k_2/2$ neighbourhood: Consider a maximal subinterval $[a', b'] \subset [a, b]$ such that $c'([a', b'])$ lies outside of the D_0 neighbourhood $B(D_0, [p, q])$ of $[p, q]$. Every point of $[p, q]$ lies in $B(D_0, \operatorname{im}(c'))$, so by connectedness there are $w \in [p, q], t \in [a, a']$, $t' \in [b, b']$ such that $d(w, c'(t)) \leq D_0$ and $d(w, c'(t')) \leq D_0$. It follows that $d(c'(t), c'(t')) \leq 2D_0$ and by Lemma 1.12(3) we have $l(c'|_{[t,t']}) \leq 2k_1D_0 + k_2$, so $\operatorname{im}(c')$ is contained in the R'-neighbourhood of $[p, q]$ and by Lemma 1.12(4) we obtain $R = R' + \lambda + \varepsilon$. □

Corollary 1.14 (Quasi-geodesic triangles are slim)**.** A geodesic metric space X is hyperbolic if and only if for every $\lambda \geq 1, \varepsilon > 0$, there exists a constant M such that every λ, ε-quasi-geodesic triangle in X is M-slim, and if X is δ-hyperbolic, then M depends only on $\delta, \lambda, \varepsilon$.

Proof. Let X be δ-hyperbolic space, $p, q, r \in X$ and $\overline{\Delta}$ be a (λ, ε)-quasi-geodesic triangle with vertices p, q, r. By Lemma 1.13 there exists a constant $R = R(\delta, \lambda, \varepsilon)$ such that all sides have Hausdorff distance of less than R to geodesics connecting the respective endpoints. Because the geodesics form a geodesic triangle $\Delta(p, q, r)$ and X is δ-hyperbolic, each geodesic is contained in the δ-neighbourhood of the union of the other two geodesics. So we have that each quasi-geodesic side is contained in the $\delta+2R$-neighbourhood of the union of the other two quasi-geodesic sides, thus $\overline{\Delta}$ is M-slim for $M = \delta+2R$.

Conversely, if there exists such a M, every geodesic triangle is M-slim since geodesics are $(1, 0)$-quasi-geodesics. □

This finally enables us to prove that hyperbolicity is an invariant of quasi-isometry.

Theorem 1.15 (Hyperbolicity is an invariant of quasi-isometry)**.** Let (X, d) and (X', d') be geodesic metric spaces and let $f : X' \to X$ be a (λ, ε)-quasi-isometric embedding. If X is δ-hyperbolic, then X' is δ'-hyperbolic, where δ' depends on δ, λ and ε.

Proof. Let Δ be a geodesic triangle in X' with sides $\gamma_1, \gamma_2, \gamma_3$ and consider the (λ, ε)-quasi-geodesic triangle in X with sides $f \circ \gamma_1$, $f \circ \gamma_2$, $f \circ \gamma_3$. By Corollary 1.14 there is a constant $M = M(\delta, \lambda, \varepsilon)$ such that this triangle is M-slim, thus for all $x \in \mathrm{im}(\gamma_1)$ there exists $y \in \mathrm{im}(\gamma_2) \cup \mathrm{im}(\gamma_3)$, such that $d(x, y) \leq M$. Since f is a (λ, ε)-quasi-isometric embedding, we have that

$$d(x, y) \leq \lambda d(f(x), f(y)) + \varepsilon\lambda \leq \lambda M + \lambda\varepsilon.$$

Analogous we see that $f \circ \gamma_2$ and $f \circ \gamma_3$ are contained in the $\lambda M + \lambda\varepsilon =: \delta'$-neighbourhood of the union of the respective other sides. Thus Δ is δ'-slim and X' is δ'-hyperbolic. □

Now we see that the term hyperbolic group is well-defined, since it does not depend on the choice of the generating set. While the value of δ might change according to Theorem 1.15, the characteristic traits of hyperbolicity are preserved.

2 The Word and Conjugacy Problem of hyperbolic groups

In this section we want to discuss the Word and Transformation problem of hyperbolic groups. Max Dehn formulated these already in 1911 ([Deh11]), long before the theory of computability emerged. The question is whether there exists an algorithm to solve these problems or not.

Problem 2.1 (The Word Problem)**.** Given a word of a group in form of a product of generators, is this word the identity?

Problem 2.2 (The Transformation Problem)**.** Given two group elements $g, h \in G$, can they be transformed into each other, i.e. exists $u \in G$ satisfying $g = uhu^{-1}$?

Another name for the Transformation Problem is *Conjugacy Problem*, since one essentially asks whether g and h are conjugate in G or not. Note that these problems are not solvable in general and vary in difficulty. In fact, the Transformation Problem contains the Word Problem as a special case, namely if w is a word in in form of a product of generators, we can ask whether or not $1 = uwu^{-1}$ i.e. $1 = w$ for some $u \in G$.

Our algorithms will depend on the length of a given input word w, i.e. the number of letters in this particular word. If $S \subset G$ is a generating set of the group G and $w = s_1 \cdots s_m$ a word in generators $S^{\pm 1}$, we shall use $|w| := m$ to denote its length, regardless of whether w is a reduced word or not. Clearly this is not the same length as defined in Section 1, e.g. $l(ss^{-1}) = l(1) = 0$, but $|ss^{-1}| = 2$.

2.1 Solving the Word Problem of hyperbolic groups

First we want to solve the Word Problem in hyperbolic groups. For this we will show that hyperbolic groups admit to a *Dehn presentation*, which allows us to apply *Dehn's Algorithm.*

Definition 2.3 (Dehn presentation)**.** A finite presentation $\langle S \mid \mathcal{R} \rangle$ of a group G is called a *Dehn presentation* if $\mathcal{R} = \left\{u_1 v_1^{-1}, \ldots, u_n v_n^{-1}\right\}$ such that:

i) $u_i = v_i$ in G

ii) $|v_i| < |u_i|$

iii) If a word w represents the identity in G then at least one of the u_i is a subword of w.

Algorithm 2.4 (Dehn's Algorithm for Solving the Word Problem)**.** Let $\langle S \mid \mathcal{R} \rangle$ be a Dehn presentation for a group G. Then we can use following algorithm to solve the Word Problem for a given word $w \in \mathcal{F}(\mathcal{S})$:

While w is not the empty word, look for a subword of the form u_i. If there is no such subword, stop, because of the third condition of the Dehn presentation, the word does not represent the identity; if u_i occurs as a subword, replace it with v_i and repeat with the shorter word obtained from w.

After at most $|w|$ iterations the word is either reduced to the empty word, i.e. $w = 1$ in G, or else it is verified that w does not represent the identity.

Now it has to be shown that if G is a hyperbolic group, then G admits to a Dehn presentation. For this it is useful to establish a local criterion for recognizing quasi-geodesics in hyperbolic spaces.

Definition 2.5 (k-Local Geodesics)**.** Let (X, d) be a metric space, $k > 0$. A path $c : [a, b] \to X$ is said to be *k-local geodesic*, if $d(c(t), c(t')) = |t - t'|$ for all $t, t' \in [a, b]$ with $|t - t'| < k$.

Proposition 2.6 (k-Local Geodesics are Quasi-Geodesics)**.** Let X be a δ-hyperbolic geodesic space and let $c : [a, b] \to X$ be a k-local geodesic, where $k > 8\delta$. Then

(1) $\operatorname{im}(c)$ is contained in the 2δ-neighbourhood of any geodesic segment $[c(a), c(b)]$ connecting its endpoints

(2) $[c(a), c(b)]$ is contained in the 3δ-neighbourhood of $\mathrm{im}(c)$

(3) c is a (λ, ε)-quasi-geodesic, where $\varepsilon = 2\delta, \lambda = (k + 4\delta)/(k - 4\delta)$.

Proof. First prove (1). Let $x = c(t)$ be a point of $im(c)$ that is at maximal distance from $[c(a), c(b)]$. First we suppose that $(t-a)$ and $(b-t)$ are greater than 4δ. Then there is a subarc of c centred at x of length strictly greater than 8δ but less then k. Let y, z be the endpoints of this arc and y', z' the points on $[c(a), c(b)]$ closest to y and z respectively. Consider a geodesic quadrilateral with vertices y, z, y', z' such that the sides $[y, z]$ and $[y', z']$ are the subarcs of c and $[c(a), c(b)]$. Dividing this quadrilateral with a diagonal, as shown in Figure 3, and applying the δ-hyperbolic criterion to each of the resulting triangles, we find w on one of the sides other than $[y, z]$ such that $d(w, x) \leq 2\delta$. If $w \in [y, y']$ then there would be a path through w joining x

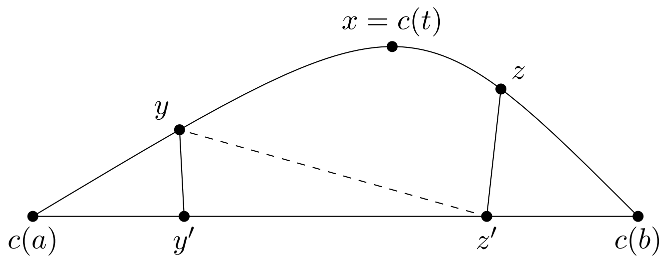

Figure 3: The geodesic quadrilateral in the proof of Proposition 2.6 (1). The dashed diagonal divides the quadrilateral so we can apply the δ-hyperbolic criterion of slim triangles.

to y' that was shorter than $d(y, y')$, thus contradicting our choice of x:

$$\begin{aligned} d(x, y') - d(y, y') &\leq d(x, w) + d(w, y') - d(y, w) - d(w, y') \\ &= d(x, w) - d(y, w) \\ &\leq d(x, w) - [d(y, x) - d(x, w)] \\ &= 2d(x, w) - d(x, y) \\ &< 4\delta - 4\delta = 0 \end{aligned}$$

Now suppose that $t - a < 4\delta$. By choosing z, z' as above we obtain a geodesic triangle with sides $[c(a), z]$, $[z, z']$, $[z', c(a)]$. It follows that there exists a $w \in [z, z'] \cup [c(a), z']$ with $d(w, x) \leq \delta$. Again, if $w \in [z, z']$, we can find a path through w joining x and z', arriving at the same contradiction to our choice of x:

$$\begin{aligned} d(x,z') - d(z,z') &\geq d(x,w) + d(w,z') - d(z,w) - d(z,w') \\ &= d(x,w) - d(z,w) \\ &\leq d(x,w) - [d(z,x) - d(x,w)] \\ &= 2d(x,w) - d(x,y) \\ &< 4\delta - 4\delta = 0 \end{aligned}$$

If both $(t-a)$ and $(b-t)$ are less than 4δ, it follows immediately that c is a geodesic and, as seen in Observation 1.10, c and $[c(a), c(b)]$ have distance of at most δ.

To prove (2), let $p \in [c(a), c(b)]$. Because of (1) every point of $\mathrm{im}(c)$ lies in the open 2δ-neighbourhood of $[c(a), p]$ or $[p, c(b)]$. Since c is connected there is $x \in \mathrm{im}(c)$ lying in both neighbourhoods. Choose $q \in [c(a), p]$ and $r \in [p, c(b)]$ such that $d(x,q) \leq 2\delta$ and $d(r,x) \leq 2\delta$. because $p \in [q,r]$ it follows that p lies in the δ-neighbourhood of $[q,x] \cup [r,x]$, thus

$$d(p, \mathrm{im}(c)) \leq d(p, [q,x] \cup [r,x]) + \max\{d(q,x), d(r,x)\} \leq 3\delta.$$

To prove (3), first note that $d(c(t), c(t')) \leq |t-t'|$ for all $t, t' \in [a,b]$, since c is a k-local geodesic. Thus in order to show that c is a quasi-geodesic one has to obtain a lower bound of $d(c(t), c(t'))$ by a linear function of $|t-t'|$. For this, divide c into subpaths of length $k' = k/2 + 2\delta$ and project the endpoints of these subarcs onto $[c(a), c(b)]$, i.e. make a choice of closest points. The lower bound is then obtained by estimating the distance between points of this sequence.

Consider a subarc of length $2k'$ and let x and y be the endpoints of this arc and let m be the midpoint of the arc; let x', y' and m' be points of $[c(a), c(b)]$ that have a distance of at most 2δ from x', y' and m' respectively. We will first show that m' lies between x' and y'.

Let x_0 (resp. y_0) be the point on $\mathrm{im}(c)$ that has a distance of 2δ from x (resp. y) in the direction of m. By δ-hyperbolicity and (1) then any geodesic triangle $\Delta(x, x', x_0)$ is contained in the 3δ-neighbourhood of x. Because of $d(x,m) = k' > 6\delta$ and choice of x_0, it follows that $d(x_0, m) > 3\delta$, meaning any such triangle lies completely outside the 3δ-neighbourhood of m, since c is k-local geodesic for $k > 8\delta$ and therefore x_0 is the point of $\Delta(x, x', x_0)$ with minimum distance to m. Similarly one can conclude that any geodesic triangle $\Delta(y, y', y_0)$ lies outside the 3δ-neighbourhood of m.

Consider the geodesic quadrilateral with vertices x', x_0, y_0 and y' and divide it by a diagonal into two triangles. Applying the slim triangle condition, it follows that m lies in the 2δ-neighbourhood of $[x_0, x'] \cup [x', y'] \cup [y', y_0]$. However $[x_0, x']$ and $[y_0, y']$ are sides of the triangles considered above and have at least distance 3δ to m, thus there exists $m'' \in [x', y'] \subset [c(a), c(b)]$ such that $d(m, m'') \leq 2\delta$.

If $m' = m''$ then it follows that m' lies between x' and y', so suppose $m' \neq m''$. By hyperbolicity of the geodesic triangle $\Delta(m, m', m'')$, any point between m' and m'' has distance of at most 3δ to m, since $d(m, m'') \leq 2\delta$ and $d(m, m') \leq 2\delta$. In particular neither x' nor y' can lie between m' and m'', thus $m' \in [x', y']$.

We express c as concatenation of $M \leq (b-a)/k'$ geodesics of length k' and a smaller piece, of length η, at the end. By the preceding argument, the projections of the endpoints of these geodesics onto $[c(a), c(b)]$ form a monotone sequence. If $p', q' \in [c(a), c(b)]$ are successive projections of points $p, q \in \mathrm{im}(c)$ one can estimate the minimum distance between p and q as follows. By (1) it holds that $d(p, p') \leq 2\delta$ and $d(q, q') \leq 2\delta$. Furthermore $d(p, q) = k' > 6\delta$ and p' is a point on $[c(a), c(b)]$ closest to p, thus we have by repeated application of the reversed triangle inequality

$$\begin{aligned} d(p', q') &\geq |d(p', p) - d(p, q')| = d(p, q') - d(p', p) \\ &\geq |d(p, q) - d(q, q')| - d(p, p') \\ &\geq |k' - 2\delta| - 2\delta = k - 4\delta. \end{aligned}$$

The triangle inequality yields that the distance from the last projection point p' of p to $c(b)$ is at least $d(p', c(b)) \geq |d(p', p) - d(p, c(b)| \geq |2\delta - \eta|$ and because p' is a point closest to p, it holds that $d(p', c(b)) \geq \eta - 2\delta$.

From the choice of M and η it follows that $b - a = Mk' + \eta$ and the concatenation of the subpaths yields

$$d(c(a), c(b)) \geq M(k' - 4\delta) + \eta - 2\delta = (b-a) - 4\delta M - 2\delta.$$

Since $M \leq (b-a)/k'$, it follows that

$$d(c(a), c(b)) \geq (b-a) - 4\delta(b-a)/k' - 2\delta = \frac{k' - 4\delta}{k'}(b-a) - 2\delta.$$

Remembering the definition of k', the claimed lower bound is established:

$$d(c(a), c(b)) \geq \frac{k/2 - 2\delta}{k/2 + 2\delta}(b-a) - 2\delta = \frac{k - 4\delta}{k + 4\delta}(b-a) - 2\delta$$

For arbitrary $t, t' \in [a, b]$ one obtains the same lower bound, by running the same argument as above for $|t - t'| > k$; if $|t - t'| \leq k$ there is nothing to prove, since c is k-local geodesic. □

Corollary 2.7. If (X, d) is a δ-hyperbolic space, there is no closed k-local geodesic subpath for $k > 8\delta$.

Proof. If $c(a) = c(b)$ then $\mathrm{im}(c)$ is contained in $B(c(a), 2\delta)$ and since c is 8δ-local geodesic it follows that $a = b$, thus c is constant. □

This Corollary allows us to find shortcuts in the Cayley graph of a hyperbolic group, which we will need in the construction of a Dehn presentation.

Lemma 2.8 (Shortcuts in hyperbolic graphs)**.** Let (X, d) be a δ-hyperbolic metric graph with unit edge lengths and δ an integer. Given any non-trivial locally-injective loop $c : [0,1]/\sim \to X$ (where $\sim$ is the equivalence relation identifying 0 and 1) beginning at a vertex, one can find $s, t \in [0,1]$ such that

(i) $l(c|_{[s,t]}) \leq 8\delta + 1$ and $c|_{[s,t]}$ is not geodesic

(ii) $c(t)$ and $c(s)$ are vertices

(iii) there exists a geodesic p joining $c(t)$ and $c(s)$ with $l(p) \leq l(c|_{[s,t]}) - 1$.

Proof. According to Corollary 2.7 X contains no closed loops which are k-local geodesics for $k = 8\delta + \frac{1}{2}$. So we can choose a non-geodesic subarc $p' = c|_{[s_0,t_0]}$ of c of length less than k and a geodesic p connecting $c(s_0)$ to $c(t_0)$. Let $s, t \in [0,1]$ such that $c(s)$ and $c(t)$ is the first respective last vertex that p passes.

We only have to show that condition (i) is satisfied, because (ii) follows from the construction and (iii) results from X having unit edge lengths, thus the length difference between two edge paths connecting the same vertices is an integer.

Let e and f be the first respective last vertex through which p' passes. Clearly, if both $e = c(s)$ and $f = c(t)$, then condition (iii) is satisfied. However, if at least one equality does not hold, we have to attend to some details.

The first case is that only one equality holds, i.e. we have the following two cases.

1.1) $e = c(s)$ and $f \neq c(t)$. Since p' had length less than k, the subarc of c connecting the vertices e to f has length of at most 8δ. Extending p' to $c(t)$ increases this length by at most 1 and makes $p' = c|_{[s,t]}$ satisfying condition i).

1.2) $f = c(t)$ and $e \neq c(s)$. This case is symmetric to case 1.1).

The second case is that $e \neq c(s)$ and $f \neq c(t)$. Here we have to consider four cases.

2.1) $d(c(s_0), e) + d(c(t_0), f) \geq 1.5$. In this case we can extend p' to $c(s)$ and $c(t)$, lengthening it by at most 0.5. Then $c|_{[s,t]} = p'$ and so condition i) is satisfied.

2.2) $d(c(s_0), e) + d(c(t_0), f) \leq 0.5$. Let $s', t' \in [0,1]$ be such that $c(s') = e$ and $c(t') = f$, then we have in this case that $l(c|_{[s',t']}) = d(c(s), c(t)) + 2$. Thus we can shorten p' such that it begins in e, moving us in the situation of the first case. See Figure 4 for an illustration of the situation in this case.

2.3) $1 < d(c(s_0), e) + d(c(t_0), f) < 1.5$. In this case we can move p' such that it begins in $c(s)$ and we are again in case 1.1).

2.4) $0.5 < d(c(s_0), e) + d(c(t_0), f) \leq 1$. In this case we can move p' such that it begins in e. Then either both ends of p' are vertices or we are in the situation of case 1.1).

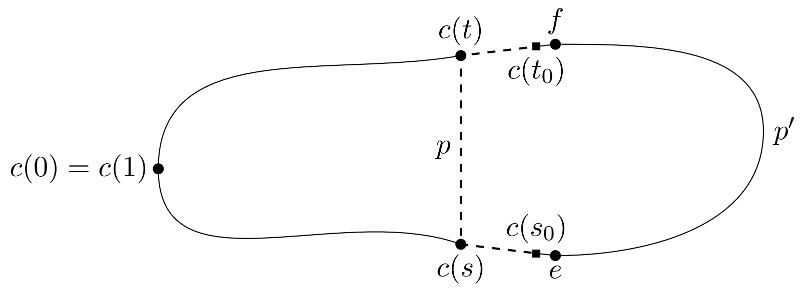

Figure 4: A shortcut in a hyperbolic graph. We are in case 2.2) in the proof of Lemma 2.8 and can shorten p' such that it begins in e.

□

Theorem 2.9 (Hyperbolic groups have a solvable Word Problem)**.** Let G be a hyperbolic group, then G admits a finite Dehn presentation as defined in Definition 2.3, hence Dehn's Algorithm applies and G has a solvable word problem.

Proof. Let the Cayley graph $\mathcal{C}_S(G)$ of G w.r.t. a finite generating set S be δ-hyperbolic (without loss of generality we may assume that δ is an integer, by increasing δ if necessary) and let k be an integer satisfying $k > 8\delta + 1$. Following the proof of the preceding Lemma, every closed loop c in $\mathcal{C}_S(G)$ (labelled with a word w which represents the identity in G) contains a subpath p' beginning and ending at vertices with $l(p') \leq k$ which is not geodesic. Let u be the subword of w labelling p'. Let p be a geodesic with same endpoints as p' as stated in Lemma 2.8 and let v be the word labelling p. The lemma states that $l(p) \leq l(p') - 1$, thus $|v| < |u|$.

Now we can obtain a Dehn presentation $\langle S \mid \mathcal{R} \rangle$ by defining $\mathcal{R}$ as the set of words $u_i v_i^{-1}$, where u_i varies over all words with length of at most k in the generators and their inverses and v_i is a word of minimal length that is equal to u_i in G.

It remains to show that $\langle S \mid \mathcal{R} \rangle$ is indeed a presentation of G. Clearly, $\mathcal{R}$ is constructed in a way such that if $w \in F(S)$ represents the identity in G, so it does in $\langle S \mid \mathcal{R} \rangle$. So let $w \in F(S)$ be a word such that $w = 1 \in \langle S \mid \mathcal{R} \rangle$. This means that w can be reduced to the empty word by applying finitely many relations $R_1, \dots R_n \in \mathcal{R}$. But these relations all have the form $u_i v_i^{-1}$

where $u_i = v_i$ in G, so $R_1, \ldots, R_n$ are relations in G and we can apply these in G to obtain the identity. □

This does not only solve the Word Problem for hyperbolic groups, but we can also see that hyperbolic groups are always finitely presentable, since they admit to a Dehn presentation. Indeed it can be shown that a group is hyperbolic if and only if it admits a finite Dehn presentation. This needs some new ideas, namely a coarse notion of area in Cayley graphs. One can show that a group is hyperbolic if and only if it satisfies a linear isoperimetric inequality, i.e. the area (in this coarse sense) covered by a loop is bounded from above by a linear function depending on the length of the loop, the perimeter of this area. If a group admits to a Dehn presentation, then one can conclude that a linear isoperimetric inequality holds. This however goes beyond the scope of this thesis and we refer the reader to [BH09] for an extensive treatment of this subject.

We will now turn to the solution of the Transformation Problem in hyperbolic groups. To this end we will use the established solution of the Word Problem in hyperbolic groups.

2.2 Solving the Conjugacy Problem in hyperbolic groups

In this section two solutions to the Conjugacy Problem of hyperbolic groups will be given. The first one is a straightforward result of hyperbolicity, however it yields a very inefficient algorithm. For this purpose it is useful to introduce another characterization of hyperbolicity. Since some further geometric consequences of hyperbolicity already were discussed, another, more efficient solution will also be presented.

Definition 2.10 (Quasi-monotone conjugacy property)**.** A group G with finite generating set S is said to have the *quasi-monotone conjugacy property (q.m.c. property)* if there is a constant $K > 0$ (the *q.m.c.-constant*), such that whenever two words $u, v \in F(S)$ are conjugate in G, one can find a word $w = a_1 \cdots a_n$ in letters of S such that

$$w^{-1}uw = v \text{ and } d(1, w_i^{-1}uw_i) \leq K \max\{|u|, |v|\}$$

for $i = 1, \ldots, n$ and $w_i = a_1 \cdots a_i$.

Algorithm 2.11 (Algorithm to determine conjugacy)**.** Let G be a group having a finite generating set $S \subset G$, a solvable Word Problem and the q.m.c. property. Let K be the q.m.c.-constant as defined above.

For each $n > 0$ let $B(n)$ be the set of words in $F(S)$ that have a length of at most n. Since G has a solvable Word Problem, given two words $v_1, v_2 \in B(n)$, one can decide if there exists $a \in S^{\pm 1}$ such that $a^{-1}v_1a = v_2$ in G; if it exists write $v_1 \sim v_2$.

Now consider the finite graph $\mathcal{G}(n)$ with vertex set $B(n)$ that has an edge joining v_1 to v_2 if and only if $v_1 \sim v_2$. By the q.m.c. property two words u and v are conjugate in G if and only if u and v lie in the same pathconnected component of $\mathcal{G}(n)$, where $n = K \max\{|u|, |v|\}$. This allows us to decide whether u and v represent conjugate elements of G.

We will show that hyperbolic groups have the q.m.c. property. To introduce the new notion of hyperbolicity, we give the following two definitions, which highlight some properties of triangles in general metric spaces.

Definition 2.12 (Gromov product)**.** Let (X, d) be a metric space and let $x \in X$. Then the *Gromov product of* $y, z \in X$ *w.r.t.* x is defined to be

$$(y \cdot z)_x = \frac{1}{2}\left(d(y, x) + d(z, x) - d(y, z)\right)$$

Definition 2.13 (Tripods in triangles)**.** If (X, d) is a metric space and one is given three points $x, y, z \in X$, then there are non-negative numbers a, b and c such that $d(x, y) = a + b$, $d(x, z) = a + c$ and $d(y, z) = b + c$. These are uniquely determined.

Proof. Let $a', b', c' \in \mathbb{R}$ positive numbers such that

$$\begin{aligned} d(x, y) &= a + b = a' + b' \\ d(x, z) &= a + c = a' + c' \\ d(y, z) &= b + c = b' + c'. \end{aligned}$$

Using the Gromov product we see that

$$\begin{aligned} a &= (y \cdot z)_x = a' \\ b &= (x \cdot z)_y = b' \\ c &= (x \cdot y)_z = c'. \end{aligned}$$

□

Let $\Delta := \Delta(x, y, z) \subset X$ be a geodesic triangle and let $a, b, c \in \mathbb{R}$ as above. Then we can compare Δ to a metric tree $T_\Delta := T(a, b, c)$ which has three vertices v_x, v_y, v_z of valence one and one vertex o of valence three, and the edges connecting o to v_x, v_y and v_z having length a, b and c respectively. This enables us to define a map $\chi_\Delta : \Delta \to T_\Delta$, which maps x, y, z to v_x, v_y, v_z in the obvious way and extends to sides of the triangle, such that the restriction to a side is an isometry; one can think of collapsing the triangle onto the tripod (see Figure 5).

Note that the preimage $\chi_\Delta^{-1}(t)$ consists of at most two points (less in the case Δ is degenerated to a line) for $t \neq o$ and of at most three points

$\{i_x, i_y, i_z\} = \chi_\Delta^{-1}(o)$, where $i_x \in [y, z]$, $i_y \in [x, z]$ and $i_z \in [x, y]$ such that

$$\begin{aligned} d(i_x, y) &= b & d(i_x, z) &= c \\ d(i_y, x) &= a & d(i_y, z) &= c \\ d(i_z, x) &= a & d(i_z, y) &= b. \end{aligned}$$

These points are uniquely determined because Δ is geodesic and we call them *internal points of* Δ.

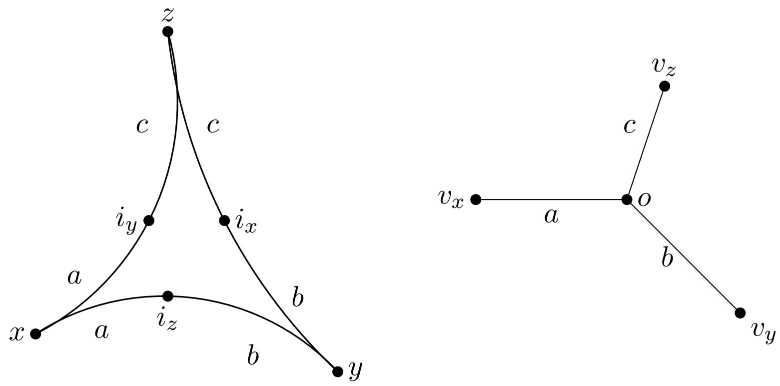

Figure 5: A geodesic triangle $\Delta = \Delta(x, y, z)$ on the left, with internal points i_x, i_y and i_z. On the right the tripod T_Δ, which is obtained from Δ under the map χ_Δ. One can think of collapsing the internal points to o.

We now introduce the concept of thin triangles, by measuring the thickness of a fiber $\chi_\Delta^{-1}(t)$ of the tripod associated to a triangle.

Definition 2.14 (Thin triangles)**.** Let Δ be a geodesic triangle in a metric space (X, d) and consider the map $\chi_\Delta : \Delta \to T_\Delta$ as defined above. Then Δ is said to be δ-*thin* for a $\delta \geq 0$, if $p, q \in {\chi_\Delta}^{-1}(t)$ implies that $d(p, q) \leq \delta$ for all $t \in T_\Delta$.

Given a geodesic triangle $\Delta(x, y, z)$ in a metric space (X, d) and a point p on one of the sides, say γ_1, the advantage of this definition over the notion of slim triangles is that we have more information over the position of a point q in the union of the other two sides, say $q \in \gamma_2$ with $d(p, q) \leq \delta$. We know that if q has the same image as p under χ_Δ, then p and q have the same distance to the vertex belonging to γ_1 and γ_2.

Clearly, a triangle that is δ-thin is also δ-slim. The following lemma shows that the converse holds.

Lemma 2.15 (Slim triangles are thin)**.** Let (X, d) be a δ-hyperbolic space. Then there exists $\delta' \geq 0$ only depending on δ, such that every geodesic triangle $\Delta := \Delta(x, y, z)$ with vertices $x, y, z \in X$ is δ'-thin.

Proof. Let i_x, i_y, i_z be the internal points of Δ.

First we show that $d(i_x, i_y) \leq 2\delta$, $d(i_x, i_z) \leq 2\delta$ and $d(i_y, i_z) \leq 2\delta$. Because Δ is δ-slim, we have that $d(i_x, [x, y] \cup [x, z]) \leq \delta$, with $[x, y]$ and $[x, z]$ being the respective sides of Δ, i.e. there exists a point $p \in [x, y] \cup [x, z]$ such that $d(i_x, p) \leq \delta$, say $p \in [x, y]$. By the reverse triangle inequality we have

$$\delta \geq d(i_x, p) \geq |d(y, p) - d(y, i_x)|.$$

Both i_x and i_y are internal points of Δ, so we have $d(y, i_x) = d(y, i_z)$ and since i_z and p lie on the same geodesic segment, we have

$$\delta \geq |d(y, p) - d(y, i_z)| = d(p, i_z).$$

Consequently $d(i_x, i_z) \leq d(i_x, p) + d(p, i_z) \leq 2\delta$. If $p \in [x, z]$ instead of $[x, y]$ we can run the same argument and obtain $d(i_x, i_y) \leq 2\delta$. Analogous we have that $d(i_z, \{i_x, i_y\}) \leq 2\delta$ and the triangle inequality yields that the distance between two internal points of Δ is at most 4δ.

Now we show that Δ is δ'-thin. Let $s \in T_\Delta$ and $p, q \in \chi_\Delta^{-1}(s)$, say $p \in [y, z]$ and $q \in [x, y]$. We have to show that $d(p, q) \leq \delta'$ for some δ'. To this end, let $c : [0, 1] \to X$ be a monotone parametrization of $[y, z]$, such that $c(0) = y$ and $c(1) = z$. Consider the geodesic triangle $\Delta_t := \Delta(x, y, c(t))$. As one varies t, the internal point of Δ_t on the side $[y, c(t)]$ varies continuously as function of t, being i_x for $t = 1$ and y for $t = 0$. Then there exists t_0 such that the internal point of Δ_{t_0} on the side $[y, c(t_0)]$ is p. Since $d(y, p) = d(y, q)$ then q is also internal point of Δ_{t_0} and we have shown in the first step that $d(p, q) \leq 4\delta$. Clearly the same argument can be applied to any other $s \in T$ and pair of points $p, q \in \chi_\Delta^{-1}(s)$. This completes the proof and shows that Δ is δ'-thin for $\delta' = 4\delta$. □

This notion of hyperbolicity now allows us to show that hyperbolic groups have the q.m.c. property.

Lemma 2.16 (Hyperbolic groups have the q.m.c. property)**.** Let G be a δ-hyperbolic group for $\delta \geq 0$ with finite generating set $S \subset G$. Then G has the quasi-monotone conjugacy property.

Proof. Let $u, v \in G$ be conjugate elements and $w = s_1 \cdots s_n$ a geodesic word (i.e. a word whose path in $\mathcal{C}_S(G)$ is a geodesic) in generators $s_i \in S^{\pm 1}$ such that $w^{-1}uw = v$, and let be $w_i = s_1 \cdots s_i$. Then we have by left-invariance of the word metric d that

$$d(1, v) = d(1, w^{-1}uw) = d(w, uw)$$
$$d(1, w_i^{-1}uw_i) = d(w_i, uw_i).$$

Consider the quadrilateral in the Cayley graph of G with vertices 1, w, u and uw (note that $uw = wv$), as well as two sides labelled with w, and two sides labelled with u and v respectively. Dividing it with the geodesic connecting 1 to $uw = wv$ we obtain two triangles (see Figure 6). Now Lemma 2.15 tells us that there is a point $p \in [1, u]$ such that $d(p, uw_i) \leq \delta'$ for $\delta' = 4\delta$. Similarly, by dividing the quadrilateral with the geodesic $[u, w]$, we obtain another set of triangles, for which we get a point p' on $[1, u]$ such that $d(p', w_i) \leq \delta'$ by the δ'-thin condition. Both p and p' lie on $[1, u]$, so their distance is at most $|u|$. The triangle inequality tells us that $d(w_i, uw_i) \leq 2\delta'|u|$. Analogous we can first divide the quadrilateral with $[u, w]$ to obtain a point $q \in [w, uw = wv]$ such that $d(uw_i, q) \leq \delta'$ and then dividing the quadrilateral with $[1, uw = wv]$ yields $q' \in [w, uw = wv]$, such that $d(w_i, q') \leq \delta'$. Again we can estimate the distance between w_i and uw_i to be at most $2\delta'|v|$. Thus G has the q.m.c. property with q.m.c.-constant $K := 2\delta'$.

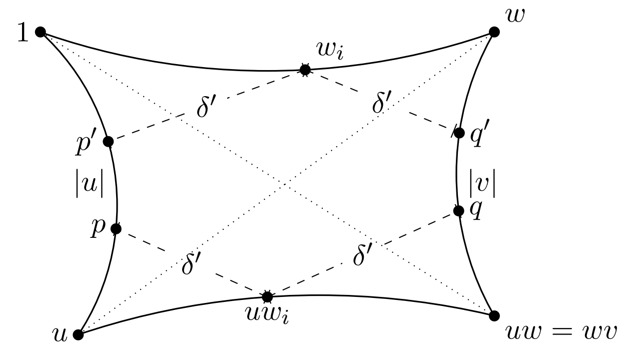

Figure 6: The quadrilateral in the Cayley graph of G and the dividing diagonals, indicated by dotted lines, for application of the thin-triangle condition in the proof of Lemma 2.16.

□

Theorem 2.17 (Hyperbolic groups have a solvable Conjugacy Problem)**.** Let G be a hyperbolic group. Then G has a solvable Transformation problem.

Proof. The preceding lemma shows that G has the q.m.c. property, hence Algorithm 2.11 can be applied. □

This gives us a solution to the Conjugacy Problem in hyperbolic groups. However, this algorithm is very inefficient, because we have to check for pathconnected components in a graph whose vertex set grows exponential with the length of the word. Using the results established in Section 1, we can obtain a more efficient algorithm. The following consequence of Lemma 1.13 and Proposition 2.6 will help us.

Corollary 2.18. Let X be a δ-hyperbolic geodesic space. There is a constant C, depending only on δ, such that if the sides of a quadrilateral $\mathcal{Q} \subset X$ are all $(8\delta + 1)$-local-geodesics, then every side of $\mathcal{Q}$ lies in the C-neighbourhood of the union of the other three sides.

Proof. Using 2.6, we see that $\mathcal{Q}$ is a (λ, ε)-quasi-geodesic quadrilateral, with ε and λ only depending on δ as stated in the proposition. Dividing $\mathcal{Q}$ with a (λ, ε)-quasi-geodesic diagonal, we obtain two quasi-geodesic triangles. By Corollary 1.14 we know that there exists M only depending on $\delta, \varepsilon, \lambda$ such that every (λ, ε)-quasi-geodesic triangle in X is M-slim. Application of the M-slim condition to the quasi-geodesic triangles yields that each side is contained in the C-neighbourhood of the union of the other three sides, for $C \geq 2M$. According to Lemma 1.13 ε and λ only depend on δ, so M depends only δ. It follows that C depends only on δ. □

The following lemma can be seen as a more strict version of the q.m.c. property, where the constant is independent of the length of two words representing conjugate elements u and v.

Lemma 2.19. Let G be a δ-hyperbolic group w.r.t. a finite generating set S. Then there is a constant $K > 0$ only depending on δ such that: If $u, v \in F(S)$ represent conjugate elements of G and if u, v and all their cyclic permutations are $(8\delta + 1)$-local geodesics, then either

(1) $\max\{|u|, |v|\} \leq K$

(2) There exists $w \in F(S)$ with length of at most K such that $w^{-1}u'w = v'$ in G, where u', v' are cyclic permutations of u and v.

Proof. Let be $w \in F(S)$ such that $wuw^{-1} = v$. Consider the geodesic quadrilateral $\mathcal{Q}$ in $C_S(G)$ with sides labelled w, u, w^{-1}, v^{-1}. By replacing u, v with cyclic permutations if necessary, we may suppose that each vertex on the top side of $\mathcal{Q}$ has at least distance $|w|$ from each vertex on the bottom (see Figure 7).

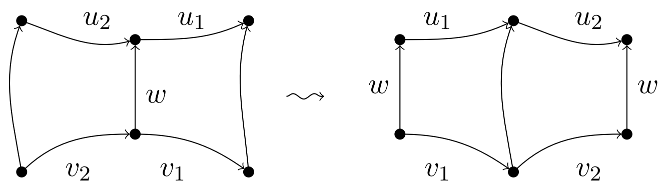

Figure 7: We can rearrange the conjugacy diagramm such that each vertex on the top has at least distance $|w|$ from each vertex on the bottom.

Consider the midpoint p of the top side, then its distance to the other sides is at most C, as shown in Corollary 2.18. If there is a point p' on the bottom side such that $d(p, p') \leq C$, then the distance of the vertices closest to p and p' would be at most $C+1$. If $|w| > C+1$ this can not happen and in this case we have bounded $|w|$ from above with $C+1$.

Suppose that p has distance of at most C to a point q on a vertical side and let x, y be the top and bottom vertices of this side. We now have that $|w| - \frac{1}{2} \leq d(p, y) \leq C + d(q, y)$ and $d(q, y) = |w| - d(x, q)$. It follows that

$$\begin{aligned} d(q, y) &= |w| - d(x, q) \leq C + d(q, y) + \frac{1}{2} \\ \Leftrightarrow d(x, q) &\leq C + \frac{1}{2}. \end{aligned}$$

The triangle inequality now yields

$$d(p, x) \leq d(p, q) + d(x, q) \leq 2C + \frac{1}{2}.$$

Because p was chosen to be the midpoint of the top side, we have that $d(p, x) = \frac{1}{2}|u|$ and we can bound u from above with $4C+1$. Similarly, if $|w| > C+1$ we have that $|v| \leq 4C+1$, so $K = 4C+1$ fulfills the statements in the lemma. $\square$

Algorithm 2.20 (Algorithm to determine conjugacy in hyperbolic groups)**.** Let G be a group that is δ-hyperbolic w.r.t. a finite generating set S. Given two words u, v in letters of $S^{\pm 1}$, look at u, v and their cyclic permutations to find non geodesic subwords with length of at most $8\delta + 1$. If such a subword is found, then replace it by a geodesic word representing the same group element in G.

Repeat this until u, v and all their cyclic permutations are $8\delta + 1$-local geodesics (working with cyclic words, this requires the application of less than $|u| + |v|$ relations from a Dehn presentation of G).

Lemma 2.19 provides a finite set Σ of words, such that u is conjugate to v in G, if and only if $w^{-1}u'w = v'$ in G for some $w \in \Sigma$. Using Dehn's algorithm, we can decide whether one of the relations is valid in G.

A possible choice for Σ would be the set of words with length of at most K with K as in the preceding lemma, together with a choice of one conjugating element for each pair of conjugate elements u_0, v_0 with $\max\{|u_0|, |v_0|\} \leq K$.

Notice that this algorithm is more efficient than the previous one, since the set of possible candidates for conjugating elements does not depend on the length of u and v.

3 Morse functions on affine cell complexes

While the Cayley graph (which is a one dimensional complex) of a group already encodes many of its geometric features, one can use higher dimensional complexes to examine a group with topological machinery. In this section the definition of the *Cayley 2-complex* or *presentation 2-complex* of a group presentation will be given. This is a space which has the given group as fundamental group and its underlying structure is a *cell complex*, which we will define below.

Definition 3.1 (Cell complex)**.** A space X constructed in the following way is called a *cell complex* or a *CW complex*. A *k-cell* e^k is a topological space that is homeomorphic to a k-disk D^k. Note that its boundary ∂e^k is homeomorphic to the $(k-1)$-sphere $\mathbb{S}^{k-1}$. We call k the *dimension* of the cell e^k. A *cell complex* or a *CW complex* X is a space obtained by gluing a collection of cells together as follows:

(1) We start with a countable set X^0, the 0-*skeleton of* X . Its points are homeomorphic to 0-disks and henceforth regarded as 0-cells.

(2) Inductively we obtain the *n-skeleton* X^n from X^{n-1} by attaching n-cells e^n_α via homeomorphisms $\varphi : \partial\mathbb{S}^{n-1} \to X^{n-1}$. This means that X^n is the quotient space of the disjoint union $X^{n-1} \bigsqcup_\alpha D^n_\alpha$ of X^{n-1} with a collection of n-disks D^n_α under the identifications $x \sim \varphi_\alpha(x)$ for $x \in \partial D^n_\alpha$. Thus as a set $X^n = X^{n-1} \bigsqcup_\alpha e^n_\alpha$ where each e^n_α is an open n-disk.

(3) One can either stop this inductive process at a finite stage, setting $X = X^n$ for some $n < \infty$, or one can continue indefinitely, setting $X = \bigcup_n X^n$. In the former case X is then a n-dimensional complex and we let $\dim(X) = n$ denote its dimension, whereas in the latter case X is said to have infinite dimension. If $\dim(X) = \infty$, X is given the weak topology: A set $A \subset X$ is open (or closed) if and only if $A \cap X^n$ is open (or closed) in X^n for each n.

This definition is adapted from [Hat02], where also some further topological properties of cell complexes are discussed. The name *CW complex* stems from two properties: On one hand, cell complexes exhibit *closure finiteness*, i.e. the closure of one cell meets only finitely other cells. On the other hand, cell complexes have the *weak topology*, which means that a set is closed if and only if it meets the closure of each cell in a closed set. Furthermore cell complexes are normal and in particular Hausdorff. This shows that cell complexes are in a way well-behaved, although the definition might seem a bit wild at first glance. We do not rely on these properties and refer the reader for the proof of these to [Hat02].

However, for introducing Morse theory, we need to endow cell complexes with an affine structure. To this end, we give the definition of a *convex cell*, then the definition of *affine cell complexes*.

Note that we refer to *affine hyperplanes* simply as *hyperplanes* and to *affine half-spaces* as *half-spaces*.

Definition 3.2 (Convex cell)**.** A *convex (polyhedral) cell* $C \subset \mathbb{R}^m$ for some $m \in \mathbb{N} \setminus \{0\}$ is a compact intersection of finitely many closed half-spaces $H_1, \ldots, H_k$. The dimension $\dim(C)$ of C is the dimension of the affine subspace spanned by C.

Each half-space can be described by a linear inequality and thus C is the solution of these inequalities. If we change some of these inequalities to equalities, we call a solution a *face of* C. Trivially a face F of C also a convex cell. The intersection of two or more faces is again a face of C. If $F = \emptyset$ or $F = C$, F is called an *improper face of* C, else F is called a *proper face of* C. Unless specified otherwise, all faces we consider are proper. If F consists of only one point, we call F a *vertex of* C.

Similar to a convex polyhedral cell, we can define a *spherical cell* $C \subset S^{m-1}$. Instead of a intersection of half-spaces, we use hemispheres and for the definition of the dimension we use the smallest dimension of a sphere containing C.

Definition 3.3 (Affine cell complex)**.** A finite-dimensional cell complex X is said to be an *affine cell complex* if it is equipped with the following structure: Let m be an integer with $m \geq \dim(X)$. For each cell $e \in X$ we are given a convex polyhedral cell $C_e \subset \mathbb{R}^m$ and a characteristic function $\chi_e : C_e \to e$ such that the restriction of χ_e to any face of C_e is a characteristic function of another cell, possibly precomposed by a partial affine homeomorphism (i.e. a restriction of an affine homeomorphism of $\mathbb{R}^m$). This basically means that the intersection of two cells is either a face of both of these cells or empty. An *admissible characteristic function for a cell* $e \in X$ is any function obtained from χ_e by precomposition with a partial affine homeomorphism. This allows us to choose isometries as admissible characteristic functions and we can think of gluing cells along faces by isometries.

The terminology is quite technical and the following might be helpful to get a better understanding of affine complexes: Let X be an affine cell complex. Given an n-dimensional cell $e \in X$, the characteristic function $\chi_e : C_e \to e$ is the following composition of continuous functions:

$$C_e \to D^n_e \hookrightarrow X^{n-1} \bigsqcup_{\alpha} e^n_\alpha \to X^n \hookrightarrow X.$$

Definition 3.4. Similarly one can define a *spherical complex*, where one uses spherical cells instead of polyhedral cells, and spherical isometries precomposed by restrictions of homeomorphisms. For convenience, we will reserve the term *convex cell* to polyhedral cells, albeit spherical cells are also convex in the spherical metric.

Definition 3.5 (Presentation 2-complex)**.** To any presentation $\langle S \mid \mathcal{R}\rangle$ of a group G one can associate the *presentation 2-complex* $K = K(S \mid \mathcal{R})$. The complex K consists of one vertex v and one edge ε_a for each $a \in S$, oriented and labelled a. Let ε_a^{-1} denote the edge ε_a traversed backwards; we assume that $\varepsilon_a^{-1} = \varepsilon_{a^{-1}}$. Thus a word $w \in F(S)$, $w = a_1 \cdots a_n$ corresponds to the edge loop that is a concatenation of the edges labelled $a_1, \ldots, a_n$. Additionally, one has for each relation $r \in \mathcal{R}$, $r = a_1 \cdots a_n \in F(S)$ one 2-cell e_r, which is attached along the loop labelled $a_1 \cdots a_n$. The map that sends the homotopy class of ε_a to $a \in G$ gives an isomorphism $\pi_1(K(S \mid \mathcal{R}), v) \cong G$.

The universal cover of this complex K is a *Cayley complex* whose 1-skeleton is the Cayley graph of the group.

Now that we can construct a 2-complex from a given presentation of a group, we give an introduction to Morse theory on cell complexes.

Definition 3.6 (Morse function)**.** Let X be an affine cell complex. A function $f : X \to \mathbb{R}$ is called a *Morse function*, if for each cell $e \in X$ $f_{\chi_e} : \mathcal{C}_e \to \mathbb{R}$ extends to an affine map $\mathbb{R}^m \to \mathbb{R}$ and is constant if and only if $\dim e = 0$, and the image of the 0-skeleton under f is discrete in $\mathbb{R}$.

If $J \subset \mathbb{R}$ is a closed, nonempty subset, we define $X_J := f^{-1}(J)$ and for a single point $t \in \mathbb{R}$ we write $X_t := X_{\{t\}}$. We call X_t the *t-level set* of the Morse function f.

To build an intuition it is useful to imagine holding a cell by one vertex and letting it hang freely. Then the height function is a Morse function restricted to this cell. See Figure 8 for an affine cell complex X and a typical Morse function f given by the height. We see three level sets, X_{t_0}, X_{t_1} and X_{t_2}, indicated by the dashed lines.

Note that t_0 and t_2 are not in the image of the 0-skeleton of X under f. We see that there is a small neighbourhood $I = (t_0 - \varepsilon, t_0 + \varepsilon)$ of t_0 and $J = (t_2 - \varepsilon', t_2 + \varepsilon')$ such that X_I is homeomorphic to $X_{t_0} \times (-\varepsilon, \varepsilon)$ and X_J is homeomorphic to $X_{t_2} \times (-\varepsilon', \varepsilon')$. However when t_0 or t_2 approaches t_1, which is in the image of the 0-skeleton of X under f, a change in homotopy occurs, since the level set X_1 through v is contractible, while X_{t_0} and X_{t_2} are not. To obtain a more precise description of this change, we need the following definitions.

Definition 3.7 (Circle-valued Morse function)**.** A function $f : X \to \mathbb{S}^1$ on an affine complex is called a *circle-valued Morse function* if it is a *cellular map* (i.e. we endow $\mathbb{S}^1$ with a cell structure by designating points on the circle as vertices and the connecting segments as edges) and lifts to a Morse function between universal covers.

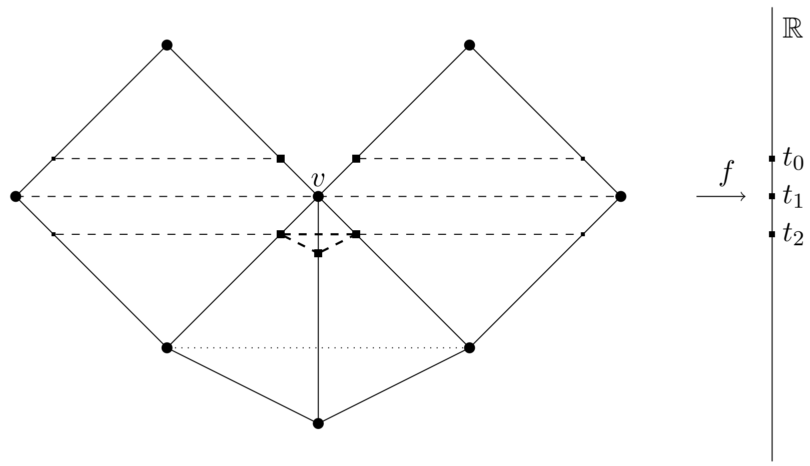

Figure 8: A cell complex X and a real valued Morse function f, given by the height of the cell. Three level sets X_{t_0}, X_{t_1} and X_{t_2} are indicated by dashed lines. Varying t_0 or t_2 by a small value does not change the homotopy type of the respective level set. A change in the homotopy type only occurs when there is a vertex in the level set, such as v.

Definition 3.8 (Links)**.** Let $C \subset \mathbb{R}^m$ be a convex n-cell. The link $\mathrm{Lk}(v, C)$ of a vertex is the collection of all unit tangent vectors at v, pointing into C. This is a spherical cell of dimension $(n-1)$. If C' is another cell, to which C is glued by an isometry f and v' is a vertex which gets identified under f with v, then the derivative of f gives us an isometry on the tangent spaces, which identifies a face of $\mathrm{Lk}(v, C)$ with a face of $\mathrm{Lk}(v', C')$. Thus in an affine complex X, the link of a vertex v is a spherical complex, which we denote with $\mathrm{Lk}(v, X)$.

This definition gives us a practical way to see how a vertex is attached or *linked* to a cell. On simplicial complexes the definition of a link is canonically, but on more general complexes this definition is not always practical. There are several ways to define the link of a vertex, however the distinguishing feature is that when the vertex is removed, one can obtain a homotopy equivalent space by coning off the link or its copy in the cell. With the given definition, we always have a natural homeomorphic copy of the link in the cell itself. We can see this copy by using the fact that the cell lives in some $\mathbb{R}^m$ and identify the tangent space with $\mathbb{R}^m$ itself, and scale down the link such that the vectors do not leave the cell.

Definition 3.9 (Ascending and descending links)**.** Suppose X is an affine cell complex and $f : X \to \mathbb{S}^1$ is a circle-valued Morse function. Choose an orientation of $\mathbb{S}^1$ which lifts to one of $\mathbb{R}$ and lift f to a map of universal covers $\tilde{f} : \tilde{X} \to \mathbb{R}$. Let $v \in X^{(0)}$ and note that the link of v in X is naturally isomorphic to the link of any lift $\tilde{v}$ in $\tilde{X}$. We say a cell $\tilde{e} \subset \tilde{X}$ contributes to the *ascending* (respectively *descending*) link of $\tilde{v} \in \tilde{e}$, if $\tilde{f}|_{\tilde{e}}$ achieves its minimum (respectively maximum) value at $\tilde{v}$. The *ascending link* $\mathrm{Lk}_{\uparrow}(v, X)$ (respectively *descending link* $\mathrm{Lk}_{\downarrow}(v, X)$) is the subset of $\mathrm{Lk}(v, X)$ naturally identified with the ascending (respectively descending) link of $\tilde{v}$.

Revisiting Figure 8, we see that the ascending link $\mathrm{Lk}_{\uparrow}(v, X)$ of the vertex v is a 0-sphere. Its copy in X_{t_0} are the two thicker square vertices above v. The level set through v is homotopy equivalent to X_{t_0} with the copy of the link coned off. Similarly, we see that $\mathrm{Lk}_{\downarrow}(v, X)$ is homeomorphic to a circle and its copy is indicated by the thick dashed lines in the figure. Again, coning off this copy of the descending link in X_{t_2} makes it homotopy equivalent to the level set X_{t_1}. We will formalize this process in the Morse Lemma, whose proof relies on the following observation.

Observation 3.10. Let $C \subset \mathbb{R}^m$ be a convex cell in Euclidean space, F, G disjoint (proper) faces with F top-dimensional. Then any strong deformation retraction from $\overline{\partial C \setminus F}$ to G extends to a strong deformation retraction from C to G. In particular $\overline{\partial C \setminus F}$ is a strong deformation retract of C.

Proof. We show that $\overline{\partial C \setminus F}$ is a strong deformation retract of C. Let $H_1, \ldots, H_{l+1}$ be the half-spaces corresponding to the top-dimensional faces of C, such that $C = \bigcap_{i=1}^{l+1} H_i$ and without loss of generality let H_{l+1} be the half-space corresponding to F. Then $C' := \bigcap_{i=1}^{l} H_i$ is also convex (albeit not necessarily a cell, but a cone) and because F is top-dimensional, $C' \setminus C$ is not empty.

Now we can construct a strong deformation retraction $f : C \times [0, 1] \to C$ of C onto $\overline{\partial C \setminus F}$. Choose a coordinate system of $\mathbb{R}^m$ such that the origin lies in $C' \setminus C$. Let $x \in C$ and e_x be the unit vector spanning the 1-dimensional subspace in which x lies. Then $e_x\mathbb{R} \cap \overline{\partial C \setminus F} = \{v_x\}$. We define f to be the following map

$$(x, t) \ \mapsto \ x + t\|v_x - x\|v_x,$$

where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^m$.

This strong deformation retraction can be thought of as the radial projection of C onto $\overline{\partial C \setminus F}$, where we chose the coordinate system of $\mathbb{R}^m$ such that the light source lies in the origin. □

We can now turn to the Morse Lemma. At first we wish to establish the fact that the homotopy type of a level set X_t does not change for small pertubations of t, if t is not in the image of the 0-skeleton of the cell complex X. Then we will prove the Morse Lemma, which describes how attaching cones to the copy of the ascending or descending link of a vertex in a level set conveys a homotopy equivalence to the level set containing the vertex. The proofs of the following lemma and the Morse Lemma are adapted from [BB97].

Lemma 3.11. Let X be a n-dimensional affine cell complex and $f : X \to \mathbb{R}$ a Morse function. If $J \subset J' \subset \mathbb{R}$ are non-empty, connected and $X_{J'} \setminus X_J$ contains no vertices of X, then $X_J \hookrightarrow X_{J'}$ is a homotopy equivalence.

Proof. For each cell $e \in X$ and each admissible characteristic function $\chi_e : C_e \to X$ we construct a strong deformation retraction $H_t^{\chi_e} = \chi_e^{-1} H \chi_e$ of $C_e \cap (f\chi_e)^{-1}(J')$ to $C_e \cap (f\chi_e)^{-1}(J)$ admitting to the following naturality properties:

(1) If X_e is precomposed by a partial affine homeomorphism h, then $H_t^{\chi_e}$ is conjugated by h, i.e. $H_t^{\chi_e h} = h^{-1} H_t^{\chi_e} h$

(2) The restriction of $H_t^{\chi_e}$ to a face of C_e is the strong deformation retraction associated to that face.

Note that $Y := X_{J'} \setminus X_J$ inherits an affine cell structure from X. Assume that $\sup J = \sup J'$ and let $t := \inf J$, $t' := \inf J'$.

Since Y contains no vertices of X, for each cell e of Y and admissible characteristic function $\chi_e : C_e \to X$ the sets $F_e := \chi_e^{-1}(X_t \cap e)$ and $G_e := \chi_e^{-1}(X_t' \cap e)$ are top-dimensional faces. This allows us for each cell $e \in Y^{(n)}$ to retract C_e onto $\overline{\partial C_e \setminus F_e}$, as shown in Observation 3.10. This is the $(n-1)$-skeleton $Y^{(n-1)}$ of Y and the retraction defines for each cell $e \in Y^{(n)}$ the strong deformation retraction $H_t^{\chi_e}|_{\mathring{C}_e}$ restricted to the interior $\mathring{C}_e := C \setminus \partial C$ of C_e, possibly conjugated with a partial affine homeomorphism h according to the first naturality property.

Repeating this process in inductively on the $(i-1)$-skeleton of Y, after having it performed on $Y^{(i)}$, gives us for each cell e and characteristic function χ_e a strong deformation retraction $H_t^{\chi_e}$ of $C_e \cap (f\chi_e)^{-1}(J')$ to $C_e \cap (f\chi_e)^{-1}(J)$. In each step we only retract the interior of a cell, leaving the boundary intact. This ensures that the strong deformation retraction of a cell is compatible with the strong deformation retraction of an adjacent cell, since the gluing happens by identifying faces, i.e. cells lying in the boundary. Thus the second naturality property is also satisfied.

For the general case observe that depending on whether $\sup J = \sup J'$ or $\inf J = \inf J'$, Y consists of either one or two disjoint cell complexes, one possibly lying above the level set of $\sup J$ (from perspective of the Morse function f) and one possibly lying below the level set of $\inf J$. Clearly we can use the same construction for cells lying above $\sup J$, with the intersection of a cell e with the $(\sup J)$-level set defining G_e and likewise the intersection of e with the $(\sup J')$-level set defining F_e.

These strong deformation retractions induce a strong deformation from $X_{J'}$ to X_J, therefore the inclusion $X_J \hookrightarrow X_{J'}$ is a homotopy equivalence. □

Proposition 3.12 (Morse Lemma)**.** If $f : X \to \mathbb{R}$ is a Morse function, $J \subset J' \subset \mathbb{R}$ are closed intervals with $\inf J = \inf J'$ and $J' \setminus J$ contains only one point r of the image of the 0-cells under f, then $f^{-1}(J')$ is homotopy equivalent to $f^{-1}(J)$ with the copies of $\mathrm{Lk}_{\downarrow}(v,X)$ (v a vertex with $f(v)=r$) coned off. A similar statement holds when $\inf J = \inf J'$ is replaced by $\sup J = \sup J'$ and $\mathrm{Lk}_{\downarrow}(v,X)$ by $\mathrm{Lk}_{\uparrow}(v,X)$.

Proof. Since $X_{J'\cap(-\infty,r]} \hookrightarrow X_{J'}$ is a homotopy equivalence by Lemma 3.11, assume that $\sup J' = r$ and $r - \varepsilon = \sup J$.

For any admissible characteristic function $\chi_e : C_e \to X$ of any other cell e we construct, again inductively on $\dim e$, a strong deformation retraction of $(f_{\chi_e})^{-1}((-\infty, r])$ onto the subset

$$(f_{\chi_e})^{-1}((-\infty, r-\varepsilon]) \cup \bigcup \{F | F \text{ is a face of } C_e \text{ with } f_{\chi_e}(F) \subset (-\infty, r]\},$$

satisfying both naturality properties from the proof of Lemma 3.11.

The construction is similar to the one of the preceding lemma, however, we have to consider what happens if a cell e has a vertex v such that $f(v) = r$. First note that if a cell $e \in X$ has $f|_e > r$ then e is disjoint from $X_{J'}$. So let $e \in Y := X_{J'} \setminus X_J$ be a cell with $f|_e \leq r$. If $f|_e < r$, then $v \notin e$ and we can, with the same notation as in Lemma 3.11, retract C_e onto G_e. However, this cannot be done if $v \in e$ with $f(v) = r$. To prove the assertion about the copies of the link coned off, note that G_e is homeomorphic to $Lk_{\downarrow}(v,e)$ and resembles the copy stated in the lemma, and thus e is homeomorphic to G_e coned off with v.

These strong deformation retractions induce a strong deformation retraction of $X_{J'}$ onto X_J with the cones attached as stated in the lemma. □

4 Morse criterion for free-by-cyclic groups

We now turn to the main theorem of this part of the thesis: A criterion given in [BRS07] for recognizing free-by-cyclic groups by examining their presentation complexes with a Morse function. Its proof models the presentation complex as a *graph of spaces*, which we will now define.

Definition 4.1 (Graph of spaces)**.** A *graph of spaces* consists of a finite graph G, a *vertex space* X_v associated to each vertex $v \in V(G)$, an *edge space* X_e associated to each $e \in E(G)$ and continuous maps $f_{\iota,e} : X_e \to X_{\iota(e)}$, $f_{\tau,e} : X_e \to X_{\tau(e)}$ for each edge e.

The *total space of the graph of spaces* is the quotient space of the disjoint union

$$\left(\bigsqcup_{v \in V(G)} X_v \right) \sqcup \left(\bigsqcup_{e \in E(G)} X_e \times [0,1] \right)$$

under the identifications $(x,0) \sim f_{\iota(e)}(x)$ and $(x,1) \sim f_{\tau(e)}(x)$ for all $x \in X_e$.

Theorem 4.2 (Recognizing free-by-cyclic groups)**.** Let $f : X \to \mathbb{S}^1$ be a circle-valued Morse function on the 2-complex X. If all ascending and descending links in X are trees, then $\pi_1(X)$ is free-by-cyclic.

Proof. We want to look at X as the total space of a graph of spaces. The circle with one vertex and one edge is a covering space of any cellular configuration of the circle, so without loss of generality we may assume that $\mathbb{S}^1$ has only one vertex and only one edge. Now we choose $\mathbb{S}^1$ consisting of one vertex v and one edge to be the underlying graph. For the vertex space we choose the preimage of the vertex $f^{-1}(v)$ and for the edge space the general point preimage $f^{-1}(\mathbb{S}^1 \setminus \{v\})$. The Morse Lemma (Proposition 3.12) states that the level set through the vertex is homotopy equivalent to the general point preimage with the ascending respective descending links coned off. These are trees and as such contractible, and since coning off a contractible space does not change its homotopy type, it follows that the basepoint preimage is homotopy equivalent to the general point preimage. Indeed, the maps from the edge space to the vertex space collapse the subgraphs corresponding to the ascending or descending link in the general point preimage, and collapsing trees is a homotopy equivalence. Therefore X is homotopy equivalent to a graph bundle over the circle and hence its fundamental group is free-by-cyclic. $\square$

4.1 Application of the Morse criterion to some examples

While the theorems and definitions elaborated in order to state and prove the Morse criterion seem to be rather daunting, its application to a given group presentation is not that technical after seeing some examples. It not only enables us to recognize free-by-cyclic groups, but its proof also yields a method to determine the automorphism in the $F \rtimes_\varphi \mathbb{Z}$ structure and in particular the rank of the free group F. We will conduct this procedure in the following examples given in [BRS07], where the first one was used to demonstrate the usage of Theorem 4.2, however no explicit description of the automorphisms where given and left as an exercise to the reader.

Example 4.3. $G = \langle a, b \mid abb = baa \rangle$. The corresponding presentation 2-complex consists of one vertex v, two 1-cells labelled a and b and a single 2-cell labelled with the relation.

Suppose we have a homomorphism $G \to \mathbb{Z}$, which maps a to A and b to B. Then the relation becomes

$$A + B + B = B + A + A \Leftrightarrow A = B.$$

Thus we can assume the homomorphism sends a and b to a generator of $\mathbb{Z}$. We can realize this homomorphism topologically by a map of the presentation 2-complex to the circle, with one base vertex and one edge (thus one 0-cell and one 1-cell). The map sends v to the basepoint of the circle and a as well as b once around the circle. We can extend this map linearly over the 2-cell, getting a circle valued Morse function, which lifts to a Morse function between universal coverings, where we take $\mathbb{R}$ to be the universal cover of the circle. The relator 2-cell lifts to the universal covering and becomes a hexagon from the perspective of the Morse function, with one vertex as its maximum point and the diametrically opposite vertex as its minimum. The two sides are labelled from bottom to top abb and baa (see Figure 9). Both the ascending and descending link are homeomorphic to segments connecting the 1-cell labelled a to the 1-cell labelled b, thus are trees. The Morse criterion for free-by-cyclic groups (Theorem 4.2) states that G is free-by-cyclic.

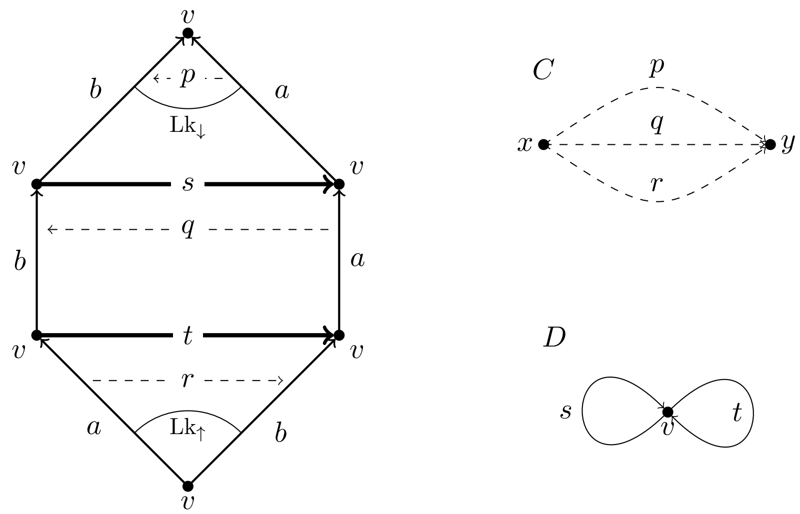

Figure 9: The relator cell from universal cover perspective for $G = \langle a, b \mid abb = baa \rangle$ in Example 4.3 on the left, with the generic point preimage C and the vertex preimage D on the right.

Looking at the basepoint preimage of the Morse function, we see the segments labelled s and t, which form a rose with two petals and base point v (D in the picture). The Morse Lemma tells us that the base point preimage D is homotopy equivalent to the generic point preimage C with the copies of the ascending respective descending link coned off. Let x denote the point where a segment of the general point preimage meets an edge labelled a and let y denote the point where a segment meets an edge labelled b. Coning off the links does not change the homotopy type, because the links are trees, so C and D are homotopy equivalent and we can conclude that $G = F_2 \rtimes_\varphi \mathbb{Z}$.

We wish to determine the automorphism φ which defines the conjugation in the semidirect product. To this end we take, under abuse of notation,

s and t to be the homotopy-equivalence classes of s and t as generators of F_2. The proof of Theorem 4.2 shows that G is the fundamental group of the graph of spaces with the circle being the underlying space, the base point preimage D being the vertex space at the basepoint and the general point preimage C being the edge space. Fixing a generator of the free group and moving around the circle gives us a homotopy equivalence, which induces a automorphism on the free group.

To obtain a precise description of this automorphism, we choose an orientation of the edges as shown in Figure 9. Looking at the general point preimage and approaching the base vertex in the circle collapses an edge in the preimage and sends other edges to s and t. We can read this homotopy equivalence off by sliding the edges in the relator cell up or down, where each direction corresponds to a direction from which we can approach the base point in the circle. We call these homotopy equivalences $H_\downarrow : C \to D$ and $H_\uparrow : C \to D$ accordingly. This results in the following homotopy equivalences, where we write e^{-1} to denote that the edge e is traversed in opposite direction:

$$H_\downarrow : p \mapsto s^{-1} \qquad H_\uparrow : p \mapsto v$$
$$q \mapsto t^{-1} \qquad q \mapsto s^{-1}$$
$$r \mapsto v \qquad r \mapsto t$$

Conversely, if we are at the base point of the circle and move in one direction into the edge space, the edges which have been collapsed by $H_\downarrow$ and $H_\uparrow$ are expanded accordingly, which gives us the following description of $H_\downarrow^{-1}$ and $H_\uparrow^{-1}$:

$$H_\downarrow^{-1} : s \mapsto rp^{-1} \qquad H_\uparrow^{-1} s \mapsto pq^{-1}$$
$$t \mapsto rq^{-1} \qquad t \mapsto rp^{-1}$$

We choose x to be the basepoint for the fundamental group of C. Then these homotopy equivalences induce isomorphisms $\pi_1(C, x) \to \pi_1(D, v)$ and $\pi_1(D, v) \to \pi_1(C, x)$ respectively. Precomposing one with the inverse of the other gives us the automorphism $\varphi : \pi_1(D, v) = F_2 \to \pi_1(D, v)$ in the semidirect product, which is topologically realised by moving a generator around the circle in one direction. Thus we arrive at the following description of φ:

$$\varphi = H_\downarrow \circ H_\uparrow^{-1} : s \mapsto s^{-1}t$$
$$t \mapsto s$$
$$\varphi^{-1} = H_\uparrow \circ H_\downarrow^{-1} : s \mapsto t$$
$$t \mapsto ts$$

Example 4.4. $G = \langle a, b \mid aba = bab \rangle$. Again the corresponding presentation 2-complex consists of one vertex v, two 1-cells labelled a and b and a single 2-cell attached along the 1-cells according to the relation.

Suppose we have a map $G \to \mathbb{Z}$ that maps a to A and b to B. Then the relation becomes

$$A + B + A = B + A + B \Leftrightarrow A = B.$$

So we can assume that the map takes a and b to a generator of $\mathbb{Z}$. We realize this map topologically by a map of the presentation 2-complex to the circle with one base vertex and one edge. As in Example 4.3, the map sends v to the basepoint and both a and b once around the circle. Extending this map linearly over the 2-cell, we obtain a circle valued Morse function, which lifts to a Morse function between universal coverings, where we take $\mathbb{R}$ to be the universal covering space of the circle. From the perspective of the Morse function, the relator 2-cell becomes a hexagon similar to that in the Example 4.3, with one side labelled aba and the other one labelled bab.

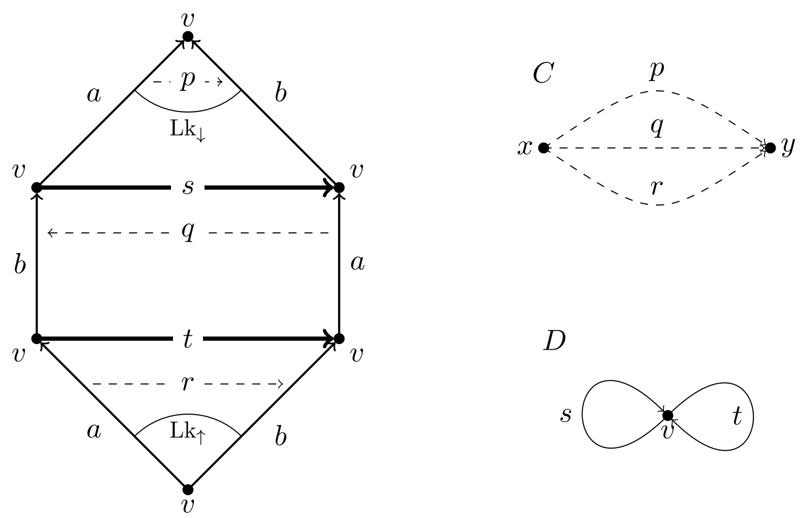

Figure 10: The relator cell from universal cover perspective for $G = \langle a, b \mid aba = bab \rangle$ in Example 4.4 on the left, with the generic point preimage C and the vertex preimage D on the right.

As we can see in Figure 10, both the ascending and descending links are homeomorphic to single segments connecting the 1-cells labelled a and b and thus are trees and G is free-by-cyclic by Theorem 4.2.

The base point preimage D of the Morse function is a rose with two petals, labelled s and t, which we take as generators of the free group. Again the Morse Lemma tells us that D is homotopy equivalent to C with the copies of the ascending respective descending link coned off — these are trees, hence C and D are homotopy equivalent and it follows that $G = F_2 \rtimes_\varphi \mathbb{Z}$. Moving the generic point preimage C towards the base point of the circle again collapses one edge, which can be read off the 2-cell as seen in Example 4.3. Again we let x denote the point where a segment in the generic point preimage meets an edge labelled a and y denote the point where an edge labelled b is met. With the orientations and notation as chosen in Figure 10, we arrive at the following description of the homotopy equivalences $H_\downarrow : C \to D$ and $H_\uparrow : C \to D$:

$$\begin{aligned} H_\downarrow : & p \mapsto s \\ & q \mapsto t^{-1} \\ & r \mapsto v \end{aligned} \qquad\qquad \begin{aligned} H_\uparrow : & p \mapsto v \\ & q \mapsto s^{-1} \\ & r \mapsto t \end{aligned}$$

The inverses of these homotopy equivalences are again given by expanding the edges that got collapsed:

$$\begin{aligned} H_\downarrow^{-1} : & s \mapsto pr^{-1} \\ & t \mapsto rq^{-1} \end{aligned} \qquad\qquad \begin{aligned} H_\uparrow^{-1} : & s \mapsto pq^{-1} \\ & t \mapsto rp^{-1} \end{aligned}$$

Precomposing one homotopy equivalence with the inverse of the other, this gives us the following description of the automorphism φ in the free-by-cyclic structure:

$$\begin{aligned} \varphi = H_\downarrow \circ H_\uparrow^{-1} : & s \mapsto st \\ & t \mapsto s^{-1} \\ \varphi^{-1} = H_\uparrow \circ H_\downarrow^{-1} : & s \mapsto t^{-1} \\ & t \mapsto ts \end{aligned}$$

The following group is also discussed in [BC07], where it serves as an example of a hyperbolic group with further interesting geometric properties, although the determination of the automorphism is not carried out in full detail.

Example 4.5. $G = \langle a, b \mid abaa = bb \rangle$ The 1-skeleton of the presentation 2-complex is again a rose with two petals, labelled a and b and vertex v, and one 2-cell attached according to the relation along a and b. Assuming we have a homomorphism $G \to \mathbb{Z}$ which takes a to A and b to B, the relation becomes

$$A + B + A + A = B + B \Leftrightarrow 3A = B.$$

So we can suppose that a maps to a generator z of $\mathbb{Z}$ and b to $3z$. We realize this map topologically by sending v to the base vertex of a circle $\mathbb{S}^1$ with one base vertex and one edge, the edge labelled a one time around the circle and the edge labelled b three times around the circle. Extending this map linearly over the circle, we get a circle valued Morse function which lifts to a Morse function between universal covers.

From the perspective of the Morse function the relator 2-cell is again a hexagon, however the edges labelled a having height one and the edges labelled b having height three. As depicted in Figure 11, the ascending and descending links are again trees, the ascending link is homeomorphic to the segment connecting a to the lower third of b, the descending link is homeomorphic to the segment connecting a to the upper third of b. So Theorem 4.2 tells us that G is free-by-cyclic.

Looking at the base point preimage D, we see a graph with three vertices, where one is v and two other located at $\frac{1}{3}$ and $\frac{2}{3}$ of the height of the segment b (labelled accordingly $\frac{1}{3}b$ and $\frac{2}{3}b$), and five edges connecting these, as depicted in Figure 12. Choosing a maximal tree in D and taking the remaining edges as generators of the free group, this allows us to determine the rank of the free group. The graph has three vertices, thus a maximal tree has at most two edges and we choose the subgraph consisting of the edges t_1 and t_2 as maximal tree T. The quotient D/T has three edges, so we can conclude that $G = F_3 \rtimes_\varphi \mathbb{Z}$.

For determining the automorphism φ we choose the homotopy classes of $p = t_3 t_1^{-1}$, $q = t_1 t_4^{-1} t_2^{-1}$ and $r = t_5 t_2^{-1}$ to be generators of $\pi_1(D, v)$. We know by the Morse Lemma that the generic point preimage C with the ascending respective descending link coned off is homotopy equivalent to D, and since these are trees, C and D are homotopy equivalent. As we can see from the proof of Theorem 4.2, the presentation 2-complex of G is homotopy equivalent to a graph of spaces with underlying space the circle with one base vertex, vertex space D and edge space C. Approaching the vertex from one direction collapses one edge and sends the others to edges of D, which can be read off by sliding the dashed lines in the relator cell up and down. The generic point preimage consists of four vertices, labelled u for the point on the edge labelled a, and x, y and z for the points on the edge labelled b, where x is the point in the lower third of the segment, y the point in the middle third of the segment and z the point in the upper third. With the orientations chosen as depicted, we can read off the homotopy equivalences $H_\downarrow : C \to D$ and $H_\uparrow : C \to D$, arriving at the following description:

$$\begin{aligned} H_\downarrow : l_1 &\mapsto t_1 \qquad & H_\uparrow : l_1 &\mapsto v \\ l_2 &\mapsto t_2 & l_2 &\mapsto t_1 \\ l_3 &\mapsto t_3 & l_3 &\mapsto t_2^{-1} \\ l_4 &\mapsto t_4 & l_4 &\mapsto t_3^{-1} \\ l_5 &\mapsto t_5 & l_5 &\mapsto t_4 \\ l_6 &\mapsto v & l_6 &\mapsto t_5 \end{aligned}$$

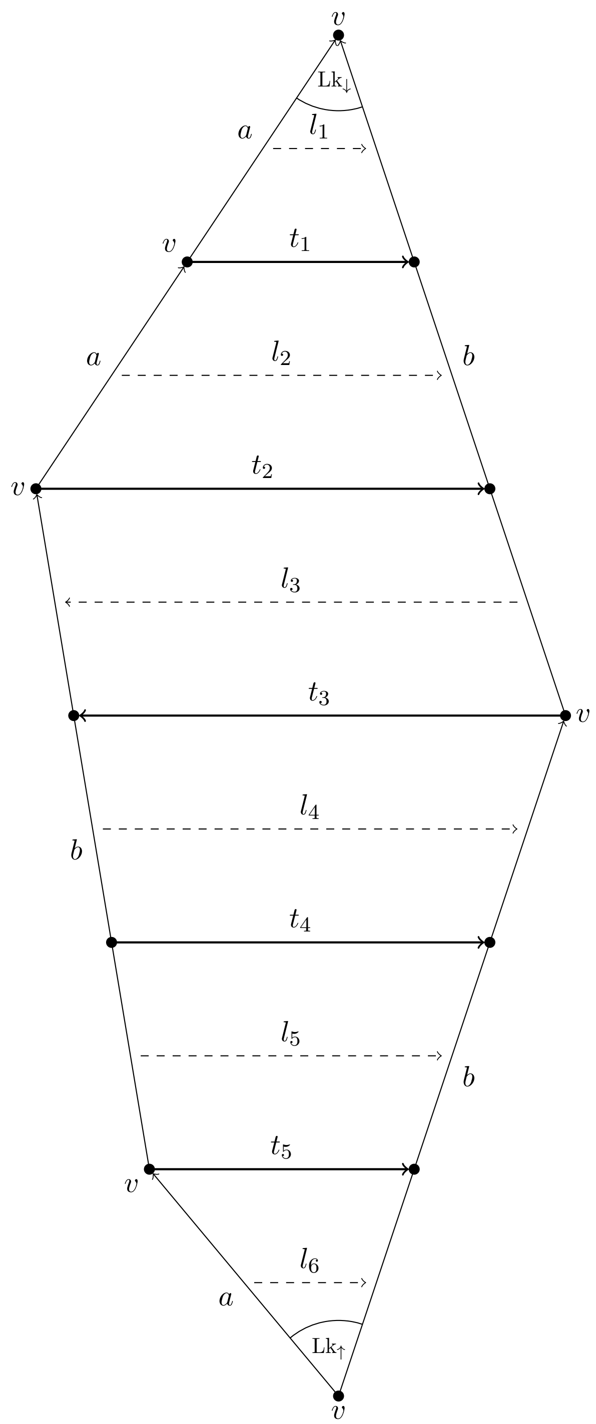

Figure 11: The relator cell of $G = \langle a, b \mid abaa = bb \rangle$ from universal cover perspective in Example 4.5.

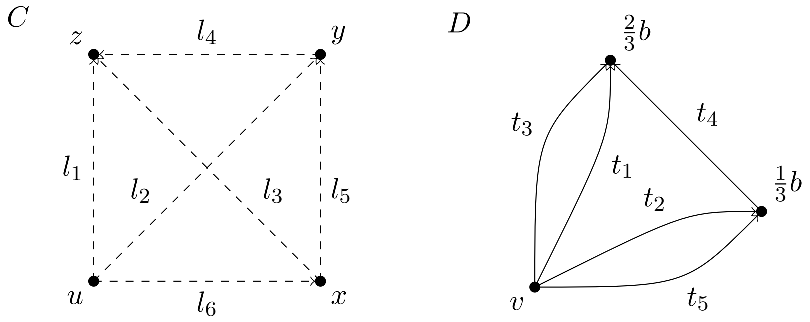

Figure 12: The generic point preimage C and the base vertex preimage D of the Morse function in Example 4.5

The inverse is given by expanding the edges that got collapsed. We are only interested in the changes on the chosen generators of our free group, so we directly look at these.

$$\begin{aligned} H_\downarrow^{-1} : p = t_3 t_1^{-1} &\mapsto l_6 l_3 l_1^{-1} \\ q = t_1 t_4^{-1} t_2^{-1} &\mapsto l_1 l_4^{-1} l_2^{-1} \\ r = t_5 t_2^{-1} &\mapsto l_6 l_5 l_2^{-1} \end{aligned} \qquad \begin{aligned} H_\uparrow^{-1} : p &\mapsto l_1 l_4^{-1} l_2^{-1} \\ q &\mapsto l_2 l_5^{-1} l_3 l_1^{-1} \\ r &\mapsto l_6 l_3 l_1^{-1} \end{aligned}$$

Finally we arrive at a description of the automorphism $\varphi : \pi_1(D, v) \to \pi_1(D, v)$ by precomposing one homotopy equivalence with the inverse of the other.

$$\begin{aligned} \varphi = H_\downarrow \circ H_\uparrow^{-1} : p &\mapsto t_1 t_4^{-1} t_2^{-1} = q \\ q &\mapsto t_2 t_5^{-1} t_3 t_1^{-1} = r^{-1} p \\ r &\mapsto t_3 t_1^{-1} = p \\ \varphi^{-1} = H_\uparrow \circ H_\downarrow^{-1} : p &\mapsto t_5 t_2^{-1} = r \\ q &\mapsto t_3 t_1^{-1} = p \\ r &\mapsto t_5 t_4 t_1^{-1} \simeq t_5 t_2^{-1} t_2 t_4 t_1^{-1} = r q^{-1} \end{aligned}$$

Example 4.6. $G = \langle a_0, \ldots, a_n \mid a_i^{a_{i+1}} = a_0 \ (1 \leq i \leq n-1), a_0^{a_1} = a_n \rangle$. The presentation 2-complex X consists of one vertex v, one 1-cell for each generator a_i labelled accordingly, one 2-cell corresponding to the relation $a_0^{a_1} = a_n$ and for each $1 \leq i \leq n-1$ a 2-cell corresponding to the relation $a_i^{a_{i+1}} = a_0$. Assuming a map $G \to \mathbb{Z}$ which takes a_i to A_i, the relations become

$$A_i = A_0 \qquad (1 \le i \le n-1)$$
$$A_0 = A_n$$

and we can assume that the a_i are taken to a generator of $\mathbb{Z}$.

Therefore each relator 2-cells becomes a quadrilateral from the perspective of the Morse function in the universal covering, with one vertex as the maximum and one as the minimum, and the quadrilateral corresponding to $a_0^{a_1} = a_n$ labelled with a_0a_1 from bottom to top on the one side and a_1a_n on the other. The cells corresponding to $a_i^{a_{i+1}} = a_0$ look similar, with the difference being that one side is labelled a_ia_{i+1} and the other $a_{i+1}a_0$.

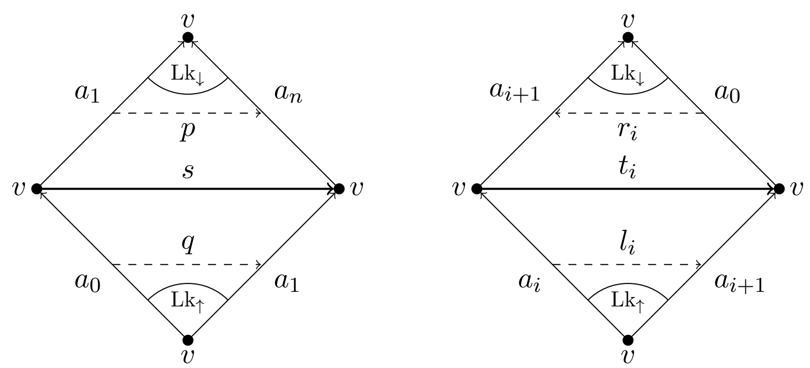

Figure 13: The relator cells of G in Example 4.6, the left one corresponding to the relation $a_0^{a_1} = a_n$, the right one corresponding to $a_i^{a_{i+1}} = a_0$

The ascending and descending links $\mathrm{Lk}_\uparrow(v, X)$ and $\mathrm{Lk}_\downarrow(v, X)$ are indicated in Figure 13. The ascending link consists of one segment connecting the edge a_0 to the edge a_1 and segments connecting a_i to a_{i+1} for each $1 \le i \le n-1$. Hence the ascending link is a tree. The descending link consists of one segment connecting a_1 to a_n and segments connecting a_{i+1} to a_0 for each $1 \le i \le n-1$. This is also a tree and we can conclude that G is free-by-cyclic. The base point preimage of the Morse function D is a rose with n petals, one labelled s as indicated in the figure, and the other ones labelled $t_1, \dots, t_{n-1}$. By the Morse Lemma, D is homotopy equivalent to C with the copies of the ascending respective descending links coned off. These are trees and we conclude that C and D are homotopy equivalent.

Again under abuse of notation, we take the homotopy classes of these edges to be the generators of the free group, hence $G = F_n \rtimes_\varphi \mathbb{Z}$ with $F_n = \langle s, t_1, \dots, t_{n-1} \rangle$.

We wish to determine automorphism φ. To this end we let x_i be the point lying on the edge labelled a_i in the generic preimage D, which is depicted in Figure 14. As in the previous examples, we slide the dashed lines corresponding to edges in the generic point preimage in the relator cells up and down, corresponding to the homotopy equivalences $H_\uparrow : C \to D$ and $H_\downarrow : C \to D$.

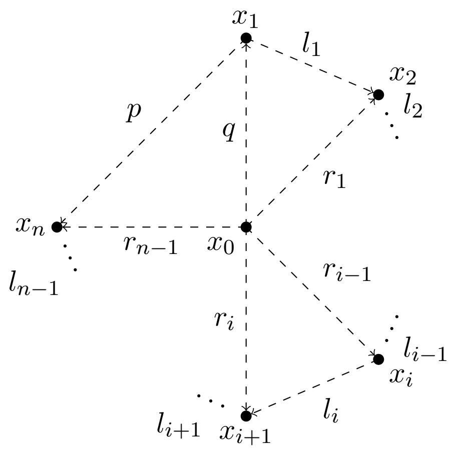

Figure 14: The generic point preimage C in Example 4.6.

With the orientations chosen in Figure 14, we obtain the following description of $H_\downarrow$ and $H_\uparrow$:

$$
\begin{aligned}
H_\downarrow :& p \mapsto s & \qquad H_\uparrow :& p \mapsto v \\
& q \mapsto v & & q \mapsto s \\
& r_i \mapsto t_i^{-1} & & r_i \mapsto v \\
& l_i \mapsto v & & l_i \mapsto t_i
\end{aligned}
$$

So $H_\downarrow$ collapses the edges q and l_i and $H_\uparrow$ collapses the edges p and r_i for each $1 \leq i \leq n-1$. The inverses are given by expanding these edges.

$$
\begin{aligned}
H_\downarrow^{-1} :& s \mapsto qpl_{n-1}^{-1} \cdots l_1^{-1} q^{-1} & \qquad H_\uparrow^{-1} :& s \mapsto qpr_{n-1}^{-1} \\
& t_i \mapsto ql_1 \cdots l_i r_i^{-1} & & t_1 \mapsto r_{n-1} p^{-1} l_1 r_1^{-1} \\
& & & t_i \mapsto r_{i-1} l_i r_i^{-1} \ (2 \leq i \leq n-1)
\end{aligned}
$$

Now we can obtain the automorphism φ by precomposing one homotopy equivalence with the inverse of the other.

$$
\begin{aligned}
\varphi = H_\downarrow \circ H_\uparrow^{-1} :& s \mapsto st_{n-1} \\
& t_1 \mapsto t_{n-1}^{-1} s^{-1} t_1 \\
& t_i \mapsto t_{i-1}^{-1} t_i & \qquad (2 \leq i \leq n-1) \\
\varphi^{-1} = H_\uparrow \circ H_\downarrow^{-1} :& s \mapsto st_{n-1}^{-1} \cdots t_1^{-1} s^{-1} \\
& t_i \mapsto st_1 \cdots t_i
\end{aligned}
$$

This concludes the examples and the treatment of the Morse criterion for free-by-cyclic groups. We have seen that it is not only possible to decide whether or not a group given by a presentation is free-by-cyclic or not, but also that we can read off the automorphism in the semidirect product. In light of the first section the question arises whether these groups have solvable Word or Transformation Problem. In fact, if the free group is of finite rank, it has been shown in [BMMV06] that the Conjugacy Problem is solvable and hence the Word Problem, as discussed in Section 2.

Acknowledgements

I want to express my gratitude towards Prof. Dr. Linus Kramer for his supervision and for providing me with the opportunity to write on the subject of hyperbolic groups, as well as his suggestion to have a closer look at the Morse criterion. Also I want to thank Daniel Skodlerack and Olga Varghese for answering my questions since my first exposure to hyperbolic geometry and Stefan Witzel, for answering my questions regarding topology and group theory. Finally, I want to thank my grandmother for her continued moral support throughout my studies.

References

[BB97] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[BC07] Noel Brady and John Crisp. CAT(0) and CAT(−1) dimensions of torsion free hyperbolic groups. *Comment. Math. Helv.*, 82(1):61–85, 2007.

[BH09] M.R. Bridson and A. Häfliger. *Metric Spaces of Non-Positive Curvature*. Grundlehren Der Mathematischen Wissenschaften. Springer, 2009.

[BMMV06] O. Bogopolski, A. Martino, O. Maslakova, and E. Ventura. The conjugacy problem is solvable in free-by-cyclic groups. *Bull. London Math. Soc.*, 38(5):787–794, 2006.

[BRS07] Noel Brady, Tim Riley, and Hamish Short. *The geometry of the word problem for finitely generated groups.* Advanced Courses in Mathematics. CRM Barcelona. Birkhäuser Verlag, Basel, 2007. Papers from the Advanced Course held in Barcelona, July 5–15, 2005.

[Deh11] M. Dehn. Über unendliche diskontinuierliche Gruppen. *Math. Ann.*, 71(1):116–144, 1911.

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[Lö11] Clara Löh. Geometric group theory, an introduction. `http://www.mathematik.uni-regensburg.de/loeh/teaching/ggt_ws1011/lecture_notes.pdf`, 2011.

Plagiatserklärung

Hiermit versichere ich, dass die vorliegende Arbeit über THE WORD PROBLEM, HYPERBOLIC GROUPS AND THE MORSE CRITERION selbstständig verfasst worden ist, dass keine anderen Quellen und Hilfsmittel als die angegebenen benutzt worden sind und dass die Stellen der Arbeit, die anderen Werken — auch elektronischen Medien — dem Wortlaut oder Sinn nach entnommen wurden, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht worden sind.

(Ort, Datum, Unterschrift)

Ich erkläre mich mit einem Abgleich der Arbeit mit anderen Texten zur Auffindung von Übereinstimmungen, sowie mit einer zu diesem Zweck vorzunehmenden Speicherung der Arbeit in eine Datenbank einverstanden.

(Ort, Datum, Unterschrift)