On CAT(0) Aspects of Geometric Group Theory and Some Applications to Geometric Superrigidity

Maxime Octave Victor Bergeron

Department of Mathematics and Statistics
McGill University, Montreal

June, 2012

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science.

Abstract

Since their popularization by Gromov in the eighties, $CAT(0)$ metric spaces of bounded curvature as defined by Alexandrov have been the locus of great progress in infinite group theory. Surveying ideas and constructions of geometric group theory, we express a bias towards groups acting on structures of this kind. As such, swiftly acquainting the reader with the theory of $CAT(0)$ spaces, we provide a variety of examples obtained by gluing together families of convex polyhedra along their isometric faces. In this context, Gromov's link condition provides a local-to-global framework for non-positive curvature. Combining this with tools from knot theory, such as the Dehn complex of an alternating knot projection, we demonstrate a result of Wise which states that the fundamental group of an alternating link complement is also the fundamental group of a non-positively curved complex. Using similar ideas, we also mention a construction of Wise relating any finitely generated group to the fundamental groups of some non-positively curved complexes. Besides providing such "explicit" constructions, we make use of tower lifts of combinatorial maps to prove Bridson and Haefliger's abstract result that every subgroup of the fundamental group of a non-positively curved two dimensional polyhedral complexes is the fundamental group of some compact non-positively curved two dimensional polyhedral complex. Then, having well established the inherent structure of $CAT(0)$ spaces, we focus on classifying their isometries, group actions upon them, and how they extend to the visual boundary. The combinatorial approach is especially effective here when we prove Haglund's result that cell-preserving isometries of $CAT(0)$ cube complexes are semi-simple. Finally, using the theory of generalized harmonic maps, we demonstrate the superrigidity result of Monod, Gelander, Karlsson and Margulis for reduced actions with no globally fixed point of irreducible uniform lattices in locally compact, compactly generated topological groups of higher rank on complete $CAT(0)$ spaces.

Résumé

Depuis leur popularisation par Gromov durant les années quatre-vingt, la théorie des espaces métriques à courbure bornée, dits $CAT(0)$, fut à la base de grandes percées dans notre compréhension des groupes infinis. Survolant des constructions de la théorie géométrique des groupes, nous portons donc une attention particulière aux actions sur les espaces $CAT(0)$ et commençons notre traité par la construction de complexes $CAT(0)$ obtenus en identifiant certaines faces isométriques d'ensembles de polyèdres convexes. Dans ce contexte, le critère du lien de Gromov nous permet de caractériser la courbure nonpositive globale de manière locale. Combinant ces idées à certaines techniques de la théorie des noeuds, nous démontrons un théorème de Wise reliant tout groupe fondamental du complément d'un entrelac alternants à un complexe de courbure nonpositive. Nous relatons aussi une construction similaire de Wise permettant de relier tout groupe présenté de manière finie au groupe fondamental d'un complexe à courbure nonpositive. Outre ces constructions concrètes, nous utilisons les tours de relèvement d'applications combinatoires afin de démontrer un théorème abstrait de Bridson et Haefliger concernant les sous-groupes de groupes fondamentaux de complexes à courbure non-positive. Ayant établi la structure des espaces $CAT(0)$, nous passons en second lieu à la classification de leurs isométries et de leurs extensions à la bordification de ces espaces. L'approche combinatoire est d'une aide particulière lorsque nous prouvons le résultat de Haglund concernant la semi-simplicité d'isométries de complexes cubiques et offre un contraste par rapport à un résultat analogue de Brisdon dans le contexte des complexes polyhédraux. Finalement, en faisant usage de la théorie des applications harmoniques généralisées, nous démontrons le résultat de superrigidité de Monod, Gelander, Karlsson et Margulis pour les actions réduites sans point fixe sur les espaces métriques $CAT(0)$ complets de réseaux uniformes et irréductibles dans des groupes de rang supérieur localement compacts engendrés par un ensemble de générateurs compact.

Acknowledgments

I would first and foremost like to thank Mikaël Pichot for providing the best guidance and feedback I could have wished for, through our countless friendly conversations.

It is with gratitude that I acknowledge Christophe Hohlweg, Olga Kharlampovich and Ada Chan for providing me with the crucial undergraduate research experiences that helped develop my mathematical abilities.

I am indebted to Dani Wise for teaching some truly insightful classes and pointing out useful references to me.

I would like to thank Marie-Hélène Goulet for helping me figure out the inner workings of Illustrator to draw the figures and Babycat for providing the external distractions necessary to make progress.

Finally, I would like thank my family who brought me this far in life and always encouraged me in my various endeavours.

Contents

1 Prologue

1.1 Outline

Metric spaces of non-positive curvature in the sense of Alexandrov [Ale51] emulate Riemannian manifolds of bounded sectional curvature. In order to generalize their elegant features to the realm of geodesic metric spaces, he introduced the notion of *upper angle*. This allows one to consider the angular excess $\delta(\triangle)$ of triangles $\triangle$ in a geodesic metric space, namely the sum of its interior angles minus the expected sum which is π, and say that such a space is non-positively curved if $\delta(\triangle) \leq 0$ for all $\triangle$. This concept is analogous in the differentiable case to the *total curvature* of a surface.

The purpose of the present survey is two-fold. Our first goal is to introduce the reader to the theory of metric spaces of non-positive curvature. As such, having laid down the basic definitions, we provide a variety of interesting examples arising mostly from gluing constructions and knot theory. We then turn our attention to the theory of groups acting on metric spaces. There, we emphasize the strong relation between the structure of non-positively curved metric spaces and the structure of groups that act on them. The techniques used throughout our exposition involve a constant interplay between analytic and combinatorial interpretations of the restrictions imposed by the curvature.

Following a similar approach that of angular excess, we introduce a general inequality to define $CAT(0)$ metric spaces in Section 2. We then use Gromov's link condition to develop some metric and combinatorial aspects of the theory of polyhedral and cube complexes of bounded curvature in Sections 3.1 and 3.2. In particular, we explore closure properties for subgroups of the fundamental group of such complexes and prove the following theorem in Section 3.3 by developing techniques of tower lifts of combinatorial maps as introduced by [Pap57] and [How81].

Theorem 1.1 ([BH99]). *Every subgroup of the fundamental group of a non-positively curved two dimensional polyhedral cell complexes is isomorphic to the* *fundamental group of some compact non-positively curved two dimensional cell complex.*

Moving away from these relatively abstract constructions, we produce a zoo of non-positively curved complexes focusing on constructions of Wise involving knots and exact sequences. In particular, it is shown in Section 4.2 that the Dehn complex of a knot projection Π is non-positively curved if and only if Π is prime and alternating, which leads up to the following theorem:

Theorem 1.2 ([Wis06]). *The fundamental group of an alternating link complement is isomorphic to the fundamental group of a compact non-positively curved cell complex.*

Finally, in section 4.3 we illustrate a modification of Rips' celebrated construction in [Rip82] by Wise which relates finitely generated groups to fundamental groups of non-positively curved complexes.

Theorem 1.3 ([Wis98]). *Given any finitely generated group G there is a short exact sequence $1 \to N \to G \to \pi_1 K \to 1$ where N is finitely generated and K is a non-positively curved cell complex obtained by gluing hyperbolic pentagons.*

Having laid down the necessary theory for metric spaces, we shift our point of view and survey some basic ideas and constructions of geometric group theory, embracing the slogan that good actions of infinite groups on metric spaces allow great insight into the group structure. Swiftly moving through standard results for groups acting on metric spaces, we prove a result of Macbeath from [Mac64] for presenting groups of homeomorphisms to conclude that a group is finitely presented if and only if it acts properly and cocompactly by isometries on a simply connected geodesic metric spaces. Since many techniques in the geometric theory of infinite groups rely on the introduction of good coarse geometric invariants, we also recall the notion of quasi-isometries and illustrate them through the Schwarz-Milnor Lemma in Section 5.8.

We then turn our attention to understanding isometries of $CAT(0)$ spaces and how they extend to the *bordification* of such spaces in Section 6.2. After defining *semi-simple* isometries analogously to the case of Riemannian manifolds, we study how they arise in $CAT(0)$ cell complexes. In section 6.3 we focus on cube complexes, giving an account of the following unpublished result of Haglund:

Theorem 1.4 ([Hag07]). *Combinatorial isometries of possibly infinite dimensional $CAT(0)$ cube complexes are semi-simple up to cubical subdivision.*

This inherently generalizes some aspects of the analogous result of Bridson for polyhedral cell complexes with finitely many types of faces.

Theorem 1.5 ([Bri99]). *Combinatorial isometries of polyhedral cell complexes of bounded curvature having only finitely many isometry types of faces are semi-simple.*

Finally, we introduce the main ideas behind geometric superrigidity and generalized harmonic maps in view of proving the recent theorems of Monod, Gelander, Karlsson and Margulis. In Section 7, it is shown in particular that :

Theorem 1.6 ([Mon06],[GKM08]). *Reduced actions with no globally fixed point of irreducible uniform lattices in locally compact, compactly generated topological groups of higher rank on complete $CAT(0)$ spaces extend continuously to the whole group.*

We wish to emphasize that while we do not prove any new results per se, the originality of the present thesis lies in the organization and presentation of the theory.

1.2 Comparison Geometry

To properly understand the geometry of geodesic metric spaces, it is convenient to introduce a notion of *angle* between geodesics issuing from a given point. Our ultimate goal will be to fruitfully compare geodesic triangles in metric spaces to similar ones in model spaces[1] M_k^2 that can be thought of as template spaces of

[1]Please consult Appendix A where standard notation and definitions are set up.

fixed curvature k. The definition of angles given by the inner product in $\mathbb{H}^n, \mathbb{E}^n$ and $\mathbb{S}^n$ can implicitly be generalized to arbitrary model spaces via the appropriate inner product. The Alexandrov angle will yield the desired notion of angle for arbitrary metric spaces.

In this light, we first define a *comparison triangle* in $\mathbb{E}^2$ for a geodesic triangle $\triangle(p,q,r)$ in a metric space X as a triangle in $\mathbb{E}^2$ with vertices $\bar{p}, \bar{q}, \bar{r}$ such that $d(p,q) = d(\bar{p},\bar{q})$, $d(p,r) = d(\bar{p},\bar{r})$ and $d(q,r) = d(\bar{q},\bar{r})$. This triangle exists, is unique up to isometry and denoted by $\bar{\triangle}(p,q,r)$ or $\bar{\triangle}$ when no confusion can arise. The interior angle of the comparison triangle at $\bar{p}$ is called the the *comparison angle* between r and q at p. We denote it by $\bar{\angle}_p(q,r)$.

Given two geodesic paths[2] $g, g' : [0,a] \to X$ issuing from the same point, consider the euclidean comparison triangle $\bar{\triangle}(g(0), g(t), g(t'))$ in $\mathbb{E}^2$ and the comparison angle $\bar{\angle}_{g(0)}(g(t), g'(t'))$.

Alexandrov Angle The *Alexandrov angle* between a pair of chosen geodesic paths g and g' issuing from a common point is the unique number $\angle(g,g')$ given by the lim sup of the comparison angles as follows:

$$0 \leq \angle(g,g') := \limsup_{t,t' \to 0} \bar{\angle}_{g(0)}(g(t), g'(t')) = \lim_{\epsilon \to 0} \sup_{0 < t,t' < \epsilon} \bar{\angle}_{g(0)}(g(t), g'(t')) \leq \pi.$$

When X is uniquely geodesic, for $p \neq x$ and $p \neq y$, the angle between the geodesic segments $[p,x]$ and $[p,y]$ is well defined and denoted $\angle_p(x,y)$. One should also note that the reason for using lim sup in Alexandrov angles is to ensure that the angular distance between geodesics issuing from a point is always a pseudometric. This is used in particular to metrize the *space of directions* at a point as mentioned in Section 2.3.

Given a triangle $\triangle$ in an arbitrary metric space X, we define *k-comparison triangles* $\bar{\triangle}^{(k)}$ and *k-comparison angles*, $\bar{\angle}^{(k)}$ in M_k^2 analogously to the $\mathbb{E}^2$ case. It is a fact that they always exist when the perimeter of $\triangle$ is less than twice the diameter[3] $2D_k$. Interestingly, the Alexandrov angle coincides with the euclidean,

[2] Always parametrized proportional to arc length.

[3] The diameter of a model space M_k^2 is defined as $D_k := \pi/\sqrt{k}$ when $k > 0$ and ∞ other-

hyperbolic and spherical angles of the three standard model spaces so we may replace $\bar{\angle}$ by $\bar{\angle}^{(k)}$ in its definition without changing its value.

1.3 Alexandrov's Lemma

The study of metric spaces of bounded curvature in the sense of Alexandrov relies heavily on constructions involving triangles. Setting the mood for what is to come, the following two technical lemma could be said to form one of the bases of this theory and will be used in Section 2.2 to characterize flat subspaces. Although the original ideas are due to Alexandrov, we follow a proof of [BH99].

Lemma 1.7 (Alexandrov's Lemma Part 1)**.** *Consider a pair of distinct points* $p, p' \in M_k^2$ *where* $k \leq 0$*. Two piecewise geodesic paths* $[p,q] \cup [q,p']$ *and* $[p,r] \cup [r,p']$ *where* $q \neq r$ *while* p *and* p' *lie on opposite sides of the geodesic line extending* $[q,r]$ *always determine a pair of geodesic triangles. We henceforth denote these triangles by* $\triangle$ *and* $\triangle'$ *corresponding to the triples of points* (p,q,r) *and* (p',q,r)*. Labelling the angles at vertices* r *of* $\triangle$ *and* $\triangle'$ *by* γ *and* γ'*, if* $\gamma + \gamma' \geq \pi$ *then*

$$d(p,r) + d(p',r) \leq d(p,q) + d(p',q) \tag{1}$$

or the path $[p,r] \cup [r,p']$ *is shortest among the two.*

Proof. Recall by Proposition A.1 that M_k^2 is a uniquely geodesic space for $k \leq 0$. It follows that there is a unique point $\tilde{p}'$ with $d(\tilde{p}', r) = d(p', r)$ such that r lies on the geodesic $[p, \tilde{p}']$. Since $\gamma + \gamma' \geq \pi$, the new angle formed also satisfies the inequality $\angle_r(q, \tilde{p}') \leq \angle_r(p', q)$. We may now apply the law of cosines in M_k^2 to deduce that $d(q, \tilde{p}') \leq d(q, p')$. Now $d(p,r) + d(r, \tilde{p}') = d(p, \tilde{p}') \leq d(p,q) + d(q, \tilde{p}') \leq d(p,q) + d(p',q)$ which concludes the proof of (1). □

Lemma 1.8 (Alexandrov's Lemma Part 2)**.** *Keeping the notation from the previous lemma as indicated in Figure 1, let* $\bar{\triangle}$ *be a new geodesic triangle obtained from the quadrilateral* (p,q,p',r) *by thinking of it as a system of four rigid bars*

wise. The bound on the perimeter is necessary to ensure the existence and uniqueness of the comparison triangle.

linked by hinges and flattening the "kink" at r. The distances between the respective vertices are preserved so this new triangle is unique up to isometry. As such, we label its three vertices by $\bar{p}, \bar{p'}$ and $\bar{q}$ and the angles at these respective points by $\bar{\beta}, \bar{\beta}'$ and $\bar{\alpha}$. Now, if we denote by $\bar{r}$ the point on the geodesic segment $[p, p']$ such that $d(p, r) = d(\bar{p}, \bar{r})$ we have:

$$\bar{\alpha} \geq \alpha + \alpha', \bar{\beta} \geq \beta, \bar{\beta}' \geq \beta' \text{ and } d(q, r) \leq d(\bar{q}, \bar{r}). \tag{2}$$

In this equation, if equality occurs anywhere it must occur everywhere. Further, such an equality occurs if and only if $\gamma + \gamma' = \pi$ or the original quadrilateral was already a geodesic triangle.

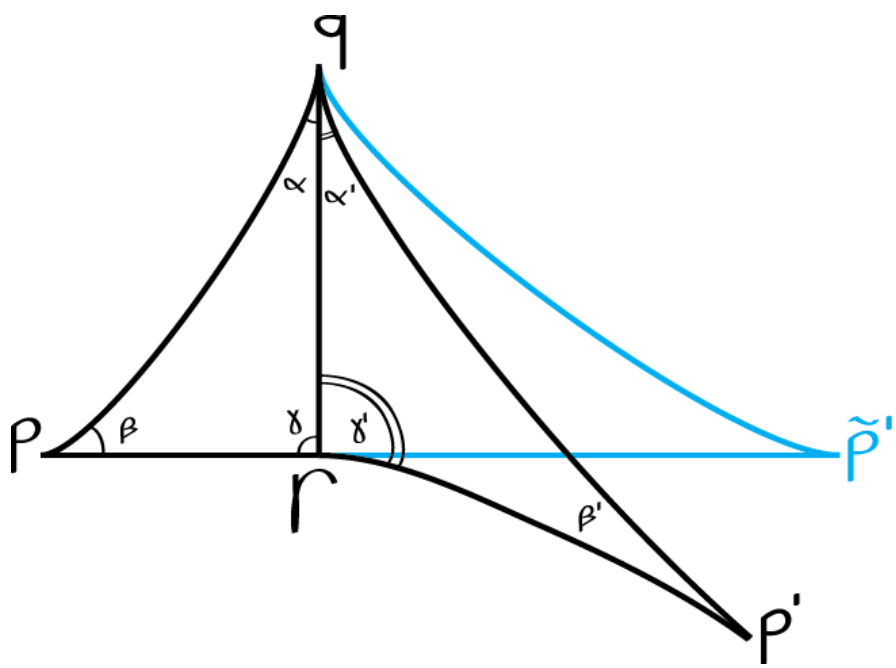

Figure 1: Alexandrov's lemma

Proof. From the proof of Lemma 1.7, we immediately obtain that $d(\bar{q}, \bar{p}') = d(q, p') \geq d(q, \tilde{p}')$. Applying the law of cosines for these two edges with the angle opposite of them in the triangles $\triangle$ and $\bar{\triangle}$ immediately yields $\beta \leq \bar{\beta}$ and

$\beta' \leq \bar{\beta}'$. Since $d(\bar{p}, \bar{p}') = d(p, r) + d(r, p') \geq d(p, p')$, we may apply the same technique to the angles opposite to the sides $\bar{p}, \bar{p}'$ and p, p' to obtain $\alpha + \alpha' \leq \bar{\alpha}$. A final application of the law of cosines yields $d(q, r) \leq d(q, \bar{r})$. In all of these cases, equality holds if and only if $\gamma + \gamma' = \pi$ in which case $\bar{\triangle}$ is isometric to the union of $\triangle$ and $\triangle'$.

□

Remark Alexandrov's Lemma holds for points in arbitrary model spaces M_k but when $k > 0$ we must additionally require the points are not too far apart or $d(p, r) + d(p', r) + d(p, q) + d(q, p') < 2D_k$ to ensure the existence of the required geodesics.

2 Bounded Curvature in Metric Spaces

We are now in a good position to define the main geometric objects of our study, metric spaces which emulate in many aspects Riemannian manifolds of non-positive sectional curvature. These metric spaces, said to be of non-positive curvature, are a special case of the notion of (locally) $CAT(k)$ spaces to be defined below. Our main reference in this section is [BH99].

2.1 Basic Preliminaries

Before proceeding, the reader unfamiliar with model spaces of a given curvature and their geodesics should read Appendix A where basic definitions and notation are set up. In the following discussion all triangles are geodesic, namely their sides are geodesic segments[4] joining their three vertices. Let $\triangle(x, y, z)$ be a triangle in a metric space X. If the perimeter of $\triangle$ is bounded by $2D_k$, we can consider the $k-$comparison triangle $\bar{\triangle} \subseteq M_k^2$. The point $\bar{x'} \in [\bar{x}, \bar{y}] \subset \bar{\triangle}$ is a comparison point of $x' \in [x, y]$ if $d(x, x') = d(\bar{x}, \bar{x'})$ and $d(x', y) = d(\bar{x'}, \bar{y})$.

The reason we are interested in comparison triangles is that since the model spaces M_k^n have constant sectional curvature k one would intuitively expect a

[4]These segments are not necessarily unique!

triangle in a space X of curvature "bounded above by k" to be more negatively curved in some sense than a similar triangle in M_k^n as illustrated in Figure 2. With this in mind, we say that a triangle $\triangle$ satisfies the $CAT(k)$ inequality or the thin triangle condition if for every $x', y' \in \triangle$ and $\bar{x'}, \bar{y'} \in \bar{\triangle}$ we have $d(x', y') \leq d(\bar{x'}, \bar{y'})$.

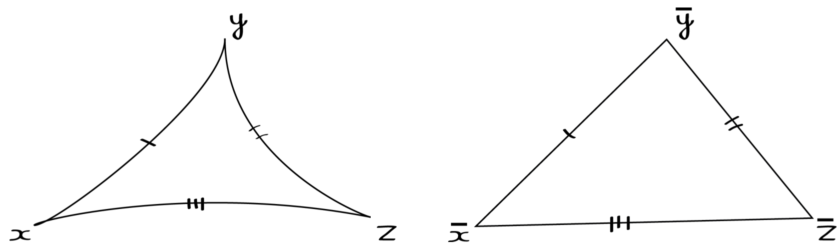

Figure 2: Thin triangle condition.

CAT(k) We say that a metric space X is $CAT(k)$ or the metric of X is $CAT(k)$ if one of the following holds:

1. If $k \leq 0$ and X is a geodesic space where all triangles satisfy the $CAT(k)$ inequality.

2. If $k > 0$ and X is D_k-geodesic with all triangles of perimeter bounded by $2D_k$ satisfying the $CAT(k)$ inequality.

These definitions are due to Alexandrov and may be found in [Ale51] under different terminology. Complete $CAT(0)$ spaces are sometimes referred to as *Hadamard spaces*, we will revisit them in Section 7. Following these definitions we use a local version of these conditions to obtain the notion of metric spaces of bounded curvature.

Definition A geodesic metric space X is said to be of curvature bounded above by k if it is locally a $CAT(k)$ space.

Remarks 1. For smooth Riemannian manifolds, having bounded curvature in the $CAT(k)$ sense is equivalent to having all sectional curvatures bounded above by k. The $CAT(k)$ condition is thus a natural generalization of sectional curvature.

2. If X is a $CAT(k)$ space then it is also $CAT(k')$ for every $k' \geq k$.

The following proposition gives us a good handle on the behaviour of such spaces. We refer the reader to [BH99] for a modern proof although it is originally due to Alexandrov.

Proposition 2.1 ([Ale51]). *Let X be a $CAT(k)$ metric space. If a pair of points x and y in X are such that $d(x,y) < D_k$ then any locally geodesic path between them coincides with the unique geodesic segment $[x,y]$. Moreover, all balls B of radius less than $D_k/2$ are convex, namely if x and y lie in B then $[x,y] \subseteq B$.*

The focus of our attention will be that of $CAT(0)$ or *non-positively curved* spaces where $D_0 = \infty$ so in the above proposition the statements are stronger than they might appear. In fact, in a $CAT(0)$ space all balls are convex and contractible. We now state two characterizations particular to this case, these are especially useful when one wants to concretely verify that a given space is $CAT(0)$ as in section 7. First, we have the *Courbure Négative* inequality of Bruhat and Tits [BT72]:

Lemma 2.2 (CN Inequality). *A geodesic metric space X is $CAT(0)$ if and only if for every triple of points p, q and r in X, given $m \in X$ such that $d(q,m) = d(r,m) = d(q,r)/2$ (i.e. m is the midpoint of $[q,r]$) we must have that*

$$d(p,q)^2 + d(p,r)^2 \geq 2d(m,p)^2 + \frac{d(q,r)^2}{2}.$$

There is also the very recent *quadrilateral condition* for such spaces due to Berg and Nikolaev and based in part on the previous Lemma.

Proposition 2.3 ([BN08]). *A geodesic metric space X is $CAT(0)$ if and only if any four points w, x, y and z of X satisfy the quadrilateral inequality*

$$d(w,y)^2 + d(x,z)^2 \leq d(w,x)^2 + d(x,y)^2 + d(y,z)^2 + d(z,w)^2.$$

The great interest of this characterization comes from the fact that it does not rely on the existence of geodesics hence allowing one to generalize the definition of $CAT(0)$ spaces to discrete spaces. We refer the reader to [Sat09] for a short and sweet proof of the result.

Example

1. All convex subsets of $\mathbb{E}^n$ are $CAT(0)$ when equipped with the induced metric.

2. $\mathbb{R}$-trees are by definition metric spaces where any two points are joined by a unique geodesic segment and the concatenation of geodesic segments always yields a geodesic. In such spaces, triangles are all "degenerate" so they are $CAT(k)$ for every k. In fact, a metric space is $CAT(k)$ for every k if and only if it is an $\mathbb{R}-$tree. As such, these spaces are sometimes called $CAT(-\infty)$.

3. Under certain hypotheses M_k-polyhedral complexes can be made to be $CAT(k)$. This is the main subject of Section 3.1.

Remark The acronym "CAT" in the definition of $CAT(k)$ spaces was originally coined by Gromov [Gro87] and stands for Cartan, Alexandrov and Topogonov. It's always more fun when cats show up in topology literature as in the nicely illustrated examples found on page two of [Wis05] and page three of [DL09]. In this vein, it would be a shame to lose the cat-like essence of the acronym when passing from english to french. Thankfully, we are in luck as Hadamard also played a role in the development of the present theory. One could therefore shamelessly insert his initial into the french acronym. *On appelle donc parfois ces espaces CHAT(0) en français.*

2.2 Convexity in $CAT(0)$ Spaces

Having laid down the basic definition of $CAT(0)$ spaces, one might wonder which features of these spaces make them so pleasant to work with. One of the most important properties from this point of view is related to the inherent convexity of their structure.

Recall that given a convex set V in a real vector space, a function $f : V \to \mathbb{R}$ is said to be convex if $f(sv_1 + (1-s)v_2) \leq (s-1)f(v_1) + sf(v_2)$ for all $s \in$

$[0,1]$. Analogously, we say that a real valued function f defined on a geodesic metric space X is *convex*[5] if for every geodesic path $g : [0,1] \to X$ parametrized proportional to arc length, the inequality

$$f(g(s)) \leq (s-1)f(g(0)) + sf(g(1))$$

holds $\forall\, s \in [0,1]$. Another way to look at this is to say that the function $f : X \to \mathbb{R}$ is convex if for all geodesics $g : I \subset \mathbb{R} \to X$ the map $f \circ g : I \to \mathbb{R}$ is convex in the traditional sense where the interval I plays the role of a convex subset of the vector space $\mathbb{R}$.

One then considers the metric as a function $d : X \times X \to \mathbb{R}$ where geodesic paths in $X \times X$ correspond to pairs of geodesic paths in X to obtain the following landmark feature of $CAT(0)$ spaces.

Proposition 2.4. *The metric of a $CAT(0)$ space is convex.*

Before proceeding with the proof, we introduce convenient notation. Given two geodesic segments $g_1, g_2 : [0,1] \to X$ parametrized proportional to arc length one can think of their *distance function* as $D_{g_1,g_2}(s) := d(g_1(s), g_2(s))$ which measures how far apart they are at a given time.

Proof. Given two geodesics as above, let us divide the convex hull of the two segments $g_1([0,1])$ and $g_2([0,1])$ into triangles $\triangle(g_2(0), g_1(0), g_1(1))$ and $\triangle'(g_2(0), g_1(1), g_2(1))$ as in Figure 3. Denoting by g_3 the geodesic segment $[g_2(0), g_1(1)]$, the $CAT(0)$ inequality for $\triangle$ shows that $d(g_1(t), g_3(t)) \leq d(\overline{g_1(t)}, \overline{g_3(t)})$ for all values of $t \in [0,1]$. Now that we have moved to the context of Euclidean geometry we immediately see that $d(\overline{g_1(t)}, \overline{g_3(t)}) \leq (1-t)d(\overline{g_1(0)}, \overline{g_3(0)})$. Combining these observations with the similar ones for $\triangle'$ yields that $D_{g_1,g_3}(t) \leq d(\overline{g_1(t)}, \overline{g_3(t)}) \leq (1-t) \cdot d(\overline{g_1(0)}, \overline{g_3(0)})$ and $D_{g_2,g_3}(t) \leq d(\overline{g_2(t)}, \overline{g_3(t)}) \leq t \cdot d(\overline{g_2(1)}, \overline{g_3(1)})$. As such, we finally obtain

$$D_{g_1,g_2}(t) \leq D_{g_1,g_3}(t) + D_{g_2,g_3}(t) \leq t \cdot d(\overline{g_2(1)}, \overline{g_3(1)}) + (1-t) \cdot d(\overline{g_1(0)}, \overline{g_3(0)})$$

[5]We invite the reader to consult [Pap05] for an overview of convexity in vector spaces and how it extends to arbitrary metric space.

which establishes the desired inequality since $d(\overline{g_2(1)}, \overline{g_3(1)}) = D_{g_1,g_2}(1)$ and $d(\overline{g_1(0)}, \overline{g_3(0)}) = D_{g_1,g_2}(0)$.

□

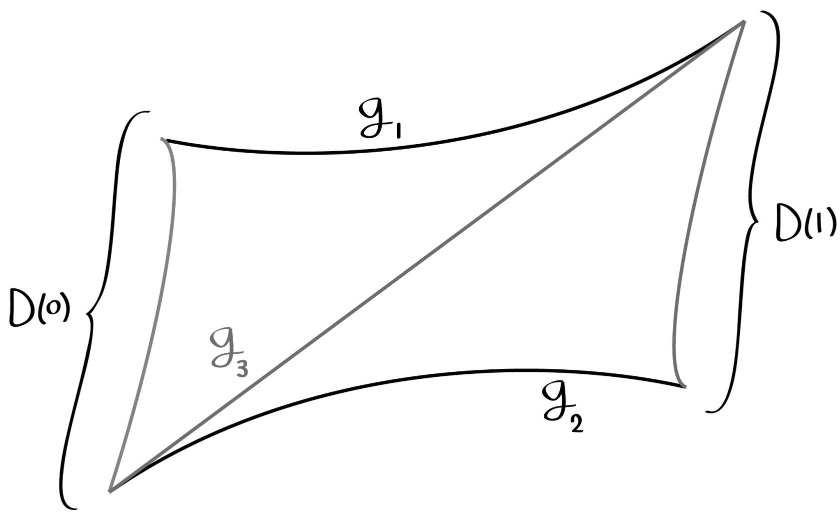

Figure 3: Convexity of the metric.

In light of the proof of this lemma, convexity of the metric simply expresses the fact that the distance between two geodesics at a given time can not exceed the distance between two comparison geodesics in euclidean space. The following key fact is a standard result for ordinary convex functions and will be used several times.

Proposition 2.5. *A locally convex function is convex and a bounded convex function is constant.*

Remark Convexity of the metric is a salient feature of non-positive curvature and in fact it is the basis for an alternative generalization of non-positive curvature to metric spaces in the sense of Busemann as developed in [Pap05]. Such spaces rely heavily on the fact that a space of infinite diameter with convex metric is necessarily uniquely geodesic.

Recall that a subset of a metric space is convex if the geodesic segment joining any two of its points lies within the subset. The following proposition stated without proof defines the useful notion of *orthogonal projection* $\pi_C : X \to C$ onto a convex subset of a $CAT(0)$ space X analogous to the similar projections frequently used in functional analysis.

Proposition 2.6 ([BH99]). *Let X be $CAT(0)$ metric space. If a convex subset $C \subset X$ is complete in the induced metric then for every $x \in X$ there is a unique point $\pi_C(x) \in C$ such that $d(x, \pi_C(x)) = \inf_{c \in C} d(x, c)$ which defines a non-expanding retraction map $\pi_c : X \to C$ homotopic to the identity.*

To better comprehend the geometry of $CAT(0)$ spaces, it is crucial to understand subspaces that are isometric to subsets of Euclidean space. The following two results are useful in characterizing "flat" subsets of $CAT(0)$ spaces, namely those that are isometric to subsets of $\mathbb{E}^2$. The first is known as the Flat Triangle Lemma and illustrates an application of Alexandrov's Lemma 1.8.

Proposition 2.7 ([Ale51]). *If $\triangle$ is a geodesic triangle in X and one of the interior Alexandrov angles of $\triangle$ is equal to the corresponding angle in some comparison $\bar{\triangle} \subset \mathbb{E}^2$ then the convex hull of $\triangle$ is isometric to the convex hull of $\bar{\triangle}$.*

Proof. The following argument follows [BH99]. Let $\triangle(p, q, r)$ be a geodesic triangle in X satisfying the above hypothesis. More precisely, suppose $\angle_p(q, r) = \bar{\angle}_p(q, r)$. We first proceed to show that for any $s \in [q, r]$, equality holds in the $CAT(0)$ condition. This is equivalent to showing that for every such s, $d(p, s) = d(\bar{p}, \bar{s})$.

To this end, fix s as above and let $\triangle'$ and $\triangle''$ be geodesic triangles with vertices (p, q, s) and (p, s, r) respectively. Let now $\tilde{\triangle}'$ and $\tilde{\triangle}''$ be their comparison triangles in $\mathbb{E}^2$ sharing the common segment joining $\tilde{p}$ to $\tilde{s}$ and such that $\tilde{q}$ and $\tilde{r}$ lie on opposite sides of $[\tilde{p}, \tilde{s}]$. By the $CAT(0)$ condition, the sum of the angles at $\tilde{r}$ in $\triangle'$ and $\tilde{\triangle}''$ must exceed π so we may apply Alexandrov's Lemma 1.8 to the pair of triangles to obtain $\angle_{\bar{p}}(\bar{q}, \bar{r}) \geq \angle_{\tilde{p}}(\tilde{r}, \tilde{s}) + \angle_{\tilde{p}}(\tilde{q}, \tilde{s})$. Since the $CAT(0)$ condition implies that the Alexandrov angle can not exceed the comparison angle we thus obtain the following chain of inequalities:

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

But now, by our original assumption, equality must hold throughout the above. We obtain in particular that $\angle_{\tilde{p}}(\tilde{r},\tilde{s}) + \angle_{\tilde{p}}(\tilde{q},\tilde{s}) = \angle_{\bar{p}}(\bar{q},\bar{r})$. A second application of Alexandrov's Lemma thus yields the desired equality $d(p,s) = d(\tilde{p},\tilde{s}) = d(\bar{p},\bar{s})$.

We are now in position to define a map $\phi : Conv(\bar{\triangle}) \to X$ from the convex hull of $\bar{\triangle}$ to the original space X which, for every $\bar{s} \in [\bar{q},\bar{r}]$ maps the segment $[\bar{p},\bar{s}]$ isometrically onto the segment $[p,s]$. We claim this map is an isometry which implies it is the desired isometry from $Conv(\bar{\triangle})$ onto $Conv(\triangle)$.

To show that it is an isometry, let $\bar{x} \in [\bar{p},\bar{s}]$ and $\bar{x'} \in [\bar{p},\bar{s'}]$ where $\bar{s}$ and $\bar{s'}$ lie on the segment $[\bar{q},\bar{r}]$. Letting $\xi = \phi(\bar{\xi})$ wherever appropriate, denote by δ_1, δ_2 and δ_3 the angles $\angle_p(q,s), \angle_p(s,s')$ and $\angle_p(s',r)$ respectively. If $\bar{\delta_i}$ represents the corresponding comparison angles in $\bar{\triangle}$, since by the first part of the proof $\triangle(\bar{p},\bar{q},\bar{s})$ is a comparison triangle for $\triangle(p,q,s)$ we must have that $\delta_1 \leq \bar{\delta_1}$. A similar argument yields $\delta_i \leq \bar{\delta_i}$ for every i. Now

$$\angle_p(q,r) \leq \delta_1 + \delta_2 + \delta_3 \leq \bar{\delta_1} + \bar{\delta_2} + \bar{\delta_3} = \angle_{\bar{p}}(\bar{q},\bar{r})$$

and by our original assumption, equality holds throughout. In particular $\delta_2 = \bar{\delta_2}$ so $d(x,x') = d(\bar{x},\bar{x'})$.

□

Appropriately, triangles described in the proposition are referred to as "flat" triangles. A generalization of the ideas in the preceding proof yields a similar condition for quadrilaterals which is the key to proving that convex hulls of asymptotic geodesics are flat strips. We refer the reader to [BH99] for a proof.

Theorem 2.8 ([Ale51]). *Given four distinct points in a $CAT(0)$ space X, if the sum of the four interior Alexandrov angles of the implicit quadrilateral ⌑ is greater* *than or equal to* 2π *then* ⌑ *is* equal *to* 2π *and the quadrilateral is in fact a flat rectangle.*

Infinite geodesics play a special role in the theory of metric spaces which will be emphasized in Section 6.1 where certain isometries will act on them as translations so they are often called *lines* or *axes*. One of the most useful relation between geodesic lines in Euclidean space is *parallelism* but it is a priori unclear how to describe such a relation in $CAT(0)$ spaces. To this end, notice that in $\mathbb{E}^n$ lines are parallel if and only if they are *asymptotic*. In other words, two lines[6] $\mathfrak{g}_1, \mathfrak{g}_2 : \mathbb{R} \to \mathbb{E}^n$ are parallel if and only if the distance $d(\mathfrak{g}_1(t), \mathfrak{g}_2(t))$ is uniformly bounded. As such, we define geodesic lines in a $CAT(0)$ space to be parallel whenever they are asymptotic. In view of the following Flat Strip Theorem found in [BH99], one might say that asymptoticity is possibly the most useful generalization of parallelism in $\mathbb{E}^n$ to arbitrary metric spaces.

Theorem 2.9 (Flat Strip Theorem)**.** *The convex hull of the union of two asymptotic geodesic lines in a* $CAT(0)$ *space is isometric to the convex hull of two parallel lines in the Euclidean plane. Appropriately, this convex hull is called a* flat strip*.*

Proof. Let $\mathfrak{g}_1, \mathfrak{g}_2 : \mathbb{R} \to X$ be two asymptotic geodesic lines in the $CAT(0)$ space X. The distance function $D_{\mathfrak{g}_1,\mathfrak{g}_2} := d(\mathfrak{g}_1(t), \mathfrak{g}_2(t))$ between the two lines is

1. Convex because X is $CAT(0)$.

2. Bounded because $\mathfrak{g}_1$ and $\mathfrak{g}_2$ are asymptotic.

so $D_{\mathfrak{g}_1,\mathfrak{g}_2}$ is a constant map. As such, we may assume without loss of generality that up to reparametrization the orthogonal projection of the first line onto the other is well behaved, namely $\pi_{\mathfrak{g}_i(\mathbb{R})}(\mathfrak{g}_j(t)) = \mathfrak{g}_i(t)$ for all $t \in \mathbb{R}$ and $(i, j) \in \{(1, 2), (2, 1)\}$. As such, consider the quadrilateral ⌑ delimited by the four points $w := \mathfrak{g}_1(t), x := \mathfrak{g}_2(t), y := \mathfrak{g}_1(t + \delta)$ and $z := \mathfrak{g}_2(t + \delta)$ where $\delta > 0$. We claim that the sum of the Alexandrov angles at the four corners of ⌑ exceeds 2π

[6] Recall that all geodesics $\mathfrak{g} : I \to X$ are parametrized such that $d(\mathfrak{g}(s), \mathfrak{g}(s')) = |s - s'|$.

so by Theorem 2.8, $\boxminus$ is a flat rectangle and the result follows. Suppose for a contradiction that the angle $\angle_x(w, z)$ in $\boxminus$ is strictly less than $\pi/2$. By definition, this means that there are points $\xi \in [x, z]$ and $\theta \in [w, x]$ such that the comparison angle $\bar{\angle}_x(\theta, \xi) < \pi/2$ in the comparison triangle $\overline{\triangle}(\theta, x, \xi)$. But now, the $CAT(0)$ inequality implies that the distance $d(w, \xi) < d(w, x)$ which contradicts the fact that $\pi_{\mathfrak{g}_2(\mathbb{R})}(w) = \pi_{\mathfrak{g}_2(\mathbb{R})}(\mathfrak{g}_1(t)) = \mathfrak{g}_2(t) = x$. It follows by symmetry that all angles are greater than or equal to $\pi/2$ and this concludes the proof.

□

The existence of a geodesic line $\mathfrak{g}$ in a $CAT(0)$ space X has strong implications on its structure. Indeed, let $X_{\mathfrak{g}}$ denote the set of geodesic lines parallel to $\mathfrak{g}$. By the Flat Strip Theorem if $\mathfrak{g}_1$ and $\mathfrak{g}_2$ are parallel to $\mathfrak{g}$ then $Conv(\mathfrak{g}_1(\mathbb{R}) \cup \mathfrak{g}_2(\mathbb{R}))$ is isometric to a flat strip consisting of other geodesic lines parallel to $\mathfrak{g}$. As such, $X_{\mathfrak{g}}$ is a convex subset and consequently is also a $CAT(0)$ space in the induced metric. Using this structure, the following product decomposition theorem is proved in [BH99].

Theorem 2.10. *Let X be a $CAT(0)$ space and consider its $CAT(0)$ subspace $X_{\mathfrak{g}}$ associated to the geodesic line $\mathfrak{g}$. If we restrict the orthogonal projection $\pi_{\mathfrak{g}(\mathbb{R})}$ to $X_{\mathfrak{g}}$ then $X_{\mathfrak{g}} \cong X^0_{\mathfrak{g}} \times \mathbb{R}$ where $X^0_{\mathfrak{g}}$ is the fibre of any chosen basepoint $x_{\mathfrak{g}} \in \mathfrak{g}(\mathbb{R})$.*

2.3 Local and Global Geometry

In this section, we collect a few constructions leading to features of metric spaces that will be used (at times implicitly) in characterizing non-positive curvature for polyhedral cell complexes.

Euclidean Cones. Cones are an important construction that generalizes the idea of the tangent space in Riemannian manifolds. Their geometry at a point carries a lot of intrinsic information about the given space as we will later see with Gromov's link condition in section 3.1. Given a metric space Y the *Euclidean Cone* $X = C_0Y$ over Y is the metric space defined as follows. As a set,

X is the quotient of $[0,\infty) \times Y$ by the equivalence relation

$$[(t,y) \sim (t',y') \text{ if } (t=t'=0) \text{ or } (t=t'>0 \text{ and } y=y')].$$

The equivalence class of (t,y) is denoted by ty and the class $(0,y)$ denoted 0 is called the *vertex of the cone*. In the definition of the metric to follow, we suggest the reader bear in mind the euclidean law of cosines as it is what everything is based upon. Let $d_\pi(y,y') := \min\{\pi, d(y,y')\}$, the distance between two points $x=ty$ and $x'=t'y'$ in X is defined by the rule:

$$d(x,x')^2 = t^2 + t'^2 - 2tt'\cos(d_\pi(y,y')),$$

It is a fact that the given formula defines a metric on $X = C_0Y$ where Y is complete if and only if X is complete.

Example If Y is the sphere $\mathbb{S}^{n-1}$, then $X = C_0Y$ is isometric to $\mathbb{E}^n$. This is one of the motivating examples behind the concept.

Remark This is a special case of the more general construction of $k-$cones $C_k(X)$ for metric spaces X obtained by modifying the metric according to the appropriate cosine law.

Space of Directions. Building upon this concept we move towards an understanding of the local geometry of a metric space X. Consider two non-trivial geodesics g and g' issuing from a given point $x \in X$. We say that g and g' define the same direction at x if the Alexandrov angle between them $\angle_x(g,g') = 0$. This establishes an equivalence relation between geodesics issuing from x whose classes are called *directions*. Equipping this space with the angular metric[7] we obtain the *space of directions at x* denoted by $S_x(X)$. We can now define the *tangent cone at x* as the Euclidean cone over $S_x(X)$, $C_0S_x(X)$.

Example Both of these concepts are natural extensions of possibly familiar ones as illustrated by the following examples:

[7]The distance between two geodesics is given by the angle between them.

1. If X is a riemannian manifold, $S_x(X)$ is isomorphic to the unit sphere in the tangent space of X at x which coincides in this case with the tangent cone.

2. If X is a polyhedral complex, $S_x(X)$ is the geometric link $Link(x, X)$.

Having defined these constructions, we now state the following theorem due to Nikolaev which characterizes curvature in the spaces of direction and tangent cones by considering *the explicit construction of geodesics in the space of directions and [proving] that the tangent cone is the four point limit of blowing ups*[8]. In fact, the result is precisely what one would intuitively expect.

Theorem 2.11 ([Nik95])**.** *If X is a metric space of curvature bounded above by k the metric completion of $S_x(X)$ is $CAT(1)$ and the metric completion of $C_0S_x(X)$ is $CAT(0)$.*

Cartan-Hadamard. We conclude this section by stating a key generalization of the Cartan-Hadamard Theorem proved in [BH99] is based on a result of [Gro87] encapsulating a major *local-to-global* aspect of the geometry of $CAT(0)$ spaces[9]. Recall that a space X equipped with a locally convex metrics is locally contractible so it has a universal cover $\tilde{X}$ on which there exists a unique induced length metric making the covering map a local isometry.

Theorem 2.12 (Generalized Cartan-Hadamard Theorem)**.** *If X is a complete, connected and locally $CAT(0)$ metric space then its universal cover $\tilde{X}$ is globally $CAT(0)$ in the induced length metric.*

Key Remark: In particular, a complete simply-connected geodesic space satisfies the $CAT(0)$ inequality *locally* if and only if it satisfies it *globally*. This is the basis of the combinatorial definition of $CAT(0)$ spaces to be introduced in Section 3.1.

[8]Please see [Nik95].

[9]The interested reader should also take a look at [AB90] for a proof under a slightly stronger hypothesis.

3 Bounded Curvature in Polyhedral Complexes

Having set up the framework of $CAT(0)$ spaces, we wish to provide concrete examples that are not Riemannian manifolds; emphasizing how pervasive these objects are in general contexts. To this end, we now turn our attention to the first main topic, M_k-polyhedral cell complexes or disjoint collections of convex polyhedra in some fixed model space glued together along isometric faces. When the choice of M_k coincides with $\mathbb{E}$, $\mathbb{H}$ or $\mathbb{S}$ we shall often say that the complex is piecewise euclidean, hyperbolic or spherical. Recall that if the set of isometry type of faces or *shapes* of such a cell complex is finite, then it is a complete geodesic metric space. The reader unfamiliar with polyhedral cell complexes and their intrinsic quotient metric should read Appendix B before proceeding.

In this section, all polyhedral cell complexes are connected and their set of shapes is finite so they are complete, geodesic length spaces.

One should keep in mind that although the set of shapes is finite, our complexes can contain infinitely many cells.

3.1 Gromov's Link Condition

The key to understanding bounded curvature in cell complexes lies in their local structure, embodying a local-to-global principle similar to the Cartan-Hadamard Theorem. With this point of view in mind it is convenient to recall the geometric link, denoted by $Link(v, X)$, of v in the complex X which coincides with the space of directions at v as defined in Section 2.3 and carries a natural cell structure induced by X. Intuitively, for a small enough $\epsilon > 0$, the link corresponds to the intersection of X with the $\epsilon-$sphere about the given point equipped with the angular metric. In the present setting, there is a useful criterion due to Gromov characterizing the existence of a metric of bounded curvature on a complex.

Link Condition Let X be an M_k-polyhedral cell complex. We say that X satisfies the *link condition* if for every vertex $v \in X^{(0)}$ the metric cell structure induced on $Link(v, X)$ is $CAT(1)$.

The *raison d'être* for this condition is that one can see with a bit of work that when $\epsilon > 0$ is small enough, the $\epsilon-$neighbourhood of a vertex $v \in X$ is isometric to the $\epsilon-$neighbourhood of the cone point[10] in $C_k(Link(v, X))$. Combining this observation with a result of [Ber83] stating that a metric space is $CAT(1)$ if and only if its $k-$cone is $CAT(k)$, one obtains the following key theorem mentioned as a *"well known and easy to prove fact"* in [Gro87][11].

Theorem 3.1. *Suppose that X is an M_k-polyhedral cell complex whose set of shapes is finite. The curvature of X is bounded above by k if and only if X satisfies the link condition.*

Using this result, one quickly deduces as indicated without proof in [Gro87] and proved in [BH99] the following convenient characterizations of bounded curvature in polyhedral cell complexes.

We state the results as definitions because that is how we will use them. If X is an M_k-polyhedral cell complex whose set of shapes is finite, there are many equivalent ways to say that it is locally or globally non-positively curved.

Definition *If $k \leq 0$ then a complex X has curvature bounded above by k if and only if it satisfies the link condition.*

A simple application of the Cartan-Hadamard Theorem then yields the global version. Notice in particular that if X has non-positive curvature its universal cover is automatically $CAT(0)$.

Definition If $k \leq 0$ then a complex X *is $CAT(k)$ if and only if it is simply connected and satisfies the link condition.*

To emphasize that this definition is far from being the only "good" one, we note that this is in turn equivalent to X not containing isometrically embedded Euclidean circles while satisfies the link condition. The same is true of requiring that X be uniquely geodesic or even that the quotient pseudometric is convex.

[10] See Section 2.3.
[11] More precisely on page 120.

Shifting our attention to the case where k is greater than zero, the requirements are slightly more technical to ensure the existence of appropriate geodesics. We only state a global definition since in practise it is only relevant to verify Gromov's link condition.

Definition *If $k > 0$ then the complex X is $CAT(k)$ if and only if it satisfies the link condition and contains no isometrically embedded circles of length less than $2\pi/\sqrt{k}$.*

There are many special results rendering a certain appeal to complexes of dimension two. Not only are they easier to visualize but they bear strong ties to group theory through tools like disc diagrams over group presentations as we will illustrate in the next few sections. When we work in dimensions less or equal to two, the cell structure induced on the link of any vertex $v \in X$ is a metric graph $Link(v, X)$ equipped with the angular metric described in Section 3.1. The vertices of this graph correspond to the $1-$cells incident to v and its the edges correspond to the corners of the $2-$cells incident to v as illustrated in Figure 4. The length of these edges is by definition the vertex angle at v between the two corresponding $1-$cells.

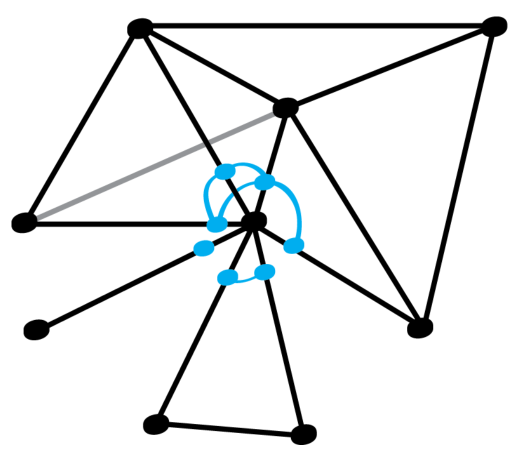

Figure 4: The figure represents the $1-$skeleton of a two dimensional simplicial complex where all triangles span a $2-$cell. The teal graph is the link of the central vertex.

Recalling our new criterions for being $CAT(k)$ we see that a metric graph is $CAT(k)$ if and only if all locally injective loops in it have length greater than or equal to $2\pi/\sqrt{k}$ and we obtain the following link condition for two dimensional complexes X mentioned in [Gro87].

Lemma 3.2 (Two Dimensional Link Condition)**.** *A two dimensional piecewise Euclidean (Hyperbolic) polyhedral cell complex X is non-positively (negatively) curved if and only if for all vertices $v \in X$, closed loops in the graph $Link(v, X)$ have length bounded below by 2π.*

After the geometry of the cells of the complex has been fixed, this criterion allows for a very combinatorial analysis of non-positive curvature in many cases. For instance, if the set of shapes of a complex X consists of a single regular Euclidean $n-$gon, edges in the metric graph corresponding to the link of a point all have length $\frac{2\pi}{n}$ so thinking of link graphs combinatorially, we simply require its girth to be bounded below by n in order to ensure non-positive curvature of X. This is the classical approach in the case of cube complexes.

3.2 Non-Positively Curved Cube Complexes

Cube complexes are polyhedral complexes all of whose cells are $n-$dimensional euclidean cubes. The set of shapes of a cube complex X can be thought of as $n-$dimensional cubes of X isometrically embedded by maps $\varphi : [-1, 1]^n \hookrightarrow X$. In fact, their inherent structure allows one to bypass much of the metric aspects of non-positive and establish a purely combinatorial definition of nonpositive curvature. Their intrinsically organized nature also facilitates the visualization of most of their properties.

Flag Complex. A simplicial complex X is a *flag complex* if all complete graphs on $n + 1$ vertices in the 1−skeleton $X^{(1)}$ span an $n-$simplex in X.

Heuristically, a complex is flag if every time one "sees" a simplex, it "is" there. The following striking results of Gromov [Gro87][12] are the base of the combinatorial

[12]Found beginning on page 122.

definition of non-positive curvature in cubulated space.

Theorem 3.3 (Gromov)**.** *A finite dimensional piecewise spherical simplicial complex where all edges have length $\pi/2$ is $CAT(1)$ if and only if it is a flag complex.*

Since the link at each vertex of a cube complex satisfies the hypothesis of the previous theorem because Euclidean cubes have angles of $\pi/2$ at every vertex, the link condition yields the following criterion which embodies the slogan that "*a cube complex is non-positively curved if and only if there are no missing cubes*".

Theorem 3.4 (Gromov)**.** *A finite dimensional cube complex is non-positively curved if and only if its link at each vertex is a flag complex.*

Remarks

1. The preceding theorem of Gromov has recently been generalized to the infinite dimensional case by M. Sageev's student Y. Algom.
2. In two dimensional cube complexes, the link at every point is a graph. Since a graph is a flag if and only if its girth is bounded below by 4 we obtain an easy criterion for non-positive curvature.
3. Criterions for cube complexes to be manifolds admitting polyhedral metrics of non-positive curvature are illustrated in [AR90].

While general cube complexes can be thought about as higher dimensional graphs, non-positively curved cube complexes are very well behaved in many aspects as they tend to emulate trees, something we will see in Section 6.3. They also tend to be a very powerful tool, in fact, work of Wise and Haglund on $CAT(0)$ cube complexes in particular Wise's monumental manuscript [Wis11] has recently allowed Agol to prove the Virtually Haken Conjecture in [AGM12].

Due to the intrinsically combinatorial nature of non-positive curvature in a cube complex X, it is often convenient to forget about the $CAT(0)$ quotient pseudometric and consider its *combinatorial metric*[13] instead. This is defined

[13] In general this metric is rather different from the $CAT(0)$ metric. For instance, geodesics are no longer unique.

by thinking of the 1−skeleton of $X^{(1)}$ as a graph where all edges have length 1 and defining a (proper) combinatorial path $\gamma = (v_0, v_1, \dots, v_n)$ as a sequence of vertices in the 0−skeleton $X^{(0)}$ where $v_i \neq v_i$. The combinatorial length of such a γ is then defined to be $l(\gamma) = n$ and the combinatorial distance between two points $x \neq y$ of $X^{(0)}$ is the combinatorial length a shortest path between them. As one would expect, shortest combinatorial paths are called combinatorial geodesics.

Much of the elegance of cube complexes stems from the fact that their properties are combinatorially encoded in a family of subspaces called hyperplanes. To define them following [Wis11] we first describe them locally. Given an $n-$cube $C \subset X$ and its isometric embedding $\varphi : [-1,1]^n \hookrightarrow X$ we define its local hyperplanes or *midcubes* as the subspaces of C obtained by restricting one of the coordinates of $[-1,1]^n$ to zero. Visually, this corresponds to slicing the cube C in two halves. For instance, an edge $[-1,1]$ has a unique midcube which corresponds to a point while a three dimensional cube $[-1,1]^3$ has three midcubes $[-1,1]^2$ isometric to two dimensional squares. One can then define the *hyperplanes* of X as its connected subspaces $\mathcal{H}$ that intersects any cube $C \subset X$ in one of its midcubes or the empty set. Intuitively, one can think of a hyperplane $\mathcal{H}$ as being defined by an initial choice of midcube $M \subset C \subset X$ and then successively "pushing out" through midcubes of adjacent cubes. This is part of the content of Sageev's key theorem stated below. We say that a cube C is *dual* to the hyperplane $\mathcal{H}$ if they intersect in a midcube of C. The *carrier* or neighbourhood of $\mathcal{H}$ is the set of its dual cubes, it is denoted by $N(\mathcal{H})$.

Theorem 3.5 ([Sag95])**.** *Let X be $CAT(0)$ cube complex.*

1. *Every midcube determines a unique hyperplane $\mathcal{H}$ which separates X into two connected components. Further the structure of $\mathcal{H}$ induced by the midcubes of X makes it a $CAT(0)$ cube complex.*

2. *The carrier $N(\mathcal{H})$ is isometric to the product $\mathcal{H} \times [-1,1]$ and is combinatorially convex as a subset of X. As such, there is always an automorphism $\sigma_{\mathcal{H}}$ fixing $\mathcal{H}$ pointwise and exchanging the endpoints of all edges dual to $\mathcal{H}$.*

3. *The combinatorial distance between two points x and y of X is equal to the number of hyperplanes separating them. In fact, a combinatorial path is a geodesic if and only if it crosses any hyperplane at most once.*

This structural result is the key to most applications of cube complexes, including the classification of their combinatorial isometries as we will see in Section 6.3. It should be noted that there is an analogous statement when X is non-positively curved obtained by replacing X with its universal cover in the appropriate places.

3.3 Tower Lifts of Combinatorial Maps

We now turn our attention to the subgroup structure of fundamental groups of polyhedral cell complexes. One would intuitively hope that if Γ is the fundamental group of a complex X of bounded curvature then subgroups of Γ can be interpreted as fundamental groups of complexes within the same class as X in some sense. It turns out that properties of this type are related to classes of complexes closed under passage to finite subcomplexes and connected covering spaces which are best expressed through the idea of towers of maps.

The concept of a *tower* originated in a paper of Papakyriakopoulos, [Pap57], where it was used to prove Dehn's Lemma. This result infers from the existence of a piecewise-linear map of a disk into a $3-$manifold the existence of and embedding corresponding to the original map on the boundary. Although it was first announced by Dehn, it was later discovered by Kneser that his proof contained a gap. To rectify the problem, Papakyriakopoulos expressed the original piecewise linear map as a sequence of inclusions and covering maps whence exploiting the inherent simplifications that occurred. The transfer of these ideas to combinatorial complexes first occurred in [How81].

In what follows let X, Y and Z be CW complexes[14] where both X and Y are compact. Given a *combinatorial*[15] map of complexes $f : X \to Z$, our goal is to decompose it as much as possible into a sequence of simpler maps. More

[14] Please consult [Hat05] for basic definitions and notations involving CW complexes.
[15] A map which sends $n-$cells to $n-$cells but isn't necessarily injective.

precisely, we wish to rewrite f as a composition $f = g \circ f'$ where $f' : X \to Y$ and $g : Y \to Z$ such that the composition maps have desirable simplifying properties.

Admissible Tower The map $g : Y \to Z$ is an admissible tower of height h if we can rewrite $g = i_0 \circ p_1 \circ i_1 \circ \ldots \circ p_h \circ i_h$ where $p_r : Z_r \to Y_{r-1}$ is a connected covering map of a compact CW complex and $i_r : Y_r \to Z_r$ is an inclusion. It is worthwhile to note that admissible towers are inherently combinatorial because they factor through covering spaces.

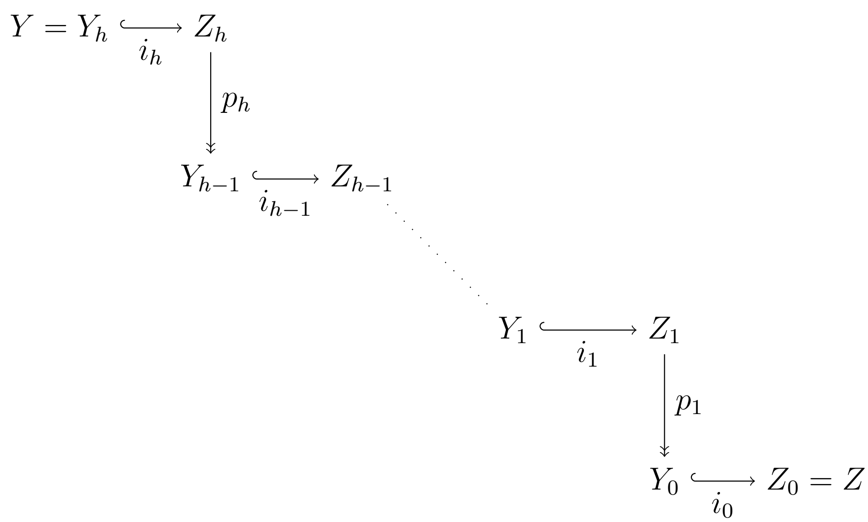

Tower Lift The map $f = g \circ f'$ is a tower lift if $f' : X \to Y$ is combinatorial and $g : Y \to Z$ is an admissible tower.

We shall henceforth denote this situation by $f : X \xrightarrow{f'} Y \xrightarrow{g} Z$ and, since we will be interested in studying the fundamental groups of spaces, we wish our tower lifts to be maximal in sense of the following lemma:

Lemma 3.6 ([How81])**.** *Consider a combinatorial map $f : X \to Z$ between connected CW complexes. If X is compact, there is a tower lift of maximal height $f : X \xrightarrow{f'} Y \xrightarrow{g} Z$ such that $f'_* : \pi_1 X \to \pi_1 Y$ is surjective.*

Proof. The proof proceeds by iteratively constructing higher tower lifts of the map f and showing that the complexity of f' decrease at each step in a quantifiable way. The process must then end in finitely many steps yielding the desired maximal tower.

Let us begin by defining the *complexity* of a combinatorial map $f : X \to Y$ to be $c(f) := |X^{(0)}| - |f(X)^{(0)}|$ or the difference in the number of $0-$cells in each CW complex. Now, setting $Y_0 := f(X)$ we trivially obtain an initial tower lift of the form

$$f : X \xrightarrow{f_0} Y_0 \xrightarrow{i_0} Z.$$

Suppose that this tower lift is not maximal, whence $f_0^* : \pi_1 X \to \pi_1 Y_0$ is not surjective. By the Galois correspondence for covering spaces and the lifting criterion there is a connected proper covering space $p_1 : Z_1 \to Y_0$ to which the map f_0 lifts to a map that we shall call f_1. Defining $Y_1 := f_1(X)$ we obtain a higher tower lift of the form $f = i_0 \circ p_1 \circ i_1 \circ f_1$ where i_1 represents the inclusion of Y_1 in Z_1.

At this point, we wish to show that $c(f_1) < c(f_0)$. To this end, let us consider the composition $(p_1 \circ i_1)$ which is surjective but not injective. As such there must be a nontrivial deck transformation σ such that $(\sigma \cdot Y_1) \cap Y_1 \neq \emptyset$. Since Y_1 is a compact cell complex and σ is a deck transformation, the intersection must be a union of closed cells and must contain a zero cell v. But now, $\sigma \cdot v \neq v$ (only the trivial deck transformation fixes a point) while $p_1(\sigma \cdot v) = p_1(v)$ so $|Y_1^{(0)}| > |Y_0^{(0)}|$ and consequently $c(f_1) < c(f_0)$.

If f_1^* is surjective, we are done. If not, we simply iterate the procedure above which must terminate because $c(f) \geq 0$ and $c(f_{i+1}) < c(f_i)$ whenever f_i^* is not surjective.

□

It should be emphasized that the preceding propositions is proved by an abstract argument and doesn't provide any explicit constructions for the next few results that rely upon it. The next proposition shows us how towers are used to exploit closure properties of classes of complexes using the idea of disc diagrams and the Cayley complex of a group presentation from geometric group theory. The reader unfamiliar with these concepts should consult Appendix D before proceeding.

Proposition 3.7 ([BH99])**.** *If Γ is a finitely presented group which injects into the fundamental group of a polyhedral cell complex K then there exists a compact two dimensional complex K_Γ such that $\pi_1(K_\Gamma) \cong \Gamma$ which can be obtained from K by successively passing to finite subcomplexes or connected covers.*

Proof. Let X be the standard $2-$complex[16] of a finite presentation

$$\Gamma \cong \langle a_1, a_2, \ldots, a_m | r_1, r_2, \ldots, r_n \rangle.$$

In other words, X has a single $0-$cell, m $1-$cells and n $2-$cells where $X^{(1)}$ consists of a bouquet of m circles labelled a_i and $X^{(2)}$ is obtained by attaching a $2-$cell along each relator r_i. We wish to show that there is a combinatorial map $f : X \to K$ inducing an injection $f_* : \pi_1\Gamma \hookrightarrow \pi_1 K$.

Indeed, once this is the case, Lemma 3.6 ensures that f admits a tower lift

$$f : X \xrightarrow{f'} K_\Gamma \xrightarrow{g} K$$

where f'_* is surjective. On the other hand, by construction, $f_* = (g_* \circ f'_*)$ is injective so f'_* must be an isomorphism $\pi_1 K_\Gamma \cong \Gamma$. Since g is an admissible tower, K_Γ is obtained from K by a sequence of passages to subcomplexes and connected covers so $K_\Gamma \in \mathbb{K}$. Since K_Γ is also the combinatorial image of a map obtained by iterative lifts of f (see the construction in the proof of Lemma 3.6), it is at most two dimensional and we are done.

To construct the desired combinatorial map $f : X \to K$, let us begin by defining f on the circles labelled a_i, so that $f(a_i)$ is mapped to a monotone parametrization of a fixed loop in K representing the homotopy class of $\phi(a_i) \in \pi_1(K, k)$ based at the zero cell corresponding to the chosen basepoint $f(v) :=$

[16] The reader unfamiliar with this construction may want to glance at Section 5.1 before proceeding.

$k \in K$. Notice that ϕ is injective so $f(a_i)$ is not nullhomotopic for any of the a_i. Now, since Γ is embedded in $\pi_1(K,k)$, every relator r_j (sequence of a_i's traversed in the corresponding order) is mapped under f to a nullhomotopic loop in K. Therefore, by van Kampen's lemma, there is a disk diagram $D_{r_j} \hookrightarrow K$ such that ∂D_{r_j} gets mapped to $f(r_j)$. If we denote by R_j the $2-$cell in X corresponding to the relator r_j, we may endow its boundary and interior with the cell structure of D_{r_j} and extend the definition of f so that $f(R_j)$ coincides with the image of the disk diagram in K and the following diagram commutes:

$$\begin{array}{ccc} R_j & = & D_{r_j} \\ \downarrow & & \downarrow \\ X & \longrightarrow & K \end{array}$$

The new cell structure induced on X by this correspondence makes f into a combinatorial map and, since f_* is injective by construction, the proof is complete.

□

3.4 Passing to Subgroups of the Fundamental Group

While subgroups of the fundamental group of a complex X of bounded curvature do not necessarily occur as fundamental groups of a subcomplex $X' \subset X$, a result similar in flavour does hold.

Theorem 3.8 ([BH99]). *If Γ is a finitely presented subgroup of the fundamental group of an M_k-polyhedral cell complex of curvature bounded above by k then Γ is the fundamental group of some compact two dimensional M_k-polyhedral complex of curvature bounded above by k in which every local geodesic can be extended to a geodesic line.*

In order to come to grasps with this theorem, let us fix $k \in \mathbb{R}$ and consider the class of complexes $\mathbb{X}(k)$ which consists of those connected M_k-complexes X whose set of shapes is finite, satisfying the link condition and whose cells are of dimension less than or equal to two. This class is clearly closed under passage to connected covering spaces since if $X \in \mathbb{X}$, any cover $\hat{X}$ must satisfy $Shapes(\hat{X}) = Shapes(X)$ and be locally isometric to X ensuring the link condition holds. On the other hand, any subcomplex $X' \subset X$ must trivially satisfy the restrictions on shapes and dimension, however if X is of dimension ≥ 3 the link condition may fail in X' as illustrated in the following example.

Example In a cube complex, the link condition holds if and only if the link at each vertex is a flag complex. In Figure 5, we see a one dimensional cube complex where the link condition fails as illustrated by the teal three cycle. However, adding a three dimensional cube behind the figure would make it non-positively curved. This reinforces the intuitive notion that a cube complex is non-positively curved if when you see the outline of a cube, it necessarily belongs to the complex.

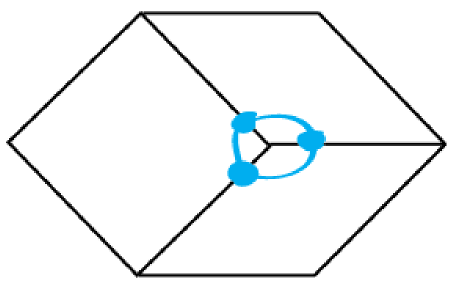

Figure 5: Cube complex

Luckily in dimension two and below, the link condition is equivalent to the absence of injective loops of length less than 2π in the link of any vertex $v \in X^{(0)}$ (which is a graph). Since $Link(v, X')$ is a subgraph of $Link(v, X)$ we see that the link condition is preserved in subcomplexes. We have proved:

Lemma 3.9. *With the induced length metric, the class of complexes* $\mathbb{X}(k)$ *is closed under passage to connected subcomplexes and connected covers.*

In view of Lemma 3.9 and Proposition 3.7, we see that if H is a subgroup of $\Gamma = \pi_1 X$ with $X \in \mathbb{X}$, there must be some $X' \in \mathbb{X}$ of dimension ≤ 2 such that $H = \pi_1 X'$. We will then have completed the proof of Theorem 3.8 once we show that X' deformation retracts onto a subcomplex with the geodesic extension property which is the content of the final lemma.

Lemma 3.10. *A compact* $2-$*dimensional* M_k-*complex satisfying the link condition collapses onto a subcomplex with the geodesic extension property.*

Before proving the lemma, let us recall the notion of an elementary collapse in a cell complex. A cell C in a combinatorial complex is called a *free face* if it is contained in the boundary of a unique higher dimensional cell C' and the intersection of the interior of C' with a small neighbourhood of C is connected. When this situation arises in a complex X we can "collapse" the free face C, meaning that there is a deformation retract from X to X' obtained by removing C and the interior of C' from X. Evidently, $\pi_1 X \cong \pi_1 X'$.

Proof. All local geodesics may be extended indefinitely in a compact M_k-polyhedral complex X of curvature $\leq k$ if and only if it does not contain a free face. Since X has only finitely many free faces, after finitely many elementary collapses we obtain the desired complex.

□

4 A Knotty Zoo of Non-Positive Curvature

Non-positively curved structures are highly natural and occur all over. A vast source of examples of non-positively curved $2-$complexes arise from constructions based on knots. These date back to Dehn's work in [Deh87] on presentations of the fundamental group of complements of knots in $\mathbb{S}^3$. This was later refined by Weinbaum [Wei71] and elegantly adapted to the non-positively curved case by Wise [Wis96]. On the other hand, one can also insert finitely generated groups into the tail end of short exact sequences containing the fundamental group of a non-positively curved cell complex. This modification of a construction of Rips' in [Rip82] by Wise is presented in the last subsection. It should be emphasized that the examples provided here are *concrete* in the sense that most of the proofs provide explicit constructions.

4.1 Basic Results for Knots

In these preliminaries we restrict our attention to the case of knots as everything can be easily generalized to collections of possibly intertwined knots called links. We follow to some extent the exposition given in [BH99] and [Rol03].

For our purposes, a *knot* $\mathcal{K}$ will be the obvious[17] equivalence class of a smooth embedding $f : \mathbb{S}^1 \hookrightarrow \mathbb{R}^3$. Having fixed a representative f, we denote by $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ its projection onto the xy plane. It is a fact that f may always be chosen such that the *regular* projection $\Pi := \pi \circ f$ is smooth with self-intersections corresponding to double points. These *double points* are paired as usual with data describing the *overpass* and the *underpass* of the knot. To make this clear in the projection $\Lambda := \Pi(\mathbb{S}^1)$, overpasses are drawn completely while underpass are drawn as missing a small neighbourhood about the corresponding double point. We refer to this altered projection as Λ'.

It is convenient to have at hand a coherent way of decomposing complicated knots into simpler ones. A natural way to do so is to consider knots as one dimensional manifolds where there is a natural decomposition operation. In that case, we say that $\mathcal{K}$ is the *connected sum*[18] of two knots $\mathcal{K}_1 \# \mathcal{K}_2$ if it is obtained by deleting a $0-$ball from both $\mathcal{K}_1$ and $\mathcal{K}_2$ and gluing together the resulting boundary spheres. In connection with this idea, a *reducing circle* for a projection Π is an embedded circle $\mathcal{C} \subset \mathbb{R}^2$ intersecting Λ in two non-double points such that both connected components of $\mathbb{R}^2 \setminus \mathcal{C}$ contain a double point as illustrated in Figure 6. In this case, we see that the corresponding knot could have been decomposed as a connected sum of two knots as in Figure 7.

[17] That is, two knots are equivalent if there is an ambient isotopy between them.
[18] The reader is invited to consult [Hat05] for a more thorough explanation.

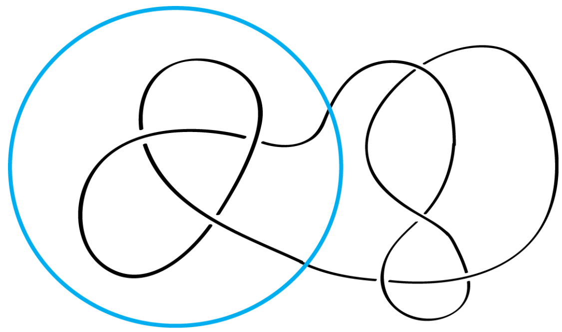

Figure 6: Reducing circle in the image of a knot projection.

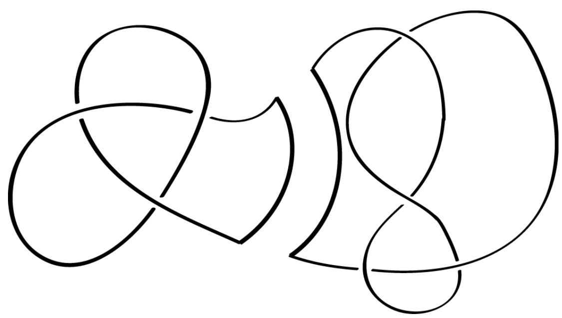

Figure 7: The resulting decomposition into prime knots.

As such, the projection Π is said to be *prime* if it does not admit a reducing circle. Finally, if when travelling along $\mathcal{K}$ in a monotone manner one encounters overpasses and underpasses in an alternating manner Π is said to be *alternating*. Knots are said to be prime or alternating if they admit a prime or alternating projection. The use of these definitions is expressed in the following two well known results, see for instance [BZ03] or [BH99].

Lemma 4.1. *Alternating knots may be decomposed as the connected sum of finitely many knots, each admitting a prime alternating projection*

As such, the following is a consequence of the Seifert van Kampen Theorem.

Corollary 4.2. *The fundamental group of the complement of an alternating knot is an amalgamated free product*[19] *along infinite cyclic groups of the fundamental groups of prime alternating knots.*

4.2 Dehn Complexes of Knot Complements

We now proceed to construct very concrete examples of non-positively curved $2-$complexes. Let $\mathcal{K}$ be a prime alternating knot and maintain the previously established notation. We begin with the following well known fact neatly proved in [Wis06] which allows one to choose a "checkerboard" colouring of the connected components of $\mathbb{R}^2 \setminus \Lambda$ for *any* knot.

Lemma 4.3 (Checkerboard Lemma)**.** *Let $\mathcal{G}$ be a planar graph where the number of edges incident to any vertex is even. Given any embedding $\mathcal{G} \hookrightarrow \mathbb{R}^2$, the connected components of $\mathbb{R}^2 \setminus \mathcal{G}$ may be coloured black and white such that no two adjacent regions are of the same colour.*

In our particular case, knot projections Λ can be viewed as an embedding in $\mathbb{R}^2$ of a graph where all vertices (double points) are incident to four edges and our chosen projection Π is alternating we may choose the following natural colouring. Given the connected components of $\mathbb{R}^2 \setminus \Lambda$: $\{A_0, A_1, \dots A_n\}$ where A_0 denotes the unbounded region, we define A_i $(i \geq 1)$ to be "white" if the anticlockwise orientation of its boundary orients its edges from overpasses to underpasses and "black" otherwise. The unbounded region A_0 is coloured by the opposite convention.

Definition The *Dehn Complex* $\mathcal{D}(\Pi)$ of the projection Π is the two dimensional cube complex with the following structure:

1. Two vertices v_+ (Top) and v_- (Bot) that we picture as lying above and below the knot projection.

2. One $1-$cell for every region A_i oriented from v_+ to v_-.

[19]See Appendix C.

3. One 2−cell ϕ for every double point $x \in \Lambda$ attached by the following map. If $A_{i_x(1)}, A_{i_x(2)}, A_{i_x(3)}$ and $A_{i_x(4)}$ are the the regions one encounters while proceeding anticlockwise around x in a small circle beginning in a white region, the attaching map of ϕ is defined by the word $A_{i_x(1)}A^{-1}_{i_x(2)}A_{i_x(3)}A^{-1}_{i_x(4)}$ where the $A_{i_x(j)}$ represent the previously defined 1−cells and the $^{-1}$ indicates it is traversed in the opposite direction.

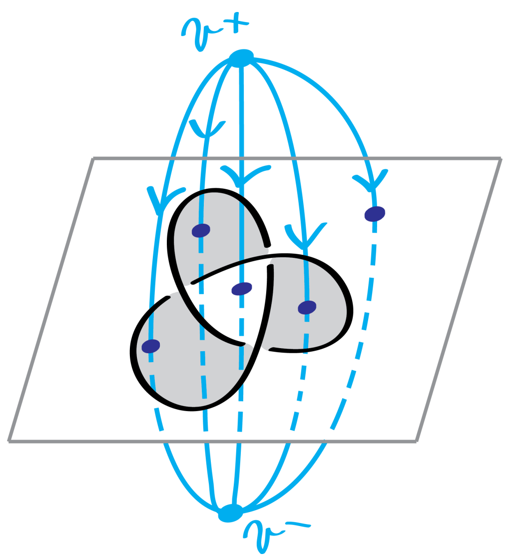

Figure 8: The 1−skeleton of the Dehn complex of a trefoil knot.

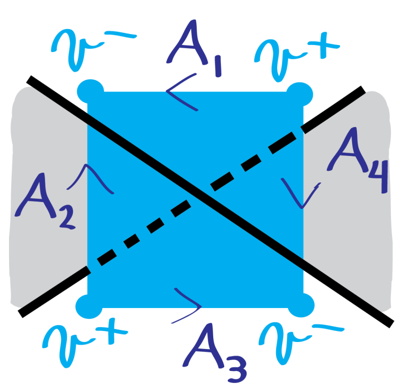

Figure 9: The attaching map of a 2−cell at a double point.

As a corollary to this construction we obtain the following *Dehn presentation* for the fundamental group of $\mathbb{R}^3 \setminus \mathcal{K}$. If we chose v_+ as our basepoint and let a_i represent the homotopy class of the loop $A_i A_0^{-1}$ the fundamental group of the complement of $\mathcal{K}$ can be expressed as

$$\langle a_0, \ldots, a_n | a_0 = 1, a_{i_x(1)} a_{i_x(2)}^{-1} a_{i_x(3)} a_{i_x(4)}^{-1} = 1 \text{ for each double point } x \rangle$$

obtained by a simple application of the Seifert-van Kampen Theorem[20].

Remark Consider Λ' as a subset of $\mathbb{R}^2$ (this is just Λ with small neighbourhoods of underpasses at crossing points deleted). Wise noticed in [Wis06] that the geometric link $Link(v_+, \mathcal{D}(\Pi))$ can be embedded in $\mathbb{R}^2 \setminus \Lambda'$ in a highly visual way. Interchanging overpasses and underpasses yields a similar result for v_-.

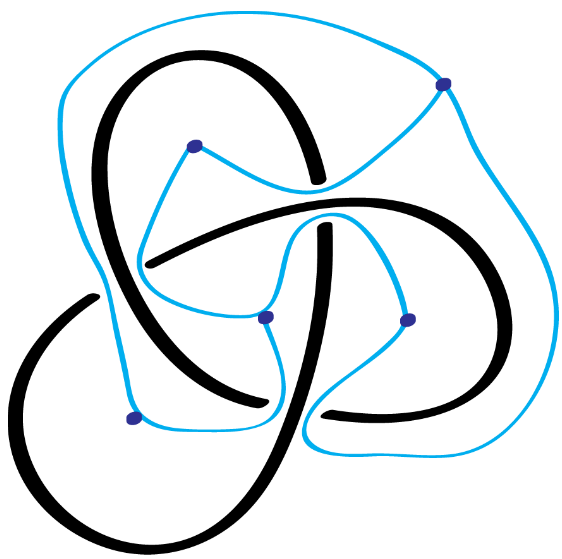

Figure 10: Link of the Dehn complex embedded in the trefoil knot.

Begin by viewing the regions $A_0, \ldots, A_n$ as sitting in $\mathbb{R}^2 \setminus \Lambda'$. The vertices of the link correspond to incident $1-$cells at v_+ so there is one for every region $A_0, \ldots A_n$. On the other hand, $1-cells$ of the link correspond to corners of $2-$cells

[20]Please see Appendix C if this is unfamiliar.

of $\mathcal{D}(\Pi)$ at v_+. As such, if we pick one point in each region to represent the vertices of the link, the $1-$cells of the link correspond to curves between adjacent regions passing through the gaps at double points and joining the corresponding vertices. This realizes $Link(v_+, \mathcal{D}(\Pi))$ as a graph in $\mathbb{R}^2 \setminus \Lambda'$. We illustrate this for the trefoil knot in Figure 10. The black lines represent Λ' and the teal graph is an embedding of the geometric link of v_+.

We henceforth metrize $\mathcal{D}(\Pi)$ by viewing it as a piecewise Euclidean complex where each $2-$cell corresponds to a unit square. In this setting we obtain the following nice class of examples:

Theorem 4.4 ([Wis06])**.** *The Dehn complex of a knot projection* Π *is non-positively curved if and only if* Π *is prime and alternating.*

Proof. Having metrized the $2-$cells of $\mathcal{D}(\Pi)$, we see that edges in $Link(v_\pm, \mathcal{D}(\Pi))$ have length $\pi/2$. As such, verifying the link condition for the complex is equivalent to ruling out paths of combinatorial length less than four.

Suppose that Π is not prime. There is by definition a reducing circle S and the components of $S \setminus \Lambda$ must lie in distinct components of $\mathbb{R}^2 \setminus \Lambda$, say A_i and A_j. Since the interior and exterior of S contain a double point, A_i and A_j must meet along at least two distinct edges in such a way that there is a combinatorial path of length two joining them in the graph mentioned in the remark. On the other hand, if we suppose that Π is not alternating, without loss of generality some edge of Λ will be an overpass at both of its endpoints when viewed from above. As such, the components of $\mathbb{R}^2 \setminus \Lambda$ meeting along this edge are once again joined by a combinatorial path of length two in $\mathbb{R}^2 \setminus \Lambda'$.

Conversely, suppose that Π is prime and alternating. Recall that the regions of $\mathbb{R}^2 \setminus \Lambda$ may be coloured black and white in such a manner that adjacent regions always have different colour. This induces a bipartite colouring on the vertices of $Link(v_\pm, \mathcal{D}(\Pi))$ implying that all closed loops in this graph must have even length. It is therefore sufficient to rule out the existence of combinatorial paths of length two but a path of length two between vertices A_i and A_j immediately contradicts the fact that the knot is alternating. □

Combining this result with Corollary 4.2 we see that the fundamental group of a knot complement is the amalgamated free product along infinite cyclic subgroups of the fundamental groups of compact non-positively curved square complexes. As it turns out, if we let $G_1 = \pi_1(X_1)$ and $G_2 = \pi_1(X_2)$ where both X_i are compact non-positively curved cell complexes of dimension at most two and glue X_1 to X_2 along a cylinder identifying the two necessary generators of $\mathbb{Z}$ in G_i then the resulting complex remains non-positively curved and coincides with $G_1 *_{\mathbb{Z}} G_2$. Modulo generalizations and precisions, we obtain the following theorem of Wise.

Theorem 4.5 ([Wis06]). *The fundamental group of the complement of an alternating link is isomorphic to the fundamental group of a compact* $2-$*dimensional piecewise Euclidean* $2-$*complex of non-positive curvature.*

Remark It is unknown whether the fundamental group of all link complements are the fundamental groups of non-positively curved cube complexes.

4.3 Algorithmic Construction of Negatively Curved $2-$Complexes

Keeping with the spirit of producing a vast array of examples of negatively curved complexes of dimension two we present a construction of Wise generating such complexes from arbitrary finitely presented groups. This is based on previous work of Rips, in particular on the proof of the following theorem:

Theorem 4.6 ([Rip82]). *Given a finitely presented group G and a constant $\lambda > 0$, there is a short exact sequence of groups*

$$1 \to N \to \Gamma \to G \to 1$$

where Γ has a finite group presentation satsifying the small cancellation[21] *condition $C'(\lambda)$ and N is a finitely generated group.*

[21] Consult [LS77] for definitions and a classical development of the theory.

Let $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ be an alphabet and consider a set of words $\mathcal{W} = \{W_1, W_2, \ldots, W_m\}$ in the letters of $\mathcal{A}$. This set of words $\mathcal{W}$ is said to contain no *2-letter repetitions* if any of the 2−letter words in the letters of $\mathcal{A}$ occur at most once as a subword of at most one of the W_i. The following lemma of Wise is known as the Good Word Lemma.

Lemma 4.7 ([Wis98]). *Given $J \in \mathbb{N}$ and an alphabet $\mathcal{A}_J = \{a_1, \ldots, a_J\}$ there is a positive*[22] *word W_J of length J^2 with no 2−letter repetitions.*

Proof. Consider the following sequence of words in $\mathcal{A}_J$: $W_1 = (a_1)$, $W_2 = (a_1a_1a_2)(a_2)$, $W_3 = (a_1a_1a_2a_1a_3)(a_2a_2a_3)(a_3)$, $W_4 = (a_1a_1a_2a_1a_3a_1a_4)(a_2a_2a_3a_2a_4)(a_3a_3a_4)(a_5)$. Proceeding in a similar manner, we define

$$W_J := (a_1a_1a_2a_1a_3a_1a_4 \ldots a_1a_J)(a_2a_2a_3a_2 \ldots a_2a_J) \ldots (a_{J-1}a_{J-1}a_J)(a_J).$$

By construction, W_J has no 2−letter repetitions and its length is equal to $\sum_{i=1}^{J}(2i-1) = J^2$ so the prophecy is fulfilled. □

This lemma turns out to be the key to the proof of the next theorem that we demonstrate following Wise's ideas.

Theorem 4.8 ([Wis98]). *Let G be a finitely presented group. There is an algorithm producing a negatively curved 2−complex K and finitely generated group N such that $\pi_1 K \cong G/N$. In others words, we have the following short exact sequence:*

$$1 \to N \to \pi_1 K \to G \to 1$$

Proof. Let $G = \langle a_1, a_2, \ldots, a_I | R_1, R_2, \ldots, R_K\rangle$ be a finite presentation. We will obtain the complex K as the standard 2−complex of a presentation extending the given one for G, "unwrapping" its relations. As such, postponing the definition of the integer J, define Γ as the group with the following presentation

$$\Gamma \cong \langle a_1, \ldots, a_I, x_1, \ldots, x_J | a_i x_j a_i^{-1} = W_{ij+}, a_i^{-1} x_j a_i = W_{ij-}, R_k = W_k\rangle$$

[22] A word where all exponents are positive.

Where $i \in \{1, \ldots, I\}$, $j \in \{1, \ldots, J\}$ and $k \in \{1, \ldots, K\}$, $W_{ij\pm}$ is a positive word of length 14 in the x_j's and W_k is a positive word of length $2||R_k|| + 8$ also in the x_j's. It is immediate from this construction that if we define N as the subgroup of Γ generated by $\langle x_1, x_2, \ldots, x_J \rangle$ that it will be normal. Further, $\Gamma/N \cong G$ so the groups fit into the short exact sequence as claimed. All that remains to be done is to equip K with a negative metric. This is where we will exploit our freedom in the choice of J using the preceding lemma and the precise length of the developed relations to divide them into pentagons.

To this end, let us metrize the 2−cells of K by subdividing them into hyperbolic regular right angled pentagons (effectively making K an M_{-1}−polyhedral complex). This procedure is illustrated in Figure 11 where the black arrows represent a_i's and the blue arrows represent x_j's. The top figure represents the conjugate relations $a_i x_j a_i^{-1} = W_{ij+}$ while the bottom shows the $R_k = W_k$ case for some relation R_k of length four.

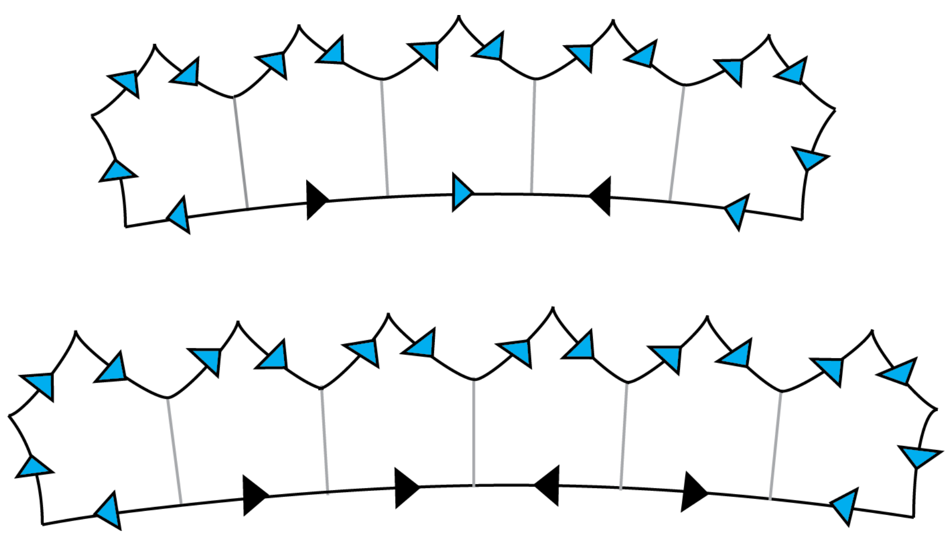

Figure 11: Hyperbolic right angled pentagon decomposition of 2−cells.

This subdivision greatly simplifies verification of Gromov's link condition (Lemma 3.2) to ensure that K is negatively curved. Indeed, notice that if we think of the 2−cells as 14-gons or $(2||R_k|| + 8)$−gons respectively, the angles at the corners of the cells are of π or $\pi/2$ (flat or sharp). The link condition then boils down to eliminating sequences of corners representing closed paths of length $< 2\pi$ in the geometric link.

Since the R_k's were chosen to be reduced, sequence of the form $a_i a_i^{-1}$ do not label any corner. Therefore any sequence of corners corresponding to a cycle in the link containing a black edge must have length exceeding 2π. We may thus restrict our attention to sequences of corners containing only blue edges. Keeping in mind that the $W_{ij\pm}$ were chosen to be positive words, all corners of blue edges must be labelled $x_j x_h$ or $x_h^{-1} x_j^{-1}$. All cycles of blue corners must therefore contain an even number of edges so we need only rule out the existence of blue cycles of combinatorial length two.

This is where the good words come in. Ensuring that there are no blue combinatorial paths of length two in the link corresponds precisely to requiring that the family of words $\{W_{ij+}, W_{ij-}, W_k$ for every $i, j, k\}$ has no 2−letter repetitions. The good word lemma provides us with a word W_J of length J^2 with this desirable property. If we choose J^2 such that it exceeds the total length of all words in the family $\{W_{ij+}, W_{ij-}, W_k$ for every $i, j, k\}$, we could then subdivide W_J to obtain a good set $\{W_{ij+}, W_{ij-}, W_k\}$ ensuring that the link condition holds. Concretely, this is done by choosing J such that is satisfies the following inequality

$$J^2 \geq 2(IJ)14 + \sum_{k=1}^{K} (2||R_k|| + 8).$$

□

Warning Groups containing finitely generated subgroups that are not finitely presented are called *incoherent*. In [Wis98], it is shown using this construction that as in Rips' result for groups satisfying small cancelation [Rip82], fundamental groups of negatively curved 2−complexes may be incoherent or have unsolvable generalized word problem.

5 Groups Acting on Metric Spaces

At this point we completely shift our point of view, no longer considering metric spaces intrinsically but using them instead to study groups acting on them. One of the first motifs in the geometric theory of groups is the idea that if we have a "nice" action of an infinite group on a metric space, structural properties of this space have strong implications on the structural properties of the group. We develop some of the tools of this theory, specializing from topological spaces to metric spaces only when necessary.

5.1 Basic Preliminaries

As usual, a group action of Γ on a space X is a homomorphism $\Gamma \to Homeo(X)$ if X is a topological space or $\Gamma \to Iso(X)$ if X is a metric space and will be denoted by $\Gamma \curvearrowright X$. The case of interest to us will be that of actions said to be *proper* and *cocompact*.

Definition An action $\Gamma \curvearrowright X$ is said to be

1. *Cocompact* if the quotient space $\Gamma \setminus X$ is compact. This is equivalent to saying that there is a compact $K \subseteq X$ such that $X = \Gamma \cdot K$.

2. *Proper* if small enough balls are properly moved. More precisely, if for every $x \in X$ there is a radius $r > 0$ such that the set $\{\gamma \in \Gamma | \gamma \cdot B(x,r) \cap B(x,r) \neq \emptyset\}$ is finite. The $\gamma \in \Gamma$ such that $\gamma \cdot U \cap U \neq \emptyset$ for some fixed open set $U \subset X$ are called $U - improper$ isometries.

3. *Geometric* if it is proper and cocompact on a metric space.

Although the above definition of a proper action is the weakest required criterion to prove the first few results presented below, one generally deals with actions in which every point has a neighbourhood mapped to disjoint sets by distinct group elements. Hatcher [Hat05] refers to these as a *covering space action* since groups acting in this way can be realized as the group of deck transformations of the covering $X \to \Gamma \setminus X$.

Lemma 5.1. *If a group Γ acts properly and cocompactly by isometries on a length*[23] *space X, then X is complete and locally compact.*

Proof. Let (x_n) be a Cauchy sequence in X. Since $\Gamma \curvearrowright X$ is proper, (x_n) can not eventually lie in a single orbit of Γ unless it is eventually constant. We may thus assume without loss of generality that consecutive points in our sequence (x_n) lie in distinct $\Gamma-$orbits. Since the action of Γ is cocompact, this sequence thus projects to a Cauchy sequence in the compact quotient $\Gamma \setminus X$ where it must converge to the orbit of some point $\overline{x} = \Gamma \cdot x$. Since the distance between points of (x_n) and the orbit $\Gamma \cdot x$ can be made arbitrarily small, it follows that (x_n) must converge to some point of the orbit. Our space X is therefore complete. On the other hand, since there is some compact $K \subseteq X$ such that $\Gamma \cdot K = X$ we also deduce that X is locally compact.

□

Corollary 5.2. *Applying the Hopf-Rinow Theorem for length spaces to the above lemma shows that X is in fact a proper*[24] *geodesic space.*

We now shift our attention to understanding groups themselves as geometric objects, a main theme in geometric group theory. A first step in this direction are the following three classical objects that one associates to a group, two of which we have already encountered without formally introducing them.

Cayley Graph. A group Γ generated by a set $\mathcal{A}$ can be obtained as a quotient of the free group on the alphabet $\mathcal{A}$, denoted by $F(\mathcal{A})$, through the natural surjection $F(\mathcal{A}) \xrightarrow{\varphi} \Gamma$. On the other hand, associated to every set of generators $\mathcal{A}$ is a (directed) graph whose vertices are the elements of Γ with an edge joining γ to γ' whenever $\gamma' = \gamma \cdot a$ for some $a \in \mathcal{A}$. The resulting graph is the *Cayley graph* of Γ with respect to the generating set $\mathcal{A}$, denoted by $\mathcal{C}_{\mathcal{A}}(\Gamma)$. This is a first striking manifestation of a group as a geometric object. Indeed, we can consider Γ as a metric graph $\mathcal{C}_{\mathcal{A}}(\Gamma)$ by setting all edges to have length equal to one. Equivalently, one can equip Γ as a set with the so called *word metric* $d_{\mathcal{A}}$

[23] See Appendix A for a definition.
[24] A metric space where closed balls are compact.

which measures the distance between γ_1 and γ_2 as the length of the shortest word in the pre-image of $\gamma_1^{-1}\gamma_2$ under the natural projection $F(\mathcal{A}) \to \Gamma$.

Cayley Complex. We now characterize a second geometric object that one can associate to a group. Consider any set of (reduced) words $\mathcal{R}$ contained in the kernel of the map $F(\mathcal{A}) \xrightarrow{\varphi} \Gamma$. Since a graph is equivalent to a one dimensional CW complex we may view the Cayley graph $\mathcal{C}_{\mathcal{A}}(\Gamma)$ as the 1−skeleton of a 2−complex $\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)$ which is obtained by attaching a 2−cells along each path in $\mathcal{C}_{\mathcal{A}}(\Gamma)$ labelled by a word contained in $\mathcal{R}$. Since the tree $\mathcal{C}_{\mathcal{A}}(F(\mathcal{A}))$ is the universal cover of $\mathcal{C}_{\mathcal{A}}(\Gamma)$ we have that $\pi_1(\mathcal{C}_{\mathcal{A}}(\Gamma)) \cong ker(\varphi)$. On the other hand, a classical application of the Seifert-van Kampen Theorem[25] shows that $\pi_1(\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)) \cong ker(\varphi)/N(\mathcal{R})$ where $N(\mathcal{R})$ is the normal closure of $\mathcal{R}$ in the free group $F(\mathcal{A})$. Recalling that whenever $N(\mathcal{R})$ coincides with the kernel of φ, the group Γ is given by the presentation $\langle\mathcal{A}|\mathcal{R}\rangle$ we have proved the following lemma.

Lemma 5.3. *The CW complex $\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)$ is simply connected $\iff$ $\Gamma \cong \langle\mathcal{A}|\mathcal{R}\rangle$.*

When either of the equivalent conditions given by the lemma hold we refer to $\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)$ as the *Cayley complex* of the presentation $\Gamma \cong \langle\mathcal{A}|\mathcal{R}\rangle$.

Standard 2-Complex. When we are given a presentation $\Gamma \cong \langle\mathcal{A}|\mathcal{R}\rangle$ there is a third natural geometric object associated to Γ. This is the so-called *standard 2-complex* of the presentation which is a CW complex $\mathfrak{S}_{\mathcal{A},\mathcal{R}}$ having a single 0−cell v, a one dimensional cells for every $a \in \mathcal{A}$ both of whose endpoints are attached to $\mathfrak{S}^{(0)}_{\mathcal{A},\mathcal{R}}v$ and a two dimensional cells for every $r \in \mathcal{R}$, each of which is attached to the graph $\mathfrak{S}^{(1)}_{\mathcal{A},\mathcal{R}}$ along the path labelled by the word it represents. A simple application of the Seifert van-Kampen Theorem shows that $\pi_1(\mathfrak{S}_{\mathcal{A},\mathcal{R}}) \cong \Gamma$ and in fact the Cayley complex $\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)$ is the universal cover of the standard 2−complex $\mathfrak{S}_{\mathcal{A},\mathcal{R}}$.

With these standard definitions at hand, we can move on to understanding the intricate relationship between a group and spaces it acts on. A neat application of the seemingly inoffensive Lemma 5.3 is the following surprising sequence of

[25]This theorem is stated in Appendix C. The reader very unfamiliar with these concepts could consult [LS77] or [Hat05].

results due to Murray MacBeath. Our proofs follow arguments of [BH99].

Lemma 5.4 ([Mac64])**.** *Suppose that X is a connected topological space and $\Gamma \curvearrowright X$ by homeomorphisms. If $U \subset X$ is a $\Gamma-$covering of X, in other words, if U is an open set such that $X = \Gamma \cdot U$ then the set of $U-$improper isometries $S := \{\gamma \in \Gamma | \gamma \cdot U \cap U \neq \emptyset\}$ generates Γ.*

Proof. Let $\langle S \rangle = H \leq \Gamma$ and denote $V := H \cdot U$ and $V' := (\Gamma \setminus H) \cdot U$. If $V \cap V' \neq \emptyset$, there must be some $h \in H$ and $h' \in \Gamma \setminus H$ such that $h' \cdot U \cap h \cdot U \neq \emptyset$. However, we then see that $h^{-1}h' \cdot U \cap U \neq \emptyset$ and so $h' \in HS \leq H$ which is a contradiction. Therefore, $V \cap V' = \emptyset$ and by connectedness of X, since V is nonempty we must have that $V' = \emptyset$. This shows that $H = \Gamma$. □

Theorem 5.5 ([Mac64])**.** *Suppose that Γ acts by homeomorphisms on a path connected and simply connected topological space X. If $U \subset X$ is a path-connected $\Gamma-$covering of X and $\mathcal{A}_S$ is an alphabet indexed by the set S of $U-$improper isometries of Γ then $\Gamma \cong \langle \mathcal{A}_S | \mathcal{R} \rangle$ where*

$$\mathcal{R} = \{a_{s_1} a_{s_2} a_{s_3}^{-1} | s_i \in S; U \cap s_1 \cdot U \cap s_3 \cdot U \neq \emptyset; s_1 s_2 = s_3 \text{ in } \Gamma\}.$$

Proof. Consider the CW complex $\mathfrak{C}_{\mathcal{A}_\mathcal{S},\mathcal{R}}(\Gamma)$ defined in Section 5.1. In light of Lemma 5.3, it suffices to show that $\mathfrak{C}_{\mathcal{A}_\mathcal{S},\mathcal{R}}(\Gamma)$ is simply connected to conclude that $\Gamma \cong \langle \mathcal{A}_S | \mathcal{R} \rangle$. Our approach to show this is by considering locally injective continuous maps $l : \partial D \to C_{\mathcal{A}_\mathcal{S}}(\Gamma)$ where D is the standard $2-$disk. If every such map can be continuously extended to a map $D \to \mathfrak{C}_{\mathcal{A}_\mathcal{S},\mathcal{R}}(\Gamma)$, it will then follow that $\mathfrak{C}_{\mathcal{A}_\mathcal{S},\mathcal{R}}(\Gamma)$ is simply connected.

Fix a base point $x_0 \in U$ and for each $s \in S$, choose a point $x_s \in U \cap s \cdot U$. Since U is path-connected, we may choose a path joining x_0 to x_s and another one joining x_s to $s \cdot x_0$. Labelling the concatenation of these paths by c_s, we obtain a continuous map $p : C_{\mathcal{A}_\mathcal{S}}(\Gamma) \to X$ as the extension of the map sending the identity $1 \mapsto x_0$ and edges emanating from the identity $a_s \mapsto c_s$. Note that since X is simply connected p actually extends to a Γ-equivariant continuous map $\mathfrak{C}_{\mathcal{A}_\mathcal{S},\mathcal{R}}(\Gamma) \to X$.

Suppose now that we have at hand a map $l : \partial D \to C_{\mathcal{A}_\mathcal{S}}(\Gamma)$ as described above. We can then consider the composition $p \circ l : \partial D \to X$ that we extend continuously (since X is simply connected) to $\Phi : D \to X$. Since D is compact and U is open, there is a finite triangulation τ of D such that given a vertex v of a triangle in τ, there is some element $\gamma_v \in \Gamma$ ensuring that Φ maps all triangles incident at v into the open set $\gamma_v \cdot U$.

Consider the triangle t with vertices labelled v_1, v_2 and v_3. By the above,

$$\emptyset \neq \Phi(t) \subseteq \gamma_{v_1} \cdot U \cap \gamma_{v_2} \cdot U \cap \gamma_{v_3} \cdot U$$

$$= \gamma_{v_1} \cdot (U \cap \gamma_{v_1}^{-1}\gamma_{v_2} \cdot U \cap \gamma_{v_1}^{-1}\gamma_{v_3} \cdot U)$$

$$= \gamma_{v_2} \cdot (\gamma_{v_2}^{-1}\gamma_{v_1} \cdot U \cap U \cap \gamma_{v_2}^{-1}\gamma_{v_3} \cdot U)$$

so labelling $s_1 := \gamma_{v_1}^{-1}\gamma_{v_2}$, $s_2 := \gamma_{v_2}^{-1}\gamma_{v_3}$ and $s_3 := \gamma_{v_1}^{-1}\gamma_{v_3}$ we have that $s_i \in S$ and $a_{s_1}a_{s_2}a_{s_3}^{-1} \in \mathcal{R}$.

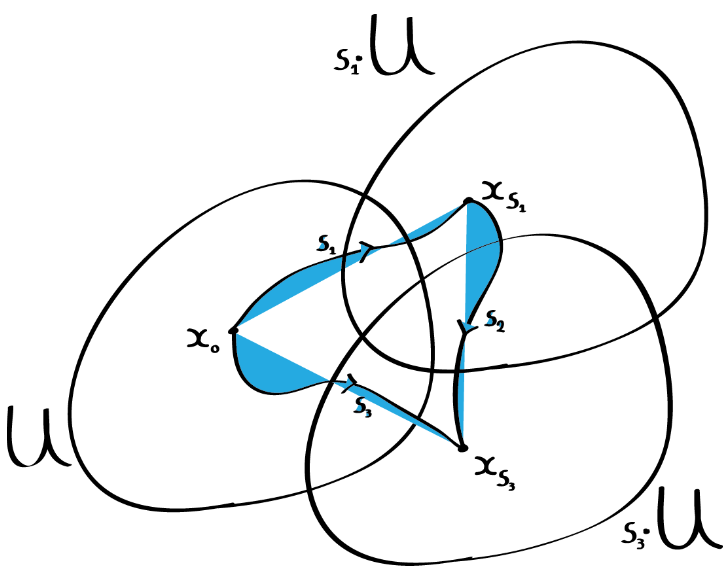

Figure 12: Presentation by generators and relations.

This observation allows us to extend the map $v \mapsto \gamma_v$ to a continuous function $\theta : 1-skeleton(\tau) \to C_{\mathcal{A}_\mathcal{S}}(\Gamma)$. Indeed, any $v \in \tau$ occurs in a triangle that we may label v_1, v_2 and v_3 so we can thus send $v_i \mapsto \gamma_{v_i}$ and edges $(v_i, v_{i+1}) \mapsto (\gamma_{v_i}, \gamma_{v_i} s_i)$ where indices are taken mod 3. To make everything compatible, if $v_i \in \partial D$ we choose γ_{v_i} to ensure that $\theta|_{\partial D}$ is a reparametrization of l.

To conclude, since the image $\theta(\partial t)$ is a labeled closed loop in $\mathcal{R} \subset C_{\mathcal{A}_S}(\Gamma)$, all such loops are boundaries of $2-$cells in $\mathfrak{C}_{\mathcal{A}_S,\mathcal{R}}(\Gamma)$. We may thus continuously extend θ to a map $D \to \mathfrak{C}_{\mathcal{A}_S,\mathcal{R}}(\Gamma)$ as desired.

□

While the above theorem is lovely, it is the following corollary that one uses in practise.

Corollary 5.6. Γ *is finitely presented* $\iff$ Γ *acts geometrically on a simply-connected geodesic metric space.*

Proof. Suppose Γ acts properly and cocompactly on a simply-connected geodesic space. Since the action $\Gamma \curvearrowright X$ is cocompact, there is some compact $K \subseteq X$ such that $\Gamma \cdot K = X$. As such, for $x_0 \in K$ and some $R > 0$, $K \subset B(x_0, R)$ and we may apply the preceding theorem with $U := B(x_0, R)$. However, by Corollary 5.2, X is proper so $\bar{U}$ is compact and Γ acts properly thus we are forced to conclude that the set $S = \{\gamma | \gamma \cdot U \cap U\}$ is finite.

Conversely, given a finitely presented $\Gamma = \langle \mathcal{A}_S | R \rangle$, the associated Cayley complex can be realized as a piecewise euclidean space in which we can embed $C_{\mathcal{A}_S}(\Gamma)$ metrized with all edges given length 1. In this case, the natural action of Γ on its Cayley graph extends to a proper cocompact action of the entire space.

□

5.2 Quasi-Isometries

In order to adequately deal with infinite groups acting on metric spaces one needs a notion that encapsulates the idea of spaces looking “essentially” the same when viewed from a distance. This is the case when there is a coarse, almost surjective, distance preserving map between them. It is with this in mind that the following definition loosens up the concept of Lipschitz maps.

Definition A map $f : X \to Y$ between metric spaces is a $(\lambda, \epsilon)-$*quasi isometric embedding* if there are constants $\lambda \geq 1$ and $\epsilon \geq 0$ such that $\forall x_1, x_2 \in X$:

$$\frac{1}{\lambda} d_X(x_1, x_2) - \epsilon \leq d_Y(f(x_1), f(x_2)) \leq \lambda d_X(x_1, x_2) + \epsilon.$$

When every point of Y lies in the $C-$neighbourhood of the image of f, we call the map a *quasi-isometry* and say that the two spaces are *quasi-isometric*. It is often the case that the constants ϵ and λ are irrelevant so we will coarsely refer to maps simply as quasi-isometries.

Remark A rather amusing instance of a quasi-isometry and its quasi-inverse is mentioned in [DK09] and arises by considering languages as "metric spaces" where the maps between them are translations. The german jewish name Schwarz was translated to russian at some point in the nineteenth century resulting in Шварц until finally the AMS decided to reconvert it to english as Švarc in the 1950′s. The sequence

$$\text{Schwarz} \to \text{Шварц} \to \text{Švarc}$$

could be informally thought about as a quasi-isometry followed by a quasi-inverse.

Definition Once a generating set $\mathcal{A}$ is chosen for a group Γ, we can define a *growth function* $\beta_{\mathcal{A}} : \mathbb{N} \to \mathbb{N}$ which associates to any integer n the number of elements in the closed ball $\overline{B(1, n)}$ about the identity in $\mathcal{C}_{\mathcal{A}}(\Gamma)$. It is a fact that $\beta_{\mathcal{A}}(n)$ is bounded by a polynomial in n for some generating set $\mathcal{A}$ if and only if it is polynomial in n for all generating sets. As such, when this is the case, we say that the group Γ has *polynomial growth*.

In view of the previous definition, we mention without proof the following striking theorem of Gromov.

Theorem 5.7 ([Gro81])**.** *A finitely generated group has polynomial growth if and only if it contains a nilpotent subgroup of finite index.*

Although Gromov's result was proved much later than the following proposition, it bears some relation to it. Indeed, the original motivation for the Schwarz-Milnor Lemma was to relate the growths of the volumes of balls in the universal cover of Riemannian manifolds X to the growth function of their fundamental group $\pi_1(X)$. It turns out that this gives rise to an interesting example of quasi-isometric embeddings of groups into spaces they act on. It was originally proved by the Russian school and later rediscovered by Milnor.

Proposition 5.8 ([Sva55],[Mil68])**.** *A group Γ acting geometrically on a length space X is finitely generated. Further, if we denote by $\mathcal{A}$ some generating set and consider $(\Gamma, d_{\mathcal{A}})$ as a metric space, the map $\gamma \mapsto \gamma \cdot x_0$ is a quasi-isometry.*

Proof. The following argument is inspired by [BH99]. Let the action $\Gamma \curvearrowright X$ be proper and cocompact an choose a compact $C \subseteq X$ such that $\Gamma \cdot C = X$. For any base point $x_0 \in X$, let $D > 0$ ensure that $B(x_0, D/3) \supset C$. Using the same argument as in the proof of Corollary 5.6, we see that Γ must be finitely generated by the set $\mathcal{A} := \{\gamma | \gamma \cdot B(x_0, D) \cap B(x_0, D) \neq \emptyset\}$.

Claim 5.9. *There is some constant $\mu > 0$ such that $d(\gamma \cdot x_0, \gamma' \cdot x_0) \leq \mu d_{\mathcal{A}}(\gamma, \gamma')$.*

Indeed, if $d_{\mathcal{A}}(\gamma, \gamma') = n$ then $\gamma^{-1}\gamma' = a_1 a_2 \ldots a_n$ for some $a_j \in \mathcal{A} \cup \mathcal{A}^{-1}$. Define $g_0 := 1$, $g_i := a_1 a_2 \ldots a_i$ and set $\mu := \max\{d(x_0, a \cdot x_0) | a \in \mathcal{A} \cup \mathcal{A}^{-1}\}$. By the triangle inequality

$$d(\gamma \cdot x_0, \gamma' \cdot x_o) = d(x_0, \gamma^{-1}\gamma' \cdot x_o) \leq \sum_{i=1}^{n} d(g_i \cdot x_0, g_{i+1} \cdot x_0)$$

and we are done since $d(g_i \cdot x_0, g_{i+1} \cdot x_0) = d(x_0, g_i^{-1} g_{i+1} \cdot x_0) = d(x_0, a_{i+1} \cdot x_0) \leq \mu$.

In view of the claim, all we have to show is that we can bound $d_{\mathcal{A}}(\gamma, \gamma')$ in terms of $d(\gamma \cdot x_0, \gamma' \cdot x_0)$. Since the action of Γ is by isometries in both metrics, it suffices to do so for $d_{\mathcal{A}}(1, \gamma)$ and $d(x_0, \gamma \cdot x_0)$. To this end, fix $\gamma \in \Gamma$ and let $c : [0, 1] \to X$ be a finite path from x_0 to $\gamma \cdot x_0$. Choose a coarsest partition of the interval $[0, 1]$ by $t_0 = 0 < t_1 < \ldots < t_{n-1} < t_n = 1$ subject to the constraint that $d(c(t_i), c(t_{i+1})) \leq D/3$. Since $\Gamma \cdot B(x_0, D/3) = X$, for every t_i there is some

γ_i such that $d(c(t_i), \gamma_i \cdot x_0) \leq D/3$. Extending the γ_i's to include $\gamma_0 := 1$ and $\gamma_n := \gamma$ we see that $d(\gamma_i \cdot x_0, \gamma_{i+1} \cdot x_0) \leq D$ ensuring that for every i, $\gamma_{i-1}^{-1}\gamma_i \in \mathcal{A}$. Labelling $a_i := \gamma_{i-1}^{-1}\gamma_i$ we obtain the following expression:

$$\gamma = \gamma_0(\gamma_0^{-1}\gamma_1)\dots(\gamma_{n-1}^{-1}\gamma_n) = a_1a_2\dots a_n.$$

Recalling that X is a length space, we may choose our curve c to have length $l \leq d(x_0, \gamma \cdot x_0) + 1$. It follows that $n \leq (d(x_0, \gamma \cdot x_0) + 1)(3/D) + 1$ since our partition into the t_i's was as coarse as possible. But since γ can be expressed as a word of length n we also obtain $d_{\mathcal{A}}(1, \gamma) \leq (d(x_0, \gamma \cdot x_0) + 1)(3/D) + 1$ and it follows that we have a quasi-isometry at hand.

□

Remark In general, a group Γ always acts by isometry on the metric space $(C_{\mathcal{A}}(\Gamma), d_{\mathcal{A}})$ by left multiplication. As such, it follows from the Schwarz-Milnor Lemma that the Cayley graphs $C_{\mathcal{A}}(\gamma)$ and $C_{\mathcal{A}'}(\Gamma)$ are quasi-isometric in the word metric. In fact, this shows that viewing a group as a metric space is well defined up to quasi-isometry.

This coarse geometric approach is quite strong as shown in this final theorem and the next section. We refer the reader to [BH99] for a proof.

Theorem 5.10. *If one of two quasi-isometric groups is finitely generated then so is the other one.*

5.3 Ends of a Space

Informally, the ends of a space may be thought of as the different ways in which one can topologically move to infinity. It is in fact a subset of the visual boundary for metric spaces, a concept we will study in Section 6.2. We proceed to give two equivalent formulations of this concept moving from the most abstract to most concrete.

Consider a compact $K \subseteq X$ and its complement $K^c := X \setminus K$. The $0-$th homotopy $\pi_0(K^c)$ counts the number of path-connected components of K^c so whenever $K_1 \subset K_2$ we get an induced map $\varphi_{i,j} : \pi_0(K_2^c) \to \pi_0(K_1^c)$. In view of this we have an inverse system of $\pi_0(K_i^c)$.

Definition The set of ends of X is defined as the inverse limit

$$\mathfrak{E}(X) := \lim_{\leftarrow i} \pi_0(K_i^c) = \{(r_i)_{i\in I} \in \prod_{i\in I} \pi_0(K_i^c) : \varphi_{i,j}(r_j) = r_i\}.$$

While this inverse limit may seem rather abstract, it can be worked out explicitly following the constructions in [Fre31]. In a general topological space X, we say that a sequence of points $x_n \to \infty$ goes to infinity if it eventually leaves any compact set. In the context of a parametrized path, this amounts to requiring that its points eventually leave any given compact set. We say that proper[26] rays $r_1, r_2 : [0, \infty) \to X$ converge to the same *end* if for every compact $K \subseteq X$, the image of both rays is eventually contained in the same path component of $X \setminus K$. This notion of convergence of rays defines an equivalence relation whose classes are denoted by $\mathfrak{e}(r)$. The set of such equivalence classes coincides with $\mathfrak{E}(X)$. To understand the topology on $\mathfrak{E}(X)$ we define convergence on ends $\mathfrak{e}(r_n) \to \mathfrak{e}(r)$ whenever given a compact $K \subseteq X$ there is a sequence of integers N_n and $m_k \in \mathbb{R}$ such that $r_n[N_n, \infty)$ and $r[N_n, \infty)$ lie in the same path-component of $X \setminus K$ for all $n > m_K$. The closed sets of $\mathfrak{E}(X)$ can now be defined as those containing all of their limit points.

Our first lemma requires the notion of a *k-path* joining points x and y of a metric space X. This is a collection of points $x = x_1, x_2, \dots x_n = y$ such that $d(x_i, x_{i+1}) \leq k$ for every i. One could say this definition is analogous to that of an $m-$string defined in Appendix B. We also find it convenient in what follows to denote by $\mathfrak{G}_{x_0}$ the set of geodesic rays based at a point $x_0 \in X$.

Lemma 5.11. *In a proper geodesic space X, there is a natural surjection $\mathfrak{G}_{x_0} \to$*

[26] A map between topological spaces is said to be *proper* if the inverse image of a compact set is compact.

$\mathfrak{E}(X)$ *for any choice of basepoint* $x_0 \in X$.

Proof. The following proof is inspired by [BH99]. First, notice that if r_1 and r_2 are proper geodesic rays in X, then $\mathfrak{e}(r_1) = \mathfrak{e}(r_2)$ if and only if for every $R > 0$ there is a $T > 0$ such that $r_1(t)$ and $r_2(t)$ are joined by a $k-$path in $X \setminus B(x_0, R)$ for every $t > T$. Indeed, given any compact $K \subset X$ there is a ball about x_0 containing it and vice versa. We may thus substitute balls for compact sets in the preceding discussion of $\mathfrak{E}(X)$. The forward implication is now clear from the definitions. To show the reverse direction, notice that if $x_1, \ldots, x_n$ is a $k-$path in $X \setminus B(x_0, R+k)$ then concatenating $[x_1, x_2]$, $[x_2, x_3]$,..., $[x_{n-1}, x_n]$ yields a path from x_1 to x_n in $X \setminus B(x_0, R)$ for any choices of geodesics. Choosing $T > 0$ to be large enough that $r_1(t)$ and $r_2(t)$ are not in $X \setminus B(x_0, R+k)$ for $t > T$ does the trick since then $r_1(t)$ and $r_2(t)$ must be in the same path component.

To complete the proof, consider a proper ray $r : [0, \infty) \to X$. Define a countable sequence of geodesic paths $c_n : [0, d_n] \to X$ joining x_0 to $r(n)$. We can extend each path c_n to be constant on (d_n, ∞). Now, since X is proper and the c_n are equi-continuous, by the Arzelà-Ascoli Theorem there is a subsequence of the c_n's uniformly convergent to a geodesic ray c. By construction, $\mathfrak{e}(c) = \mathfrak{e}(r)$ and we are done. □

Example Let T_3 be a rooted metric tree of valence 3 where the length of an edge separated from the root by n vertices is defined to be $1/2^{n+1}$. Here, $\mathfrak{E}(T_3)$ can be vividly pictured as the uncountably many infinite paths one can take starting at the root without backtracking. In fact, $\mathfrak{E}(T_3) \simeq C$ where C is the cantor set.

Indeed, recall that C can be characterized as the set of points in the interval $[0, 1]$ admitting a ternary expansion using only the digits 0 and 2 and, as such, $C \simeq \{0, 1\}^{\mathbb{N}}$. Since T_3 is connected, by the preceding lemma it is enough to consider all equivalence classes of locally injective proper rays based at the root to characterize $\mathfrak{E}(T_3)$. But then, we have a clear bijection between such equivalence classes and $\{0, 1\}^{\mathbb{N}}$ since $\mathfrak{e}(r)$ is uniquely characterized by the choice of one of the two possible directions at each vertex after the root. This correspondence yields a homeomorphism because the condition on convergent sequences of ends defining closed sets translates into a condition on convergent sequences of points yielding a bijection between closed sets in both spaces.

It is interesting to note that, by construction, T_3 embeds inside the unit disk as shown in Figure 13.

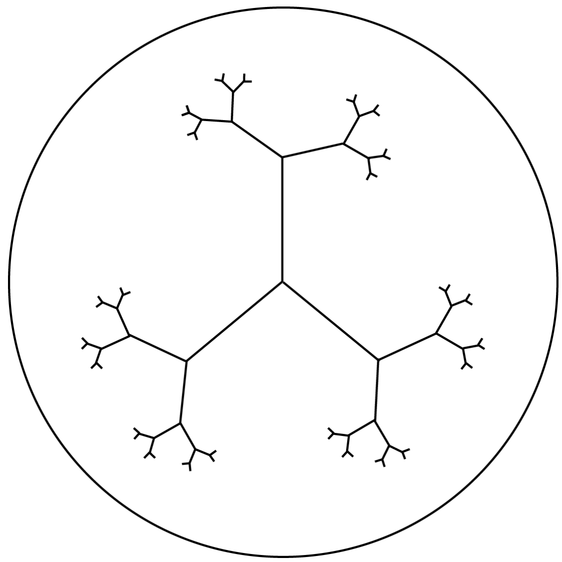

Figure 13: T_3

In light of this new concept, we also obtain another well known quasi-isometry invariant. Indeed, suppose that we have a quasi-isometry $f : X \to Y$ between proper geodesic spaces. Given any proper ray r in X, one can consider the proper ray

$$f_*(r) := [f(r(0)), f(r(1))], [f(r(1)), f(r(2))], \ldots, [f(r(n)), f(r(n+1))], \ldots$$

obtained by concatenating geodesic segments in Y. Since f is a quasi-isometry, the map $f_{\mathfrak{e}} : \mathfrak{E}(X) \to \mathfrak{E}(Y)$ given by $\mathfrak{e}(r) \mapsto \mathfrak{e}(f_*(r))$ is a well defined homeomorphism.

Proposition 5.12. *Every quasi-isometry between proper geodesic spaces induces a homeomorphism between their ends.*

Applying these constructions to the case of groups, we can "unambiguously" define the *ends of a group* Γ as

$$\mathfrak{E}(\Gamma) := \mathfrak{E}(C_{\mathcal{A}}(\Gamma)).$$

Whenever $\mathcal{A}$ is finite, this space is fairly well understood. The following landmark pair of theorems are due to Hopf and Stallings.

Theorem 5.13 ([Hop43])**.** *If* Γ *is a finitely generated group then* $\mathfrak{E}(\Gamma)$ *is compact and has either* $0, 1, 2$ *or infinitely many elements. In the finite case,* $\mathfrak{E}(\Gamma)$ *has* 0 *elements if and only if* Γ *is finite and two elements if and only if* Γ *contains* $\mathbb{Z}$ *as a finite index subgroup.*

Proof. We only prove that if $\mathfrak{E}(\Gamma)$ is finite it has 0, 1 or 2 elements following an argument found in [BH99]. Let $\mathcal{C}_{\mathcal{A}}(\Gamma)$ be the Cayley graph of Γ with respect to some finite generating set. The action of Γ on itself by left multiplication extends to an action of Γ on C by (quasi) isometries. By Lemma 5.12 we thus obtain a homomorphism $\Gamma \to Homeo(\mathfrak{E}(C))$.

Assume that $\mathfrak{E}(C)$ is finite. Since Γ is infinite, the above homomorphism has a nontrivial kernel, H which must have finite index in Γ. Suppose for the contrapositive that we have three distinct ends e_0, e_1 and e_2. Let r_1 and r_2 be proper geodesic rays such that $r_1(0) = r_2(0)$ and $\mathfrak{e}(r_i) = e_i$. As $[\Gamma : H] < \infty$, there must be some positive constant μ with the property that every vertex of C lies in the μ neighbourhood of some element of H. This ensures the existence of a proper ray r_0 with the property that $\mathfrak{e}(r_0) = e_0$, $d(r_0(n), 1) \geq n$ and $r_0(n) \in H$ for every $n \in \mathbb{N}$.

Define $\gamma_n := r_0(n)$ and let $\rho > 0$ be such that $r_1([\rho, \infty))$ and $r_2([\rho, \infty))$ lie in different path components of $C \setminus B(1, \rho)$. It follows that for $t, t' > 2\rho$,

$$d(r_1(t), r_2(t) > 2\rho$$

since the path joining the two points must cross through $B(1, \rho)$. Notice now that $\gamma_n \in H$ implies that $\mathfrak{e}(\gamma_n \cdot r_i) = \mathfrak{e}(r_i)$ for $i = 1, 2$.

Fix $n > 3\rho$ and let i range over $1, 2$. We must then have that $\gamma_n \cdot r_i(0) = \gamma_n$ which lies in a different path component of $C \setminus B(1, \rho)$ than $r_i([\rho, \infty))$. Since translation by γ_n doesn't effect ends and $d(\gamma_n, 1) \geq n > 3\rho$, $\gamma_n \cdot r_i$ must pass through $B(1, \rho)$ for some $t > 2\rho$ as in Figure 14. When $t > 2\rho$ is such that $\gamma_n \cdot r_i$ lies in $B(1, \rho)$ for $i = 1, 2$, their distance is bounded by the ball so

$$d(\gamma_n \cdot r_1(t), \gamma_n \cdot r_2(t)) \leq 2\rho.$$

Since γ_n is an isometry, this contradicts the first inequality established above. It follows that C can only have $0, 1$ or 2 ends. □

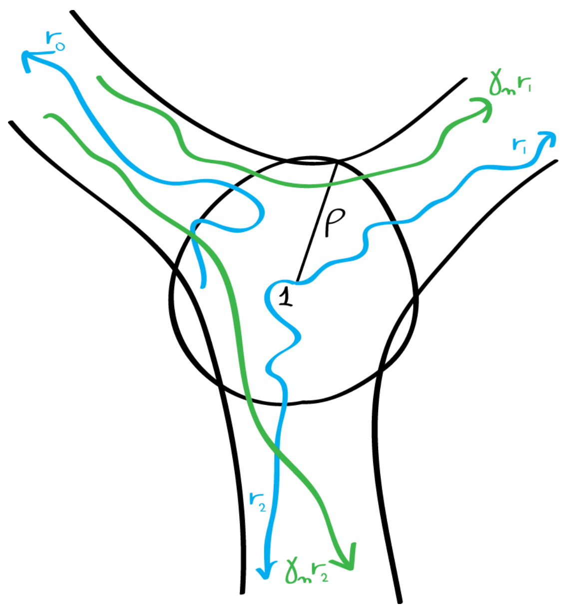

Figure 14: Three ends

We conclude by mentioning without proof the following theorem of Stallings.

Theorem 5.14 ([Sta68])**.** *Γ has infinitely many ends if and only if it splits as an amalgamated product*[27] *$\Gamma_1 *_C \Gamma_2$ or an HNN extension $\Gamma_1 *_C$ where C is a finite group with $[\Gamma_1 : C] \geq 3$ and $[\Gamma_2 : C] \geq 2$.*

[27]Please see Appendix C for the relevant definitions.

6 Groups Acting on $CAT(0)$ Spaces

Having set up the basic preliminaries concerning group actions, we now return to the realm of metric spaces of bounded curvature. It turns out that groups acting on $CAT(0)$ spaces by isometries share many of the elegant properties of isometry groups of non-positively curved Riemannian manifold.

6.1 Types of Isometries

Associated to any isometry σ of a $CAT(0)$ space X is the so called *displacement function*, $d_\sigma : X \to \mathbb{R}_{\geq 0}$ defined by the rule $d_\sigma(x) := d(\sigma \cdot x, x)$ and *translation length* $\delta_\sigma := \inf\{d_\sigma(x) : x \in X\}$. The set of points at which this translation length is attained is denoted by $Min(\sigma) := \{x : d_\sigma(x) = \delta_\sigma\}$. This notation will be used heavily throughout the other sections.

Remark It should be emphasized that $Min(\sigma)$ is convex since if two points x and y belong to $Min(\sigma)$ the geodesic segment $[x, y]$ must also belong to $Min(\sigma)$ because isometries preserve distances.

In analogy with the classification of isometries of riemannian manifolds we say that an isometry of a $CAT(0)$ space is *Elliptic* if it has a fixed point, *Hyperbolic* if d_σ attains a strictly positive minimum and *Parabolic* if $Min(\sigma) = \emptyset$.

Definition Isometries that fall into the first two categories or those where $Min(\sigma) \neq \emptyset$ are *semi-simple*. The parabolic isometries are in many aspects the "rogue" ones.

Example

1. Isometries of $\mathbb{E}^n$ are all of the form $x \mapsto Ax + b$ where $A \in O(n)$ and $b \in \mathbb{R}^n$ so they are all semi-simple.
2. Isometries of $\mathbb{H}^n$ are not necessarily semi-simple. Consider a rotation about a point at infinity in the Poincarré disc model of $\mathbb{H}^2$.

We now proceed to develop the structure intrinsic to metric spaces admitting semi-simple isometries of either type. Our first tool is the fact that the *circumcentre* of a bounded subset of $CAT(0)$ spaces is well defined. Recall that the *radius* of a bounded set Y, $\mathfrak{r}_Y := \inf_{r \in \mathbb{R}}\{Y \subseteq B(x, r) : x \in X\}$.

Proposition 6.1 ([Bro89]). *Suppose that X is a complete $CAT(0)$ space and that $Z \subseteq X$ is a bounded set of radius $\mathfrak{r}_Z$. There is a unique point $\mathfrak{c}_Z \in X$ such that $Y \subseteq \overline{B}(\mathfrak{c}_Z, \mathfrak{r}_Z)$. This is the circumcentre of Z.*

Proof. Let us consider a sequence of points $\mathfrak{c}_n \in X$ such that $d(\mathfrak{c}_n, Z) \xrightarrow{n\to\infty} \mathfrak{r}_Z$. An application of the $CN-$inequality[28] to a fixed pair of points $\mathfrak{c}_n, \mathfrak{c}_m \in X$ along with an arbitrary $z \in Z$ implies that

$$d^2(\mathfrak{c}_n, \mathfrak{c}_m) \leq 2 \cdot \left[d^2(\mathfrak{c}_n, Z) + d^2(\mathfrak{c}_m, Z) - 2 \cdot \mathfrak{r}_Z^2\right] \tag{3}$$

so the sequence $(\mathfrak{c}_n)$ is cauchy and converges to a circumcentre $\mathfrak{c}_Z$. On the other hand, equation (3) holds for arbitrary pairs of points in X so a circumcentre is unique. □

The idea of using the circumcentre of a set acted upon by a group to find a fixed point is central; for instance, it was applied by Élie Cartan in the context of Lie groups in his *Leçons sur la géométrie des espaces de Riemann* and in the case of groups acting on Euclidean buildings by Bruhat and Tits in [BT72]. This has the following consequences for groups with bounded orbits in $CAT(0)$ spaces:

Lemma 6.2 (Detecting Semi-Simple Isometries). *If a group acts on a complete $CAT(0)$ space with a bounded orbit then the action necessarily has a fixed point. Therefore, an isometry of a complete $CAT(0)$ space X with a bounded orbit is elliptic. In fact, if the isometry σ^n is elliptic for $n \neq 0$ then σ is elliptic. Similarly, if the isometry σ^n is hyperbolic for $n \neq 0$ then σ is hyperbolic.*

Proof. The circumcentre of the bounded orbit is well defined and preserved by the action of the group so it is a fixed point. If σ^n fixes a point $x \in X$ then the orbit of x under the action of σ has at most n points. We refer the reader to [BH99] for a proof of the last statement which follows more easily from Theorem 6.3. □

[28] Lemma 2.2.

While hyperbolic isometries do not fix points, they stabilize their higher dimension analogues, a family of geodesic lines. The following theorem is crucial for understanding hyperbolic isometries.

Theorem 6.3 (Axes of Hyperbolic Isometries)**.** *If X is a $CAT(0)$ space then σ is hyperbolic if and only if there are geodesic lines stabilized by σ where it operates by translation of length $\delta_\sigma > 0$. These geodesic lines are parallel and their union coincides with $Min(\sigma) \cong Y \times \mathbb{R}$ where $\sigma(y,t) = (y, t + \delta_\sigma)$.*

Definition Geodesic lines on which a hyperbolic isometry acts by non-trivial translation are called its *axes*.

Proof. The following argument is inspired by [BH99]. Suppose that σ stabilizes a geodesic line γ where it operates by non-trivial translation. By definition, γ is a complete convex set so we have $\delta_\sigma = \delta_{\sigma|_\gamma}$ where $\sigma|_\gamma$ is the restriction of σ to γ. Therefore, if σ acts by non-trivial translation on γ, it is necessarily hyperbolic. Suppose conversely that σ is hyperbolic. We claim that the union over all integers $n \in \mathbb{Z}$ of the geodesic segments $[\sigma^n x, \sigma^n + 1]$ is an axis for σ. A first step in this direction is to show that

$$[x, \sigma x] \cup [\sigma x, \sigma^2 x] = [x, \sigma^2 x]. \tag{4}$$

Denoting by $\frac{x+\sigma x}{2}$ the midpoint of the segment $[x, \sigma x]$, equation (4) is equivalent to showing that $d(\frac{x+\sigma x}{2}, \sigma(\frac{x+\sigma x}{2})) = 2d(x, \frac{x+\sigma x}{2})$. Recall that $Min(\sigma)$ is a convex subset of X so $\frac{x+\sigma x}{2} \in Min(\sigma)$ and consequently $d(\frac{x+\sigma x}{2}, \sigma(\frac{x+\sigma x}{2})) = d(x, \sigma x) = 2d(x, \frac{x+\sigma x}{2})$ where the last equality follows by definition of the midpoint. The result now follows because local geodesics are geodesics in $CAT(0)$ spaces, see Proposition 2.1.

Let us now fix parametrizations $r, r' : \mathbb{R} \to X$ for two different axes of σ. We have already shown that $\sigma r(t) = r(t+\delta_\sigma)$ while $\sigma r'(t) = r'(t+\delta_\sigma)$ so $d(r(t), r'(t)) = d(r(t+\delta_\sigma), r'(t+\delta_\sigma))$. As such, the distance function between the two geodesics is periodic which implies that it is bounded. Recalling that the $CAT(0)$ metric is convex allows us to deduce that the axes lie a constant distance apart so they are parallel. Finally, since $Min(\sigma)$ is convex and axes are parallel, applying Theorem 2.10 we obtain the decomposition $Min(\sigma) \cong Y \times \mathbb{R}$.

□

The technique to produce an axis in the preceding theorem should be kept in mind during the proof of the main theorem in Section 6.3. Some classes of hyperbolic isometries have even stronger implications on the structure of a metric space and deserve a special name.

Definition A *Clifford translation* is an isometry σ or a metric space X such that $Min(\sigma) = X$. In other words, they are those isometries which shift every point of x by an equal distance. Simple examples of spaces with many Clifford translations are Banach spaces. These special isometries admit a surprising amount of additional structure and the existence of non-trivial Clifford translation has strong consequences on the properties of a metric space. In fact, we shall have to rull them out in Section 7 to establish superrigidity results.

Let $\sigma \neq 1$ and τ be Clifford translations. By virtue of Theorem 6.3 there is a splitting $X = Y \times \mathbb{R}$ such that $\sigma(y, r) = (y, r + \delta_\sigma)$ so we may write $\sigma = 1 \times t_\sigma$ where t_σ denotes translation by δ_σ. The action of τ on $Y \times \mathbb{R}$ must then be of the form $\tau' \times t_\tau$ where τ' is a Clifford translation of Y so $\tau\sigma = \sigma\tau = \tau' \times t_\tau t_\sigma$ and we see that the Clifford isometries form an abelian group $\mathbf{H}$.

We can now define an action of $\mathbb{R}$ on $\mathbf{H}$ by letting $\lambda \in \mathbb{R}$ send the isometry σ as above which maps $(y, r) \mapsto (y, r + \delta_\sigma)$ to the isometry $(y, r) \mapsto (y, r + \lambda\delta_\sigma)$. Further, the norm on $\mathbf{H}$ which maps σ to δ_σ satisfies the parallelogram law so $\mathbf{H}$ is a vector space equipped with an inner product. In fact, the following theorem proved in [BH99] holds.

Theorem 6.4. *The group of Clifford translations* $\mathbf{H}$ *of a complete* $CAT(0)$ *space* X *is a Hilbert space. Further,* X *splits as a product* $Y \times \mathbf{H}$ *preserved by* $Iso(X)$.

Without providing the proof, let us simply mention how the above splitting arises. Since X is complete, its closed subsets are also complete so the projection[29]

[29] Defined in Proposition 2.6.

$\pi : X \to \mathbf{H} \cdot x_0$ onto the closed convex orbit $\mathbf{H} \cdot x_0$ of a fixed point $x_0 \in X$ is well defined. We can then consider the splitting $\pi^{-1}(x_0) \times \mathbf{H}$ and the map $\theta : \mathbf{H} \times \pi^{-1}(x_0) \to X$ defined by the rule $(\sigma, \xi) \mapsto \sigma(\xi)$ is an isometry. It should be mentioned that $\mathbf{H}$ is analogous to the Euclidean de Rham factor in Riemannian geometry.

6.2 Visual Boundary and Bordification

The bordification of a complete $CAT(0)$ metric space encapsulates in a sense what the space looks like from the point of view of an internal observer. It is obtained by attaching to a space its visual boundary which consists of the set of points at infinity and topologizing this union in a coherent way.

The following constructions emulate in a sense the study of the geometry of geodesics in non-positively curved spaces by Hadamard in [Had98] and Busemann in [Bus55]. However, the first definition of a visual boundary for arbitrary metric spaces was given in [EO73] and we give an account similar to [BH99]. One could also consult [Hot97] or [Pap05] for the perspective of Busemann non-positively curved spaces. The same ideas are applied in the context of hyperbolic groups by Gromov in [Gro87].

Throughout this section we assume that X is a complete $CAT(0)$ space, also called a Hadamard space.

Recall that in a $CAT(0)$ space, two geodesic rays are parallel[30] if and only if they are asymptotic. It is the existence and uniqueness of parallel rays emanating from different points of a metric space that allow us to characterize its boundary at infinity.

Proposition 6.5 ([BH99])**.** *Given two points x and y in a complete $CAT(0)$ space X, once a geodesic ray based at x is specified, there is a unique parallel geodesic ray based at y.*

[30] See Theorem 2.9.

Sketch of the Proof. Suppose that we are given x and y as in the statement and write $r : [0, \infty[\to X$ for the geodesic ray based at x. Given any $n \in \mathbb{N}$ let γ_n be a parametrization of the unique geodesic segment $[y, r(n)]$ joining y to the point $r(n)$. For any fixed $s \in [0, 1]$, we have a sequence of points $\gamma_n(s)$ which is Cauchy by an application of the $CAT(0)$ inequality. Since X is complete, we have that $\gamma_n(s) \to r'(s)$ for some unique point $r'(s) \in X$ and we can define $r' : [0, \infty[\to X$ as the pointwise limit of the geodesics $s \to r'(s)$. The ray r' is unique because geodesic rays are parallel if and only if they are a bounded distance apart. If $r'' : [0, \infty[\to X$ was another ray emanating from y parallel to r, its initial distance from r' would be zero implying $r' = r''$.

□

In view of this proposition, we can unambiguously define the *visual boundary* ∂X as the set of geodesic rays based at some point x_0 modulo parallelism. A good example to keep in mind is the disc model for the hyperbolic plane $\mathbb{H}^2$ where $\partial\mathbb{H}^2$ coincides with the boundary of the model disc. The main purpose for this definition is to consider the "closure at infinity" $\overline{X} = X \cup \partial X$ which contains the missing fixed point of parabolic isometries. This new space $\overline{X}$ is called the *bordification* of X. Heuristically, $\overline{X}$ corresponds to what an individual travelling within X would experience it as. So far, these spaces only consist of a set of points so we need a useful topology on it. The following realization of $\overline{X}$ from [EO73] allows one to elegantly define the so called *cone topology.*

Once a basepoint $x_0 \in X$ has been fixed, there is a system of closed balls $\overline{B_\epsilon} := \overline{B(x_0, \epsilon)}$ where $0 < \epsilon < \infty$. Since closed balls in $CAT(0)$ spaces are bounded, complete and convex, we have a projection $\pi_\epsilon : X \to \overline{B_\epsilon}$ as defined in Proposition 2.6. Concretely, if $x \in X$, there is a unique segment $[x_0, x]$ and if $[x_0, x] \cap \overline{B_\epsilon} = [x_0, x']$ then $\pi_\epsilon(x) = x'$. Restricting this map to closed balls yields projections $\pi_{\epsilon,\epsilon'} : \overline{B_\epsilon} \to \overline{B_{\epsilon'}}$ wherever $\epsilon \geq \epsilon'$ so the closed balls form an inverse system and we may take its inverse limit

$$\lim_{\substack{\longleftarrow \\ \epsilon}} \overline{B_\epsilon} = \{(r_\epsilon)_{\epsilon \in [0,\infty[} \in \prod_{\epsilon \in [0,\infty[} \overline{B_\epsilon} : \pi_{\epsilon,\epsilon'}(r_\epsilon) = r_{\epsilon'}\}.$$

We equip it with the inverse limit topology, namely the one making all maps $\lim_{\leftarrow_\epsilon} \overline{B_\epsilon} \to \overline{B_\epsilon}$ continuous. Concretely, $\lim_{\leftarrow_\epsilon} \overline{B_\epsilon}$ correspond to the set of geodesics[31] $r : [0, \infty[\to X$ based at x_0 such that $\pi_{\epsilon,\epsilon'}(r(\epsilon)) = r(\epsilon')$ equipped with the topology of uniform convergence on compact sets. In other words, a sequence of geodesics r_n in the inverse limit converges to a ray r in $\overline{X}$ if for all compact sets $K \subset [0, \infty[$ the restrictions $r_n|_K$ converge uniformly to $r|_K$ as functions.

There is a natural identification $\eta : \overline{X} \to \lim_{\leftarrow_\epsilon} \overline{B_\epsilon}$ sending points $x \in X$ to the geodesic $r : [0, \infty[\to X$ who's image is the segment $[x_0, x]$ and sending points $\xi \in \partial X$ to the unique geodesic ray $r : [0, \infty[\to X$ based at x_0 in the equivalence class of ξ. This identification yields the desired *cone topology* on $\overline{X}$ and the inclusion of X into $\overline{X}$ under this topology is a homeomorphism onto a dense set.

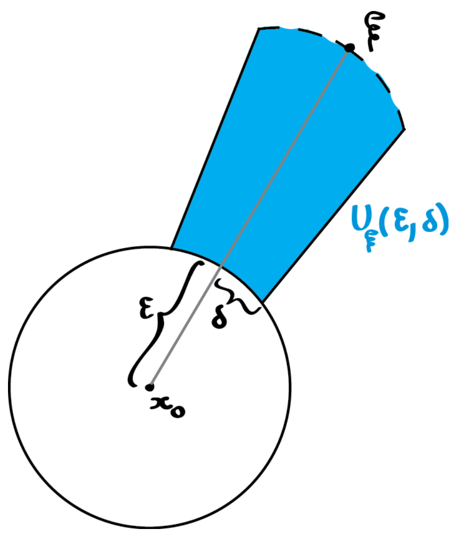

Figure 15: A neighbourhood of a point at infinity.

A convenient basis for the topology on $\overline{X}$ is given by two families of sets. First we have open balls about points in $x \in X$. Then, if $\xi \in \partial X$ corresponds in the inverse limit construction to the geodesic ray $r : [0, \infty[\to X$ parametrized such that $r([0, \epsilon]) = r([0, \epsilon]) \cap \overline{B(x_0, \epsilon)}$, a basis[32] for its neighbourhoods is given by the

[31] The maps are either geodesic rays or for some $t > 0$, $r(t') = r(t)$ for all $t' \geq t$.

[32] This follows by general considerations of the compact-open topology as explicated in

sets of the form $U_\xi(\epsilon,\delta) := \{x \in \overline{X} : d(x,x_0) > \epsilon, d(\pi_\epsilon(x), r(\epsilon)) < \delta\}$ as shown in Figure 15.

Proposition 6.6 ([BH99])**.** *The cone topology on $\overline{X} = X \cup \partial X$ doesn't depend on the choice of the basepoint $x_0 \in X$.*

Sketch of Proof. Let $\overline{X}$ be topologized as above with respect to the basepoint $y_0 \neq x_0$. The cone topology is induced from the inverse limit topology so it suffices to show that the projections $\pi_\epsilon^{x_0} : \overline{X} \to \overline{B(x_0,\epsilon)}$ are continuous. This holds for any fixed $\epsilon > 0$ if for all $\theta \in \overline{X}$ and for every neighbourhood V of $\pi_\epsilon^{x_0}(\theta)$ there is a neighbourhood U of θ such that $\pi_\epsilon^{x_0}(U) \subset V$. If $\theta \in X$, this is clear but if $\theta \in \partial X$ it requires some work.

By definition, θ correspond uniquely to two geodesic ray $r_{y_0}, r_{x_0} : [0,\infty[\to X$ based at y_0 and x_0 respectively and parametrized in a way that for $\lambda \in \mathbb{R}_{\geq 0}$, $r_\cdot([0,\lambda]) = r_\cdot([0,\lambda]) \cap \overline{B(\cdot,\lambda)}$. As such, we will abuse notation and write statements such as $\pi_\epsilon^{x_0}(r_{x_0})$ instead of $\pi_\epsilon^{x_0}(\theta)$ in what follows. It suffices to check the continuity condition on basis elements so suppose we are given an open ball of radius δ about $\pi_\epsilon^{x_0}(\theta)$. For every $\lambda > 0$ there is a basis element of the cone topology (induced by the basepoint y_0) on $\overline{X}$ of the form

$$U_\theta(\lambda, \delta/3) = \{x \in \overline{X} : d(x,y_0) > \lambda, d(\pi_\lambda^{y_0}(x), r_{y_0}(\lambda)) < \delta/3\}.$$

All that remains to be verified is that if $x \in U_\theta(\lambda,\delta/3)$ then $d(\pi_\epsilon^{x_0}(x), \pi_\epsilon^{x_0}(\theta)) = d(\pi_\epsilon^{x_0}(x), \pi_\epsilon^{x_0}(r_{x_0})) < \delta$. But using the triangle inequality we see that

$$d(\pi_\epsilon^{x_0}(x), \pi_\epsilon^{x_0}(r_{x_0})) \leq$$

$$d(\pi_\epsilon^{x_0}(x), \pi_\epsilon^{x_0}\pi_\lambda^{y_0}(x)) + d(\pi_\epsilon^{x_0}\pi_\lambda^{y_0}(x), \pi_\epsilon^{x_0}\pi_\lambda^{y_0}(r_{y_0})) + d(\pi_\epsilon^{x_0}\pi_\lambda^{y_0}(r_{y_0}), \pi_\epsilon^{x_0}(r_{x_0}))$$

where each summand on the right hand side is bounded above by $\delta/3$. For the first and last summands this is a consequence of the proof of Proposition 6.5 while for the middle summand it follows directly from the definition of $U_\theta(\lambda,\delta/3)$. The

[Mun00].

situation is heuristically illustrated in Figure 16 for the convenience of the reader. □

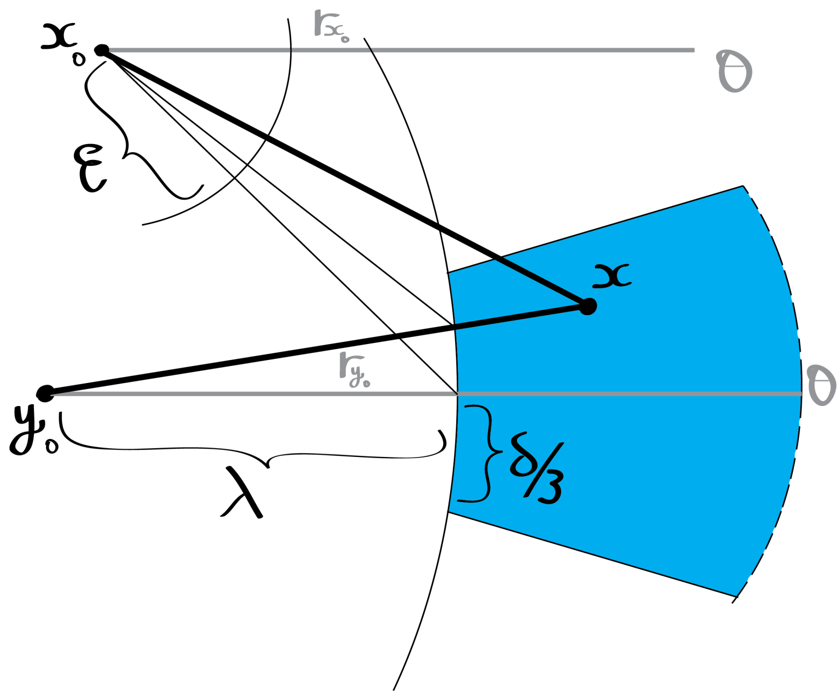

Figure 16: Visualizing the triangle inequality.

Remarks 1. The independence of the cone topology from the choice of a basepoint should be a hint that there is a functorial description lurking nearby. Indeed, it is shown in [BH99] following a construction of Gromov that if we denote the complete space of real valued continuous functions on X modulo additive constants by C_*X, then one can describe $\overline{X}$ as the closure of the inclusion $X \hookrightarrow C_*X$ sending a point x to the map $y \mapsto d(x, y)$.

2. It should be noted that the visual boundary continuously surjects onto the ends of a metric space as defined in Section 5.3.

Now, if $\sigma : X \to X$ is an isometry, then σ sends rays r based at a point x to rays $\sigma(r)$ based at the point $\sigma(x)$ so we obtain the following corollary.

Corollary 6.7. *Let X be a complete $CAT(0)$ space. For every isometry $\sigma \in Iso(X)$, there is a natural extension $\overline{\sigma}$ to $\overline{X}$ which is a homeomorphism.*

Example Recall from a previous example that a rotation $\sigma \in Iso(\mathbb{H}^2)$ about a point at infinity in the hyperbolic plane is parabolic since it rotates all geodesic lines and has no fixed point in $\mathbb{H}^2$. However, if we consider $\overline{\sigma} \in Homeo(\overline{\mathbb{H}^2})$ it has a fixed point in $\partial\mathbb{H}^2$.

Remark We have equipped the visual boundary with the cone topology but since it was built from a metric space, one might ask if there is also a natural way to metrize it. There are in fact several ways to do this, the most intuitive of which is perhaps to define the distance between ξ_1 and ξ_2 in ∂X as the supremum over all points $x \in X$ of the angle $\angle_x(\xi_1, \xi_2)$. One can then endow the visual boundary with the so-called *Tits metric* which is the length metric d induced by this angular distance. It turns out that the resulting space called the Tits boundary $\partial_T X$ encodes the geometry of flats in X. One should be warned however that in general $\partial_T X$ is not homeomorphic to ∂X with the cone topology.

6.3 Isometries of Cube Complexes

In this section, we illustrate an aspect in which $CAT(0)$ cube complexes behave analogously to trees. More particularly, we give an account of Haglund's result that automorphisms of $CAT(0)$ cube complexes are semi-simple (up to cubical subdivision). All of the results proved here come from [Hag07]. We invite the reader to review the contents of Section 3.2 for a definition of non-positively curved cube complexes and preliminary results.

All distances expressed in this section are combinatorial, as such, all vertices of cube complexes considered lie in the $0-$skeleton.

Due to the combinatorial nature of the metric, slight modifications must be imposed on the definitions pertaining to isometries introduced in the Section 6.1. To this end, we clarify that given a cube complex X we will be working with *automorphisms* $\sigma : X \to X$ which are the bijections sending $n-$cubes to $n-$cubes. It should be emphasized that an automorphism is necessarily an isometry in both the combinatorial and $CAT(0)$ metrics. As such, they are often called combinatorial or cellular isometries in the literature.

Given an automorphism σ of a cube complex X we have, as before, the associated displacement function $d_\sigma(x) = d(x, \sigma(x))$ along with the translation length $\delta_\sigma = \inf_{x \in X^{(0)}} d_\sigma(x)$ where the only difference is that d now represents the combinatorial distance and that the infimum is taken over vertices instead of arbitrary points. Recall that in the general case, an isometry f of a $CAT(0)$ space was said to be semi-simple if it was elliptic (had a fixed point) or hyperbolic (had no fixed point but its displacement function attained the translation length) and in the latter case f necessarily acted by its translation length on an infinite geodesic called an axis. Analogously, we say that an automorphism $\sigma \in Aut(X)$ of a cube complex X is (combinatorially) *semi-simple* if either it fixes a vertex $v \in X^{(0)}$ or it acts by translations of length δ_σ on a set of infinite combinatorial geodesics that we will also call its *axes*. The key technical obstruction to showing that all automorphisms are semi-simple is that some automorphisms act as glide reflections along certain hyperplanes. This is illustrated in the following definition.

Stable Action Let $\mathcal{H}$ be a hyperplane in a cube complex X and recall that such a hyperplane separates X in two connected components. We say that $\sigma \in Aut(X)$ has an inversion along $\mathcal{H}$ if it interchanges these two components. The automorphism σ is then said to act *stably* if for any $n \geq 0$, σ^n does not have an inversion along any hyperplane.

For instance, if one considers the cube complex consisting of a single euclidean square and a rotation σ of $\pi/2$ then σ does not have an inversion but σ^2 does hence the action is not stable. In fact, the action is not semi-simple since its fixed point doesn't lie in the $0-$skeleton and there are no infinite geodesics. Fortunately, the requirement to eliminate such pathologies is not really an obstruction. Indeed, one can readily pass to a cubical subdivision which is defined analogously to the barycentric subdivision[33] in the case of polyhedral complexes but where the sub-cells are cubes instead of simplices as shown in Figure 17. For the square mentioned above, the subdivision results in four squares with the fixed point now lying in the $0-$skeleton implying semi-simplicity. The following lemma is straightforward.

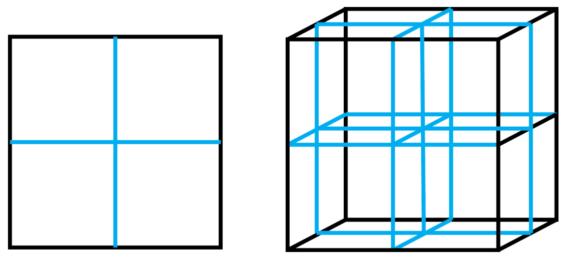

Figure 17: The cubical subdivision of a two and three dimensional cube.

Lemma 6.8 ([Hag07]). *If σ is an automorphism of the cube complex X, it acts stably on the cubical subdivision of X.*

Working towards the main result, recall that we can view a (connected and undirected) graph $\mathcal{G}$ as a one dimensional cube complex in which case an automorphism is a bijection of its vertices which preserves adjacencies. The following abstract result for graphs will be used to understand automorphisms of general $CAT(0)$ cube complexes.

Theorem 6.9 (Thomas Elsner). *Let γ be an infinite combinatorial geodesic in a graph $\mathcal{G}$. If $\sigma \in Aut(\mathcal{G})$ preserves γ then γ must contain a fixed point, a pair of consecutive vertices exchanged by σ or be an axis of the action of σ.*

Proof. Without loss of generality, we may assume that σ is not trivial. Let the vertices of γ be indexed by the integers $(p_i)_{i\in\mathbb{Z}}$ so that its edges are of the form (p_i, p_{i+1}). Since $\sigma(\gamma) = \gamma$ there is a bijection $\phi : \mathbb{Z} \to \mathbb{Z}$ such that $\sigma(p_n) = p_{\phi(n)}$ so accounting for the fact that σ preserves adjacencies, $|\phi(n+1) = \phi(n)| = 1$ and we may write $\phi(n) = \theta + \epsilon n$ where $\epsilon \in \{\pm 1\}$ determines the presence or absence

[33]See Appendix B.

of reflections. Indeed, when $\epsilon = -1$, when θ is even, $\sigma(p_{\theta/2}) = p_{\theta/2}$ so we find a fixed point and when θ is odd, $\sigma(p_{\frac{\theta-1}{2}}) = p_{\frac{\theta+1}{2}}$ so we find an inversion exchanging two consecutive vertices of γ. It should be noted that in both of these cases a "reflection" occurs in γ but it only occurs along a hyperplane in the odd case. When $\epsilon = 1$, it is clear that σ acts on γ by translations of length θ so we need only show that $\theta = \delta_\sigma$ which follows by comparing the distances $d(x, \sigma^n(x))$ and $d(p_0, \sigma^n(p_0))$ using the triangle inequality.

□

Since an automorphism σ of a cube complex X induces an automorphism of its $1-$skeleton which is a graph and the exchanging of a pair of adjacent vertices in the above result implies the existence of an inversion we immediately obtain the following corollary which sheds some light on the similarity between trees and $CAT(0)$ cube complexes.

Corollary 6.10 ([Hag07]). *Let X be a $CAT(0)$ cube complex with $\sigma \in Aut(X)$ acting stably without inversions. If σ preserves an infinite combinatorial geodesic γ then it is an axis of σ. Further, σ has the same translation length on each axis and any axis of σ is also an axis of σ^n where translation is by $\delta_{\sigma^n} = n \cdot \delta_\sigma$.*

Finally, the following lemma of [Hag07] is the essential key to finding the axes of a stable automorphism which doesn't fix a vertex. This is where the real work behind the main theorem is hidden.

Lemma 6.11 (Axis Building Lemma, [Hag07]). *Let X be a $CAT(0)$ cube complex, suppose that $\sigma \in Aut(X)$ acts stably*[34] *and let $x \in X^{(0)}$ be a point where $\delta_\sigma = d_\sigma(x)$. Then, for all $n \geq 0$ we must have $d(x, \sigma^n(x)) = n \cdot \delta_\sigma$.*

Proof. Let $Min(\sigma)$ be the set of vertices $x \in X^{(0)}$ such that $\delta_\sigma = d_\sigma(x) = d(x, \sigma(x))$. Since the combinatorial metric is discrete, $Min(\sigma) \neq \emptyset$ and we may assume without loss of generality that $\delta_\sigma > 0$.

[34] It is easy to see that this hypothesis is necessary by considering the example following the definition.

Suppose for a contradiction that the following subset of $X^{(0)}$ is nonempty:

$$\Xi := \{x \in Min(\sigma) \text{ with } d(x, \sigma^n(x)) < n \cdot \delta_\sigma \text{ for some } n\}. \tag{5}$$

Choose a point $x \in \Xi$ such that the value of n for which it is a counterexample is minimal. Let now γ_0 be a fixed geodesic segment joining x to $\sigma(x)$ and let

$$\gamma = \gamma_0 \cup \sigma\gamma_0 \cup \ldots \cup \sigma^{n-1}\gamma_0$$

be the concatenation of translates of γ_0 by powers of σ which results in a path from x to $\sigma^n(x)$. Since $x \in \Xi$, γ is not a geodesic and there are[35] some hyperplanes of X that cross it twice. By minimality of n, any such hyperplane $\mathcal{H}$ must cross γ before $\sigma(x)$ and after $\sigma^{n-1}(x)$ so it divides $\gamma = \gamma_A \cup e_A \cup \gamma_{\mathcal{H}} \cup e_B \cup \gamma_B$ into three disjoint subpaths and two edges dual to $\mathcal{H}$. Denoting the subpath which doesn't contain x or $\sigma^n(x)$ by $\gamma_{\mathcal{H}}$ we may choose $\mathcal{H}$ such that the length of $\gamma_{\mathcal{H}}$ is minimal and consequently is a geodesic.

Since the carrier of a hyperplane is convex and both endpoints of the geodesic $\gamma_{\mathcal{H}}$ lie in $N(\mathcal{H})$ we have that $\gamma_{\mathcal{H}} \subset N(\mathcal{H})$ so there is a combinatorial geodesic $\overline{\gamma_{\mathcal{H}}}$ parallel and opposite to $\gamma_{\mathcal{H}}$ along $\mathcal{H}$. By the existence[36] of inversions along hyperplanes there is a vertex $y \in \Xi$ such that $\sigma(y) \in \overline{\gamma_{\mathcal{H}}}$ is opposite to $\sigma(x)$. The inherent symmetry in $N(\mathcal{H})$ allows us to further deduce that the path from y to $\sigma(y)$ given by the edge (x, y) followed by the path γ_A along with the necessary initial segment of $\overline{\gamma_{\mathcal{H}}}$ is a geodesic that we denote by γ'. It may be helpful to consult Figure 18. In fact, by minimality of n we now have a geodesic $\gamma'' = \gamma' \cup \sigma(\gamma') \cup \sigma^2(\gamma') \cup \ldots \cup \sigma^{n-2}(\gamma')$.

Therefore, the hyperplane $\mathcal{K}$ dual to the edge (x, y) separates y from $\sigma(y)$ and y from $\sigma^{n-1}(y)$. Notationally we will write $y/_{\mathcal{K}}/\sigma(y)$ and $y/_{\mathcal{K}}/\sigma^{n-1}(y)$. Let us assume that $y/_{\mathcal{K}}/\sigma^{n-1}(x)$, this will be shown to be true in the claim below. Applying σ we obtain $\sigma(y)/_{\sigma(\mathcal{K})}/\sigma^n(x)$ which is the same as saying $\sigma(y)/_{\mathcal{H}}/\sigma^{n-1}(x)$ which is a contradiction because $\mathcal{H}$ would have to cross the geodesic $\sigma^{n-1}[(x), \sigma(x)]$

[35]See Theorem 3.5 part 3.
[36]See Theorem 3.5 part 2.

twice. As such, it only remains to prove the following claim:

Claim 6.12. $y/_{\mathcal{K}}/\sigma^{n-1}(x)$

Proof of the Claim. If this were not the case then $\sigma^{n-1}(x)/_{\mathcal{K}}/\sigma^{n-1}(y)$ so $\mathcal{K}$ would be dual to the edges (x, y) and $(\sigma^{n-1}(x), \sigma^{n-1}(y))$ which implies that $\mathcal{K} = \sigma^{n-1}(\mathcal{K})$. However, combining this with $y/_{\mathcal{K}}/\sigma^{n-1}(y)$ would imply that σ has an inversion along $\mathcal{K}$, a contradiction. □

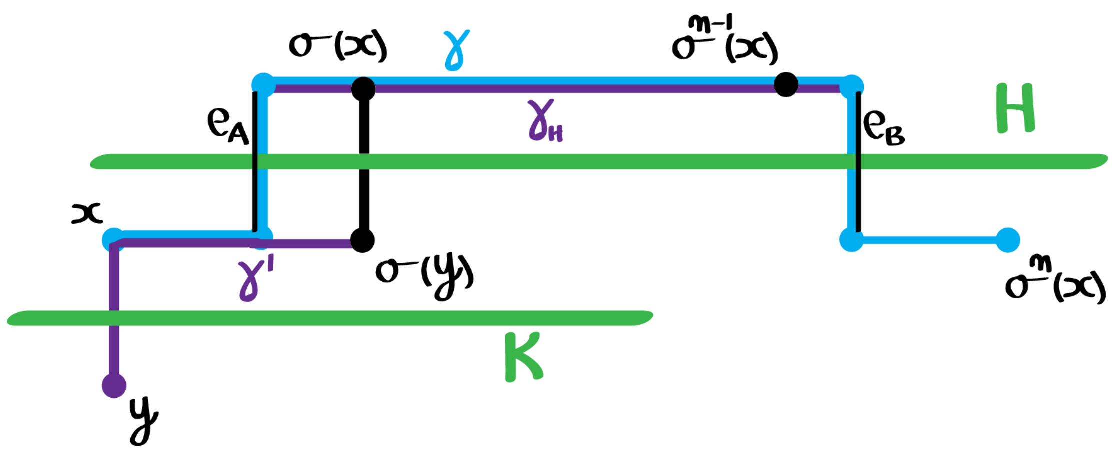

Figure 18: The various combinatorial paths and hyperplanes.

This completes the proof of the theorem. □

The reason for the previous lemma's name is the following procedure analogous to the proof of Theorem 6.3: Suppose that σ is an automorphism of a $CAT(0)$ cube complex X which acts stably and doesn't have a fixed point. Since we are working in the combinatorial metric, δ_σ is a strictly positive integer and must be attained for some point x, namely $\delta_\sigma = d_\sigma(x) > 0$. Consider now a fixed combinatorial geodesic γ_0 joining x to $\sigma(x)$. The Axis Building Lemma tells us that $d(x, \sigma^n(x)) = n\delta_\sigma$ so for any n the concatenation of geodesic segments

$$\sigma^{-n}\gamma_0 \cup \ldots \cup \sigma^{-1}\gamma_0 \cup \gamma_0 \cup \sigma\gamma_0 \ldots \cup \sigma^n\gamma_0$$

is a geodesic segment and taking the union over all n yields an infinite combinatorial geodesic γ preserved by σ. By construction, Corollary 6.10 applies and we have found an axis for σ. We have proved the following:

Theorem 6.13 ([Hag07])**.** *If X is a $CAT(0)$ cube complex and $\sigma \in Aut(X)$ acts stably then it is combinatorially semi-simple, namely it fixes a point $x \in X^{(0)}$ or it acts by translations of length δ_σ on an infinite combinatorial geodesic.*

This theorem has some neat application. For instance it forbids many groups from acting without fixed points on $CAT(0)$ cube complexes as we will now show.

Definition We say that a subgroup $H \leq G$ is *distorted* if there is a sequence of elements $h_n \in H$ indexed by $\mathbb{N}$ such that the following three conditions hold:

1. $d_G(1, h_n) < d_G(1, h_{n+1})$
2. $\lim_{n\to\infty} d_G(1, h_n) = \infty$
3. $\lim_{n\to\infty} \frac{d_G(1,h_n)}{d_H(1,h_n)} = 0$

where d_G and d_H represent the word metrics in G and H for a chosen set of generators.

For instance, in the Baumslag-Solitar group $BS(m,n) = \langle a, b | ba^m b^{-1} = a^n \rangle$ the subgroup $\langle a \rangle$ is distorted whenever $m \neq n$. In light of this concept from geometric group theory we have the following corollary to Haglund's result.

Corollary 6.14 ([Hag07])**.** *If Γ is a group containing a distorted cyclic subgroup H then every action of Γ by automorphisms on a $CAT(0)$ cube complex has a fixed point.*

Proof. Let us suppose that the distorted cyclic subgroup H is generated by an element $\gamma_0 \in \Gamma$ and that we have fixed a set of generators $\mathcal{S} := \{s_1, s_2, \dots s_k\}$ for Γ. If Γ acts on the $CAT(0)$ cube complex X, we may assume (by passing to a cubical subdivision if necessary) by the argument following the axis building lemma that $\delta_{\gamma_0^n} = n\delta_{\gamma_0}$. If we take a geodesic decomposition $\gamma_0^n = s_1 s_2 \dots s_{k_n}$ in $C_{\mathcal{S}}(\Gamma)$ we must have that $\delta_{\gamma_0^n} \leq k_n \max_{s\in\mathcal{S}} \delta_s$ and consequently $n \cdot \delta_{\gamma_0} \leq k_n \max_{s\in\mathcal{S}} \delta_s$. But now because H is distorted we have

$$\frac{k_n}{n} = \frac{d_\Gamma(1, \gamma_0^n)}{d_H(1, \gamma_0^n)} \xrightarrow{n\to\infty} 0$$

which shows that $\delta_{\gamma_0} = 0$ so γ_0 has a fixed point. □

6.4 Isometries of Polyhedral Cell Complexes

The semi-simplicity of the isometries of $CAT(0)$ cube complexes should be compared with the somewhat analogous result of Bridson which applies to $CAT(0)$ polyhedral cell complexes X whose set of shapes is finite. In this context, the automorphisms of cell complexes are those isometries preserving the cell structure. In other words, they are the cellular isometries $\sigma : X \to X$ sending cells isometrically onto cells. It should be noted that requiring the number of shapes to be finite implicitly impose the restriction that the complex be finite dimensional. As such, Haglund's result could be thought of as some kind of a generalization of the above theorem to the infinite dimensional case for cube complexes

Theorem 6.15 ([Bri99]). *If X is a $CAT(0)$ polyhedral cell complex whose set of shapes is finite then every automorphism of X is semi-simple.*

Sketch of the Proof. Let σ be an automorphism of X and recall that by definition of the translation length, there is a sequence of points (x_n) in X such that $d_\sigma(x_n) \to \delta_\sigma$. Since the set $Shapes(X)$ is finite, by passing to a subsequence, we may assume that there is some fixed $\hat{S} \in Shapes(X)$ such that $x_n \in \varphi_n(\hat{S})$ for all n where $\varphi_n : \hat{S} \to S_n$ represents an embedding of the shape $\hat{S}$ into X. We now use the classic trick and consider the pullback sequence $\hat{x}_n := \varphi_n^{-1}(x_n)$ which lies in the compact shape $\hat{S}$ and passing to yet another subsequence, we may assume that $\hat{x}_n \to \hat{x}$. Now, by the triangle inequality

$$\delta_\sigma \leq d(\sigma\varphi_n(\hat{x}), \varphi_n(\hat{x})) \leq d(\sigma \cdot \varphi_n(\hat{x}), \sigma \cdot x_n) + d(\sigma \cdot x_n, x_n) + d(x_n, \varphi_n(\hat{x})) \quad (6)$$

where the first summand tends to zero because $d(x_n, \varphi(\hat{x})) \leq d_{\hat{S}}(\hat{x}_n, \hat{x})$, the second summand tends to δ_σ and the last summand tends to zero. One can show that the set of numbers $d(\sigma\varphi_n(\hat{x}), \varphi_n(\hat{x}))$ is a discrete subset of $\mathbb{R}$ using the finiteness of $Shapes(X)$ along with the fact that σ maps cells isometrically onto cells. Therefore, there must be some n for which $d(\sigma\varphi_n(\hat{x}), \varphi_n(\hat{x})) = \delta_\sigma$ so the isometry is semi-simple.

□

The interest in this kind of result for general polyhedral complexes is that, as shown in [Bri99] it provides fixed points for some actions of lattices upon them.

7 Geometric Superrigidity

We now turn our attention to a last and rather different aspect of $CAT(0)$ geometry, the world of rigidity. The basic scheme behind any *superrigidity*[37] result is that given a topological group G along with a subgroup $\Gamma \leq G$ for which there is a morphism $f : \Gamma \to H$ into a third group H then, under suitable hypothesis, there is a unique extension of f to a morphism $\hat{f} : G \to H$.

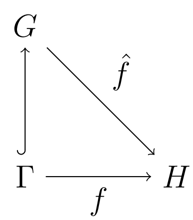

The case we are most interested in is that of *geometric superrigidity* where $H = Isom(X)$ for X a complete $CAT(0)$ space also called a Hadamard space. Here, the hypothesis on H and f can be recast into conditions on the implicit action $\Gamma \curvearrowright X$.

[37]This term was coined by Mostow in view of Margulis' results at the International Congress of Mathematicians held in Vancouver in 1974.

All actions of groups on metric spaces in this section are by isometries on complete $CAT(0)$ spaces.

We also reemphasize the crucial properties of a $CAT(0)$ space X that will be used heavily throughout this section:

1. *The $CAT(0)$ metric is convex.*

2. *Balls in a $CAT(0)$ space are uniformly convex: given $r, \delta > 0$ there exists an $\epsilon = \epsilon(r, \delta) > 0$ such that for any closed ball of radius r centred about a point $x \in X$, if y_1 and y_2 lie within the ball and $d(y_1, y_2) \geq \delta$ we must have $d(x, \frac{y_1+y_2}{2}) \leq r - \epsilon$.*

The second point emphasizes the idea that, like in euclidean space, the midpoint of the geodesic segment joining two distant points in a given ball must lie "deep" within the ball. The focus of our attention will be on a very specific type of lattices but the interested reader should also glance at [Bur94]'s short *Survey of Rigidity Properties of Group Actions on $CAT(0)$-Spaces*.

The constructions and results of this section originate from analogous ones in [GKM08] where they were carried out in the context of so-called Busemann non-positively curved spaces (metric spaces where the distance function is convex) assuming uniform convexity as an axiom. See [Pap05] for a more detailed approach to this weaker notion of nonpositive curvature.

For the remainder of this section, we let X denote a complete $CAT(0)$ spaces.

7.1 Generalized Harmonic Maps

Harmonic maps $\varphi : M \to N$ between Riemannian manifolds are critical points of the Dirichlet energy functional. Heuristically, one can think of a harmonic transformation φ as requiring a minimal amount of "energy" to deform M into its image $\varphi(M) \subseteq N$. Following [GKM08] and [Jos97], we generalize this notion of energy to the case of equivariant maps between metric spaces with respect to a group action. Such critical maps will later allow us to see that certain orbit maps for subgroup actions on $CAT(0)$ spaces are continuous, facilitating their extension to the whole group.

Recall that if G is a locally compact topological group, there is a unique (up to scaling) countably additive, translation invariant measure defined on the Borel algebra generated by all compact subsets of G. This is the *Haar* measure[38] defined on G that we will henceforth denote by μ.

Irreducible and Uniform Lattices If a discrete subgroup $\Gamma \leq G$ has finite co-volume and compact quotient, namely $\mu(G/\Gamma) < \infty$ and G/Γ is compact, we say that Γ is a *uniform lattice* in G. If Γ is a uniform lattice in $G = G_1 \times G_2 \times \ldots \times G_n$, we say that it is *irreducible* if $\Gamma(\prod_{j \neq i} G_j)$ is dense in G for any $1 \leq i \leq n$. In particular, this means that the projection onto the $i-$th coordinate $\pi_i(\Gamma)$ is dense in G_i for any i.

Let us now specialize to the case where we are given a finitely generated irreducible uniform lattice in a compactly generated locally compact group $\Gamma \leq G_1 \times G_2 = G$ with an action $\Gamma \curvearrowright X$ by isometries that we wish to extend to an action of G. The first step in this direction will be the construction of an associated $CAT(0)$ space of equivariant functions $\varphi : G \to X$ induced by $\Gamma \curvearrowright X$. Since Γ is uniform, there is a measurable relatively compact Ω such that $G = \Gamma \cdot \Omega$. Compact sets have finite Haar measure so we may renormalize μ so that (Ω, μ) becomes a probability space.

Space of Equivariant Maps Denote by $L^2(\Omega, X)$ the set of $\Gamma-$equivariant measurable maps $\varphi : G \to X$ subject to the L^2 condition that $\int_\Omega d^2(\varphi(\omega), x_0) d\mu < \infty$ for some fixed $x_0 \in X$. Since μ is finite, this condition does not depend on the choice of x_0 by the triangle inequality in X. Recalling that $\Gamma \cdot \Omega = G$ and all φ satisfy $\gamma \cdot \varphi(\omega) = \varphi(\gamma \cdot \omega)$ we abuse notation and make no real distinction between $\varphi : \Omega \to X$ and $\varphi : G \to X$. We give $L^2(\Omega, X)$ the metric $\rho^2(\varphi, \psi) = \int_\Omega d^2(\varphi(\omega), \psi(\omega)) d\mu$.

[38] Please consult [FF99] for a development of this theory.

The following lemma characterizes geodesics in $L^2(\Omega, X)$.

Lemma 7.1 ([KS93],[Jos97])**.** *Let* $\varphi_0, \varphi_1 \in L^2(\Omega, X)$ *and define a family of geodesics* $\{\sigma_\omega : [0,1] \to X\}_{\omega\in\Omega}$ *such that* $\sigma_\omega(0) = \varphi_0(\omega)$ *and* $\sigma_\omega(1) = \varphi_1(\omega)$. *The family of maps* $\varphi_t(\omega) := \sigma_\omega(t)$, $t \in [0,1]$ *is a geodesic joining* φ_0 *to* φ_1 *in* $L^2(\Omega, X)$.

Proof. We verify that $\rho^2(\varphi_0, \varphi_t) = \int_\Omega d^2(\varphi_0(\omega), \varphi_t(\omega))d\mu = \int_\Omega t^2 d^2(\varphi_0(\omega), \varphi_1(\omega))d\mu = t^2\rho^2(\varphi_0, \varphi_1)$ from which the result follows. □

Lemma 7.2 ([Jos97])**.** *If* X *is a complete* $CAT(0)$ *space then* $L^2(\Omega, X)$ *is also* $CAT(0)$ *and complete.*

Proof. The completeness of $L^2(\Omega, X)$ is straightforward. To see that it is $CAT(0)$ notice by Lemma 7.1 that given $f, g, h \in L^2(\Omega, X)$, the triangle comparison property, or the CN inequality holds pointwise for $f(\omega), g(\omega)$ and $h(\omega)$ so integrating yields the desired result. □

At this point, we are in a good position to define an energy functional on $L^2(\Omega, X)$. However, since we will want to minimize it we need conditions to ensure that it will be finite and attain a minimum[39]. To this end, let $\langle\Sigma\rangle = \Gamma$ be a finite generating set and $\langle K\rangle = G_1$ be a compact generating set. We introduce a notion of "size" for an element $g_1 \in G_1$ by the map $h : G_1 \to \mathbb{R}$, $h(g_1) := e^{-d_K(1,g_1)^2+1}$ where d_K is the word metric on G_1. This function along with the L^2 condition will ensure that energy is finite. We further assume that the displacement function $d_\Sigma \to \infty$ as $x \to \infty$ where $d_\Sigma = \max_{\sigma\in\Sigma} d(\sigma \cdot x, x)$ and x eventually leaves any bounded set. This notion is described as non-evanescence in [Mon06] and is equivalent to the fact that there are no fixed points in the boundary ∂X when X is a proper metric space[40]. This condition is used to ensure that energy will attain its infimum.

[39] A more general outline for this kind of construction is indicated in Section 7.4

[40] A metric space where closed balls are compact.

Definition The $G_1 - Energy$ [41] of a function $\varphi \in L^2(\Omega, X)$ is defined to be

$$E(\varphi) = \int_{\Omega \times G_1} h(g_1) d^2(\varphi(\omega), \varphi(\omega g_1)) = \int_{(\Gamma \backslash G) \times G_1} h(g_1) d^2(\varphi(g), \varphi(g g_1)).$$

Note that it corresponds to the total amount of resulting "stretch" in the image under φ of the right action $G_1 \curvearrowright \Gamma \backslash G$ proportional to the "size" of the elements of G_1.

We immediately highlight a few properties of energy $E : L^2(\Omega, X) \to \mathbb{R}$.

1. E is convex by convexity of the $CAT(0)$ metric and Lemma 7.1.
2. E is finite by the definition of h and using the triangle inequality in X combined with compactness of Ω.
3. E is continuous.
4. E is G_2-invariant from the right i.e. $E(\varphi) = E(\varphi(\cdot g_2))$.

In fact, letting $M = \inf\{E(\varphi) : \varphi \in L^2(\Omega, X)\}$ we can finally define the key ingredient of our construction.

Definition We say that φ is *harmonic* if $E(\varphi) = M$.

Theorem 7.3 ([GKM08], Theorem 3.2)**.** *There exists a harmonic map.*

Sketch of the proof. Let the norm of $\varphi \in L^2(\Omega, X)$ be $||\varphi|| := \rho(\varphi, x_0)$ where x_0 is the map constantly equal to x_0. Defining a family of maps φ_n satisfying

1. $E(\varphi_n) \leq M + \frac{1}{n}$
2. $||\varphi_n|| \leq \inf\{||\varphi|| : E(\varphi) \leq M + \frac{1}{n}\} + \frac{1}{n}$

one can show using the fact that $d_\Sigma \to \infty$ and the uniform convexity of X that they are uniformly bounded. This then leads to the fact that they form a cauchy sequence. Finally, since $L^2(\Omega, X)$ is complete and E is continuous we obtain a harmonic map. □

[41] See [Jos97] for a more general notion of energy for maps between metric spaces.

Remark Notice that since energy is G_2-invariant if we are given a harmonic function φ, its translate $\varphi(\cdot g_2)$ is also harmonic.

Lemma 7.4 (Harmonic Functions Are G_1-Parallel). *If φ and ψ are harmonic, then the geodesic segment $[\varphi(g), \varphi(gg_1)]$ is parallel to $[\psi(g), \psi(gg_1)]$.*

Sketch of proof. This follows from the convexity of the $CAT(0)$ metric along with the minimality of $E(\varphi)$. □

The following lemma, implicit in the proof of the main theorem of [Mon06] is key to the next results.

Lemma 7.5 (Key Lemma). *Fixing any $g_1 \in G_1$, the function $g \mapsto d(\varphi(g), \varphi(gg_1))$ is essentially constant. The same holds when we fix any $g_2 \in G_2$.*

Proof. Recall that in a $CAT(0)$ space, a geodesic segment $[a, b]$ is parallel to $[x, y]$ if and only if $[a, x]$ is parallel to $[b, y]$ as shown in [Bus48]. Combining this with Lemma 7.4 we have that $d(\varphi(gg_2), \varphi(gg_1g_2)) = d(\varphi(g), \varphi(gg_1))$. On the other hand, $d(\varphi(g), \varphi(gg_1)) = d(\gamma \cdot \varphi(g), \gamma \cdot \varphi(gg_1)) = d(\varphi(\gamma g), \varphi(\gamma gg_1))$ since φ is $\Gamma-$equivariant and Γ acts by isometry on X. But now, since Γ is irreducible, the right action $G_2 \curvearrowright \Gamma \setminus G$ is ergodic and we obtain the result. □

Remark This is the first place where the irreducibility of the lattice Γ is used, all previous constructions, in particular the construction of the space $L^2(\Omega, X)$ hold in a more general setting.

Corollary 7.6. *Harmonic maps are essentially continuous.*

Sketch of proof. The Key Lemma 7.5 shows that $d(\varphi(g), \varphi(gg_i)) = \rho(\varphi, \varphi(\cdot g_i))$ where $i = 1, 2$ so showing that the (right) action of G on $L^2(\Gamma \backslash G, X)$ is continuous we obtain that φ is essentially continuous. □

It follows that by changing the value of a harmonic map on a set of measure zero we obtain a continuous harmonic map. As such, in the rest of this section we assume harmonic maps to be continuous.

7.2 Reduced Actions on $CAT(0)$ Spaces

The following definition is a substitute introduced by Monod which corresponds to Zariski density in the context of Lie groups.

Definition We say the action $\Gamma \curvearrowright X$ is *reduced* if there is no unbounded closed convex proper subset $Y \subsetneq X$ such that the Hausdorff distance

$$d_{\mathcal{H}}(Y, \gamma \cdot Y) = \inf_{\epsilon}\{N_\epsilon(Y) \subset \gamma \cdot Y \text{ and } N_\epsilon(\gamma \cdot Y) \subset Y\}$$

is finite for all $\gamma \in \Gamma$.

The reason for introducing this definition is because it has the following strong consequence for the action $\Gamma \curvearrowright X$ outlined in Section 2.5 of [GKM08][42].

Lemma 7.7. *If X is a complete $CAT(0)$ space with no non-trivial Clifford translations and $\langle\Sigma\rangle = \Gamma \curvearrowright X$ is* reduced *with no globally fixed point where $|\Sigma| < \infty$, then*

1. *The displacement map $d_\Sigma \to \infty$ as $x \to \infty$ where $d_\Sigma = \max_{\sigma\in\Sigma} d(\sigma \cdot x, x)$.*
2. *There is no non-empty closed convex proper $\Gamma-$invariant subset of X.*[43]

Sketch of proof. The function d_Σ is convex being the maximum of convex functions and the absence of Clifford translations implies that it is not constant. Since the proper sub-level sets $\{x \in X : d_\Sigma(x) \leq \delta\}$ for any $\delta > 0$ are convex and of bounded Hausdorff distances from each other under the action of Γ they must be bounded because $\Gamma \curvearrowright X$ is reduced so we must have that $d_\Sigma \to \infty$.

Using the uniform convexity of X, the fact that $d_\Sigma \to \infty$ implies by convexity of d_Σ, uniform convexity of X and an application of Zorn's Lemma that there exists a minimal closed convex $\Gamma-$invariant subset $Y \subseteq X$. If $Y \neq X$, it must be bounded because $\Gamma \curvearrowright X$ is reduced and Since X is $CAT(0)$, it has a well defined circumcentre[44] which is a fixed point. □

[42] Compare with Lemma 63 of [Mon06].
[43] In the terminology of [GKM08], the action is then said to be $C-$minimal.
[44] See Proposition 6.1.

Theorem 7.8 ([GKM08], [Mon06])**.** *Let Γ be an irreducible uniform lattice in a locally compact and compactly generated group $G = G_1 \times G_2 \times \ldots \times G_n$ and let X be a complete $CAT(0)$ space with no nontrivial Clifford translations. If $\Gamma \curvearrowright X$ is a reduced action with no global fixed points then it extends to a continuous action $G \curvearrowright X$ which factors through one of the G_i's.*

We give a rough sketch of the argument to be fully carried out in the next section:

Sketch of the proof. We will only prove the case where $n = 2$, the general case follows by induction. Let $\varphi : G \to X$ be a continuous harmonic function. If $\varphi(G_2)$ is bounded, we may take φ to be G_2-invariant using the fact that $\Gamma \curvearrowright X$ is reduced. On the other hand, if $\varphi(G_2)$ is unbounded, we may take φ to be G_1-invariant using the fact that there are no nontrivial Clifford translations. Without loss of generality, φ is continuous and G_2-invariant. We use φ it to show that the $\Gamma-$orbits in X are continuous when Γ is endowed with the G_1-topology. The $G-$action can then be defined by

$$g \cdot x := \lim_{\pi_1(\gamma) \to \pi_1(g)} \gamma \cdot x.$$

□

(Counter)-Example [Mon06] Consider the following scenario where all groups are discrete: $G = (\mathbb{Z}/2\mathbb{Z} \ltimes \mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z} \ltimes \mathbb{Z})$ along with its index two irreducible subgroup $\Gamma = \mathbb{Z}/2\mathbb{Z} \ltimes (\mathbb{Z} \oplus \mathbb{Z})$. Here $\Gamma \curvearrowright \mathbb{R}$ where $\mathbb{Z}$ acts by translation and $\mathbb{Z}/2\mathbb{Z}$ acts by reflections but the action does not extend to G. ($\mathbb{R}$ has many nontrivial Clifford translations!)

The assumption that $\Gamma \curvearrowright X$ is reduced is rather strong. Using the techniques of generalized harmonic maps [GKM08] obtained the following result of group action extensions into the visual boundary ∂X of a Busemann non positively curved space without the reduced action assumption. This was first proved by [Mon06] for $CAT(0)$ spaces omitting the use of generalized harmonic maps and generalizing instead a splitting result for the the space X itself found in [BH99].

Theorem 7.9 ([GKM08], [Mon06]). *Let Γ be an irreducible uniform lattice in a locally compact and compactly generated group $G = G_1 \times G_2 \times \ldots \times G_n$ and let X be a complete proper*[45] *$CAT(0)$ space. If $\Gamma \curvearrowright X$ without a global fixed point, then there is a closed, invariant subset $\emptyset \neq \mathcal{L} \subseteq \partial X$ on which the action extends continuously to G and factors through one of the G_i's.*

Remark It is shown in [Mon06] that Theorem 7.9 actually implies and generalizes a celebrated theorem of Margulis for uniform lattices found in [Mar91].

7.3 Proof of Theorem 7.8

Proof. Let $\Gamma \leq G_1 \times G_2$ be as in the statement of Theorem 7.8.

The first trick of this proof revolves around an iterative construction of closed convex hulls in uniquely geodesic metric spaces. Let Y be a subset of a uniquely geodesic metric space. Defining $Y_0 := Y$ and $Y_n := \{\frac{x+y}{2} : x, y \in Y_{n-1}\}$ we see that $Y_n \supset Y_{n-1}$ (taking x=y) and that

$$\overline{Conv(Y)} = \overline{\cup_{i=1}^{\infty} Y_n}$$

since any $z \in Conv(Y)$ must occur as the midpoint of a pair of points in Y_n for some n. The use of this construction is apparent when one recalls that since energy is convex, if φ and ψ are harmonic, the convex combination $\frac{\varphi+\psi}{2}$ remains harmonic.

Let $\varphi : G_1 \times G_2 \to X$ be harmonic. As we have seen, we can (and will) assume harmonic maps to be continuous. We proceed by splitting up the problem into cases as to whether $\varphi(G_2)$ is bounded or not to obtain a G_i-invariant harmonic map.

1. $\varphi(G_2)$ is bounded:

   Consider the orbit of φ under the (right) action of G_2, namely the set $G_2 \cdot \varphi := \{\varphi(\cdot g_2) : g_2 \in G_2\}$ in which all maps are harmonic. The Key Lemma

[45] A metric space where closed balls are compact.

7.5 ensures that the distance between any two functions in $G_2 \cdot \varphi$ is bounded so the circumcentre[46] of $\overline{Conv(G_2 \cdot \varphi)}$ is a well defined G_2-invariant map φ_0 which remains harmonic by our iterative construction of $\overline{Conv(G_2 \cdot \varphi)}$.

2. $\varphi(G_2)$ is unbounded:

   Letting $Y := \overline{Conv(\varphi(G_2))}$, we claim that $d_{\mathcal{H}}(\gamma \cdot Y, Y) < \infty$ for all $\gamma \in \Gamma$. Indeed, if $\gamma = (\gamma_1, \gamma_2)$ we see that $\gamma \cdot \varphi(g_2) = \varphi(\gamma_2 g_2 \gamma_1)$ since φ is $\Gamma-$equivariant so $d(\gamma \cdot \varphi(g_2), \varphi(\gamma_2 g_2)) = d(\varphi(\gamma_2 g_2 \gamma_1), \varphi(\gamma_2 g_2))$ is constant by the Key Lemma 7.5. Since $\Gamma \curvearrowright X$ is reduced, we conclude that $X = Y$.

   Recalling our iterative construction for closed convex hulls we have $X = \overline{\cup_{n=0}^{\infty} Y_n}$ where points of $\cup_{n=0}^{\infty} Y_n$ can be described as $\psi(1)$ where $\psi \in Conv(G_2 \cdot \varphi)$ is harmonic. To see this, notice that $\psi(\cdot) = \varphi(\cdot g_2)$ and write elements of Y_n as nested combinations of elements of $\varphi(G_2)$. Now, for any $g_1 \in G_1$, the segments $[\psi(1), \psi(g_1)]$ and $[\psi'(1), \psi'(g_1)]$ where $\psi, \psi' \in Conv(G_2 \cdot \varphi)$ are parallel by Lemma 7.4. As such, the map $\psi(1) \mapsto \psi(g_1)$ extends to a Clifford translation on the set $\overline{\{\psi(1) : \psi \in conv(G_2 \cdot \varphi\}}$. But $Y \subset \overline{\{\psi(1) : \psi \in conv(G_2 \cdot \varphi\}} = X$ so because X has no non-trivial Clifford translations $\varphi(g_2) = \varphi(g_1 g_2)$ for all g_2. Letting g_1 vary we see that φ is G_1-invariant.

Without loss of generality, we can now assume that $\varphi : G \to X$ is a continuous, harmonic, $\Gamma-$equivariant and G_2-invariant map. We would like to define the extension of the action $\Gamma \curvearrowright X$ by the rule

$$g \cdot x := \lim_{\pi_1(\gamma) \to \pi_1(g)} \gamma \cdot x \tag{7}$$

so we need to ensure that the limit makes sense. To do so, consider the *orbit map* of a given point $x \in X$, $O_x : \Gamma \to X$ defined by the rule $\gamma \mapsto \gamma \cdot x$.

Claim 7.10. *If we equip Γ with the topology[47] induced from G_1 by the projection map $\pi_1 : \Gamma \to G_1$ then O_x is continuous for all $x \in \varphi_0(G)$.*

[46] See Proposition 6.1.
[47] This topology is not necessarily Hausdorff.

To see this, let g_x be such that $\varphi_0(g_x) = x$ and notice that since φ_0 is G_2-invariant we may assume that $g_x \in G_1$. Using $\Gamma-$equivariance of φ_0, $\gamma \mapsto \gamma \cdot x = \gamma \cdot \varphi_0(g_x) = \varphi_0(\gamma \cdot g_x)$ so we can factor O_x.

$$\Gamma \xrightarrow[\cdot g_x]{} G_1 \xrightarrow[\varphi_0]{} X, \qquad O_x : \Gamma \to X$$

Since φ_0 is continuous and $\cdot g_x : \Gamma \to G_1$ is continuous for the given topology, we conclude as claimed that O_x is continuous when Γ is endowed with the G_1-topology.

Now, keeping in mind that Γ acts by isometries we see that the set

$$\{x \in X : O_x \text{ is continuous when } \Gamma \text{ is equipped with the } G_1 \text{ topology}\}$$

is nonempty, closed, convex and $\Gamma-$invariant so it must be equal to X by Lemma 7.7. Notice that since O_x is continuous when Γ is endowed with the G_1-topology, the action $\Gamma \curvearrowright X$ factors through $\pi_1(\Gamma)$ so we may think of Γ as a dense subgroup of G_1 because it is irreducible in G. As such, using the continuity of O_x we may define the action of $g \in G$ by the Equation (7) above where the action factors through G_1.

□

7.4 The General Method

To conclude this section we outline some ideas behind the existence of generalized harmonic maps.

Given a $CAT(0)$ (whence uniformly convex) space X, the nonnegative function $d_y^2 : X \to \mathbb{R}^+$ defined by the rule $d_y^2(x) := d^2(x, y)$ is strictly convex so for all geodesics $\sigma : [0,1] \to X$ and for all $\epsilon > 0$ there is a $\delta > 0$ such that

$$d_y^2(\sigma(1/2)) \geq \frac{1}{2} d_y^2(\sigma(0)) + \frac{1}{2} d_y^2(\sigma(1)) - \delta \implies d(\sigma(0), \sigma(1)) < \epsilon. \tag{8}$$

To any function $E : X \to \mathbb{R} \cup \{\infty\}$ we can associate the Moreau-Yosida approximations $E^{(n)} := \inf_{x\in X}\{nE(x) + d_y^2(x)\}$, $n \in \mathbb{N}$. The following Lemma and Theorem are proved in [Jos97] in the more general context of any metric space admitting a function similar to d_y^2 in the sense that it satisfies (8), this "convexity" condition is key to their proofs.

Lemma 7.11. *If X is a complete $CAT(0)$ space and $EX \to \mathbb{R} \cup \{\infty\}$ is convex, (lower semi-) continuous and not constantly equal to ∞ then for all $n \in \mathbb{N}$ there is a unique $x_n \in X$ such that $E^{(n)} = nE(x_n) + d_y^2(x_n)$.*

Using this lemma, one readily obtains the following theorem which produces under suitable hypothesis a minimum of the function $E : X \to \mathbb{R} \cup \{\infty\}$, the key to finding harmonic maps. This result should be compared with Theorem 7.3.

Theorem 7.12. *If the sequence $(d_y^2(x_n))_{n\in\mathbb{N}}$ is bounded as $n \to \infty$ then (x_n) converges to a minimizer of E as $n \to \infty$.*

Now, if we have a lattice $\Gamma \leq G$ which acts $\Gamma \curvearrowright X$ we already know that for any finite probability space Ω and a complete $CAT(0)$ space X, the space $L^2(\Omega, X)$ of $\Gamma-$equivariant maps is also complete and $CAT(0)$. In view of the previously introduced general tools, we can define several different energy functionals tailored to any given situation. Our first example is the G_1-energy used in the previous section,

$$E(\varphi) = \int_{(\Gamma\backslash G)\times G_1} h(g_1) d^2(\varphi(g), \varphi(gg_1))$$

but more generally, for any positive symmetric continuous function $h : G \times G \to \mathbb{R}^+$ we can define an energy functional as:

$$E(\varphi) = \int_{\Delta(\Gamma)\backslash G\times G} h(g_1, g_2) d^2(\varphi(g_1), \varphi(g_2))$$

where $\Delta(\Gamma) = \{(\gamma, \gamma) : \gamma \in \Gamma\}$ or any other sensible variant of these definitions. In order to construct generalized harmonic maps by applying [Jos97]'s theorem we need only ensure that:

1. At least one map $\varphi \in L^2(\Omega, X)$ has finite energy.

2. The functions $(d_y^2(x_n))_{n\in\mathbb{N}}$ are uniformly bounded.

While condition (1) is intimately linked to the choice of the function h and the nature of $\Gamma \leq G$, condition (2) usually follows from some kind of "irreducibility" assumption for the action $\Gamma \curvearrowright X$ which implies a condition similar to $d_\Sigma \to \infty$ used in the previous sections. There are a variety of available criterions that work in particular situations allowing the use of harmonic maps to solve many different problems, see for instance [DO85], [Don87], [JY90], [KS93], [Jos97], [Pan06] and of course [GKM08].

8 Concluding Remarks

The present survey of non-positive curvature in geometric group theory was far from exhaustive and there are many paths that one could pursue from here. To conclude, without burdening ourselves with technicalities, we wish to highlight a few directions in which applications of the geometric methods developed could be fruitful.

A first consideration arises by trying to imagine how results like Theorem 7.8 could be applied to spaces obtained by gluing constructions similar to those studied in Section 3.1. It turns out that there is an entirely analogous theory of combinatorial harmonic maps as developed by Wang, Gromov, Izeki and Nayatani in [Wan98], [Wan00], [Gro03], [IN05] and summarized by Pansu in [Pan06]. One simply considers a complex as a simplicial complex (by passing to an appropriate subdivision) and thinks of each edge of the $1-$skeleton as a coil spring. The energy of a mapping from a group into the complex is then simply the resulting potential energy of the coil springs, and mappings which result in equilibrium positions are called harmonic. Finding a class of concrete examples to which analogous abstract results would apply could yield useful insights. This discretized combinatorial approach to energy and harmonic maps might be a good place to start.

These ideas fit into a superrigidity research program which consists in finding the minimal necessary conditions on a group Γ such that an action of this group on a complete $CAT(0)$ space X would either have a fixed point or stabilize a convex subset which is isometric to X up to rescaling metrics on its non-trivial factors. It is unknown, to the best of our knowledge, whether there is an infinite group which does not act without a fixed point on any complete $CAT(0)$ space. This would be a good place to apply a theorem similar to Theorem 7.8 but the stringent requirements, that of a reduced action and non-evanescence, severely limit its applications. Further investigations will be necessary.

On a different topic, during the course of his study of manifolds in 1912, Dehn encountered three fundamental algorithmic problems of group theory as studied from the perspective of group presentations:

1. *The Word Problem* which requires determining if a word on an alphabet of generators for a group is trivial.

2. *The Conjugacy Problem* which requires determining whether two elements of a group are conjugates.

3. *The Isomorphism Problem* which requires determining if two groups given by two presentations are isomorphic.

It was discovered in the $1950'$s that these problems are undecidable for generic finitely presented groups. However, such algorithms are known to exist for many classes of groups and, in particular, for so-called *hyperbolic* groups as defined by Gromov in [Gro87]. Without going into too much detail, we say a finitely generated group $\Gamma \cong \langle \mathcal{A} | \mathcal{R} \rangle$ is *hyperbolic* if there is some $\delta \geq 0$ such that every geodesic triangle $\triangle$ in $\mathcal{C}_{\mathcal{A}}(\Gamma)$ is $\delta-$thin, meaning that each of its sides lies in the δ-neighbourhood of the other two. While the solutions to the word and conjugacy problem for hyperbolic groups have been known for some time[48], the isomorphism problem in this class was solved quite recently by [DG11] following the initial ideas of [Sel95] that were extended in [DG08].

[48]See [BH99] for proofs and other references.

Since this other global feature of negative curvature is in a sense dual to the local one of Alexandrov, one can wonder which relations exist between groups acting geometrically on $CAT(0)$ spaces and hyperbolic groups. For instance, it is a theorem of [Gro87] that a finitely generated group Γ acting properly and cocompactly by isometries on a $CAT(0)$ space X is hyperbolic if and only if X doesn't contain any isometric embeddings of $\mathbb{E}^2$. On the other hand, according to [Bri07] and to the best of our knowledge, it is currently unknown wether every hyperbolic group acts properly and cocompactly by isometries on a $CAT(0)$ or $CAT(-1)$ space. One might also independently wonder if the Isomorphism Problem is solvable for groups acting geometrically on $CAT(0)$ spaces and in the affirmative, if techniques behind the solutions in hyperbolic groups could be of any help. Working in this direction, one could first try to use the additional structure of those groups acting geometrically on $CAT(0)$ cube complexes to understand the problem in this class of groups.

In a similar vein, recall the results mentioned in Section 3.3 expressing a finitely presented subgroup of the fundamental group of a non-positively curved complex as the fundamental group of some other non-positively curved complex obtained by successively passing to connected covers and finite subcomplexes. It would be interesting to replace the abstract constructions used to prove the result by concrete ones which would yield an algorithm for finding the new complex. A concrete understanding of this process could yield greater insight into the structure of the finitely presented subgroups of groups acting on $CAT(0)$ spaces in general.

Finally, one could remark that we have only touched a few of the known examples of $CAT(0)$ spaces other than the gluing constructions that we mentioned and actual Riemannian manifolds of bounded sectional curvature. There are many other constructions that one could study and it would also be interesting to find entirely new ones.

A Model Spaces

Following [BH99] we provide a brief account of the material necessary to understand model spaces. Most of our study takes place in the realm of metric spaces (X, d) that we will denote by X whenever there is no ambiguity with respect to the metric. In view of generalizing to such spaces many of the concepts of differential geometry, our main concern is the notion of "shortest paths" between points. For points $x, y \in X$, a *geodesic path*[49] joining x to y is a map $c : [0, 1] \subset \mathbb{R} \to X$ such that $c(0) = x, c(1) = y$ and $d(c(t), c(t')) = |t - t'|$ for all $t, t' \in [0, 1]$. Images of such paths are referred to as *geodesic segments* and denoted $[x, y]$. If there is a (unique) geodesic joining any two points of x, we say X is a (uniquely) *geodesic space*. A subset $C \subset X$ is called *convex* if every pair of points $x, y \in C$ can be joined by a geodesic segment $[x, y] \subset C$.

Given a curve $c : [a, b] \to X$, we define its *length* as the supremum over all finite partitions $[t_0, t_1] \cup [t_1, t_2] \ldots \cup [t_{n-1}, t_n]$ of the geodesic segment $[a, b] \subset \mathbb{R}$ of the sum of the distances $d(c(t_i), c(t_{i+1}))$. Those curves having finite length in this sense are said to be *rectifiable*. With this notion at hand, a metric space X is said to be a *length space* if $d(x, y)$ coincides with the infimum of the length of all possible curves joining x to y. One should keep this in mind in Section B where we "rig" our construction of a metric on cell complexes to make them into length spaces. In this sense, we are equipping them with the "inner metric".

Throughout this work, geodesic metric spaces will be frequently compared to a standard set of *model spaces* M_k^n of constant sectional curvature k and dimension n which can be roughly divided into three distinct classes according to whether k is zero, positive or negative. Such spaces can be defined as complete $1-$connected Riemannian manifolds of constant sectional curvature k but we approach them from the purely metric point of view. In the three main cases mentioned, the standard models are Euclidean n-space with the usual metric $\mathbb{E}^n = M_0^n$, the n-sphere with the "angular" metric $\mathbb{S}^n = M_1^n$ and the hyperboloid model $\mathbb{H}^n = M_{-1}^n$ with the "Minkowski" metric. These metrics arise from considering the given

[49]Notice that this notion is stronger than the differential geometric one.

spaces as subsets of $\mathbb{R}^{n+1}$ with the appropriate inner product.

In the case of $\mathbb{S}^n \subset \mathbb{R}^{n+1}$ the metric is obtained by associating to two points $x, y \in \mathbb{S}^n$, the smallest value of the two possible angles between them. Formally, this is done by saying that we take the unique solution of the equation $\cos(d(x,y)) = (x,y)$ with the restriction that $0 \leq d(x,y) \leq \pi$ where (x,y) is the usual inner product in $\mathbb{R}^{n+1}$. Similarly, if we view $\mathbb{H}^n \subset \mathbb{R}^{n+1}$ as the hyperboloid model it coincides with the upper sheet of the set of points $x \in \mathbb{R}^{n+1}$ with $\langle x, x\rangle = -1$ where $\langle x, y\rangle := -x_{n+1}y_{n+1} + \sum_{i=1}^{n} x_i y_i$ is the Minkowski inner product. We can then define the distance analogously to the spherical case as the unique solution of the equation $\cosh(d(x,y)) = -\langle x, y\rangle$ with $d(x,y) \geq 0$.

Model spaces when $k < 0$ and $k > 0$ are respectively obtained by scaling the metrics of $M^n_{-1} = \mathbb{H}^n$ and $M^n_1 = \mathbb{S}^n$ by a factor of $\frac{1}{\sqrt{|k|}}$ to obtain M^n_k. One of the reasons we will favour those M^n_k where $k \leq 0$ or $CAT(0)$ spaces is the following attractive property which is familiar in the euclidean or hyperbolic case and follows easily from the scaling of the metric.

Proposition A.1. *For $k \leq 0$, M^n_k is a uniquely geodesic metric space where all balls are convex.*

There is a similar result when $k > 0$, for pairs of points x, y such that $d(x,y) < \pi/\sqrt{k}$ and closed balls of radius $< \pi/(2\sqrt{k})$. The quantity $D_k := \pi/\sqrt{k}$ when $k > 0$ and ∞ when $k \leq 0$ is frequently referred to for such spaces and called the *diameter*.

The interested reader may consult [BH99] for an elaborate discussion of model spaces and their isometries or [Ber03] for a friendly approach to the Riemannian point of view.

B Polyhedral Cell Complexes

The standard examples of metric spaces of a given curvature are Riemannian manifolds but the constructions carried out below yield a much more exotic class of length spaces that often are not manifolds. In order to maintain transparency of the ideas, following [BH99] we first work in the context of simplicial complexes and introduce a canonical subdivision process to carry over results in the general setting of M_k-polyhedral cell complexes[50].

The construction of M_k-simplicial complexes is entirely analogous to that of simplicial complexes in Euclidean space. An $n-$simplex $S \subset M_k^n$ is the convex hull of $n+1$ points in general position. The *faces* or cells $F \subset S$ are the convex hulls of subsets of the initial set of $n+1$ vertices and the *interior* of S, $int(S)$, is the set of its points that do not lie on any face. To build an M_k-simplicial complex we fix a disjoint family of M_k-simplices $\{S_\lambda \subset M_k^{n_\lambda} : \lambda \in \Lambda\}$ and glue them together along isometric faces of our choice to obtain a connected[51] complex X. In analogy with the classical definition of CW complexes[52], we denote by $X^{(i)}$ the set of faces of X spanned by at most $i+1$ points. This is the $i-$skeleton of X. The following definition is the key to understanding the properties of the metric structure we will endow X with.

Shapes We denote by $Shapes(X)$ a fixed choice of representatives called *shapes* for the set of different isometry types of faces of X. As such, for every face F of a simplex $S \subset X$ there is an associated $\hat{F} \in Shapes(X)$ corresponding to the unique chosen representative for faces of X isometric to it. Since simplices are also faces, we can think of $\hat{S} \in Shapes(X)$ as lying in an abstract copy of M_k^n above X and equipped with a family of isometric embeddings $\varphi_S : \hat{S} \to X$ sending the shape $\hat{S}$ to any given simplex $S \subset X$ isometric to it.

Each simplex $S \subset X$ inherits a metric $d_{\hat{S}}$ from the corresponding $\hat{S} \in Shapes(X)$ defined by pulling pairs of points in S back to $\hat{S}$. As such, we can endow X with the quotient pseudometric and under suitable hypothesis turn it into a length space as outlined below. In what follows we refer to this pullback metric as the *local* metric.

[50]These will also be called M_k- polyhedral complexes, polyhedral complexes, cell complexes or even complexes when the meaning is evident from the context.

[51]In general, this is not required but since we are interested in geodesic metric spaces, disconnected complexes are rather uninteresting.

[52]Consult [Hat05] for a definition and basic properties.

To define a global metric, we need a coherent notion of path in X between arbitrary points allowing us to make sense of the "distance" between them. At this point, one would intuitively think of the distance between points in X as the length of the shortest path joining them where the length of such paths would be given by adding the successive distances travelled within fixed simplices. With this in mind, an *m-string* in X from x to y is an ordered sequence of $m+1$ points in X with the property that successive pairs of points in the sequence lie in a common simplex. The length $l(\Sigma)$ of an $m-$string Σ can now be defined as the sum of the local distances between its successive pairs of points. Finally, as one would intuitively expect, the *quotient pseudometric*[53] on X is defined by the formula

$$d(x,y) := \inf\{l(\Sigma) | \Sigma \text{ is a string from } x \text{ to } y\}$$

where if there is no such string, we set $d(x,y) := \infty$. We henceforth ubiquitously endow X with this pseudo-metric that will turn out to be a length metric when $Shapes(X)$ is a finite set.

It is unfortunate but true that the quotient pseudometric is not a true metric in general, namely it can happen that $d(x,y) = 0$ while $x \neq y$. For instance, this happens if we consider a graph with two vertices and infinitely many edges joining them who's length tends to zero. To capture the essence of the property that cause this pathology, recall that the *star* of a point x is defined as

$$star(x) := \{int(S) : S \subset X : x \in S\}$$

and notice that these pathological examples arise when the distance (measured in the local metric) from some point $x \in X$ to $\overline{star(x)} \backslash star(x)$ is zero. More formally, we encapsulate this quantity as follows:

$$\epsilon(x) := \inf\{\epsilon(x,S) | S \subset K \text{ is a simplex containing } x\},$$

[53] We allow ourselves to define it in this way since it coincides with the general definition of quotient pseudo metrics. In [BH99] it is first called the intrinsic pseudometric.

where

$$\epsilon(x, S) := \inf\{d_{\hat{S}}(x, F) | F \text{ is a face of } S \text{ and } x \notin T\}.$$

If $S = \{x\}$ simply define $\epsilon(x, S)$ as ∞. When $\epsilon(x)$ is bounded away from zero, points $y \in X$ that lie very close to x with respect to the pseudo-metric must lie in a simplex containing x. Showing that this condition holds uniformly will allow us to draw a strong conclusion on the metric as shown by the proof of the following key theorem.

Theorem B.1. *If X is an M_k-simplicial complex with finitely many isometry types of faces then it is a complete geodesic length space when equipped with the intrinsic quotient pseudo-metric.*

The proof of this theorem exploits the finiteness property of $Shapes(X)$ to give a uniform lower bound on $\epsilon(x)$ as x ranges over the points of X. The following pivotal lemma indicates the second step of this reasoning.

Lemma B.2. *Suppose that $\epsilon(x)$ is uniformly bounded away from zero and let $x, y \in X$ be two points such that $d(x, y) < \epsilon(x)$. Then, any simplex S containing y necessarily contains x and $d_{\hat{S}}(x, y) = d(x, y)$. In particular, the intrinsic pseudometric is a metric and (X, d) is a length space.*

Proof. Since $d(x, y) < \epsilon(x)$, $\exists$ an $m-$string Σ joining x to y such that $l(\Sigma) \leq \epsilon(x)$. However, from $l(\Sigma) \leq \epsilon(x)$, it follows that the first three points of Σ lie in a common simplex S, so we can replace Σ by a shorter $m-$string, omitting the second point. Continuing in this manner, one eventually obtains a $1-$string Σ consisting of two points $\{x, y\}$ such that $d_{\hat{S}}(x, y) = d(x, y)$. □

Proof of the Theorem. By virtue of the lemma, to obtain the length space it suffices to show that $\epsilon(x) > 0$ for every $x \in X$. Consider $\hat{x} \in \hat{S}$ and define a quantity $\epsilon(\hat{x})$ as the minimum distance $d(\psi(\hat{x}), \hat{F}')$ over all fixed faces $\hat{F}' \subset \hat{S}' \in Shapes(X)$ and isometric embeddings $\psi : \hat{S} \hookrightarrow \hat{S}'$ onto a face of $\hat{S}'$ whose interior is disjoint from the face $\hat{F}' \subset \hat{S}'$. Since there are only finitely many shapes, $\epsilon(\hat{x}) > 0$ and if $x = \varphi_S(\hat{x})$ it follows that $\epsilon(x) \geq \epsilon(\hat{x}) > 0$ as desired.

We now know $\epsilon(x) \neq 0$. Suppose however that there is a sequence $(x_n) \in X$ such that $\epsilon(x_n) \to 0$. This can not happen since it would require shapes of arbitrarily small size contradicting the fact that $Shapes(K)$ is finite so in fact, $\epsilon(x) > c > 0$ for some constant c. To show that X is complete, consider a cauchy sequence (y_n). It must then hold that for some $N > 0$, for all $n, m \geq N$ we have $d(y_n, y_m) < c$. It then follows that for all $n \geq N$, $y_n \in S$, some fixed simplex of X. Such a simplex is complete so the sequence converges. □

In fact, the resulting space is also geodesic as shown in [BH99] and the result can be extended to arbitrary polyhedral complexes with a certain amount of work through the process of *barycentric subdivision* we will describe shortly:

Theorem B.3 ([BH99])**.** *Let X be an M_k-polyhedral cell complex equipped with the intrinsic quotient pseudo-metric. If the set of shapes of X is finite then it is a complete geodesic metric space*[54]*.*

An M_k-polyhedral cell complex is defined in essentially the same way as an M_k-simplicial complex. In this case, the building blocks are M_k-convex polyhedral cells $P \subset M_k^n$ defined as the convex hull of a finite sets of points[55] in M_k^n. The faces of P are its non-empty intersections with various hyperplanes of the ambient space such that P is contained in one of the closed half-spaces they define. As usual, the 0 dimensional ones are its vertices. To construct an M_k-polyhedral cell complex we start with a disjoint family of M_k-convex polyhedral cells $\{P_\lambda \subset M_k^{n_\lambda} : \lambda \in \Lambda\}$ and glue them together along isometric faces of our choice to obtain a connected complex X. All notions previously defined in the simplicial case cary over to polyhedral complexes with minor adjustments. In fact, given a polyhedral complex, we can give it the structure of a simplicial complex using the following procedure.

Given any convex polyhedral cell $P \subset M_k^n$, there is a point $b_P \in P$ fixed by all isometries of the cell called the barycentre of the vertices of P. The *first*

[54]Recall that we require complexes to be connected!

[55]If $k > 0$, the points need to lie in an open ball of radius $D_k/2$ in order to ensure the existence of required geodesics.

barycentric subdivision of P is the M_k-simplicial complex P' defined as follows: given a strictly increasing sequence of faces $F_0 \subset F_1 \subset \ldots \subset F_n$, there is a geodesic simplex corresponding to the convex hull of the barycentres of the F_i's. The collection of these simplices impose the necessary structure on P.

If we apply this process individually to the cells of X, an M_k-polyhedral cell complex, we obtain its first barycentric subdivision X'. This X' is a polyhedral complex whose cells are all geodesic simplices but it is not always a simplicial complex as the intersection of two simplexes may only be the union of faces. To remedy this, we repeat the subdivision process a second time and obtain the *second barycentric subdivision* X'' which always has the structure of a simplicial complex.

Example Begin with a cylinder depicted as a square where a pair of opposite edges are identified. This is a very simple polyhedral complex consisting of a single cell, namely the square. Notice that in the first barycentric subdivision, the two identified triangles intersect at two disjoint vertices (three dots in the picture) when in a simplicial complex they should intersect on a single face. This complication does not occur in the second barycentric subdivision as the identified edges are "far enough apart" as illustrated in Figure 19.

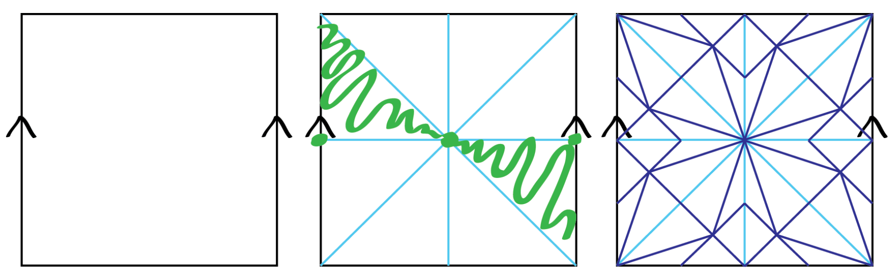

Figure 19: Barycentric subdivision in a cylinder.

C A Word on Amalgamated Products

Due to their appearance in a few key places in the text, we place a few words on amalgamated free products and Higman-Neumann-Neumann (HNN) extensions based on the treatment given in [LS77].

Given two groups Γ_1 and Γ_2 with presentations $\Gamma_1 = \langle \mathcal{A}_1 | \mathcal{R}_1 \rangle$ and $\Gamma_2 = \langle \mathcal{A}_2 | \mathcal{R}_2 \rangle$, one can form their *free product* which is concretely defined as the disjoint union of their presentations $\Gamma_1 * \Gamma_2 = \langle \mathcal{A}_1 \sqcup \mathcal{A}_2 | \mathcal{R}_1 \sqcup \mathcal{R}_2 \rangle$. This construction instantiates the coproduct in the category of groups and readily generalizes to concepts mentioned above as follows. Suppose we are given subgroups $H_1 \leq \Gamma_1$ and $H_2 \leq \Gamma_2$ along with an isomorphism $\phi : H_1 \to H_2$. The *free product with amalgamation* of Γ_1 and Γ_2 along the subgroups H_1 and H_2 by the isomorphism ϕ is defined by the presentation

$$\Gamma_1 *_{H_1 \simeq H_2} \Gamma_2 = \langle \mathcal{A}_1 \sqcup \mathcal{A}_2 | \mathcal{R}_1 \sqcup \mathcal{R}_2, h_1 = \phi(h_1) \text{ for every } h_1 \in H_1 \rangle.$$

More generally, if we are given a group H along with a family of monomorphisms $\phi_\lambda : H \hookrightarrow \Gamma_\lambda$, the amalgamated free product of the Γ_λ along the isomorphic images of the group H is defined as the quotient of $*_\lambda \Gamma_\lambda$ by the normal closure of the subgroup generated by elements of the form $\phi_\lambda(h)\phi_{\lambda'}(h)^{-1}$. Notice that for every λ there is a natural inclusion of Γ_λ into the amalgamated product.

The construction of HNN extensions is similar in flavour to the amalgamated product. Let $\Gamma = \langle \mathcal{A} | \mathcal{R} \rangle$ be a group along with two subgroups H_1 and H_2 isomorphic via the map $\phi : H_1 \to H_2$. The HNN extension of Γ relative to the subgroups H_1 and H_2 related by the map ϕ is defined concretely as

$$\Gamma *_H := \langle \mathcal{A}, t | \mathcal{R}, \phi(h) t^{-1} h t \ \forall \ h \in H_1 \}$$

where the letter $t \notin \mathcal{A}$ is called the *stable letter* and H is an abstract group isomorphic to the *associated subgroups* H_1 and H_2.

In fact, both constructions can be thought of as being part of Stalling's unifying concept of a bipolar structure [Sta68]. Informally, they both involve two subgroups related by an isomorphism but in the first case the subgroups are in different groups while in the second they both lie in the same one hence one might say the first case is "disconnected" while the later is "connected". To give both constructions a concrete realization we illustrate their relation to topology in two key examples found in [LS77] but first we state the following crucial theorem referring the reader to [Hat05] for a proof.

Theorem C.1 (Seifert-van Kampen)**.** *Let X be a topological space with a distinguished basepoint x_0 and suppose that it can be written as the union of open sets $X = \cup_i A_i$ where $x_0 \in A_i$ and A_i is path connected for each i. Denote by φ_{ij} the homomorphism of fundamental groups $\pi_1(A_i \cap A_j, x_0) \to \pi_1(A_i, x_0)$ induced by the inclusion. If $A_i \cap A_j \cap A_k$ is path connected for all choices of i, j and k then $\pi_1(X) \cong *_i \pi_1(A_i)/\langle \varphi_{ij}(\sigma)\varphi_{ji}(\sigma)^{-1} : \sigma \in \pi_1(A_i, x_0) \cap \pi_1(A_j, x_0)\rangle$.*

Example Let X and Y be path-connected topological spaces.

1. Let U be an open path-connected subspace of X homeomorphic by a function h to the open path-connected subspace $V \subset Y$. Choose a base point $u \in U$ for the fundamental group $\pi_1(X)$ and a corresponding one $f(u) \in V$. If we assume that the homomorphisms $\pi_1(U) \to \pi_1(X)$ and $\pi_1(V) \to \pi_1(Y)$ induced by the inclusion map are injective then the homeomorphism h induces an isomorphism $h_* : \pi_1(U) \to \pi_1(V)$. A simple application of the Seifert-van Kampen Theorem implies that the space Z obtained by identifying U with V has fundamental group $\pi_1(Z) = \pi_1(X) *_H \pi_1(Y)$ where $H \simeq \pi_1(U) \simeq \pi_1(V)$.

2. Similarly, if we suppose that in the above scenario U and V are both subspaces of X and we attach a handle to X from U to V the Seifert-van Kampen Theorem show that for the resulting space Z, the fundamental group $\pi_1(Z)$ is the HNN extension $\pi_1(X)*_H$.

D Disc Diagrams

Disc diagrams are central to much of geometric group theory as developed in [LS77] and we will use them in several instances. The following elegant definition of the concept can be found in [Wis11].

Disc diagram A *disc diagram* or *van Kampen diagram* D is a compact, connected, simply connected and contractible combinatorial 2−complex equipped with a chosen embedding $D \hookrightarrow \mathbb{R}^2$. Thinking of $\mathbb{S}^2$ as $\mathbb{R}^2 \cup \infty$ we define the boundary path of the disc diagram $\partial_p D$ as the attaching map of the 2−cell containing the point at infinity which is missing to give $\mathbb{S}^2$ a cell structure.

Most of the time, one refers to a disc diagram in a complex X which is a combinatorial map $D \to X$. The key to using disc diagrams lies in the following result known as van Kampen's Lemma for which our elegant statement is once again taken from [Wis11]. Although the result is originally due to van Kampen, we refer the reader to [LS77] for a modern proof.

Lemma D.1 ([VK33])**.** *Let X be a CW complex. A closed combinatorial path $P \to X$ is nullhomotopic if and only if there is a disc diagram $D \to X$ with $P \simeq \partial_p D$ such that the following diagram commutes:*

$$\begin{array}{ccc} \partial_p D & \longrightarrow & D \\ \| & & \downarrow \\ P & \longrightarrow & X \end{array}$$

E Notation

We enumerate here various frequently used notations for the convenience of the reader, more or less in the order in which they appear. They are also defined in the text where appropriate.

1. $d(x, y)$ and $B(x, \epsilon)$: distance function for metric spaces and their open balls.

2. $[x, y]$: geodesic segment between points x and y.

3. $Iso(X)$ and $Homeo(X)$: isometry and homeomorphism groups.

4. M^n_k,D_k, $\mathbb{E}^n$, $\mathbb{H}^n$ and $\mathbb{S}^n$: model spaces of constant curvature k and their diameter followed by Euclidean, hyperbolic and spherical space.

5. $\bar{\triangle}$, $\bar{\angle}_p(q,r)$: comparison triangle and angle, usually in euclidean space.

6. $D_{g_1,g_2}(s)$: the distance between two geodesics g_1 and g_2 parametrized proportionally to arc length at time s.

7. $\Gamma = \langle \mathcal{A} | \mathcal{R} \rangle$, $C_{\mathcal{A}}(\Gamma)$, $\mathfrak{C}_{\mathcal{A},\mathcal{R}}(\Gamma)$, $\mathfrak{S}_{\mathcal{A},\mathcal{R}}$: the group Γ given by a presentation on the generators $\mathcal{A}$, the Cayley graph, Cayley complex and standard $2-$complex of the presentation.

8. $\Gamma \curvearrowright X$: the group Γ acting on the space X by isometry or homemorphism.

9. $C_0(X)$, $S_x(X)$: Euclidean cone and space of directions.

10. $Link(x,X)$ is the geometric link of x in the cell complex X.

11. $Shapes(X)$, $d_{\hat{S}}$: the set of isometry classes $\hat{S}$ of faces S of a cell complex X and the local metric of such a complex, d is reserved for the quotient length metric.

12. $f : X \xrightarrow{f'} Y \xrightarrow{g} Z$: tower lift of f.

13. $\mathcal{K} = \mathcal{K}_1 \# \ldots \# \mathcal{K}_n$, $\mathcal{D}(\Pi)$: a knot decomposed as a connected sum and the Dehn complex of a knot projection Π.

14. D, ∂D: euclidean disc and its boundary path.

15. $d_\sigma : X \to \mathbb{R}_{\geq 0}$, δ_σ, $Min(\sigma)$: displacement function and translation length of σ followed by the set of points at which σ attains its translation length.

16. ∂X, $\overline{X}$: The visual boundary and bordification of X.

17. $d_{\mathcal{H}}(A,B)$: Hausdorff distance between subsets A and B of a metric space.

18. $L^2(\Omega, X)$, $\rho^2(\varphi, \psi)$: space of equivariant maps and its distance.

19. d_Σ : the displacement map $\max_{\sigma\in\Sigma} d(\sigma \cdot x, x)$.

References

[AB90] S.B. Alexander and R.L. Bishop. The hadamard-cartan theorem in locally convex metric spaces. *Enseign. Math*, 36:309–320, 1990.

[AGM12] I. Agol, D. Groves, and J. Manning. The virtual haken conjecture. *Arxiv preprint arXiv:1204.2810*, 2012.

[Ale51] A.D. Aleksandrov. A theorem on triangles in a metric space and some of its applications. *Trudy Matematicheskogo Instituta im. VA Steklova*, 38:5–23, 1951.

[AR90] IR Aitchison and JH Rubinstein. An introduction to polyhedral metrics of non-positive curvature on 3-manifolds. *Geometry of Low-Dimensional Manifolds*, 2:127–161, 1990.

[Ber83] VN Berestovskii. Borsuk's problem on the metrization of a polyhedron. In *Soviet Math. Dokl*, volume 27, pages 56–59, 1983.

[Ber03] M. Berger. *A panoramic view of Riemannian geometry*. Springer Verlag, 2003.

[BH99] M.R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319. Springer Verlag, 1999.

[BN08] ID Berg and IG Nikolaev. Quasilinearization and curvature of aleksandrov spaces. *Geometriae Dedicata*, 133(1):195–218, 2008.

[Bri99] M.R. Bridson. On the semisimplicity of polyhedral isometrics. *Proceedings-American Mathematical Society*, 127:2143–2146, 1999.

[Bri07] M. Bridson. Problems concerning hyperbolic and cat (0) groups. *Preprint*, 2007.

[Bro89] K.S. Brown. *Buildings*. Springer New York etc., 1989.

[BT72] F. Bruhat and J. Tits. Groupes réductifs sur un corps local. i. données radicielles valuées. *Publ. math. IHES*, 41(5):251, 1972.

[Bur94] M. Burger. Rigidity properties of group actions on cat (0)-spaces. In *Proceedings of the International Congress of Mathematicians*, volume 1, page 2, 1994.

[Bus48] H. Busemann. Spaces with non-positive curvature. *Acta Mathematica*, 80(1):259–310, 1948.

[Bus55] H. Busemann. *The geometry of geodesics*, volume 6. Academic press, 1955.

[BZ03] G. Burde and H. Zieschang. *Knots*, volume 5. de Gruyter, 2003.

[Deh87] M. Dehn. *Papers on group theory and topology: Translated and introduced by John Stillwell.* Springer, 1987.

[DG08] F. Dahmani and D. Groves. The isomorphism problem for toral relatively hyperbolic groups. *Publications mathematiques de l'IHES*, 107(1):211–290, 2008.

[DG11] F. Dahmani and V. Guirardel. The isomorphism problem for all hyperbolic groups. *Geometric And Functional Analysis*, pages 1–78, 2011.

[DK09] C. Drutu and M. Kapovich. Lectures on geometric group theory. *Preprint (spring 2009 revision)*, 2009.

[DL09] T.K. Dey and K. Li. Cut locus and topology from surface point data. In *Proceedings of the 25th annual symposium on Computational geometry*, pages 125–134. ACM, 2009.

[DO85] K. Diederich and T. Ohsawa. Harmonic mappings and disc bundles over compact kahler manifolds. *Publ. Res. Inst. Math. Sci. Kyoto Univ*, 21:819–833, 1985.

[Don87] S.K. Donaldson. Twisted harmonic maps and the self-duality equations. *Proceedings of the London Mathematical Society*, 3(1):127–131, 1987.

[EO73] P. Eberlein and B. O'Neill. Visibility manifolds. *Pacific Journal of Mathematics*, 46(1):45–109, 1973.

[FF99] G.B. Folland and GB Folland. *Real analysis: modern techniques and their applications*, volume 361. Wiley New York, 1999.

[Fre31] H. Freudenthal. Über die enden topologischer räume und gruppen. *Mathematische Zeitschrift*, 33(1):692–713, 1931.

[GKM08] T. Gelander, A. Karlsson, and G.A. Margulis. Superrigidity, generalized harmonic maps and uniformly convex spaces. *Geometric and Functional Analysis*, 17(5):1524–1550, 2008.

[Gro81] M. Gromov. Groups of polynomial growth and expanding maps. *Publications Mathématiques de l'IHÉS*, 53(1):53–78, 1981.

[Gro87] M. Gromov. Hyperbolic groups. *Essays in group theory*, 8:75–263, 1987.

[Gro03] M. Gromov. Random walk in random groups. *Geometric and Functional Analysis*, 13(1):73–146, 2003.

[Had98] J. Hadamard. Les surfaces à courbures opposées et leurs lignes géodésique. *J. Math. Pures Appl.*, 4:27–73, 1898.

[Hag07] F. Haglund. Isometries of cat (0) cube complexes are semi-simple. *Arxiv preprint arXiv:0705.3386*, 2007.

[Hat05] A. Hatcher. *Algebraic topology*. 2005.

[Hop43] H. Hopf. Enden offener räume und unendliche diskontinuierliche gruppen. *Commentarii Mathematici Helvetici*, 16(1):81–100, 1943.

[Hot97] P.K. Hotchkiss. The boundary of a busemann space. *Proceedings of the American Mathematical Society*, 125(7):1903–1912, 1997.

[How81] J. Howie. On pairs of 2-complexes and systems of equations over groups. *J. Reine Angew. Math*, 324:165–174, 1981.

[IN05] H. Izeki and S. Nayatani. Combinatorial harmonic maps and discrete-group actions on hadamard spaces. *Geometriae Dedicata*, 114(1):147–188, 2005.

[Jos97] J. Jost. *Nonpositive curvature: geometric and analytic aspects*. Birkhauser, 1997.

[JY90] J. Jost and S.T. Yau. *Harmonic maps and group representations*. Univ., Sonderforschungsbereich 256, 1990.

[KS93] N.J. Korevaar and R.M. Schoen. Sobolev spaces and harmonic maps for metric space targets. *Comm. Anal. Geom*, 1(4):1–99, 1993.

[LS77] R.C. Lyndon and P.E. Schupp. *Combinatorial group theory*, volume 89. Springer Verlag, 1977.

[Mac64] AM Macbeath. Groups of homeomorphisms of a simply connected space. *The Annals of Mathematics*, 79(3):473–488, 1964.

[Mar91] G.A. Margulis. *Discrete subgroups of semisimple Lie groups*, volume 17. Springer, 1991.

[Mil68] J. Milnor. A note on curvature and fundamental group. *J. Diff. Geom*, 2(1):1–7, 1968.

[Mon06] N. Monod. Superrigidity for irreducible lattices and geometric splitting. *Journal of the American Mathematical Society*, 19(4):781–814, 2006.

[Mun00] J.R. Munkres. Topology, 2nd, 2000.

[Nik95] I. Nikolaev. The tangent cone of an aleksandrov space of curvature k. *manuscripta mathematica*, 86(1):137–147, 1995.

[Pan06] P. Pansu. Superrigidité géométrique et applications harmoniques. *Arxiv preprint math/0612736*, 2006.

[Pap57] C.D. Papakyriakopoulos. On dehn's lemma and the asphericity of knots. *Proceedings of the National Academy of Sciences of the United States of America*, 43(1):169, 1957.

[Pap05] A. Papadopoulos. *Metric spaces, convexity and nonpositive curvature*, volume 6. European Mathematical Society, 2005.

[Rip82] E. Rips. Subgroups of small cancellation groups. *Bulletin of the London Mathematical Society*, 14(1):45, 1982.

[Rol03] D. Rolfsen. *Knots and links*. American Mathematical Society, 2003.

[Sag95] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, 3(3):585, 1995.

[Sat09] Takashi Sato. An alternative proof of berg and nikolaevs characterization of cat(0)-spaces via quadrilateral inequality. *Archiv der Mathematik*, 93:487–490, 2009. 10.1007/s00013-009-0057-9.

[Sel95] Z. Sela. The isomorphism problem for hyperbolic groups i. *Annals of Mathematics*, pages 217–283, 1995.

[Sta68] J.R. Stallings. On torsion-free groups with infinitely many ends. *The Annals of Mathematics*, 88(2):312–334, 1968.

[Sva55] AS Svarc. A volume invariant of coverings.(russian). In *Dokl. Akad. Nauk SSSR (NS)*, volume 105, pages 32–34, 1955.

[VK33] E.R. Van Kampen. On some lemmas in the theory of groups. *American Journal of Mathematics*, 55(1):268–273, 1933.

[Wan98] M.T. Wang. A fixed point theorem of discrete group actions on riemannian manifolds. *J. Differential Geom*, 50(2):249–267, 1998.

[Wan00] M.T. Wang. Generalized harmonic maps and representations of discrete groups. *Communications in analysis and geometry*, 8(3):545–564, 2000.

[Wei71] CM Weinbaum. The word and conjugacy problems for the knot group of any tame, prime, alternating knot. In *Proc. Amer. Math. Soc*, volume 30, pages 22–26, 1971.

[Wis96] D. Wise. Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups. *Princeton thesis*, 1996.

[Wis98] D.T. Wise. Incoherent negatively curved groups. *Proceedings of the American Mathematical Society*, 126(4):957–964, 1998.

[Wis05] D.T. Wise. The coherence of one-relator groups with torsion and the hanna neumann conjecture. *Bulletin of the London Mathematical Society*, 37(5):697, 2005.

[Wis06] D.T. Wise. Subgroup separability of the figure 8 knot group. *Topology*, 45(3):421–463, 2006.

[Wis11] D. Wise. The structure of groups with a quasiconvex hierarchy, 186 pages, preprint (2011) downloaded on september 9. In *2011 from the conference webpage for the NSFCBMS conference 3-Manifolds, Artin Groups and Cubical Geometry, from August*, pages 1–5, 2011.