

UNIVERSITÀ DI PISA



CORSO DI LAUREA MAGISTRALE IN MATEMATICA

Non smoothable locally CAT(0) spaces

17 ottobre 2014

TESI DI LAUREA MAGISTRALE

Candidato
**Kirill
Kuzmin**

Relatore
**Dott.
Roberto Frigerio**

Controrelatore
**Prof.
Carlo Petronio**

ANNO ACCADEMICO 2013/2014

La filosofia è scritta in questo grandissimo libro che continuamente ci sta aperto innanzi a gli occhi (io dico l'universo), ma non si può intendere se prima non s'impara a intender la lingua, e conoscer i caratteri, ne' quali è scritto. Egli è scritto in lingua matematica, e i caratteri son triangoli, cerchi, ed altre figure geometriche, senza i quali mezzi è impossibile a intenderne umanamente parola; senza questi è un aggirarsi vanamente per un oscuro laberinto.

- Galileo Galilei, *Il Saggiatore*

Introduction

In this dissertation we will talk about manifolds of curvature bounded above in the sense of CAT, an approach originally formalized by Aleksandrov in [Ale51]. The concept is general, but we are interested in seeing how much freedom it leaves in the manifold setting. Let us introduce some terminology to set the problem.

A geodesic is an isometric embedding of an interval into a metric space. We consider geodesic metric spaces, where for any pair of points there is a geodesic having them as extremes. Then, given 3 points in a geodesic space, we can construct a geodesic triangle having them as vertices, and define a concept of curvature bounded above by κ by comparing triangles in the space with triangles having sides of the same length in 2-dimensional model spaces of curvature κ ; that is, the Euclidean space for $\kappa = 0$, a suitably rescaled sphere for $\kappa > 0$ and a suitably rescaled hyperbolic space for $\kappa < 0$. A space is locally CAT(κ) if every point has a neighbourhood which is CAT(κ).

The (locally) CAT(κ) condition allows us to describe the curvature of a large class of metric spaces, of which polyhedral metric complexes are notable representatives: these are spaces made up by polyhedra attached via isometries. If we restrict our attention to Riemannian manifolds, they are locally CAT(κ) if and only if all of their sectional curvatures are less or equal than κ . We may ask ourselves, however, if for manifolds the requirement on a distance function compatible with the topology to be induced by a Riemannian metric is restrictive when talking about curvature. The answer is yes; and we will in fact exhibit a smooth closed 4-manifold that supports a locally CAT(0) distance function but is not homeomorphic to any Riemannian manifold of non positive curvature.

We will follow the approach presented in [DJL12]. The manifold will be the geometric realization of a cubical complex, i.e. polyhedral complexes made up by euclidean cubes. Cubical complexes are a very important tool in the geometry of non positively curved spaces; an example is provided by the recent proof of the Virtual Fibration Conjecture. Their power comes from the fact that the locally CAT(0) condition on them can be verified in purely combinatorial terms. In particular, we have to look at a neighbourhood of a vertex, which is a cone on a simplicial complex called *link* at that vertex.

We begin the exposition with a review of definitions on curvature of metric spaces. Then we construct, for a given complete CAT(0) space X , its boundary at infinity ∂X . We continue by describing a natural structure of metric space on finitely generated groups, and an equivalence relation called *quasi isometry* useful in this context. Quasi isometry invariants become more evident when seen through a construction allowing us to see the

space “from the infinity”, namely the asymptotic cone. With the strength of this construction, we define hyperbolic groups, hyperbolic spaces, and generalizations of these concepts: relatively hyperbolic spaces and groups, and spaces with isolated flats. Hyperbolic spaces and spaces with isolated flats share a nice behaviour of their boundary at infinity when talking about quasi isometries which are *quasi equivariant*, i.e. equivariant with respect to a cocompact properly discontinuous action of a group G ; in particular, such quasi-isometry induces a homeomorphism between the boundaries at infinity.

We then pass to examine polyhedral complexes and prove the facts we need in the sequel, including the description of the curvature in combinatorial terms we have already cited.

The cubical complex of the main theorem’s thesis will be constructed starting from a link at any of its vertices. To be a 4-manifold, this link has to be a triangulation of \mathbb{S}^3 with certain properties. We will define these properties and construct the required triangulation.

Finally, we pass to describing the manifold. There is a standard way, introduced by Davis in [Dav08], to create a cubical complex once known the link we want its vertices to have. This construction also allows to give a concise description of the fundamental group of the manifold, which we will call P . If P was homeomorphic to a smooth Riemannian manifold M of non positive curvature, the structure of the triangulation and thus of the fundamental group of both manifolds would allow us to conclude that the respective universal covers \tilde{P} and \tilde{M} would have both an isometrically embedded euclidean plane, with the embedding totally geodesic in the smooth case, and that the boundaries at infinity, topologically \mathbb{S}^3 in both cases, would be homeomorphic, with the homeomorphism taking the boundary of the flat found in \tilde{P} to the boundary of the flat found in \tilde{M} . This will turn to be impossible: the boundary of the flat is a knot in \mathbb{S}^3 , but it will be non trivial in the former case and trivial in the second.

Contents

Introduction	4
1 Curvature of metric spaces	7
1.1 Basic definitions	7
1.2 Constructions in non-positive curvature	10
1.3 Tools from Geometric Group Theory	13
1.4 Isolated flats condition	28
2 Polyhedral metric complexes	39
2.1 Polyhedral complexes	39
2.2 Links	42
3 Triangulations of S^3	48
3.1 Preliminary definitions	48
3.2 The 600-cell	48
3.3 The triangulation with a knot	50
4 The manifold	53
4.1 The Davis complex	53
4.2 The desired manifold	55
Bibliography	59

1 Curvature of metric spaces

We are first going to recover some basic definitions on metric spaces and on how one can define a concept of curvature bounded above on a particular class of them. Then we are going to explore some recent developments in the theory of spaces of non positive curvature.

1.1 Basic definitions

We recover quickly some notations we will use throughout the text. The distance function on a metric space will be often denoted by d , and we will call with the same letter the distance of a point x from a subset A , which is the infimum over a in A of $d(x, a)$. The R -neighbourhood of a subset of a metric space is the set of all points with distance strictly less than R from that subset; we will call closed R -neighbourhood the analogous concept where the distance can be equal or less than R . The Hausdorff distance between two subsets is the infimum of R such that the first is contained in the R -neighbourhood of the second, and vice versa. It fulfils the axioms of a distance, but it may be infinite for unbounded sets.

The diameter of a metric space is the supremum of the distances between its points. The open balls of centre x_0 and radius r will be denoted with $B(x_0, r)$ and the closed ones with $\overline{B}(x_0, r)$. A local isometry is a function between metric spaces which restricts to an isometry on a collection of open sets covering the domain.

Definition 1.1.1 (Geodesics). *Let X be a metric space. A geodesic in X is an isometric embedding of a real interval in X .*

If a geodesic is defined on a closed interval, it is said to be a geodesic segment, and images of the endpoints of the interval are the endpoints of the geodesic.

A geodesic defined on a half line will be called geodesic ray, whilst a geodesic line is defined on the whole \mathbb{R} .

A closed geodesic is an isometric embedding of a suitably rescaled copy of \mathbb{S}^1 into the space.

A local geodesic is a locally isometric immersion of an interval or of a suitably rescaled \mathbb{S}^1 into the space.

For brevity, we will often designate the image of a geodesic with the same name. We will do the same for paths.

Talking about paths, if we have a suitable definition of length for a collection of them, we may define a pseudo-distance on the space, i.e. a function with the same properties of distance except it may be 0 even between different

points. The length has to be non negative and additive, i.e. the concatenation of paths must be a path of the collection and its length must be equal to the sum of the lengths of then original paths, and for every pair of points in the space there must be a path of the collection between them. In this setting, we define the pseudo distance of a pair of points to be the infimum of the lengths of paths between them. This becomes a distance if we can argue somehow that the length of any path between two different points must be bounded below by a positive constant depending on the points.

The typical examples are the integral definition of length of piecewise \mathcal{C}^1 paths in a Riemannian manifold and the induced length of continuous paths in a metric space. We recover it briefly:

Definition 1.1.2 (Length of a continuous curve). *Let (X, d) be a metric space and $[a, b]$ a real interval (the definition adapts as well to other intervals and to \mathbb{S}^1) and let $\mathcal{P} = \{a = t_0, t_1, \dots, t_n = b\}$ be a partition of $[a, b]$. The length of a curve $\alpha: [a, b] \rightarrow X$ is*

$$\sup_{\mathcal{P}} \sum_{i=1}^n d(\alpha(t_i), \alpha(t_{i-1})).$$

The induced distance is called *induced length metric* and leads to a distance greater or equal than the original one; on subspaces it is not to be confused with the induced distance, which is the plain restriction of the distance function of the space. If the induced length metric coincides with the original distance, we say that the metric space is a *length space*.

The metric spaces which we will consider in this text will usually have many geodesics.

Definition 1.1.3 (Geodesic metric space). *A metric space X is said to be geodesic if for every x and y in X there is a geodesic segment in X with endpoints x and y . We will indicate one such segment with the notation xy .*

X is said to be uniquely geodesic if such segment is unique.

A geodesic space is obviously a length space. Geodesics allow us to define convexity in geodesic metric spaces. A subspace is said to be *convex* if any geodesic between its points lies entirely in the subspace.

We will now describe the simplest metric spaces, which we will use for comparisons and further definitions; they are called *model spaces*.

Definition 1.1.4 (Model spaces). *Let κ be a real number and $n \geq 2$ be an integer. The model space of curvature κ and dimension n , denoted by \mathbb{M}_{κ}^n , is the unique complete simply connected Riemannian manifold of dimension n and all sectional curvatures equal to κ . In particular,*

- If $\kappa < 0$, the hyperbolic space \mathbb{H}^n with distance function multiplied by $\frac{1}{\sqrt{-\kappa}}$;
- If $\kappa = 0$, the euclidean space \mathbb{E}^n ;
- If $\kappa > 0$, the sphere \mathbb{S}^n with distance function multiplied by $\frac{1}{\sqrt{\kappa}}$.

All model spaces are uniquely geodesic, except for $\kappa > 0$; in that case the geodesic is unique only for points closer than the diameter of the space, equal to $\frac{\pi}{\sqrt{\kappa}}$. The cases we will mostly use throughout this text are $\kappa = 0, -1, 1$.

Given 3 points x, y, z in a geodesic metric space X , a *geodesic triangle* or simply *triangle* with vertices x, y, z is the union of three segments xy, yz, zx ; we will denote it with the notation Δxyz .

Definition 1.1.5 (CAT(κ) space). *Let κ be a real number, and let \bar{d} be the distance function in \mathbb{M}_κ^2 . A geodesic metric space (X, d) is CAT(κ) if for every triangle Δxyz in X , with the additional hypothesis that*

$$d(x, y) + d(y, z) + d(z, x) < \frac{2\pi}{\sqrt{\kappa}}$$

if $\kappa > 0$, and for every choice of points p on the side xy and q on the side xz , the inequality $d(p, q) \leq \bar{d}(\bar{p}, \bar{q})$ holds, where \bar{p} and \bar{q} are points on the sides $\bar{x}\bar{y}$ and $\bar{x}\bar{z}$ of the (unique up to isometry) triangle $\Delta \bar{x}\bar{y}\bar{z}$ in \mathbb{M}_κ^2 with $d(x, y) = \bar{d}(\bar{x}, \bar{y})$ and so on, such that $d(x, p) = \bar{d}(\bar{x}, \bar{p})$ and $d(x, q) = \bar{d}(\bar{x}, \bar{q})$.

The triangle $\Delta \bar{x}\bar{y}\bar{z}$ of the definition is called *comparison triangle* for Δxyz .

The CAT(κ) condition just provided is global. Let us exhibit a local definition, which gives a good generalization of the concept of a manifold with curvature bounded above by κ .

Definition 1.1.6 (Locally CAT(κ) space). *A metric space is said to be locally CAT(κ), or of curvature bounded above by κ , if every point has a neighbourhood which is CAT(κ).*

A CAT(0) space will be shortly called of *non-positive curvature*.

The following fact, which is proven in many classic texts like [BH99, Appendix II.1A], is the starting point of the problem here exposed and solved.

Theorem 1.1.7. A Riemannian manifold with its distance function has riemannian curvature less or equal than κ on every 2-subspace of every tangent space if and only if it is a locally CAT(κ) space.

The purpose of this exposition is to prove that for a smooth manifold endowed with a distance function compatible with the topology the request that this distance derives from a Riemannian structure is actually a restricting condition, when talking about the curvature. The precise statement is the following:

Theorem 1.1.8. There is a closed smooth 4-manifold endowed with a distance function that makes it a locally CAT(0) space but which is not homeomorphic to any Riemannian manifold of non positive curvature.

The dimension 4 is the least possible for a counterexample. In dimension 1 there are only two connected manifolds: \mathbb{R} and \mathbb{S}^1 . The classification theorem for compact surfaces implies that a compact surface supports a locally CAT(0) distance function if and only if it supports a Riemannian metric of non positive curvature. In dimension 3 the question is subtler, but the fact that a compact locally CAT(0) manifold supports a Riemannian metric of non positive curvature is still true [DJL12, Proposition 1]. The proof uses heavily Thurston Geometrization Conjecture, which now a Theorem thanks to Perel'man's work, which is, as far as we know, the best substitute for a classification theorem in the 3-dimensional setting.

1.2 Constructions in non-positive curvature

We now exhibit some standard constructions in the setting of CAT(0) spaces. We first need the following definition.

Definition 1.2.1 (Asymptotic rays). *Two rays c, c' defined on $[0, +\infty)$ in a CAT(0) space (X, d) are said to be asymptotic if the function*

$$t \mapsto d(c(t), c'(t))$$

is bounded.

Note that in general, if two geodesics γ, γ' in a CAT(0) space (X, d) are defined on intervals $[a, b]$ and $[a', b']$ respectively, then the function defined on $[0, 1]$ and sending s to $d(\gamma(a + s(b - a)), \gamma'(a' + s(b' - a)))$ is convex [BH99, Proposition II.2.2]. It follows in particular that two asymptotic rays with the same starting point are necessarily the same.

Definition 1.2.2 (Boundary at infinity). *Let X be a complete CAT(0) space. The boundary at infinity of X , denoted by ∂X , is the set of all rays in X quotiented by the equivalence relation of being asymptotic.*

We shall also denote by \overline{X} the set $X \cup \partial X$. We will see that it has a topological structure which coincides with the original one when restricted to X . Note that an isometry takes asymptotic rays in asymptotic rays, so it extends naturally to a function on \overline{X} , which we will see to be a homeomorphism with respect to the topology of \overline{X} . To move ourselves towards this result we first need the following

Definition 1.2.3 (Angle). *Let (X, d) be a metric space and $\gamma_1: [0, a_1] \rightarrow X$, $\gamma_2: [0, a_2] \rightarrow X$ be geodesics with $\gamma_1(0) = \gamma_2(0)$. The angle between γ_1 and γ_2 is defined recurring to the law of cosines:*

$$\angle(\gamma_1, \gamma_2) = \limsup_{(t_1, t_2) \rightarrow (0, 0)} \arccos \left(\frac{t_1^2 + t_2^2 - d(\gamma_1(t_1), \gamma_2(t_2))}{2t_1 t_2} \right).$$

It is easy to see that the angle is subadditive, that the angle between a geodesic and itself is 0 and that, if γ is a geodesic defined in a neighbourhood of 0, the angle between the two opposite parts of γ with respect to 0 is π . Any other angle has a value between 0 and π . It is a matter of calculation that the angle between two geodesics on a Riemannian manifold is the euclidean angle in the tangent space, with scalar product given by the actual Riemannian metric between the two tangent vectors to the geodesics [BH99, Proposition II.1A.7]. It can also be proven [BH99, Proposition II.1.7] the following characterization of CAT(κ) spaces:

Lemma 1.2.4. Let κ be a real number. A geodesic metric space X is CAT(κ) if and only if in any geodesic triangle in X the angles at the vertices are less or equal than the corresponding angles of a comparison triangle.

In particular, if $\kappa = 0$, the sum of the angles in any geodesic triangle of a CAT(0) space is less or equal than π . This turns useful in the proof of the following

Lemma 1.2.5. For every x_0 in a complete CAT(0) space (X, d) and for every element of ∂X there is a geodesic ray $c: [0, +\infty) \rightarrow X$ with $c(0) = x_0$ representing it.

Proof. Let γ be a geodesic ray representing an element of ∂X and let y_0 be its endpoint. Define c_n to be the geodesic segment $x_0\gamma(n)$. By the triangle inequality, when we fix t the sequence $(c_n(t))_{n \in \mathbb{N}}$ is defined for every n big enough. We claim that it is a Cauchy sequence, so by completeness it has a limit, and by continuity of the distance the function

$$c(t) = \lim_{n \rightarrow +\infty} c_n(t)$$

is a geodesic. The function $d(c(t), \gamma(t))$ is convex and therefore it is always less or equal than $d(x_0, y_0)$ by construction and continuity of the distance function.

So it remains to prove the Cauchy property. We suppose $m > n$ are non negative integers, big enough that $c_i(t)$ is defined for any $i \geq n$. The angle at $\gamma(n)$ in the comparison triangle $\Delta_{x_0 y_0 \gamma(n)}$ can be made arbitrarily small if n is big enough because of the triangle inequality. The corresponding angle in X is less or equal than it, but its sum with the angle at $\gamma(n)$ in $\Delta_{x_0 \gamma(n) \gamma(m)}$ is π , so the latter can be made arbitrarily close to π . The corresponding angle in $\Delta_{x_0 \gamma(n) \gamma(m)}$ is greater or equal than the angle in the space, so the angle α at \bar{x}_0 can be made arbitrarily small. The distance $d(c_n(t), c_m(t))$ is less or equal than the distance between $\overline{c_n(t)}$ and $\overline{c_m(t)}$, which is equal to $2t \sin(\frac{\alpha}{2})$, and α can be made arbitrarily small by choosing a suitable n , like we have just seen. \square

To describe the topology of \bar{X} we need the following fact [BH99, Proposition II.2.4]:

Lemma 1.2.6. Let X be a CAT(0) space and B a convex subset which is complete with respect to the induced distance function. Then there is a continuous (in fact, 1-Lipschitz) retraction $\rho: X \rightarrow B$ such that for every x in X , its image $\rho(x)$ is the unique point of B at minimal distance from x . Furthermore, if x and y in X are such that $\rho(x) \neq \rho(y)$, then the two geodesics joining the points to their images are disjoint.

Consider a complete CAT(0) space X and choose a point x_0 in it. Then the closed ball $\bar{B}(x_0, r)$ is a complete convex subset of X , so there is a retraction ρ_r like previously stated. If $r' > r$ we have that $\rho_r \circ \rho_{r'} = \rho_r$. In fact, for x in X with $d(x_0, x) \geq r$ the projection $\rho_r(x)$ is the intersection of the unique geodesic between x_0 and x with the boundary of $\bar{B}(x_0, r)$; it can be proven by contradiction using the triangle inequality. But then the behaviour of the composition is easily proven to be the one we described. We can then consider the collection $(\bar{B}(x_0, r))_{r \in [0, +\infty)}$ together with maps $\rho_{r', r}: \bar{B}(x_0, r') \rightarrow \bar{B}(x_0, r)$, which are simply the restrictions of ρ_r defined before. These maps form an inverse system of topological spaces. Call \tilde{X} the inverse limit.

Lemma 1.2.7. Let (X, d) be a complete CAT(0) metric space. \tilde{X} has a natural bijection with \bar{X} defined before, such that the restriction of the topology of \tilde{X} to X coincides with the one induced by the distance.

Proof. Elements of the inverse limit are functions $c: [0, +\infty) \rightarrow X$ such that for every $r' > r$ one has $\rho_r(c(r')) = c(r)$. Using the previous considerations, there are two cases:

- $c(t)$ coincides with the parametrization of the geodesic segment x_0x for some x when $t \leq d(x_0, x)$, and then it is constant;
- c is a geodesic ray.

In the first case, we identify c with x , in the second we identify c with its class in ∂X . This is a bijection with \overline{X} thanks to the previous considerations.

The topology restricted to X is the standard one because internal parts of the balls cover X . \square

Let us describe the neighbourhoods of the boundary points in \overline{X} . Choose a geodesic ray c issuing from x_0 and an open set U in X where it passes. Consider all the possible pairs (\tilde{c}, t) where \tilde{c} is a geodesic ray issuing from x_0 and t is a real number such that $\tilde{c}(t) \in U$. The points $\tilde{c}(t')$ for which exists a pair (\tilde{c}, t) of this type with $t < t'$ form an open neighbourhood of the class of c in \overline{X} ; furthermore these sets, which we will call $A(U, x_0)$, along with open sets of X , form a basis for the topology of \overline{X} . Note that these new open sets, when intersected with X , give open sets of X . With this description in mind, we can prove that the topology does not depend from x_0 .

Lemma 1.2.8. Let y_0 be another point in X and construct the topology on \overline{X} following the previous Lemma but starting from y_0 . Then it is the same topology as the one constructed starting from x_0 .

For a proof, see [BH99, Proposition II.8.8].

Let us now explore the behaviour of \overline{X} under extensions of isometries of X .

Lemma 1.2.9. Let X be a complete CAT(0) space and $f: X \rightarrow X$ an isometry. Then the induced function $\overline{f}: \overline{X} \rightarrow \overline{X}$ is a homeomorphism.

Proof. It is clear that, along with old open subsets of X , we can restrict our attention only to open sets of the type $A(B(c(t), r), x)$, with x in X , c geodesic issuing from x , r and t a positive real numbers and B the open ball. An isometry sends a set of this type in another set of this type. \square

1.3 Tools from Geometric Group Theory

Let G be a finitely generated group. It is possible to make it a metric space, or to construct a geodesic metric space starting from it using a set of generators. We will see that the constructions we make will not depend, in a precise sense, from the finite set of generators chosen.

We begin with defining an equivalence relation between metric spaces that is useful in this context.

Definition 1.3.1 (Quasi-isometry). Let (X, d) , (X', d') be metric spaces, $K \geq 1$, $C \geq 0$, $D \geq 0$ real numbers. A function $f: X \rightarrow Y$ is said to be a (K, C, D) -quasi-isometry, or simply quasi-isometry, when the constants are clear or inessential, if

- For every x, y in X , the following inequality holds:

$$\frac{1}{K}d(x, y) - C \leq d'(f(x), f(y)) \leq Kd(x, y) + C;$$

- Y is contained in the D -neighbourhood of $f(X)$ (quasi-surjectivity).

Two metric spaces are said to be quasi-isometric if there is a quasi-isometry between them.

It is easy to verify that quasi-isometry is an equivalence relation. In particular, any quasi isometry $f: X \rightarrow Y$ has a *quasi inverse*, i.e. a function $g: Y \rightarrow X$ which is a quasi isometry too, perhaps with different parameters, and such that $f \circ g(y)$ is uniformly close to y and $g \circ f(x)$ is uniformly close to x . The composition of quasi isometries is a quasi isometry too.

If we do not assume quasi surjectivity the function will be called a *quasi isometric embedding*. The typical case is a quasi isometric embedding of an interval of \mathbb{R} ; this will be called *quasi geodesic* and it typically arises when composing geodesics in the domain of a quasi isometry f with f .

Now we can make constructions using groups.

Definition 1.3.2 (Cayley graph). Let G be a finitely generated group and S some finite set of generators. The Cayley graph of G relative to S , or simply Cayley graph of G when the set of generators is clear or inessential, denoted by $\mathcal{C}(G, S)$ or $\mathcal{C}(G)$ is a graph having G as set of vertices and an arc between x and y in G if and only if $x^{-1}y$ is an element of S or an inverse of an element of S .

The Cayley graph is endowed with a distance function which is simply the path distance when length 1 is given to the arcs. The restriction of such distance to the set of vertices is the *word metric* on G (relatively to S). The word metric has values in the set of non negative integers and the distance between elements x and y of G is given by the length of the shortest string of elements of S or of their inverses representing $x^{-1}y$.

Clearly the previous definitions depend on S . However, this dependence vanishes if we consider the metric spaces up to quasi isometry.

Lemma 1.3.3. Let S and S' be finite set of generators for a group G . Then the identity of G is a quasi isometry when we consider the word metrics given by S and S' . Furthermore, the relative Cayley graphs are quasi isometric.

Proof. There is a positive integer K such that every element of S can be written like a string of at most K elements of S' or their inverses and vice-versa. Then the identity of G is clearly a $(K, 0, 0)$ -quasi-isometry between the two word metrics. The Cayley graph relative to a set of generators is in an obvious way quasi-isometric to the set of vertices; then we conclude by composition of quasi-isometries. \square

We will now define a construction defined on metric spaces which becomes useful when talking about quasi-isometry invariants.

We begin with a construction from infinite combinatorics.

Definition 1.3.4 (Ultrafilter). *An ultrafilter on an index set I is a family $\mathcal{U} \subset \wp(I)$ such that:*

- $I \in \mathcal{U}, \emptyset \notin \mathcal{U}$;
- \mathcal{U} is closed under intersections and supersets;
- If $A \cup B \in \mathcal{U}$, then either $A \in \mathcal{U}$ or $B \in \mathcal{U}$.

A principal ultrafilter is one consisting of subsets containing a fixed element. We will avoid them; in fact we will only consider non principal ultrafilters on \mathbb{N} . Non principal ultrafilters on infinite sets do exist, if we assume the Axiom of Choice.

Given an ultrafilter \mathcal{U} on I and an I -sequence $(x_i)_{i \in I}$ of elements of a set X , we will say that a property P of elements of X holds \mathcal{U} -almost everywhere on $(x_i)_{i \in I}$ if the set of indices i such that $P(x_i)$ holds is in \mathcal{U} . If X is a topological space, we say that the \mathcal{U} -limit of a sequence $(x_i)_{i \in I}$ is l , and write $\mathcal{U} - \lim x_i = l$, if for every open set A containing l , \mathcal{U} -almost every x_i belongs to A . In compact spaces \mathcal{U} -limits always exist; they are unique in Hausdorff spaces.

Let us define the object of our interest.

Definition 1.3.5 (Asymptotic cone). *Let (X, d) be a metric space, \mathcal{U} a non principal ultrafilter on \mathbb{N} , $(d_n)_{n \in \mathbb{N}}$ a sequence of non negative real numbers which is not \mathcal{U} -a.e. bounded by any real number, $(x_n)_{n \in \mathbb{N}}$ a sequence of points in X . Consider the set C of all sequences $(y_n)_{n \in \mathbb{N}}$ of elements of X such that the sequence*

$$\left(\frac{d(x_n, y_n)}{d_n} \right)_{n \in \mathbb{N}}$$

is \mathcal{U} -a.e. bounded by some real number. Define an equivalence relationship \sim in C by declaring two such sequences equivalent if

$$\mathcal{U} - \lim \frac{d(x_n, y_n)}{d_n} = 0.$$

The set $\mathcal{C}(\mathcal{U}, (x_n), (d_n), X) = C/\sim$ is said asymptotic cone of X relative to \mathcal{U} and (d_n) , or simply $\mathcal{C}X$, asymptotic cone of X , if the other parameters are clear or inessential. This set has a metric space structure with distance function given by

$$d([y_n], [y'_n]) = \mathcal{U} - \lim \frac{d(y_n, y'_n)}{d_n},$$

which always exists and is unique because of the previous considerations.

We introduce the following model theoretic construction to give a somewhat alternate view on asymptotic cones. We follow here [CK90].

Definition 1.3.6 (Ultrapower). *Let X be any structure, i.e. set with operations and relations, and let \mathcal{U} be a non principal ultrafilter on a set of indices I . The \mathcal{U} -ultrapower of X , or simply ultrapower of X when the ultrafilter is clear or inessential, denoted by *X , is the set of all I -sequences of elements of X quotiented by the relation $(x_i)_{i \in I} \sim (y_i)_{i \in I}$ if $x_i = y_i$ for \mathcal{U} -almost every $i \in I$.*

We denote by $[x_i]$ the equivalence class of $(x_i)_{i \in I}$. Having fixed the ultrafilter, a function $f: X \rightarrow Y$ induces a function ${}^*f: {}^*X \rightarrow {}^*Y$ defined on $[x_i]$ to be $[f(x_i)]$, and a relation $\mathcal{R} \subseteq X \times X$ gives a relation ${}^*\mathcal{R} \subseteq {}^*X \times {}^*X$ such that ${}^*\mathcal{R}([x_i], [y_i])$ if and only if $\mathcal{R}(x_i, y_i)$ for \mathcal{U} -a.e. i . There is a canonical immersion of X in *X given by the classes of constant sequences.

When talking about a metric space (X, d) the structure consists of:

- The set X ;
- \mathbb{R} with the whole of its structure;
- A function $d: X \times X \rightarrow \mathbb{R}$ respecting the hypotheses of distance.

Associated to a given structure X we define the superstructure $V(X)$ the following way:

Definition 1.3.7 (Superstructure). *Let X be a set. Suppose for simplicity no element of X contains other elements of X . The superstructure $V(X)$ on X is the union $\bigcup_{n \in \mathbb{N}} V_n(X)$, where the V_n 's are recursively defined as follows:*

- The set $V_0(X)$ is X ;
- The set $V_n(X)$ is the powerset of $V_{n-1}(X)$.

The definition given in [CK90] is slightly different, in fact there V_n is the union on $i \leq n$ of the V_i 's we have defined here.

Note that on a superstructure is naturally defined a membership relation \in . Superstructures allow us to take in exam particular subsets of an ultrapower of X , particular subsets of subsets, and so on. The precise construction goes as follows. Suppose I is an index set and \mathcal{U} is a non principal ultrafilter on I . We first take $Y = {}^*X$, and, for every x in X , call *x the class in *X of the sequence constantly equal to x ; this is the immersion of X in *X we already told about. Define, like we already did, $Y = V_0(Y) = {}^*X$ to be the set of all classes of sequences of elements of X .

If A is a subset of X , we call *A the subset of *X made up of all classes of sequences $[x_i]$ such that $x_i \in A$ for \mathcal{U} -a.e. i in I ; we can say, without lost of generality, that *A is the set of classes of sequences of elements of A . If $(A_i)_{i \in I}$ is a sequence of subsets of X , we denote by $[A_i]$ and call the class of the sequence (A_i) the subset of *X made up of classes of sequences $[x_i]$ such that $x_i \in A_i$ for \mathcal{U} -a.e. i in I . We then define $V_1(Y)$ like the set of all classes of sequences in $V_1(X)$. The sets in $V_1(Y)$ are said to be *internal (sub)sets* of *X . It is important to stress that internal sets are not all the subsets of *X .

Recursively, if A is a subset of $V_n(X)$, we call *A the subset of V_n made up of classes of sequences of elements of A . If $(A_i)_{i \in I}$ is a sequence of elements of $V_n(X)$, then its class $[A_i]$ is a subset of $V_n(Y)$ made up of sequences $(B_i)_{i \in I}$ of elements of $V_n(X)$ such that $B_i \in A_i$ for \mathcal{U} -a.e. i in I . Finally, $V_{n+1}(Y)$ is the set of all classes of sequences in $V_n(X)$.

Definition 1.3.8 (Non standard universe). *Let X be a set like in the previous Definition and $V(X)$ be its superstructure. A non standard universe is a triple $(V(X), V(Y), *)$ where $V(Y)$ is the union on n in \mathbb{N} of the $V_n(Y)$'s just constructed and $*$: $V(X) \rightarrow V(Y)$ is the function, called transfer map, just defined.*

Note that we can say, more concisely, that $V_n(Y) = {}^*V_n(X)$.

The power of the concept of non standard universe comes from their property to transport certain predicates on $V(X)$ to $V(Y)$ and vice versa.

Theorem 1.3.9 (Transfer principle). *Let $(V(X), V(Y), *)$ be a non standard universe. A first order formula on $V(X)$ with constants c_1, \dots, c_k is satisfied in $V(X)$ if and only if the same formula with constants respectively substituted by ${}^*c_1, \dots, {}^*c_k$ is satisfied in $V(Y)$.*

For a proof, see [CK90, Theorem 4.4.5].

In fact, the previous theorem is a generalization of the classic Łoś's theorem in model theory, which establishes the same result for X and *X only.

When we take the ultrapower of a metric space, which we may call mimicking the standard nomenclature *hypermetric space*, it comes endowed with a *hyperdistance* function $*d: X \times X \rightarrow *\mathbb{R}$. An ultrapower of reals is said a *hyperreal field*.

A hyperreal field is an ordered field, superfield of the reals, with operations and order constructed like before, and has infinite elements, greater in absolute value than any natural, and infinitesimal ones, less in absolute value than any $\frac{1}{n}$ for n natural but distinct from 0. Every finite, i.e. non infinite hyperreal t has a unique standard part $\text{std}(t)$, which is a real number such that $t - \text{std}(t)$ is infinitesimal.

By the Transfer Principle, the hyperdistance fulfils the hypotheses of distance, except it has hyperreal values. Let us explore further properties of ultrapowers of metric spaces:

Lemma 1.3.10. Let (X, d) be a geodesic metric space. Then for any index set I and any non principal ultrafilter \mathcal{U} in I the ultrapower $*X$ is also geodesic.

Proof. Let $p = [p_i]$ and $q = [q_i]$ be points in $*X$. The hyperdistance between them is $D = [d_i]$, where $d_i = d(p_i, q_i)$. Let γ_i be a geodesic between p_i and q_i . We have that $[[0, d_i]]_{i \in I} = [0, D] \subseteq *\mathbb{R}$. Therefore, if $t = [t_i]$ is a hyperreal less or equal than D , the point $\gamma_i(t_i)$ is defined on \mathcal{U} -a.e. i in I . Then the function γ sending t to $[\gamma_i(t_i)]$ is a geodesic in $*X$. \square

The previous proof is actually a particularization of the proof of the Transfer Principle; we outlined it for more clarity. The following proposition is subtler.

Lemma 1.3.11. Let (X, d) be a uniquely geodesic metric space. Then for any index set I and any non principal ultrafilter \mathcal{U} in I the ultrapower $*X$ is uniquely geodesic.

The problem may come from the fact that a function defined on (a subset of) $*\mathbb{R}$ with values in $*X$ is not necessary obtained from a sequence of functions on \mathbb{R} with values in X like in the proof of the previous Lemma. The Transfer Principle tells us only that geodesics of this type are unique between two fixed points. So we need a different proof.

Proof. Let $p = [p_i]$ and $q = [q_i]$ be points in $*X$, γ_i the geodesic between p_i and q_i and γ the geodesic between p and q found in the previous Lemma. Suppose there is another geodesic β between them. It passes from a point

$r = [r_i]$ not in the image of γ , and therefore r_i is not in the image of γ_i for \mathcal{U} -a.e. i in I . Because of the uniqueness of geodesics in X we have that

$$d(p_i, r_i) + d(r_i, q_i) > d(p_i, q_i)$$

for \mathcal{U} -a.e. i in I and so

$${}^*d(p, r) + {}^*d(r, q) > {}^*d(p, q),$$

but then β could not be a geodesic, absurd. \square

Spaces which respect the CAT (0) condition are uniquely geodesic, and so we already know how geodesics in *X are made. The following fact becomes thus easy to prove.

Lemma 1.3.12. If X is a CAT (0) metric space, then any ultrapower *X is a CAT (0) metric space, with the comparison triangles are taken in ${}^*(\mathbb{R}^2) = ({}^*\mathbb{R})^2$.

From the transfer principle follows also the following version of the Lemma 1.2.6 for ultrapowers of CAT (0) spaces.

Lemma 1.3.13. Let X be a complete CAT (0) metric space, and let *X an ultrapower of X . Let C be the collection of closed subsets of X and A a subset of *X in *C . Then there is a retraction ρ of X on A with the same properties of the Lemma 1.2.6.

We return now to asymptotic cones. Note that the hyperreal represented by (d_n) in the definition of asymptotic cone is infinite. With this in mind, a rephrasing of definition of the asymptotic cone $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$ is of a subset of *X consisting of all points $[y_n]$ such that

$$\frac{{}^*d([x_n], [y_n])}{[d_n]}$$

is finite, quotiented by the relation that identifies pairs of points having that ratio infinitesimal, with distance function given by

$$d([y_n], [z_n]) = \text{std} \left(\frac{{}^*d([y_n], [z_n])}{[d_n]} \right).$$

We will call $[x_n]$ the *center* of the cone and $[d_n]$ the *radius*, despite it is not a true radius.

The main interest in asymptotic cones lies in the fact they annihilate the finite errors comparing in the definition of quasi-isometry. The exact, straightforward, statement is the following

Lemma 1.3.14. Let X and Y be metric spaces and $f: X \rightarrow Y$ a (K, C, D) -quasi-isometry. Let \mathcal{U} be an ultrafilter on \mathbb{N} , $(x_n)_{n \in \mathbb{N}}$ a sequence in X and $(d_n)_{n \in \mathbb{N}}$ a real sequence. Then $*f$ brings points of the subset of $*X$ that defines the cone $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$ in the subset of $*Y$ that defines the cone $\mathcal{C}(Y, \mathcal{U}, (f(x_n)), (d_n))$, when two such points are identified in the former cone their images are identified in the latter and the function $\mathcal{C}(f, \mathcal{U}, (x_n), (d_n))$, or simply $\mathcal{C}(f)$, induced at the quotient is a K -bi-Lipschitz homeomorphism between $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$ and $\mathcal{C}(Y, \mathcal{U}, (f(x_n)), (d_n))$.

We will now define particular types of geodesic spaces which often arise when taking asymptotic cones of groups, i.e. of their Cayley graphs.

Definition 1.3.15 (Real tree). *A topological space X is said to be a real tree if for any pair of points x, y in X there is a unique, up to reparametrization, topological embedding of $[0, 1]$ into X with endpoints x and y .*

If a real tree has a distance function, it naturally has a structure of geodesic metric space. Note that a real tree is $\text{CAT}(\kappa)$ for every real κ .

Definition 1.3.16 (Tree graded space). *A topological space X is said to be tree-graded respect to a family \mathcal{P} of subspaces called pieces if:*

- *Every element of \mathcal{P} is path connected;*
- *Two different pieces have at most one point in common;*
- *The image of every topological embedding of \mathbb{S}^1 is contained in one piece.*

The definition usually given in literature [HK05, Definition 2.1.3] requires X to be a complete geodesic space and pieces to be convex. We see that adding the first of these hypotheses ensures the second.

Lemma 1.3.17. Let X be a tree graded space in the topological sense. Suppose furthermore the topology is given by a distance function d that makes X a complete geodesic metric space. Then the pieces are convex.

Proof. Take two distinct points x and y in the same piece P . They are joined by a path $\alpha: [0, 1] \rightarrow X$ entirely contained in P and also by a geodesic $\gamma: [0, d(x, y)] \rightarrow X$. We can suppose without loss of generality α is injective. This claim will be the subject of a technical Lemma of which we defer the proof.

Suppose that the image of γ is not entirely contained in the image of α . The set of t in $[0, d(x, y)]$ such that $\gamma(t)$ is not contained in the image of α is

open. Take one such t and consider the endpoints $a < b$ of the maximal open interval containing t and mapped by γ outside the image of α . Then $\gamma(a)$ and $\gamma(b)$ lie in the image of α . It follows that $\gamma([a, b])$ together with the subpath of α between $\gamma(a)$ and $\gamma(b)$ form a simple closed loop, which stays thus in the same piece, the one containing the image of α and hence x and y . By arbitrariness of t all of the image of γ is contained in this piece. \square

We will implicitly use an argument similar to the one used in the proof of the previous Lemma whenever we want to produce a simple closed curve from two injective curves with the same endpoints but having different images. To actually have injective curves, we may make use of the following

Lemma 1.3.18. Let X be a T1 topological space and $\alpha: [0, 1] \rightarrow X$ be any continuous path in it. Then there is an injective continuous path in X with endpoints $p = \alpha(0)$ and $q = \alpha(1)$.

Proof. Consider the set Z of all continuous paths in X between p and q defined on $[0, 1]$. Consider the following order relation on Z : if γ_1 and γ_2 are in Z , then $\gamma_1 \leq \gamma_2$ if and only there is an open set $U \subseteq (0, 1)$ such that:

- The paths γ_1 and γ_2 coincide outside U ;
- The path γ_2 is locally constant in U .

To prove the antisymmetry, if $\gamma_1 \leq \gamma_2$ using a set U and $\gamma_2 \leq \gamma_1$ using a set V it is easy to see that γ_1 and γ_2 coincide outside $U \cap V$, and from locally constant property they coincide on $U \cap V$ too. If $\gamma_1 \leq \gamma_2$ with respect to a set U_1 and $\gamma_2 \leq \gamma_3$ with respect to a set U_2 it can be seen that $\gamma_1 \leq \gamma_3$ with the set $U_1 \cup U_2$.

Consider a chain $C = \{\gamma_i\}_{i \in I}$ in Z , where I is a set of indices. We claim that for any t there is an index i such that for any $\gamma_j \geq \gamma_i$ the value of $\gamma_j(t)$ coincides with $\gamma_i(t)$. This is immediate for any t which is not involved in any open set defining the order between a pair of elements of C . Otherwise, if $\gamma_{i_1} \leq \gamma_{i_2}$ with a set U containing t , then $i = i_2$ suffices. Then we can define a path β that takes on t its definitive value. Making the same distinction on t we used for the proof of existence, the continuity of β follows as well.

The Zorn Lemma gives us a maximal element $\tilde{\alpha}$. We claim that if there are $0 \leq t < s \leq 1$ such that $\tilde{\alpha}(t) = \tilde{\alpha}(s)$, then $\tilde{\alpha}$ takes the same value on any point between t and s . Otherwise, there would be a greater element in Z constructed by taking $U = (t, s)$ and making it in this interval be constantly $\tilde{\alpha}(t) = \tilde{\alpha}(s)$. But then it can be seen that defining the equivalence relation \sim on $[0, 1]$ that identifies points on which $\tilde{\alpha}$ takes the same values, the quotient $[0, 1] / \sim$ is homeomorphic to an interval (perhaps a point), thanks

to the hypothesis on the topology. The induced map $\tilde{\alpha}/\sim: [0, 1]/\sim \rightarrow X$ is injective. \square

Note that a real tree is tree graded respect to the collections of its singletons so tree graded spaces provide, in a sense we will explore further, a generalisation of real trees.

We will now give some definitions about actions of a group on a space by isometries, including a condition which will be essential in the proof of the main theorem.

From now on metric spaces in exam will be *proper*, i.e. closed balls are compact subsets. Note that proper spaces are complete.

Definition 1.3.19 (Geometric action). *Let X be a proper geodesic metric space and G a group acting on it by isometries. We say that the action of G is geometric if it is properly discontinuous and cocompact, i.e. the quotient space of X by the action is compact.*

The prototypical example of a geometric action is that of the fundamental group of a compact geodesic metric space on its universal cover. Geometric actions have a clear interpretation in terms of quasi-isometries.

Lemma 1.3.20 (Milnor-Švarc Lemma). *Let (X, d) be a proper geodesic metric space and G a group geometrically acting on it. Then G is finitely generated by a set S and for any x_0 in X the function f sending an element g of G to $g \cdot x_0$ is a quasi-isometry between G with the word metric d_S given by S and (X, d) .*

Proof. Quasi-surjectivity is an immediate consequence of the cocompactness. In particular, if R is the (finite) diameter of the quotient, taken any x_0 in X we can be sure that the projection $X \rightarrow X/G$ is surjective on $\overline{B}(x_0, R)$ or, said otherwise, the translates of this ball through the G action cover X .

The set of elements g in G such that $g \cdot x_0 \in \overline{B}(x_0, 2R)$ is finite for the properly discontinuous property of the action. Call S_0 the set of such elements. If $S_0 = G$ we are done. Otherwise, let

$$r = \inf_{g \in G \setminus S_0} d(x_0, g \cdot x_0) - 2R > 0$$

and $S = S_0 \setminus \{1\}$. We claim that S generates G and that for any h_1, h_2 in G the inequality $d_S(h_1, h_2) \leq \frac{1}{r} d(h_1 \cdot x_0, h_2 \cdot x_0) + 1$ holds.

Take g in G and a geodesic $\gamma: [0, D = d(x_0, g \cdot x_0)] \rightarrow X$ between x_0 and $g \cdot x_0$. Let

$$N = \left\lfloor \frac{D}{r} \right\rfloor + 1.$$

For every i in $\{1, \dots, N - 1\}$, let g_i be such that

$$d\left(g_i \cdot x_0, \gamma\left(\frac{iD}{N}\right)\right) \leq R;$$

by construction we can always find a g_i with this properties; also set $g_0 = 1$ and $g_N = g$. We have that $d(x_0, g_{i-1}^{-1}g_i \cdot x_0) = d(g_{i-1} \cdot x_0, g_i \cdot x_0) < 2R + r$ and thus $g_{i-1}^{-1}g_i \in S$. But then g can be represented with a string of at most N elements of S and $N \leq \frac{D}{r} + 1$. The thesis follows by taking $g = h_2^{-1}h_1$.

Vice versa, consider a point $x = h \cdot x_0$ in the orbit of x_0 and an element $s \in S$. By hypothesis, $2R \geq d(x_0, s \cdot x_0) = d(h \cdot x_0, hs \cdot x_0)$. Then from the triangle inequality for any h_1, h_2 in G we have that

$$d(h_1 \cdot x_0, h_2 \cdot x_0) \leq 2Rd_S(h_1, h_2)$$

holds.

The two inequalities together ensure the quasi isometry property. \square

If the group G already has a finite generating set S' , by the transitivity of the quasi isometry relation and Lemma 1.3.3 we have that $\mathcal{C}(G, S')$ and X are quasi isometric too via one of the maps from the previous lemma.

We pass now to some definitions using asymptotic cones.

Definition 1.3.21 (Hyperbolic space). *A complete metric space X is said to be Gromov hyperbolic, or simply hyperbolic in this dissertation, if any asymptotic cone of it is a real tree.*

In the literature hyperbolic spaces are defined in a different way [Gro87]; it can be proven that the classical definitions are equivalent to the one we use here [Dru02, Proposition 3.A.1]. The definition we use has the advantage to behave in a clear way under quasi-isometries.

Lemma 1.3.22. If two complete metric spaces X and Y are quasi-isometric and X is hyperbolic then Y is hyperbolic too.

Proof. A quasi isometry between spaces becomes a bi-Lipschitz homeomorphism between corresponding asymptotic cones, like stated in 1.3.14. The notion of real tree is clearly invariant under homeomorphism. \square

The proof can be enriched if we consider geodesic spaces. If X is geodesic, any asymptotic cone of it is geodesic too, so the unique topological path between two points in the cone is a geodesic with the right parametrization. Then the corresponding path in the corresponding cone of Y is a geodesic too, up to reparametrization, thanks to the bi-Lipschitz property.

Note that the standard hyperbolic space is hyperbolic in the sense we have just defined. To prove this, we need another definition of hyperbolic space, the one usually found in literature. We say that a complete geodesic metric space is δ -hyperbolic if in any geodesic triangle every side is included in the δ -neighbourhood of the union of the other two.

Lemma 1.3.23. Let X be a δ -hyperbolic geodesic metric space. Then X is hyperbolic.

Proof. Fix an ultrafilter \mathcal{U} on \mathbb{N} , sequences $(x_n)_{n \in \mathbb{N}}$ in X and $(d_n)_{n \in \mathbb{N}}$ in \mathbb{R} and consider $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$. The hypothesis of the Lemma are true also for the \mathcal{U} -ultrapower *X because of the transfer principle. But then the asymptotic cone is also geodesic, and every triangle in it fulfils the same hypothesis on δ . By dividing all distances by the infinite $[d_n]$ we have that every geodesic triangle in $\mathcal{C}(X)$ is in fact a tripod: every side is contained in the union of the other two.

It remains then to prove that this characterises real trees. The argument is a bit long and rather technical, so we leave it for a separate lemma. \square

The converse of the previous statement is also true, but we will not use this fact. Let us establish the technical Lemma announced before.

Lemma 1.3.24. Let (X, d) be a geodesic metric space such that every geodesic triangle is a tripod. Then X is a real tree, in the topological sense.

Proof. Note that a space like in the hypotheses is uniquely geodesic.

If there are two points in the space joined by more than one topological path we can find, by taking appropriate subpaths of these two, a simple closed non constant loop in X . With this hypothesis, we will prove that there is a triangle which is not a tripod.

It will also be sufficient to find a quadrilateral $x_1x_2x_3x_4$ with geodesic sides that intersect only if they are consecutive and only at extremes: this way x_1 and x_3 cannot form a tripod with both x_2 and x_4 .

Starting with the simple closed non constant loop we have, we take a geodesic between two points of it and, up to throwing away some common pieces, we can suppose there are two points, the geodesic γ joining them and another simple curve α having them as extremes, and this two curves intersect only at extremes. We suppose for simplicity, without lost of generality, that both are defined on $[0, 1]$

For every point p of α there is a nearest point q to it on γ by compactness. The geodesic joining p and q intersects γ only in q , otherwise there would be a nearer point to p on γ . Then q is unique: if there were another point q' on γ at the same distance from p , there would be a triangle non tripod having

for sides the piece of γ between q and q' , and pieces of geodesics joining q and q' to p from these points to their first intersection.

So there is a function $f: [0, 1] \rightarrow \gamma[0, 1]$ associating to t the nearest point on γ to $\alpha(t)$. This function cannot be constant, and thus neither locally constant, on $(0, 1)$ because near 0 it takes values near $\gamma(0)$ and near 1 values near $\gamma(1)$. So there is a t_0 in $(0, 1)$ such that in any of its neighbourhoods there is a point where f takes a value different from $f(t_0)$. Let $L = d(\alpha(t_0), f(t_0))$. Take t_1 in $(0, 1)$ such that $d(\alpha(t_0), \alpha(t_1)) < \frac{L}{3}$ and $d(\alpha(t_1), f(t_1)) > \frac{2}{3}L$, which we can do by continuity of the distance function, with the additional property that $f(t_0) \neq f(t_1)$, which can be assured by construction. If the geodesics $\alpha(t_0)f(t_0)$ and $\alpha(t_1)f(t_1)$ intersect, we find a triangle non tripod, otherwise a piece of them, the piece of γ between $f(t_0)$ and $f(t_1)$ and a piece of the geodesic $\alpha(t_0)\alpha(t_1)$ form a quadrilateral of the type pointed out in the beginning of the proof. \square

The proof that \mathbb{H}^n is hyperbolic relies then on the fact that it is δ -hyperbolic.

Lemma 1.3.25. \mathbb{H}^n is δ -hyperbolic for every $\delta > \log(\cot(\frac{\pi}{8}))$, for example $\delta = 1$.

Proof. Every triangle in \mathbb{H}^n lies in a 2-subspace, so it suffices to prove the thesis for \mathbb{H}^2 . Take a geodesic triangle there and an internal point p and trace the three geodesics from p passing by the vertices. The thesis for the original triangle descends from the thesis for the ideal triangle with extremes the points on $\partial\mathbb{H}^2$ corresponding to these three geodesics: the original triangle lies inside the new one and a point on one its side is closer to the union of the other two than to the union of the two new sides asymptotic to the geodesic from p and passing by the opposite vertex.

There is only one ideal triangle in \mathbb{H}^2 up to isometry: we consider, in the half plane model, the one of vertices -1 , 1 and ∞ , and prove that the side between -1 and 1 lies in the δ -neighbourhood of the union of vertical sides.

For a positive real r , the r -neighbourhood of a vertical line is boarded by oblique (euclidean) lines making an angle α with the vertical and departing from the same ideal finite point. The angle α is such that

$$\int_{\frac{\pi}{2}}^{\frac{\pi}{2}+\alpha} \frac{1}{\sin \theta} d\theta = r,$$

and for $\alpha = \frac{\pi}{4}$, which is the adequate value for containing the third side in this case, $r = \log(\cot(\frac{\pi}{8}))$. \square

This implies immediately the following

Corollary 1.3.26. If ε is a positive real number, then every complete metric space satisfying the CAT $(-\varepsilon)$ property is hyperbolic.

Having turned groups into metric spaces, the following definition is natural.

Definition 1.3.27 (Hyperbolic group). *A finitely generated group G is said to be hyperbolic if there is a finite set of generators S for which the Cayley graph $\mathcal{C}(G, S)$ is hyperbolic.*

The definition we gave is by Lemma 1.3.14 a quasi-isometry invariant, so being a hyperbolic group is, by Lemma 1.3.3, independent from the generating set. The considerations we made previously allow us to conclude the following

Lemma 1.3.28. If ε is a positive real number, every group acting geometrically on a complete geodesic CAT $(-\varepsilon)$ space is hyperbolic.

In particular, fundamental groups of compact locally CAT $(-\varepsilon)$ spaces are hyperbolic.

We are going now to introduce a subtler definition. We begin by returning to ultrapowers. Given a collection \mathcal{A} of subspaces of a metric space (X, d) , we already know how to define ${}^* \mathcal{A}$. Take an asymptotic cone $\mathcal{C}(X)$ of X having x as centre and R as radius. If Z is a subspace of ${}^* X$, we call *intersection of Z with $\mathcal{C}(X)$* the set of points z of Z such that $\frac{{}^* d(x, z)}{R}$ is finite. When we later quotient the intersection of Z with the cone by the relation that identifies points with hyperdistance infinitesimal when divided by R , we obtain a subset of $\mathcal{C}(X)$ which we call the *projection of Z to the cone*. We can then define $\mathcal{C}(\mathcal{A})$ as the collection of nonempty projections of sets in ${}^* \mathcal{A}$ to the cone. Note that different sets may lead to the same projection.

Alternatively, if we have an asymptotic cone $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$, we can consider sequences $(A_n)_{n \in \mathbb{N}}$ of elements of \mathcal{A} such that

$$\frac{d(x_n, A_n)}{d_n}$$

is bounded and sequences (y_n) representing elements of the cone and such that $y_n \in A_n$ for \mathcal{U} -a.e. n in \mathbb{N} : these sequences determine a subset $[A_n]$ of the cone, and we call $\mathcal{C}(\mathcal{A}, \mathcal{U}, (x_n), (d_n))$, or simply $\mathcal{C}(\mathcal{A})$ the collection of all possible $[A_n]$.

Let us apply these considerations to the following

Definition 1.3.29 (Relatively hyperbolic space). *Let X be a complete metric space and \mathcal{A} a family of subspaces such that any two distinct members have*

infinite Hausdorff distance. The space X is said to be relatively hyperbolic with respect to \mathcal{A} if every asymptotic cone $\mathcal{C}(X)$ is tree-graded with respect to $\mathcal{C}(A)$.

Take now a finitely generated group G and consider a finite collection $\mathcal{H} = \{H_1, \dots, H_n\}$ of finitely generated subgroups. Let $\mathcal{A}(\mathcal{H})$ be the collection of all left cosets of subgroups in \mathcal{H} .

Definition 1.3.30 (Relatively hyperbolic group). *A group G is said to be relatively hyperbolic with respect to \mathcal{H} if it is a relatively hyperbolic space with respect to $\mathcal{A}(\mathcal{H})$.*

A hyperbolic group is considered to be hyperbolic relatively to its trivial subgroup. In this case $\mathcal{C}(\mathcal{A})$ is the set of all singletons of $\mathcal{C}(X)$. This convention is made even if the cosets have not infinite Hausdorff distance from each other.

Group actions on CAT(0) metric spaces have additional features that make their structure richer and more interconnected with that of the group itself. We begin with a geometric definition.

Definition 1.3.31 (Flat). *Let $k \geq 2$ be an integer. A k -flat in a CAT(0) space X is the image of an isometric embedding of \mathbb{E}^k into X .*

One classical result that links the algebra of the fundamental group of a compact locally CAT(0) space to its geometry is the following:

Theorem 1.3.32 (Flat Torus Theorem). *Let k be a positive integer. Suppose that the fundamental group of a compact CAT(0) space X with universal cover \tilde{X} has a subgroup G isomorphic to \mathbb{Z}^k . Then there is a k -flat in \tilde{X} invariant under the action of G and with quotient by such action a flat torus locally isometrically immersed in X .*

If X is a smooth manifold of non-positive curvature, the embeddings and immersions of the theorem are totally geodesic in the Riemannian sense.

For a proof in the CAT(0) setting, see [BH99, Theorem II.7.1]. The proof in the smooth case can be found in the original papers [GW71], [LY72].

Another classical result on flat spaces is concerned with their group of isometries.

Theorem 1.3.33 (Bieberbach Theorem). *Let G be a group acting geometrically on \mathbb{E}^n . Then G has a finite index subgroup isomorphic to \mathbb{Z}^n that acts geometrically by translations. Furthermore, the translations associated to a basis of this subgroup form a basis for the underlying vector space \mathbb{R}^n .*

The thesis of the theorem is expressed saying that G is virtually free abelian of rank n .

1.4 Isolated flats condition

We have already seen the strict connection between hyperbolic groups, hyperbolic spaces and complete $\text{CAT}(\kappa)$ spaces, with κ strictly negative. In the more general $\text{CAT}(0)$ setting the large scale structure of the space, i.e. properties invariant under quasi-isometry, can be far more complicated. The situation is easier to deal with in the smooth setting, of course. Another condition that simplifies and partially rigidifies the geometry is the following

Definition 1.4.1 (Isolated flats). *Let X be a proper $\text{CAT}(0)$ space, G a group acting geometrically on it, and let \mathcal{F} be a collection of flats in X . Suppose that:*

- \mathcal{F} is G -invariant;
- For any pair F_1, F_2 of flats in \mathcal{F} and any positive R , the flat F_2 is not contained in the R -neighbourhood of F_1 (we say that there are no two parallel flats in \mathcal{F}).
- X is relatively hyperbolic with respect to \mathcal{F} .

Then X has the isolated flats property with respect to \mathcal{F} .

We adopt the convention that a $\text{CAT}(0)$ space which is also hyperbolic has isolated flats. This is because being hyperbolic is the same thing as being relatively hyperbolic with respect to singletons, which we may think to as “0-dimensional flats”.

Before proceeding, we make an observation involving ultrapowers and asymptotic cones. The definition of isolated flats makes use of these concepts through the relative hyperbolicity. Elements of ${}^*\mathcal{F}$ are hyperflats in the ultrapower, meaning that they are isometric embeddings of ultrapowers of an \mathbb{E}^n constructed with the same index set and the same ultrafilter. At the level of asymptotic cones, they became regular flats. Hyperflats in ${}^*\mathcal{F}$ are obviously internal sets, and r -neighbourhoods of internal sets, with r hyperreal, are internal sets too.

Lemma 1.4.2. Let $F_1 \neq F_2$ be hyperflats in ${}^*\mathcal{F}$. Then they are distinct in any asymptotic cone, i.e. if their intersection with a cone $\mathcal{C}(X)$ is nonempty, then their projections to the cone determine two distinct flat pieces in $\mathcal{C}(X)$.

Proof. Suppose by contradiction there are hyperflats F_1 and F_2 that are equal in some cone $\mathcal{C}(X)$ of centre x and radius R . This means that the intersection of R -neighbourhoods of F_1 and F_2 has a diameter D which is infinite when divided by R . Take another asymptotic cone $\mathcal{C}'(X)$, having centre in the

intersection of these neighbourhoods and radius D . In this cone, F_1 and F_2 become two flats having diameter of the intersection 1, which is absurd because of relative hyperbolicity, in particular the tree graded property. \square

We will now state and prove some technical lemmas about spaces with isolated flats. The hypothesis we will assume from now on is: let X be a proper CAT(0) with a geometric action of a group G and having isolated flats with respect to a family \mathcal{F} . We begin with proving that the definition we gave implies another definition of isolated flats, the one usually found in the literature.

Lemma 1.4.3. In the above hypotheses on X we have:

1. Every flat of X is contained in an R -neighbourhood of a flat in \mathcal{F} ;
2. For every $r > 0$ there is $D(r) > 0$ such that for any pair of flats in \mathcal{F} the diameter of the intersection of their r -neighbourhoods is less than D .

The last condition says precisely the flats of \mathcal{F} are isolated in the Hausdorff topology of convergence on bounded subsets (see [HK05, Definition 2.1.1] for a definition).

Proof. Suppose X is relatively hyperbolic with respect to \mathcal{F} . Take a flat F in X . ${}^*F \subseteq {}^*X$ is a hyperflat. Take an asymptotic cone $\mathcal{C}(X, \mathcal{U}, (x_n), (d_n))$ such that it intersects *F : the projection on the cone is a flat which we will call $\mathcal{C}(F)$. Take a simple closed curve in $\mathcal{C}(F)$: by relative hyperbolicity, it must be contained in a piece of $\mathcal{C}(\mathcal{F})$. By taking arbitrarily large triangles, the whole of *F is contained in a piece of $\mathcal{C}(\mathcal{F})$. That means F is contained in a neighbourhood of a flat of \mathcal{F} , otherwise the proposition “for every r and for every flat of \mathcal{F} there is a point of F at distance greater than r from that flat” would become, in the ultrapower, “for every flat in ${}^*\mathcal{F}$ there is a point in *F with distance more than $[d_n]$ from it”, i.e. no $\mathcal{C}(F)$ could not be contained in any flat piece of the corresponding $\mathcal{C}(\mathcal{F})$.

It remains to prove the isolation condition. The converse, i.e. the proposition “there is r such that exist pairs of flats in \mathcal{F} with intersection of their r -neighbourhoods having arbitrarily large diameter” would become in the ultrapower “there is a pair of distinct flats in ${}^*\mathcal{F}$ with the intersection of their r -neighbourhoods having diameter greater or equal than an infinite hyperreal \tilde{D} ”. Taking a point of this intersection as the center of the asymptotic cone and D as the radius, we obtain a contradiction for the relative hyperbolicity, because the two flats are distinct in the cone by the previous Lemma. \square

The characterization we just found allows us to conclude some results on the structure of the space.

Lemma 1.4.4. \mathcal{F} is locally finite, i.e. every compact set in X intersects only finitely many elements of \mathcal{F} .

Proof. It suffices to prove the property for closed balls. Choose $r > 0$ and x_0 in X . By hypothesis, if two flats F_1, F_2 in \mathcal{F} intersect both $\overline{B}(x_0, r)$ then there are no points in the intersection of r -neighbourhoods of F_1 and F_2 at distance $D(r)$ from x_0 , where D is defined in the Lemma 1.4.3. For every flat in \mathcal{F} intersecting $\overline{B}(x_0, r)$ take a point in the (non empty) intersection of $\partial\overline{B}(x_0, R)$ with it: the balls of radius r centered in these points define disjoint open sets in $\partial\overline{B}(x_0, R)$. By properness of X , these points are finite. \square

Lemma 1.4.5. The action of G on \mathcal{F} has finitely many orbits and stabilizers of elements of \mathcal{F} lie in finitely many conjugacy orbits.

Proof. Let K be a compact set whose translates by G cover X . It intersects only finitely many flats in \mathcal{F} , and the number of orbits of the action is less or equal that that finite number. The thesis on the stabilizers easily follows. \square

Lemma 1.4.6. The stabilizer of a flat F in \mathcal{F} acts geometrically on F .

Proof. Fix a compact set K whose G -translates cover X . Consider a subset $\{g_i\}_{i \in I}$ of G , with I a set of indices, such that for every i in I the set $g_i \cdot K$ intersects F . By hypothesis, $\{g_i \cdot K\}_{i \in I}$ cover F .

If for some i, j in I the flats $g_i^{-1} \cdot F$ and $g_j^{-1} \cdot F$ coincide, then $g_j g_i^{-1}$ is in $H = \text{Stab } F$, and thus $g_j \in H g_i$. By Lemma 1.4.4 the g_i 's lie in a finite number of right cosets of H . Furthermore, for any i, j the sets $\bigcup_{h \in H} h g_i \cdot K$ and $\bigcup_{h \in H} h g_j \cdot K$ lie at finite Hausdorff distance, less or equal than the Hausdorff distance between $g_i \cdot K$ and $g_j \cdot K$. Take then representatives g_{i_1}, \dots, g_{i_m} for the right cosets of H intersecting $\{g_i\}_{i \in I}$ and a $K' \supseteq K$ compact and large enough so that $g_{i_1} \cdot K'$ contains all of $g_{i_2} \cdot K, \dots, g_{i_m} \cdot K$: by construction, the H -translates of K' cover F . \square

Like already seen in the case of hyperbolic spaces, the geometric structure of spaces with isolated flats reflects in the geometric structure of groups acting geometrically on them. We begin by observing that a maximal free abelian subgroup H in G stabilizes a flat F_H of the same dimension as the rank of H by the Flat Torus Theorem 1.3.32. This flat is parallel to a flat F in \mathcal{F} for the previous considerations, and F is stabilized by H too thanks to [BH99, Theorem II.6.8]; in fact, parallel flats determine a product of a flat by an interval stabilized by the same subgroup. Then F must have the

same dimension as F_H , otherwise the stabilizer of F would contain an abelian subgroup containing H of rank higher than H . So a maximal free abelian subgroup of G determines the flat of \mathcal{F} it stabilizes.

Vice versa, a flat of \mathcal{F} is stabilized, according to Bieberbach Theorem 1.3.33, by a virtually free abelian subgroup H of the same rank as the dimension of the flat. Similarly to the previous reasoning, the subgroup found is maximal because no pair of flats in \mathcal{F} are parallel. So a flat in \mathcal{F} determines the virtually free abelian subgroup it is stabilized by. We have then found a one-to-one correspondence between maximal virtually free abelian subgroups of G of rank at least 2 and flats of \mathcal{F} .

Lemma 1.4.7. Let X be a CAT(0) space and G a group acting geometrically on it. Then G is relatively hyperbolic with respect to a collection of virtually free abelian subgroups of rank at least two if and only if X has isolated flats with respect to a G -invariant family \mathcal{F} of flats.

Proof. Suppose X has isolated flats with respect to the family \mathcal{F} . The stabilizers of the elements of \mathcal{F} lie in finitely many conjugacy classes of subgroups by Lemma 1.4.5, and are virtually free abelian. Choose a set of representatives H_1, \dots, H_m . If H_i stabilizes F and g is an element of G , then $gH_i g^{-1}$ stabilizes $g \cdot F$. So a quasi isometry induced by the action associates left cosets of the H_i 's to flats in \mathcal{F} . Finite index subgroups are quasi isometric to the groups they are contained in and X is relatively hyperbolic with respect to \mathcal{F} , so, by homeomorphism of corresponding asymptotic cones, G is relatively hyperbolic with respect to $\{H_1, \dots, H_m\}$.

Suppose vice versa that G is relatively hyperbolic with respect to a finite collection of virtually free abelian subgroups $\{H_1, \dots, H_m\}$. A finite index free abelian subgroup of each one stabilizes a flat F_i by the Flat Torus Theorem 1.3.32. We have already seen that if two such flats were parallel, they would be in fact stabilized by the same subgroup, so the associated flats are not parallel. We have already seen that translates of these flats are stabilized by conjugates of the H_i 's and that a quasi isometry induced by the action associates left cosets of H_i 's to the translates of flats they stabilize. Two corresponding asymptotic cones of X and G are homeomorphic; so the fact that G is relatively hyperbolic with respect to $\{H_i\}$ implies X has isolated flats with respect to the G -translates of $\{F_i\}$. \square

We now want to study the behaviour of spaces with isolated flats under quasi isometries. Remember first that if we have the action of a group G on two sets X and Y , we say that a function $f: X \rightarrow Y$ is G -equivariant if for every x in X and every g in G the equality $f(g \cdot x) = g \cdot f(x)$ holds. In the

category of metric spaces with quasi isometries as morphisms we consider G -quasi-equivariant quasi isometries (briefly G -q.e.q.i.) f such that the distance between $f(g \cdot x)$ and $g \cdot f(x)$ is uniformly bounded. Note that a quasi inverse of a G -q.e.q.i. is automatically a G -q.e.q.i.

An example is readily provided by two proper metric spaces X and Y with a geometric action of a group G on both. Remember that G acts on itself by left multiplications and that this action is by isometries on any Cayley graph. Then, for any x_0 in X and y_0 in Y , the maps $g \mapsto g \cdot x_0$ and $g \mapsto g \cdot y_0$ are G -q.e.q.i.'s, and the quasi isometry defined like the composition of one of these two with a quasi inverse of the second is G -quasi-equivariant too.

We concentrate now on spaces with isolated flats. Remember that this property is invariant under quasi isometries. The fact we want to prove is the following

Lemma 1.4.8. Let X and Y be proper CAT(0) spaces with isolated flats with respect to a geometric action of a group G , let $f: X \rightarrow Y$ be a G -q.e.q.i. and γ be the geodesic between p and q in X . Then there is a constant λ depending only on the spaces, the group and f such that the Hausdorff distance between the image of $f \circ \gamma$ and the image of the unique geodesic between $f(p)$ and $f(q)$ is less or equal than λ .

Furthermore, if γ is a geodesic ray in X issuing from p , there is a geodesic ray β in Y issuing from $f(p)$ such that $f \circ \gamma$ and β have Hausdorff distance less or equal than λ .

First part of the proof. We begin the proof for a very particular case, namely the one where X and Y are both \mathbb{E}^n and G is \mathbb{Z}^n acting by translations. This case falls into the hypothesis because \mathbb{E}^n has isolated flats with respect to $\{\mathbb{E}^n\}$. Chosen basepoints x_0 and $y_0 = f(x_0)$ in X and Y respectively, f brings $g \cdot x_0$ in a point uniformly close to $g \cdot y_0$. It follows that f is uniformly close to an affine map, and then the thesis follows immediately, because affine maps bring geodesics into geodesics. \square

Before continuing with the proof, note that an action of a group G on a space X induces an action of *G on *X , where the ultrapowers are made on the same index set and the same ultrafilter. In particular, a G -q.e.q.i f between X and Y becomes *f , which is a *G -q.e.q.i. with the same constants, between *X and *Y . For example, the action by translations of \mathbb{Z}^n on \mathbb{E}^n induces an action by translations of an ultrapower $({}^*\mathbb{Z})^n = {}^*(\mathbb{Z}^n)$ on the corresponding ultrapower ${}^*\mathbb{E}^n$. The proof we just made tells us that in the large scale, a \mathbb{Z}^n -q.e.q.i between two \mathbb{E}^n 's becomes an affinity between corresponding asymptotic cones.

In the more general case of a space X having isolated flats with respect to a family \mathcal{F}_X , we already know that there is only a finite number of G -orbits in \mathcal{F}_X . When we have $f: X \rightarrow Y$ as in the hypothesis, it brings flats in \mathcal{F}_X near to flats in \mathcal{F}_Y , the family with respect to which Y has isolated flats, and the correspondence is bijective, mediated by free abelian subgroups of G . So restricting f to a flat in \mathcal{F}_X and then projecting, using Lemma 1.2.6, on the corresponding flat in \mathcal{F}_Y , gives a map uniformly close to an affinity. Different flats in \mathcal{F}_X are transformed via different affinities this way, but there is only a finite number of them, up to isometry. In fact, if F_0 is a flat in \mathcal{F}_X corresponding to F'_0 in \mathcal{F}_Y and $f|_{F_0}$ later composed with the projection on F'_0 is close to an affinity φ , then, taken a g in G , the flat $g \cdot F_0$ in X corresponds to the flat $g \cdot F'_0$ in Y , and $f|_{g \cdot F_0}$ is uniformly close to $g \circ \varphi \circ g^{-1}$. The conclusion follows from the finiteness of G -orbits in \mathcal{F}_X .

These considerations extend directly to ultrapowers. The only observation to make here is that ${}^*\mathcal{F}_X$ has only a finite number of *G -orbits. In fact, if I is the index set and \mathcal{U} is the ultrafilter, an I -sequence F_i of flats in \mathcal{F}_X induces a subdivision of I in sets of flats F_i in the same G -orbit. Exactly one of these sets is in \mathcal{U} , so we may suppose without loss of generality that all the flats of the sequence lie in the same well defined G -orbit.

Take two sequences $(F_i)_{i \in I}$ and $(F'_i)_{i \in I}$ such that the flats of both lie in the same G -orbit. There are elements g_i in G , such that $g_i \cdot F_i = F'_i$. But then, in the ultrapower, we have that $[g_i] \cdot [F_i] = [F'_i]$. The finite number of G orbits of flats in X implies thus a finite number of *G -orbits of hyperflats in *X .

Second part of the proof. Step 1: Suppose, to fix the notation, that f is a (K, C) -quasi-isometry, and take a geodesic $\alpha: [0, a] \rightarrow X$. The geodesic α' between $f(\alpha(0))$ and $f(\alpha(a))$ is a closed convex subset of a proper CAT(0) space, so we can apply the Lemma 1.2.6 to get a 1-Lipschitz retraction ρ of Y on it. It follows by taking subdivisions of the domain of α in arbitrarily small pieces that for any point in the image of α' there is a point in the image of $\rho \circ f \circ \alpha$ at distance at most C from it. Then, to prove the thesis it suffices to prove that the distance of points on $f \circ \alpha$ from α' is uniformly bounded independently of α .

Step 2: Suppose by contradiction that for any positive real λ there are geodesics in X whose image through f has points at distance greater than λ from the geodesic in Y joining the images of the endpoints. This means that in *X there is a geodesic α such that there are points in the image of ${}^*f \circ \alpha$ at infinite distance from the geodesic α' in *Y joining the images of the endpoints.

We have already seen in the Lemma 1.3.11 that images of the geodesics

are internal subsets in the ultrapower. So, by Transfer Principle it is well defined the infinite hyperreal Λ , the supremum of the distances of points in $*f \circ \alpha$ from α' . Take a point $r = *f \circ \alpha(t)$ at distance at least $\frac{\Lambda}{2}$ from α' and consider the asymptotic cone $\mathcal{C}(X)$ of centre $\alpha(t)$ and radius Λ , the corresponding asymptotic cone $\mathcal{C}(Y)$ and the K -bi-Lipshitz homeomorphism $\mathcal{C}(f)$ between them induced by f . The intersection of α' with $\mathcal{C}(Y)$ projects to a geodesic $\tilde{\alpha}'$ in the cone, whilst the intersection of $*f \circ \alpha$ projects to a continuous injective K -bi-Lipschitz curve, both defined on an interval of a suitable asymptotic cone of \mathbb{R} , which is still \mathbb{R} . The projection of r , which we will call \tilde{r} , stays in the image of the latter curve, at a distance at least $\frac{1}{2}$ from $\tilde{\alpha}'$.

Step 3: Now we distinguish two cases. In the first one, the endpoints of $*f \circ \alpha$ both lie in $\mathcal{C}(Y)$. Call $\tilde{\alpha}: [0, \tilde{a}] \rightarrow \mathcal{C}(Y)$ the projection on the cone of this curve and let $(t_1, t_2) \subset [0, \tilde{a}]$ be the maximal interval whose image contains \tilde{r} and does not intersect the image of $\tilde{\alpha}'$. Then $\tilde{\alpha}([t_1, t_2])$ along with the piece of $\tilde{\alpha}'$ between $\tilde{\alpha}(t_1)$ and $\tilde{\alpha}(t_2)$ forms a simple closed curve, that stays thus on a single flat piece \tilde{F} of $\mathcal{C}(Y)$. The function $\mathcal{C}(f)$ is an affinity when restricted to the pre image of \tilde{F} because of the first part of the proof. But then the image of $\tilde{\alpha}|_{[t_1, t_2]}$ should coincide with the image of the geodesic between two endpoints, which is the piece of $\tilde{\alpha}'$ between these two points, absurd.

Step 4: The other case is when at least one of the endpoints of $*f \circ \alpha$ lies outside $\mathcal{C}(Y)$. In this case $\tilde{\alpha}'$ becomes a geodesic ray or line in the cone and $\tilde{\alpha}$ a continuous K -bi-Lipschitz embedding of \mathbb{R} or of a half line into the cone. If there is an interval $[t_1, t_2]$ whose endpoints are sent by $\tilde{\alpha}$ into the image of $\tilde{\alpha}'$ and the internal part outside it, the configuration may be dealt with like in the first case. So we may suppose that $\tilde{\alpha}$ and $\tilde{\alpha}'$ intersect at most once in a closed segment inside the cone.

Take then a point \tilde{s} on $\tilde{\alpha}'$ that stays outside this (maybe empty) intersection; this is also true for points near \tilde{s} . The point \tilde{s} is represented by some point s on α' in $*Y$, and, by the argument at the beginning of the second part of the proof, called ρ the projection on α' there is a point in the image of $\rho \circ *f \circ \alpha$ at distance at most C from s . This means that, taking the projection $\tilde{\rho}$ on $\tilde{\alpha}'$ the point \tilde{s} is in the image of $\tilde{\rho} \circ \tilde{\alpha}$, like any other point in $\tilde{\alpha}'$.

Take a point s' on α' such that its projection to the cone \tilde{s}' is near \tilde{s} on $\tilde{\alpha}'$, but different from it, and still outside the intersection. We can suppose s comes before s' on α' . Connect each of \tilde{s} and \tilde{s}' with a geodesic to a pre image via $\tilde{\rho}$ on $\tilde{\alpha}$. We can suppose each of these geodesics intersects $\tilde{\alpha}$ only once in the endpoint. Thus the piece of $\tilde{\alpha}$ between these two endpoints, the

two geodesics and the piece of $\tilde{\alpha}'$ between \tilde{s} and \tilde{s}' form a simple closed curve, which then stays on a single flat piece of $\mathcal{C}(Y)$.

Step 5: This piece is the projection on the cone of a well defined hyperflat F in ${}^*\mathcal{F}$. To fix the ideas, suppose that α' is parametrized on the hyperreal interval $[0, A]$ and that $\alpha'(t_0) = s$. It is well defined, for the previous considerations on internal sets, the hyperreal t' , the supremum of positive hyperreals t less or equal than t_0 for which the distance of $\alpha'(t)$ from F is greater than Λ . Take the asymptotic cone $\mathcal{C}'(Y)$ of centre $\alpha'(t')$ and radius Λ , and call F' the non empty projection of F on this cone. For the other points, we will continue to use the convention to put a tilde over a point to indicate its projection to the cone. The symbol $\tilde{\alpha}$ will be again used for the projection of ${}^*f \circ \alpha$ to the cone, and $\tilde{\alpha}'$ analogously for α'

We search for an intersection of $\tilde{\alpha}$ and $\tilde{\alpha}'$ in F' . If $t' = 0$ and $\widetilde{\alpha'(0)}$ stays in F' we are done. Otherwise,

$$\widetilde{\alpha'(t')} = \tilde{\alpha}'(\tau)$$

is at a positive distance less or equal than 1 from F' . We claim that in $\tilde{\alpha}'([\tau, \tau + 3])$ there is a point in F' . If it were not, by construction the distance of every point in $\tilde{\alpha}'([\tau, \tau + 3])$ from F' is less or equal than 1. Consider the two geodesics η' and η'' from $\tilde{\alpha}'(\tau)$ and $\tilde{\alpha}'(\tau + 3)$ to the points u' and u'' nearest to them on F' . These point are different, therefore the pieces of η' and η'' from u' and u'' respectively until their first intersection with $\tilde{\alpha}'$, the piece of $\tilde{\alpha}'$ between these two intersections and the geodesic, lying on F' , between u' and u'' form a simple closed curve, so it stays on a single flat piece in $\mathcal{C}'(Y)$, which must be F' for the tree graded property, absurd.

Step 6: Now we found a point s'' on α' which stays at a hyperdistance from F infinitesimal when divided by Λ , and such that the points of $\tilde{\alpha}'$ preceding \tilde{s}'' are not on F' . By the convexity of the distance function on geodesics all the points of the piece of α' between s'' and s' have a hyperdistance from F infinitesimal when divided by Λ ; in particular, the points on $\tilde{\alpha}'$ after \tilde{s}'' and close enough to it stay on F' . If \tilde{s}'' stays on $\tilde{\alpha}$ we are done.

Suppose by contradiction this is not true. Then there are no intersections of $\tilde{\alpha}'$ and $\tilde{\alpha}$ neither near \tilde{s}'' . Take then two points \tilde{s}_0 and \tilde{s}_1 on $\tilde{\alpha}'$ in this neighbourhood of \tilde{s}'' , the first before \tilde{s}'' and the second after. Similarly to the argument in Step 4 we obtain a simple closed curve containing the piece of $\tilde{\alpha}'$ between \tilde{s}_0 and \tilde{s}_1 and a piece of $\tilde{\alpha}$, which thus stays in a single flat piece of $\mathcal{C}'(Y)$: it must be F' for the tree graded property but cannot be F' for the choice of \tilde{s}_0 .

Step 7: When we interpret this configuration in the ultrapower, we obtain three points v_1, v_2 and v_3 in *Y such that the first stays on F , the second




Figure 1: Lemma 1.4.8, Step 6 of the second part of the proof

on α' before s and the third on $*f \circ \alpha$ and whose distances are infinitesimal when divided by Λ . We repeat the same argument to find t'' , the infimum of hyperreals t in $[t_0, A]$ for which $\alpha'(t)$ stays at a greater distance from F than Λ and then three points w_1 , w_2 and w_3 with analogous properties, with w_2 on α' but after s . Furthermore, the choice of s ensures us that $*d(v_2, w_2)$ is not infinitesimal when divided by Λ .

Let \bar{F} be the hyperflat in $*X$ corresponding to F in $*Y$. We already know that $*f|_{\bar{F}}$ composed with the projection on F is uniformly close to an affine map, where the uniformity is extended on all hyperflats in $*X$. From the quasi-isometry we also know that pre images of v_3 and w_3 stay at a distance from \bar{F} which is infinitesimal when divided by Λ . By projecting them on \bar{F} and repeatedly using quasi-isometry, we obtain that the piece of $*f \circ \alpha$ between v_3 and w_3 has a Hausdorff distance from the geodesic between the two points which is infinitesimal when divided by Λ , and this geodesic as at a Hausdorff distance with the same property from the piece of α' between v_2 and w_2 ; here we also used the convexity of distance on geodesics. But this is absurd because we supposed there was a point of $*f \circ \alpha$ between v_3 and w_3 at a distance at least $\frac{\Lambda}{2}$ from α' .

Step 8: Finally, let γ be a geodesic ray in X issuing from p . Consider the geodesic $*\gamma$ defined on positive hyperreals with values in $*X$. The first part of the proof, along with the Transfer Principle, tells us that, taken an infinite positive real T , the image of $*f \circ *\gamma$ and the geodesic between $*f(*p)$ and $*f(*\gamma(T))$ have Hausdorff distance less than λ . By taking the standard parts of points at finite distance from $*f(*p) = f(p)$ we obtain a geodesic ray issuing from $f(p)$ at Hausdorff hyperdistance less than λ from the image of $f \circ \gamma$. \square

The previous lemma is true in hyperbolic spaces too. The proof is simpler thanks to the more rigid structure of real trees in comparison with tree graded spaces.

By construction, every geodesic ray in Y with the same properties of the

one found in the end of the previous lemma, except perhaps the endpoint, is asymptotic to the ray we already constructed. This is a fundamental ingredient in the proof of the following

Lemma 1.4.9. Suppose X, Y, G and f are as in the previous lemma. Then there is a G -equivariant homeomorphism $\partial f: \partial X \rightarrow \partial Y$ such that, if a sequence $(x_n)_{n \in \mathbb{N}}$ in X converges to a boundary point \bar{x} in the topology of \bar{X} , then $(f(x_n))_{n \in \mathbb{N}}$ converges to $\partial f(\bar{x})$ in \bar{Y} .

Proof. Define ∂f on the class of a geodesic ray γ to be the class of the geodesic ray β constructed in the proof of the previous Lemma. This function is well defined for the considerations we made, and clearly, if \tilde{f} is a G -quasi-equivariant quasi inverse of f , then $\partial \tilde{f} = (\partial f)^{-1}$ by the quasi-isometry property. The G -equivariance of ∂f follows easily too. So it remains to prove the continuity; the continuity of the inverse is identical.

To fix the notations, suppose f is a (K, C) -quasi-isometry and λ is the constant found in the previous Lemma. Choose p in X and select a ray γ issuing from it. Consider the ray β in Y issuing from $f(p)$ and representing the class of $\partial f([\gamma])$. It suffices to prove that, chosen positive reals t and ε , the preimage of the open set $A(B(\beta(t), \varepsilon), f(p)) \cap \partial Y$ containing the class of β , with A defined in the considerations preceding the Lemma 1.2.8, is a neighbourhood of the class of γ in ∂X .

Choose then s large enough to ensure that $\frac{s\varepsilon}{t} - 2C - 4\lambda$ is positive and call $\delta = \frac{1}{2K} \left(\frac{s\varepsilon}{t} - 2C - 4\lambda \right) > 0$. Define $r = K(s + C)$. We claim that $\partial f(A(B(\gamma(r), \delta), p) \cap \partial X) \subseteq A(B(\beta(t), \varepsilon))$, which suffices. Indeed, let $s' \geq s$ be $d(f(p), f(\gamma(r)))$. We know that there is a point in the image of β at a distance at most λ from $f(\gamma(r))$. The distance of this point from $f(p)$ lies between $s' - \lambda$ and $s' + \lambda$. From the triangle inequality follows that $d(f(\gamma(r)), \beta(s')) \leq 2\lambda$. If a geodesic γ' represents a class in $A(B(\gamma(r), \delta), p) \cap \partial X$, then a point $\gamma'(r')$ on it stays at distance less than δ from $\gamma(r)$ and thus $d(f(\gamma(r)), f(\gamma'(r'))) < K\delta + C$; furthermore there is a point on the geodesic β' issuing from $f(p)$ and representing $\partial f([\gamma'])$ with distance less or equal than λ from $f(\gamma'(r'))$. Combining the triangle inequality and an argument similar to the one already used in this proof we obtain that $d(\beta(s'), \beta'(s')) < 4\lambda + 2K\delta + 2C$. By the CAT(0) inequality and by construction, we have $d(\beta(t), \beta'(t)) < \varepsilon$.

The last assertion follows by construction and the properties of quasi isometries. \square

Note that, by definition, if F is a k -flat in X of the family \mathcal{F} involved in the isolated flats definition, it is stabilized by a virtually abelian subgroup of G of rank k , and its boundary $\partial F \subseteq \partial X$ given by the classes of geodesic rays

contained in F , homeomorphic to \mathbb{S}^{k-1} , is brought by ∂f in the boundary of a flat in Y stabilized by the same subgroup.

2 Polyhedral metric complexes

This section is devoted to polyhedral complexes, specifically simplicial and cubical ones, to the natural metric space structure on them and to a proof of the classical theorem by Gromov which provides necessary and sufficient conditions for them to be CAT(κ) spaces.

2.1 Polyhedral complexes

We begin with some basic definitions on polyhedral complexes.

Definition 2.1.1 (Polyhedron). *Let κ be a real number and n be an integer equal or greater than 2. A polyhedron of dimension n , or n -polyhedron of curvature κ is a compact subset of \mathbb{M}_κ^n having non empty internal part and being the intersection of a finite number of closed half spaces.*

We will also call a real closed bounded interval 1-polyhedron and a single point 0-polyhedron.

Note that polyhedra of dimension 0 and 1 do not have a well defined curvature, but for sake of brevity we will not treat them separately when talking about collections of polyhedra, some having higher dimension and thus a curvature, usually the same for all of them, and some having dimension 0 or 1. By definition, polyhedra are convex subsets of the relative model spaces.

Polyhedra have faces, which are polyhedra too.

Definition 2.1.2 (Face). *Let σ be the boundary of one of the half spaces involved in the definition of a polyhedron P , such that $\sigma \cap P$ has nonempty internal part in σ . A face of codimension 1 is $\sigma \cap P$ for one such σ .*

If $n = 1$, a face is an endpoint of the interval. A 0-polyhedron has no faces.

Note that a face of codimension 1 of a polyhedron of dimension n and curvature κ is a polyhedron of dimension $n - 1$ and curvature κ . So we can also say that a face of codimension 1 has dimension $n - 1$.

We can then define faces of codimension 2 and dimension $n - 2$, and so on, until 1-dimensional faces, which we will call *arcs* or *edges*, and 0-dimensional ones, which we will call *vertices*.

If P is a polyhedron, we will denote by ∂P the union of its faces and by $\text{Int } P$ the internal part $P \setminus \partial P$. Note that $\partial \{\text{point}\} = \emptyset$.

The examples of polyhedra which will be used mostly throughout this text are simplices and cubes.

Denoting by I the real interval $[0, 1]$, a *standard euclidean cube* of dimension n is simply I^n . Its faces are products of some factors I and some factors corresponding to one endpoint of I on the relative coordinate, in some order, and are euclidean cubes too.

To define a simplex of dimension $n \geq 2$ and curvature κ begin with choosing $n + 1$ points in \mathbb{M}_κ^n in general position, i.e. not lying on a hyperplane. For every choice of n points among them, consider the half space having for boundary the unique subspace passing for them, and containing the remaining point. Intersect all these subspaces to obtain the simplex.

Note that the combinatorics of such metric n -simplex relatively to its faces is isomorphic to that of the abstract n -simplex, but metric simplices have an additional structure, so when constructing metric simplicial complexes later we must take in account that structure. To have more freedom in operating in that setting, we define the *regular* metric n -simplex of curvature κ having the additional property that the group of isometries of \mathbb{M}_κ^n has a subgroup that permutes the vertices of the simplex in every possible way. Note that a regular n -simplex of curvature κ is determined up to isometry by the length of its 1-faces, which has to be less than some positive constant depending on the dimension if $\kappa > 0$. We will always use an upper bound of $\frac{\pi}{2\sqrt{\kappa}}$ for the length in that case, which will always suffice.

After defining polyhedra, we may pass to assemble them into polyhedral complexes. We begin with the topological definition.

Definition 2.1.3 (Polyhedral complex). *A polyhedral complex is a topological space X together with a collection of embeddings $\{f_\alpha: P_\alpha \rightarrow X\}_\alpha$, where P_α are polyhedra, such that:*

- *Every $x \in X$ is contained inside one and only one $f_\alpha(\text{Int } P_\alpha)$;*
- *For every α and every immersion $\iota: \sigma \rightarrow P_\alpha$ of a face into P_α , the composition $f_\alpha \circ \iota$ is an embedding $f_\beta: P_\beta \rightarrow X$ of the collection.*
- *Given two distinct polyhedra $f_\alpha: P_\alpha \rightarrow X$ and $f_\beta: P_\beta \rightarrow X$ either $f_\alpha(P_\alpha) \cap f_\beta(P_\beta) = \emptyset$, or the pre image of the intersection via f_α is a single face of P_α and analogously for f_β .*

When no confusion arises, we will call polyhedron also the image of one of the f_α in the definition. The following result tells us how we can define on a polyhedral complex a pseudo-distance only in terms of its combinatorics and metrics on its faces.

Lemma 2.1.4. Let X be a polyhedral complex defined by $\{f_\alpha: P_\alpha \rightarrow X\}_\alpha$, and let κ be a real number. Suppose every polyhedron has a metric structure

deriving from it being of curvature κ such that whenever the image of a P_β coincides with the image of a face of a P_α , the function $f_\alpha^{-1} \circ f_\beta$ is an isometric embedding. Then there is a unique pseudo-distance on X such that the induced length metric on every polyhedron coincides with the standard one.

Proof. Let x and y be points in X . A sequence $x_0 = x, x_1, \dots, x_m = y$ of points in X such that every pair $\{x_{i-1}, x_i\}$ of consecutive elements lie in the same polyhedron (P_i, d_i) is said to be a *string* between x and y . It is now easy to see that the function

$$d(x, y) = \inf \sum_{i=1}^m d_{P_i}(x_{i-1}, x_i)$$

is well defined and meets the required properties. \square

A polyhedral complex endowed with this pseudo-distance is said to be a *polyhedral metric complex*. Note that f_α become 1-Lipschitz maps but are not, in general, isometric embeddings. We want to isolate a condition that makes the defined function a true distance and the embeddings local isometries.

Take a point x in X . If (P, d) is a polyhedron whose image contains x , define $\varepsilon_P(x) = \inf d_P(x, F)$, with F ranging over the faces of P not containing x ; if (P, d) is a point, and thus x is its image, we define this quantity being $+\infty$. Then take $\varepsilon(x) = \inf_P \varepsilon_P(x)$.

Lemma 2.1.5. If for every x in X we have $\varepsilon(x) > 0$, then X is a metric space and the embeddings of polyhedra are local isometries.

Proof. We will prove the following: fix x in (X, d) . If y is such that $d(x, y) < \varepsilon(x)$, then every polyhedron (P, d_P) containing y contains x and $d(x, y) = d_P(x, y)$, which obviously suffices.

Take then a string $(x_0 = x, x_1, \dots, x_m = y)$ of length less than $\varepsilon(x)$. We claim that the string obtained by omitting x_1 has shorter length. In fact there is a polyhedron P_2 containing both x_1 and x_2 and by definition of $\varepsilon(x)$ and hypothesis on the length it also contains x . But then

$$d_{P_2}(x_0, x_2) \leq d_{P_2}(x_0, x_1) + d_{P_2}(x_1, x_2).$$

Proceeding by induction we get to the thesis. \square

Let us examine separately polyhedral metric complexes we are mostly interested in.

Definition 2.1.6 (Cubical complex). *A cubical complex is a polyhedral complex where all the faces are standard euclidean cubes.*

Definition 2.1.7 (Simplicial metric complex). *A simplicial metric complex is a polyhedral complex where all the faces are metric simplices with the same curvature.*

Cubical complexes and simplicial metric complexes made up of isometric regular simplices meet the requirements of Lemma 2.1.5 and are thus true metric spaces.

Note that single polyhedra are in an obvious way polyhedral metric complexes with the new distance function coinciding with the old one.

Lemma 2.1.4 allows us to construct a metric complex by gluing polyhedra. We begin by taking the disjoint union of a collection of polyhedra $\{P_\alpha\}_\alpha$ of various dimensions, but of the same curvature, and we let \tilde{f}_α be the natural embeddings. Then we take a collection of pairs of polyhedra in the set such that one polyhedron of the pair isometrically embeds in the second as one of its faces. We require that no polyhedron can be identified with more than one face of another one, and that, if two distinct polyhedra σ_1 and σ_2 embed each in two other polyhedra P_1 and P_2 , then either one of the σ embeds in the second as a face, or they both embeds as faces in a third polyhedron σ_3 which embeds as a face in both P_1 and P_2 . We require furthermore that, if a polyhedron Q_1 embeds as a face in Q_2 and Q_2 embeds as a face in Q_3 , then Q_1 embeds as a face in Q_3 via the composition of the two embeddings. Consider the quotient \sim of the disjoint union by these isometries. We apply the previous lemma to the quotient space with $f_\alpha = \sim \circ \tilde{f}_\alpha$ to get a pseudo-distance function. We say that the complex is obtained by *gluing polyhedra along identifications of faces*.

Note that despite single polyhedra are complete metric spaces, we did not at any point claim that general polyhedral complexes, even when they are metric spaces, are complete too, and this is in fact false. There are subtler arguments that allow us to get to such conclusion under certain hypotheses, but we will not mention them here because the polyhedral complexes we will use will be either compact, or obtained by lifting the polyhedral structure of a compact polyhedral complex to its universal cover.

2.2 Links

In this subsection we will define a concept of link similar to the simplicial one, but suitable for dealing with polyhedral complexes. We begin with defining it for single polyhedra.

Definition 2.2.1. *Let P be a polyhedron, thought as a subset of a model space \mathbb{M} and v a vertex of P . The link of P at v , which we will indicate by $\text{lk}_P v$, is the subset of $T_v \mathbb{M}$ consisting of all the unit vectors tangent to some of the geodesics between v and another point of P .*

Note that the link at a vertex of an n -polyhedron is an $n - 1$ -polyhedron of curvature 1. Note also that in both cases of a simplex and a cube, the link at a vertex is a simplex of 1 dimension less.

To an isometric embedding of a face σ , having v as a vertex, into a polyhedron P corresponds an isometric embedding of $\text{lk}_\sigma v$ in $\text{lk}_P v$. Therefore, the following definition makes sense.

Definition 2.2.2 (Link of a vertex). *Let K be a polyhedral metric complex and v a vertex in it. The link of K at v , denoted by $\text{lk}_K v$ is the polyhedral metric complex obtained by gluing the links at v of the polyhedra containing v along the isometries given by the inclusion of the faces into the corresponding polyhedra.*

In the simplicial context St and $\overset{\circ}{\text{St}}$ will indicate the closed and the open star and we will call simplicial link their set difference, to avoid any confusion with the metric link, or simply link, we have just defined.

Note that, when the hypotheses of Lemma 2.1.5 are satisfied every vertex has a neighbourhood homeomorphic to the cone on its link. We want therefore study cones from a metric point of view to understand more deeply the structure of polyhedral complexes. The letter κ will always denote a real number thought as the curvature of the objects we want to define.

We now recall the law of cosines in \mathbb{M}_κ^2 with distance function d : let $\triangle xyz$ be a geodesic triangle there, and α the angle at x . Let $a = d(y, z)$, $b = d(x, y)$ and $c = d(x, z)$. Then:

$$\cosh(a\sqrt{-\kappa}) =$$

$$= \cosh(b\sqrt{-\kappa}) \cosh(c\sqrt{-\kappa}) - \sinh(b\sqrt{-\kappa}) \sinh(c\sqrt{-\kappa}) \cos(\alpha)$$

if $\kappa < 0$,

$$a^2 = b^2 + c^2 - 2bc \cos(\alpha)$$

if $\kappa = 0$ and

$$\cos(a\sqrt{\kappa}) = \cos(b\sqrt{\kappa}) \cos(c\sqrt{\kappa}) + \sin(b\sqrt{\kappa}) \sin(c\sqrt{\kappa}) \cos(\alpha)$$

if $\kappa > 0$.

Let now (Y, d) be a metric space and d_π be the distance on Y defined by $\inf\{d, \pi\}$. We construct the topological cone CY on Y by taking

$$Y \times [0, +\infty) / \sim$$

if $\kappa \leq 0$ and

$$Y \times \left[0, \frac{\pi}{2\sqrt{\kappa}}\right] / \sim$$

if $\kappa > 0$, where \sim identifies all points with real coordinate 0 to a point said *vertex*, and denote the class of (y, t) by ty . We will now describe a distance function on CY , depending on κ .

Definition 2.2.3 (κ -cone). *The κ -cone on Y , denoted by $C_\kappa Y$, is CY where the distance between two points $t_1 y_1$ and $t_2 y_2$ is t_2 if $t_1 = 0$ and vice versa, and:*

$$\operatorname{arccosh} \left(\frac{\cosh(t_1 \sqrt{-\kappa}) \cosh(t_2 \sqrt{-\kappa}) - \sinh(t_1 \sqrt{-\kappa}) \sinh(t_2 \sqrt{-\kappa}) \cos(d_\pi(y_1, y_2))}{\sinh(t_1 \sqrt{-\kappa}) \sinh(t_2 \sqrt{-\kappa}) \cos(d_\pi(y_1, y_2))} \right)$$

if $\kappa < 0$,

$$\sqrt{t_1^2 + t_2^2 - 2t_1 t_2 \cos(d_\pi(y_1, y_2))}$$

if $\kappa = 0$ and

$$\arccos(\cos(t_1 \sqrt{\kappa}) \cos(t_2 \sqrt{\kappa}) + \sin(t_1 \sqrt{\kappa}) \sin(t_2 \sqrt{\kappa}) \cos(d_\pi(y_1, y_2)))$$

if $\kappa > 0$.

In the particular case where Y is \mathbb{S}^{n-1} the space $C_\kappa Y$ is isometric to \mathbb{M}_κ^n if $\kappa \leq 0$ and to a half space of \mathbb{M}_κ^n otherwise. But in the general case, we first have to show that the definition makes sense.

Lemma 2.2.4. The above defined function is a distance.

Proof. The only non trivial thing to prove is the triangle inequality. Let d and d_π be the two previously defined distances on Y and $t_x x, t_y y, t_z z$ three points in $C_\kappa Y$. If one of them is the vertex, the thesis follows from the triangle inequality in \mathbb{M}_κ^2 . If $d(x, y) + d(y, z) < \pi$, the set $\{x, y, z\}$ with the defined distances embeds into \mathbb{M}_κ^3 so the thesis follows again by triangle inequality.

So remains the case where $d(x, y) + d(y, z) \geq \pi$. In $(\mathbb{M}_\kappa^2, \bar{d})$ choose points $\tilde{v}, \tilde{x}, \tilde{y}$ and \tilde{z} such that $\bar{d}(v, x) = t_x$ and analogously for y and z , the angle at v between \tilde{x} and \tilde{y} is less or equal than $d_\pi(x, y)$, the angle at v between \tilde{y} and \tilde{z} is less or equal than $d_\pi(y, z)$ and the angle at v between \tilde{x} and \tilde{z} is equal to $d_\pi(x, z)$. The thesis follows from triangle inequality in \mathbb{M}_κ^2 and monotonicity of the length of a side of a triangle as a function of the opposite angle. \square

Having defined the distance function on the cone, we may pass to examine its curvature. First, let us describe when it makes sense to talk about it, i.e. when the cone is geodesic.

Lemma 2.2.5. If (Y, d) is geodesic, $C_\kappa Y$ is geodesic.

Proof. Take two points $t_x x$ and $t_y y$ in $C_\kappa Y$. If one of the points is the vertex, say the first, it can be seen by triangle inequality and construction that the unique geodesic between them is $\gamma: [0, t_y] \rightarrow C_\kappa Y$ given by $\gamma(t) = ty$.

In the other case, let $L = d(x, y)$. If $L \geq \pi$, similarly to the previous case the unique geodesic between $t_x x$ and $t_y y$ is the concatenation of the geodesics from the vertex to the two points. Otherwise, let $\beta: [0, L] \rightarrow Y$ be a geodesic between x and y . Consider in \mathbb{M}_κ^2 a triangle with one vertex v and the two sides exiting from v having other extremes \bar{x} and \bar{y} , length t_x and t_y respectively and forming an angle L . Let L' be the length of the third side and denote by $\bar{\gamma}: [0, L'] \rightarrow \mathbb{M}_\kappa^2$ its parametrization as a geodesic beginning in \bar{x} . Let $a: [0, L'] \rightarrow [0, L]$ denote the angle formed by $v\bar{x}$ and $v\bar{\gamma}(t)$ and $b: [0, L'] \rightarrow \mathbb{R}$ the distance between v and $\bar{\gamma}(t)$, for t in $[0, L']$. Then the curve $t \mapsto b(t)\beta(a(t))$ defined on $[0, L']$ is a geodesic between $t_x x$ and $t_y y$. \square

So, if Y is geodesic, we can try to study the curvature of a κ -cone on it in the sense of CAT.

Lemma 2.2.6 (Berestovskij Theorem). If Y is CAT(1), then the cone $C_\kappa Y$ is CAT(κ).

For a proof, see [BH99, Theorem II.3.14].

Theorem 2.2.7. Let K be a polyhedral metric complex respecting the hypotheses of 2.1.5 with faces of curvature κ . If the link at every vertex is a CAT(1) space, then K is locally CAT(κ).

Proof. The thesis follows immediately from 2.2.6 for a neighbourhood of a vertex. It remains to note that any other point x of K has a neighbourhood isometric to an open ball in a neighbourhood of a vertex of a polyhedron containing x . \square

If we want a global statement, we need to assume something on the structure of the geodesics. We state the following result only for curvature 1.

Lemma 2.2.8. Let K be a polyhedral metric complex with faces of curvature 1. If the link at every vertex is CAT(1) and there are no closed geodesics of length strictly less than 2π , then K is CAT(1).

For a proof, see [BH99, Theorem II.5.4].

Let us now concentrate specifically on cubical complexes. Note that the link at the vertex of an n -cube is a spherical, i.e. of curvature 1, regular $n - 1$ -simplex with all arcs of length $\frac{\pi}{2}$. Then the link at the vertex of a cubical complex is a simplicial complex of curvature 1 and arcs of length $\frac{\pi}{2}$. The metric link of such a complex is again a complex of the same type and it is not only combinatorially equivalent, but also isometric to the link defined like $\mathring{\text{St}} \setminus \mathring{\text{St}}$. So, the metric of the simplicial complexes that arise by cubical complexes can be described in purely combinatorial terms. We begin with the following

Lemma 2.2.9. Let K be a simplicial metric complex with spherical faces having edges of length $\frac{\pi}{2}$ and let γ be a geodesic in it. If v is a vertex of K then $\mathring{\text{St}}(v) \cap \gamma$ has length π .

Proof. $\mathring{\text{St}}v$ is $C_1 \text{lk}_K v$ without the base. If γ passes through v then the thesis is obvious, otherwise consider the collection of geodesic segments exiting from v , passing through point of $\mathring{\text{St}}(v) \cap \gamma$ and having length $\frac{\pi}{2}$. This allows us to establish an isometry between a convex subset of $\mathring{\text{St}}v$ and a convex subset of the northern hemisphere of \mathbb{S}^2 with v going to the north pole. The curve γ must go to a geodesic, hence a half maximal circle, which leads to the thesis. \square

We now define a combinatorial property of simplicial complexes.

Definition 2.2.10 (Flag simplicial complex). *A simplicial complex is said to be flag if every subset of vertices spanning the 1-skeleton of a simplex spans that simplex.*

Note that a flag complex is determined by its 1-skeleton. Now we can prove the following

Lemma 2.2.11. If a locally finite simplicial complex K is flag, then giving it a structure of metric complex with spherical faces with sides of length $\frac{\pi}{2}$ leads to a CAT(1) space.

Proof. We have to prove that the link at every vertex is a CAT(1) space and that there are no closed geodesics of length strictly less than 2π . If we prove the latter for the whole complex, by previous considerations it will be true also for the link at every vertex, which is again a simplicial complex of the same type, and then for every link at the vertex of this complex and so on by induction which finishes because of locally finiteness. Then we can prove

the thesis also by induction on the dimension of the complex, because a link of a flag complex is still flag.

Take then a geodesic γ in K and consider the family of vertices v such that $\mathring{\text{St}}(v)$ intersects γ . By the previous Lemma, it suffices to prove that there are two vertices of the family having disjoint stars. If there are not, then any pair of vertices of the family is joined by an arc in K , and thus the vertices span a simplex because K is flag. But a simplex cannot contain a closed geodesic. \square

We are now ready to prove the main result of this section.

Theorem 2.2.12. Let K be a locally finite cubical complex. If the link at every vertex is a flag simplicial complex, then K is locally CAT(0).

Proof. This follows immediately by reassembling Theorem 2.2.7 and Lemma 2.2.11. \square

3 Triangulations of \mathbb{S}^3

The main purpose of this section is to exhibit triangulations of \mathbb{S}^3 , having some properties defined below, which will turn useful in the following sections.

3.1 Preliminary definitions

We begin by defining some properties of the simplicial complexes of our interest.

Definition 3.1.1 (Subcomplex spanned by a set of vertices). *Let V be a subset of the set of vertices of a simplicial complex K . The subcomplex spanned by V is the set of all simplices in K that have vertices in V .*

Definition 3.1.2 (Full subcomplex). *A subcomplex of a flag simplicial complex is said to be full if it coincides with the subcomplex spanned by its vertices.*

Definition 3.1.3 (Square). *A square is a subcomplex of a simplicial complex K , with 4 vertices v_0, v_1, v_2 and v_3 , 4 arcs between v_i and $v_{i+1 \bmod 4}$, for $i = 0, 1, 2, 3$ and such that there is no arc between v_0 and v_2 or between v_1 and v_3 in K .*

Note that a square is a full subcomplex, isomorphic to a flag triangulation of \mathbb{S}^1 .

The main result of this section is the following

Theorem 3.1.4. *Let L be any knot in \mathbb{S}^3 . There is a flag triangulation of \mathbb{S}^3 which contains only one square, such that the isotopy type of its immersion in \mathbb{S}^3 is precisely L .*

3.2 The 600-cell

The 600-cell, which we will call for brevity K_{600} , is a regular triangulation of \mathbb{S}^3 . This means that, fixed any pair of tetrahedra, including the same taken twice, and any bijection between the vertices of the first and the second, there is a unique simplicial automorphism of K_{600} extending this bijection. The simplicial link at every vertex is an icosahedron. Thus K_{600} is flag: a complete graph in the 1-skeleton lies entirely in the closed star of a vertex, which is a cone over an icosahedron, which is flag.

Furthermore, the 600-cell has no squares: two adjacent arcs share the vertex of a cone on an icosahedron: it is easy to see that the other two




Figure 2: Subdivision of a part of the boundary of K_{543} .

extremes are either joined by an arc of the icosahedron, or cannot be joined both to another point.

Take a tetrahedron σ in K_{600} and consider all the tetrahedra in it that do not intersect σ . The boundary of this 3-manifold, topologically a 3-ball, is like that of a tetrahedron in which every face is subdivided according to Figure 2. Call this complex K_{543} ; being 543 the number of tetrahedra in it.

For any simplicial complex X realising a triangulation of a 3-manifold, possibly with boundary, we can subdivide each tetrahedron in it by substituting it with K_{543} coherently with the description of the boundary we just made. Call X^* the obtained triangulation.

Lemma 3.2.1. X^* is a flag simplicial complex without squares.

Proof. Subdivide the vertices in two families: the ones that are internal to the old tetrahedra of X and the ones on their boundaries.

Take a complete graph in the 1-skeleton of X^* . We may suppose none of its vertices is internal, otherwise it would be the vertex of a cone on the icosahedron, which is flag.

If one of the vertices lies on a facet of an old tetrahedron, it with the adjacent boundary vertices form a cone on a pentagon, which is flag, so we may suppose all vertices of the complete subgraph lie on old edges. The midpoint of an edge with its adjacent edge vertices form a cone on two points, the extremes of the edge, and this is flag too. So the only remaining case is that of subgraphs made only by the old vertices, but they are isolated in the new triangulation.

Suppose now there is a square in X^* . If two of its vertices lie in the same old tetrahedron, we are done because K_{543} has no squares. Observing the triangulation induced on the facets of the old tetrahedra, we may as well

conclude there cannot be squares there. So the only remaining possibility is that a hypothetical square has a vertex inside one of the old tetrahedrons, the adjacent ones on its boundary and the fourth vertex somewhere else.

Two vertices on the boundary of K_{543} that are joined to the same internal vertex are either joined themselves, which concludes, or are the two vertices of a pair of triangles in the boundary having a side in common not lying on that side. In the latter case the two triangles have to be, furthermore, in the simplicial link of the same internal vertex, the first in the square. We see, observing the complement of K_{543} in K_{600} that this is possible only if the two vertices lie on different faces of an original tetrahedron. The original triangulation was simplicial, so it follows that in X^* the two vertices cannot be joined to the same vertex. \square

We call the subdivision provided by the previous Lemma *flag-no-square* (FNS) subdivision of a simplicial complex.

3.3 The triangulation with a knot

We now want to prove the Theorem 3.1.4. We begin by constructing a triangulation of a neighbourhood of a square. It will be further subdivided in order to guarantee the flag property.

In \mathbb{R}^3 take the region $0 \leq x \leq y \leq 1$, $0 \leq z \leq 1$; this is a triangular prism. The sets $z \leq x \leq y$, $x \leq z \leq y$, $x \leq y \leq z$ provide a triangulation of this prism. Take three of them and cyclically identify the face $x = y$ of the previous one with the face $x = 0$ of the following one, keeping the orientation on the z -axis: this way we have a triangulation of a larger triangular prism, which we will call block, with an arc joining the centre of the top face (the union of the faces $z = 1$ of the original prisms) with the centre of the bottom face (union of the $z = 0$ faces).

Take 4 blocks and cyclically identify the top face of one with the bottom face of the following one. This way we get a triangulation of a solid torus, having a square in its core. In this triangulations there are 36 tetrahedrons, 24 of which are the join of a triangle on the boundary with an internal vertex, and 12 are the join of an arc of the boundary with an arc in the core. Note that in this triangulation every internal vertex is adjacent only to vertices coming from its original block, except from the internal vertex in an adjacent block.

Take a knot in \mathbb{S}^3 (the argument works as well if we take a link). Denote by \mathbb{D}^2 the disk of radius 1 in \mathbb{R}^2 , and with $\widetilde{\mathbb{D}}$ the disk with radius $\frac{1}{2}$. Take a regular neighbourhood of it and identify it with $\mathbb{D}^2 \times \mathbb{S}^1$. Triangulate $\widetilde{\mathbb{D}}^2 \times \mathbb{S}^1$

the way we have just described and call it T . Then triangulate simplicially $\mathbb{S}^3 \setminus \text{Int } \mathbb{D}^2 \times \mathbb{S}^1$: we can always do this with a 3-manifold.

It remains to triangulate the remaining $\mathbb{S}^1 \times \mathbb{S}^1 \times [0, 1]$ region in a coherent way to the already present triangulation on the boundary. The boundary are two parallel triangulated tori; the triangulations have a common subdivision. Triangulate using this subdivision the layer $\mathbb{S}^1 \times \mathbb{S}^1 \times \{\frac{1}{2}\}$. It remains to connect the triangulation of this layer with the triangulations on the toric boundary components.

The problem we want to solve is the following. Consider a triangulation S of a surface F and a subdivision S' of the triangulation. Then it is possible to triangulate $F \times [0, 1]$ such that the triangulation on $F \times \{0\}$ is S and on $F \times \{1\}$ is S' . Fix an auxiliary total order $<$ on the vertices of S , take σ a triangle of S and suppose the vertices of σ are $v_0 < v_1 < v_2$. Take first the cone on the subdivision of σ in S' with vertex $v_0 \times \{0\}$, then the cone with vertex $v_1 \times \{0\}$ on the subdivision induced on the join of $v_0 \times \{0\}$ and the arc $v_1 v_2 \times \{1\}$ by the previous coning. In $\sigma \times [0, 1]$ remains thus only a tetrahedron with vertices $v_0 \times \{0\}$, $v_1 \times \{0\}$, $v_2 \times \{0\}$ and $v_2 \times \{1\}$. Repeat this for every triangle in S . The triangulations induced on prisms given by the triangles agree on common faces thanks to the total order we gave to the vertices.

Call X the triangulation of $\mathbb{S}^3 \setminus T$ we have just obtained. Take the FNS subdivision of X obtaining a triangulation X^* . This leads to a subdivision of the boundary of T : every triangle in it is subdivided the way we saw when talking about FNS subdivisions. Then, to obtain a subdivision T^* of T too, we proceed as follows:

- If a tetrahedron of T is a join of a triangle on the boundary with an internal vertex, we subdivide it like the cone on this triangle;
- If a tetrahedron of T is a join of an arc on the boundary with an arc in the core, the arc on the boundary is divided in two, which leads to a natural division in two of the tetrahedron.

This subdivision does not introduce new vertices in the internal part of T . In the end, we will have a triangulation of \mathbb{S}^3 with a square describing a knot in a prescribed isotopy class. We now prove the following

Lemma 3.3.1. The above described triangulation of \mathbb{S}^3 is flag and has as unique square the one describing the knot.

Proof. Take a complete graph in the 1-skeleton. We can suppose it has at least one vertex among the internal vertices of T , otherwise the whole thesis descends directly from Lemma 3.2.1.

Note that an internal vertex is still adjacent only to vertices in its block and to the internal vertex in an adjacent block. So a complete graph on 3 vertices or more lies in the same block: we can see, observing the boundary, that the triangulation obtained on a block is flag.

Suppose there are other squares in the triangulation. Like before, one of the vertices must be an internal vertex of T . If there are two or more of them, the square must lie into T , which is not possible unless it is the already known square. Then we see, using the structure of the FNS subdivision, that the other extremes of two arcs exiting from an internal vertex are either adjacent, or cannot be adjacent both to the same other vertex. \square

4 The manifold

We are now moving towards the construction of the manifold which will allow us to prove the main theorem.

4.1 The Davis complex

Let K be a simplicial complex and $v(K)$ its set of vertices; let $I = [0, 1]$. Consider the euclidean cube $I^{v(K)}$. We have already seen that its faces are a product of I on some coordinates and of one endpoint of I on some other coordinates. Let us call the *type* of the face the subset of $v(K)$ corresponding to the I factors in it.

Definition 4.1.1 (Davis complex of a simplicial complex). *Let K and $v(K)$ be like above. The Davis Complex associated to K , which we will indicate by P_K , is a cubical complex, subcomplex of $I^{v(K)}$, consisting of all those faces whose type spans a simplex in K .*

The next fact follows immediately from the construction.

Lemma 4.1.2. Let K be a simplicial complex, P_K the associated Davis complex, and fix a vertex v in it. The map associating to a simplex σ in K the link at v in the unique cube in P_K having for type the vertices of σ and passing by v is a simplicial isomorphism between K and $\text{lk}_{P_K} v$.

Consider for simplicity the vertex $w = \{0\}^{v(K)}$ of P_K . The previous lemma tells us that the subset of P_K consisting of points with all coordinates strictly less than 1 is homeomorphic to an open cone on K with vertex w . So, if the geometric realization of K is \mathbb{S}^{n-1} , then P_K is a manifold: the said set is homeomorphic to an open ball, and all analogously defined sets centered in various vertices cover P_K . Furthermore, if K is a PL triangulation of \mathbb{S}^{n-1} , there is a diffeomorphism of it with the differentiable \mathbb{S}^{n-1} that makes the embeddings of simplices smooth. So the previous cone structure can be made a diffeomorphism with an open ball of \mathbb{R}^n making thus P_K a smooth manifold.

Let us now describe the fundamental group of P_K . Begin with noting that $(\mathbb{Z}/2\mathbb{Z})^{v(K)}$ acts by reflections on P_K ; here a reflection r_i is a map acting like $t \mapsto 1 - t$ on the coordinate i and leaving unaltered the others. Let us call W_K the group consisting of all possible liftings of actions of $(\mathbb{Z}/2\mathbb{Z})^{v(K)}$ to the universal cover of P_K , and let p be the covering map.

The homomorphism $p \circ: W_K \rightarrow (\mathbb{Z}/2\mathbb{Z})^{v(K)}$ is obviously surjective. The kernel consists by definition of all cover automorphisms, which is precisely

$\pi_1 P_K$. So we have a short exact sequence

$$1 \rightarrow \pi_1 P_K \rightarrow W_K \rightarrow (\mathbb{Z}/2\mathbb{Z})^{v(K)} \rightarrow 1,$$

in particular $\pi_1 P_K$ is a finite index subgroup of W_K . It remains only to understand which group it is. Let us first explore the following definition.

Definition 4.1.3 (Right angled Coxeter group). *Let K be a simplicial complex. The right angled Coxeter group associated to K , denoted by Γ_K is a group with set of generators $\{x_i\}_i$ in bijection with the set of vertices $\{v_i\}_i$ of K , each having order two, and a relation $x_i x_j = x_j x_i$ if and only if v_i and v_j are joined by an arc in K .*

In fact Γ_K depends only on the 1-skeleton of K .

Now we can state and prove the following

Lemma 4.1.4. The group W_K is precisely the right angled Coxeter group Γ_K associated to K .

Proof. Fix a vertex w in P_K and one preimage \tilde{w} in the universal cover \widetilde{P}_K to use as basepoints. Call the i -th arc at w the unique arc having the factor I at the i -th coordinate and passing by w : it is stabilized by r_i . Denote by x_i the lift of r_i stabilizing the lift of the i -th arc passing by \tilde{w} . We claim that the x_i generate W_K . In fact, \widetilde{P}_K has a cubical structure inherited by that of P_K . We prove that the group G generated by x_i acts transitively on the set of vertices, which suffices, along with the fact that G projects surjectively on $(\mathbb{Z}/2\mathbb{Z})^{v(K)}$ in the above mentioned exact sequence.

We prove the claim by induction on the combinatorial distance of a vertex of \widetilde{P}_K from \tilde{w} , i.e the minimum number of arcs in a path from \tilde{w} to the final vertex. Clearly, x_i 's allow us to bring \tilde{w} in any vertex at distance 1 from it.

Suppose that G can bring \tilde{w} in any vertex at distance n from it. Let u be such a vertex and g an element of G such that $g \cdot \tilde{w} = u$. But then using $g \circ x_i$ we can get to any vertex near u departing from \tilde{w} . Varying u , we prove the thesis for all vertices at combinatorial distance $n + 1$ from \tilde{w} .

Let us now explore the relations. Being x_i^2 a lift of the identity that stabilizes an arc by \tilde{w} , it is the identity itself. With a similar argument, it is clear that an arc between v_i and v_j in K , which means a 2-cube on coordinates i and j in P_K implies that $x_i x_j = x_j x_i$.

There are no other relations. Suppose to the contrary there is one. The proof we made of the fact that the action of W_K is transitive on vertices tells us in fact that that the 1-skeleton of \widetilde{P}_K is the Cayley graph

$$\mathcal{C} \left(W_K, \{x_i\}_{i \in v(K)} \right).$$

So, if we take a string representing the identity element in W_K , it represents a closed path in the Cayley graph. But a closed path in the simply connected \widetilde{P}_K bounds a disc made up of 2-cubes, which represent already known relations. \square

4.2 The desired manifold

We now begin to restrict our point of view to particular complexes K . The non positive curvature enters now in scene.

Lemma 4.2.1. If the complex K is flag, P_K is locally CAT (0), and thus \widetilde{P}_K is CAT (0).

Proof. This follows immediately from Theorem 2.2.12 and Cartan-Hadamard for CAT (0) spaces [BH99, Theorem II.4.1]. \square

Under certain hypotheses, immersions of simplicial complexes imply immersions of relative Davis complexes.

Lemma 4.2.2. Let K be a flag simplicial complex and L a full subcomplex. The immersion of L in K naturally induces a locally isometric embedding of P_L in P_K .

For a proof, see [Dav08, Appendix I.6].

The previous fact has consequences at the level of fundamental groups.

Lemma 4.2.3. Let X be a locally CAT (0) compact space, Y a compact geodesic space and $f: Y \rightarrow X$ a locally isometric immersion. Then the induced map at the level of fundamental groups is injective.

Proof. A local isometry takes local geodesics to local geodesics. Take a non trivial element α in $\pi_1 Y$. There is a representative in its equivalence class which is a local geodesic. In fact, take a sequence of length minimizing curves in the class of α parametrized proportionally to arclength: by Ascoli-Arzelà Theorem the limit exists, being non trivial it is non constant and, being the space geodesic, it is a local geodesic γ .

Consider a lift of $f \circ \gamma$ at the universal cover \widetilde{X} : it is a local geodesic and hence a non constant geodesic: this means that $\pi_1 f(\alpha)$ is non trivial. \square

Remember that in Theorem 3.1.4 we obtained a triangulation of \mathbb{S}^3 containing a unique square representing a prefixed class of isotopy of knots. In the 3 dimensional setting, the triangulation is automatically PL. We can use it to provide the manifold which will solve our problem.

Theorem 4.2.4. Let K be the triangulation provided by Theorem 3.1.4 when we take a non trivial knot. Then P_K is a smooth 4-manifold such that the natural metric complex distance function is locally CAT(0), but not homeomorphic to any smooth manifold of non positive curvature.

We continue with proving some global properties of such \widetilde{P}_K . From now on, this notation will refer to our manifold and \widetilde{P}_K will refer to its universal cover.

Lemma 4.2.5. The universal cover of the manifold of the above theorem is PL-equivalent to the standard \mathbb{R}^4 , and the boundary at infinity of the universal cover is homeomorphic to \mathbb{S}^3 .

Proof. This follows from Stone's Theorem, proven in [Sto76], which is the PL version of the classic smooth Cartan Hadamard Theorem, that affirms that the universal cover of a complete non positively curved smooth n -manifold is diffeomorphic to \mathbb{R}^n . \square

It follows immediately that \widetilde{P}_K is in fact diffeomorphic to the standard \mathbb{R}^4 .

Observe now more closely the square in K . The Davis complex of a square is the standard 2-dimensional torus. By Lemma 4.2.3 there is a \mathbb{Z}^2 in $\pi_1 P_K$, which is a finite index subgroup of the Coxeter group Γ_K . So Γ_K has a subgroup isomorphic to \mathbb{Z}^2 which can be described explicitly in terms of the natural generators. Let v_1, v_2, v_3 and v_4 be the vertices of the square, in this order, and x_i the respective generators of Γ_K . A \mathbb{Z}^2 subgroup in $\pi_1 P_K$ is generated by $x_1 x_3 x_1 x_3$ and $x_2 x_4 x_2 x_4$; it is an index 4 subgroup of a \mathbb{Z}^2 in Γ_K generated by $x_1 x_3$ and $x_2 x_4$.

Lemma 4.2.6. Γ_K is relatively hyperbolic with respect to the subgroup generated by $x_1 x_3$ and $x_2 x_4$, which is isomorphic to \mathbb{Z}^2 .

Proof. This follows from [Cap07]. The main observation in proving the Lemma is that there are no squares in the triangulation that share one side or more. \square

Being the fundamental group of P_K a finite index subgroup of Γ_K , it is quasi isometric to it, thus it is relatively hyperbolic with respect to the abelian subgroup of rank 2 given by the intersection of the subgroup of Γ_K we already told about with $\pi_1 P_K$.

Now we can pass to prove the Theorem 4.2.4. Remember that in a complete smooth manifold M of non positive curvature the universal cover \widetilde{M} is always diffeomorphic to the tangent space at any point via the exponential

map. It follows then immediately that $\partial\widetilde{M}$ is homeomorphic to the unit sphere in any tangent space via the map that associates to a tangent vector the geodesic ray departing from the point with that tangent. If such a manifold has a totally geodesically immersed submanifold N , its universal cover \widetilde{N} totally geodesically embeds in \widetilde{M} . If we take now p in \widetilde{N} , the unit sphere $\mathbb{S}_p\widetilde{M}$ in $T_p\widetilde{M}$ has embedded an unknotted sphere of the appropriate dimension given by its intersection with $T_p\widetilde{N}$, because the immersion of \widetilde{N} is obviously locally flat. At the boundary at infinity, the fact that \widetilde{N} is totally geodesic says that we have in fact a homeomorphism of pairs

$$\left(\mathbb{S}_p\widetilde{M}, \mathbb{S}_p\widetilde{N}\right) \cong \left(\partial\widetilde{M}, \partial\widetilde{N}\right)$$

given by the same map described before. So, in the latter pair, the second sphere is unknotted in the first, too.

Proof. Suppose that there is a smooth 4-manifold M of non positive curvature homeomorphic to P_K . The homeomorphism induces an isomorphism on fundamental groups, which we will call G from now on, and hence a G -equivariant quasi isometry between universal covers \widetilde{P}_K and \widetilde{M} . For the relative hyperbolicity of G , we have that the boundaries at infinity of \widetilde{P}_K and \widetilde{M} are homeomorphic by a G -equivariant homeomorphism, as we have seen in Lemma 1.4.9.

By the smooth Flat Torus Theorem there is an \mathbb{E}^2 totally geodesically embedded in \widetilde{M} . By the reasoning we made, $\partial\mathbb{E}^2$ is a trivial knot in the 3-sphere $\partial\widetilde{M}$.

In P_K , the torus is not locally flat by construction. The homeomorphism between the boundaries at infinity of \widetilde{P}_K and \widetilde{M} would take the boundary of the corresponding flat F to an unknot. If we could manage to prove that ∂F is not unknotted in $\partial\widetilde{P}_K$ we would have finished. In fact, we will prove that the complement has a non abelian fundamental group, whilst the complement of a (normal neighbourhood of an) unknot in \mathbb{S}^3 is a solid torus and thus has fundamental group \mathbb{Z} .

Remember that $\partial\widetilde{P}_K$ is defined like an inverse limit and that, for a small positive r and for a vertex v of \widetilde{P}_K the ball $\overline{B}(v, r)$ is isometric to a neighbourhood of the vertex in the cone $C_0 \text{lk}_{\widetilde{P}_K} v$. There is a continuous projection $\rho_r: \partial\widetilde{P}_K \rightarrow \partial\overline{B}(v, r)$. We will use the fact that this projection is, in fact, a proper homotopy equivalence [FL04].

Let U be the complement of the knot in $\text{lk}_{\widetilde{P}_K} v$, identified with the boundary of the ball of radius r like above. Then $U_\infty = \rho_r^{-1}(U)$ has the same fundamental group as the complement of a non trivial knot, which is, by

[Pap57], non abelian. We want to prove that the inclusion $U_\infty \hookrightarrow \partial\widetilde{P}_K \setminus \partial F$ is injective at the level of fundamental group, which suffices.

This latter fact is proven in [DJL12, 5, Fact 4] and relies on the observation that if we have a knotted solid torus in \mathbb{S}^3 and a simple closed curve in the torus homotopic to its core, maybe further knotted, then the immersion of the complement of the torus in the complement of the curve is injective at the level of fundamental groups. \square

References

- [Ale51] Aleksandr Danilovich Aleksandrov. A theorem on triangles in a metric space and some of its applications. *Trudy Matematicheskogo Instituta Imeni Steklova*, 38:5–23, 1951.
- [BH99] Martin Robert Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 1999.
- [Cap07] Pierre-Emmanuel Caprace. Buildings with isolated subspaces and relatively hyperbolic coxeter groups. *arXiv preprint math/0703799*, 2007.
- [CK90] Chen Chung Chang and Howard Jerome Keisler. *Model Theory*. Studies in Logic and the Foundations of Mathematics. Elsevier Science, 1990.
- [Dav08] Michael Davis. *The geometry and topology of Coxeter groups*, volume 32. Princeton University Press, 2008.
- [DJL12] W. Michael Davis, Tadeusz Januszkiewicz, and Jean-François Lafont. 4-dimensional locally $\text{cat}(0)$ -manifolds with no riemannian smoothings. *Duke Mathematical Journal*, 161(1):1–28, 01 2012.
- [Dru02] Cornelia Druțu. Quasi-isometry invariants and asymptotic cones. *International Journal of Algebra and Computation*, 12(01n02):99–135, 2002.
- [FL04] F. Thomas Farrell and Jean-François Lafont. Involutions of negatively curved groups with wild boundary behavior. *arXiv preprint math/0405261*, 2004.
- [Gro87] Mikhael Gromov. *Hyperbolic groups*. Springer, 1987.
- [GW71] Detlef Gromoll and Joseph A. Wolf. Some relations between the metric structure and the algebraic structure of the fundamental group in manifolds of nonpositive curvature. *Bulletin of the American Mathematical Society*, 77(4), July 1971.
- [HK05] G. Christopher Hruska and Bruce Kleiner. Hadamard spaces with isolated flats. *Geometry & Topology*, 9:1501–1538, 2005.

- [LY72] H. Blaine Lawson and Shing-Tung Yau. Compact manifolds of non-positive curvature. *Journal of Differential Geometry*, 7:211–228, 1972.
- [Pap57] Christos 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.
- [Sto76] David A. Stone. Geodesics in piecewise linear manifolds. *Transactions of the American Mathematical Society*, 215:1–44, 1976.