Geometric structures on negatively curved groups and their subgroups

Samuel Brown

A thesis submitted in partial fulfilment for the award of the degree of

Doctor of Philosophy

Department of Mathematics

University College London

Submitted August 2016

Revised January 2017

Declaration

I, Samuel Brown, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

....

Abstract

In this thesis, we investigate two explicit families of geometric structures that occur on hyperbolic groups. After recalling some introductory material, we begin by giving an overview of the theory of *special cube complexes*, with a particular focus on properties of subgroups of hyperbolic special groups. We then describe an explicit algorithm, based on Stallings' notion of *folding* for graphs, to construct a local isometry between cube complexes that represents the inclusion of a subgroup $H \subset G$, and show that this terminates if and only if the subgroup is quasiconvex. This provides a potential method by which quasiconvexity for various subgroups could be verified.

In the second part of the thesis, we investigate another family of geometric structures: negatively curved simplicial complexes. We show that groups satisfying a "uniform" $C'(1/6)$ small cancellation condition have such a structure, and then move on to prove a gluing theorem (with cyclic edge groups) for these complexes. Using this theorem, we extend the family of groups known to be $\mathrm{CAT}(-1)$ to include hyperbolic limit groups, hyperbolic graphs of free groups with cyclic edge groups, and more generally hyperbolic groups whose JSJ components are 2-dimensionally $\mathrm{CAT}(-1)$.

Primary Supervisor: Dr Henry Wilton

Acknowledgements

Studying for a PhD has been an immensely rewarding experience, but never an easy one. There are many people without whose support I could not have completed this thesis, and I am grateful for this opportunity to express my thanks.

Most importantly, I would like to thank my supervisor, Dr Henry Wilton. It has been an enormous pleasure to spend the last four years working with Henry, and his mathematical guidance, moral support, and boundless optimism have kept me going throughout the last four years. I could not have hoped for a better experience of supervision.

More broadly, I would like to express my gratitude to the wider geometric group theory community for welcoming me into their fold at many enjoyable conferences and seminars during my PhD. I am particularly grateful to those at the mathematics departments at Oxford and Cambridge, for their warm welcome to me during many visits over the years.

My experience as a PhD student has been greatly enhanced by the work of both the academic and the administrative staff in the UCL Maths Department, and so I would like to express my gratitude to them. I would also like to take the opportunity to thank the EPSRC for funding my research, and all those who have inspired and nurtured my passion for mathematics throughout my education.

On a less academic note, my life as a PhD student would have been far poorer without the friendship and entertainment offered by my colleagues in the KLB postgraduate office at UCL. They have helped keep me sane, and provided regular welcome distractions from my research. I wish them all the very best for the remainder of their PhD's, and into the future. I would also like to thank all my other friends, both in London and further afield, for their support over the last four years.

Lastly, I would like to thank my family for continuing to inspire me; and my partner Louise, whose unending love and friendship mean more to me than I can adequately express.

Contents

Chapter 1

Introduction

Indisputably one of the most outstanding achievements of twentieth or twenty-first century mathematics has been Bill Thurston's geometrization program for three-dimensional manifolds. First proposed by Thurston in the 1980s, the geometrization conjecture states that every 3-manifold can be decomposed by cutting along spheres and tori into pieces which admit one of precisely eight different geometries. Perhaps the richest and most interesting of these eight options is *hyperbolic* geometry, which describes manifolds whose universal cover is hyperbolic 3-space $\mathbb{H}^3$. It turns out that any compact 3-manifold which is *aspherical* (contains no essential spheres) and *atoroidal* (contains no essential tori) is hyperbolic. This is known as Thurston's *hyperbolization theorem*, proved by Thurston in the special case of Haken manifolds (see [Thu86] and subsequent papers) and in general by Perelman as part of his resolution of the geometrization conjecture (see [Per02]). What the hyperbolization theorem hints at is some sort of underlying "genericity" of hyperbolic geometry; to ensure that a 3-manifold is hyperbolic, it suffices simply to rule out certain "obvious" ways in which it could fail to be so. The deep underlying fact is that this prevalence of negative curvature extends well beyond the realm of 3-dimensional geometry.

Motivated by the geometry of hyperbolic manifolds, Mikhail Gromov defined in [Gro87] the notion of a *word–hyperbolic* (or simply *hyperbolic*) group. A hyperbolic group shares some characteristics with the fundamental group of a hyperbolic manifold; in particular, it possesses an intrinsic notion of negative curvature. However the class of hyperbolic groups is much more general than the class of fundamental groups of hyperbolic 3-manifolds, and they are central objects of study in geometric group theory.

As with hyperbolic manifolds, hyperbolic groups are known to have a certain genericity; Gromov himself observed that (roughly speaking) in group presentations with a fixed generating set and q relations, the probability that the corresponding group is hyperbolic goes to 1 as the lengths of the relations go to infinity. However, a direct analogue of the hyperbolization theorem for hyperbolic groups remains elusive. One possible such analogue is proposed in the following question.

Question 1.0.1. Let Γ be a group which is type F_∞ and contains no Baumslag–Solitar subgroups. Is Γ hyperbolic?

It is worth unpicking this question to understand exactly how it relates to the hyperbolization theorem. A group is said to have *type* F_∞ if it has a classifying space with finitely many cells in each dimension; in particular, it is the fundamental group of an aspherical CW complex (with finite n-skeleta), and so this hypothesis is analogous to the "aspherical" hypothesis in the hyperbolization theorem. A *Baumslag–Solitar* group is a group with presentation $\langle a, b \mid a^{-1} b^m a = b^n \rangle$; these groups are all known not to be hyperbolic, and ruling them out is analogous to the "atoroidal" hypothesis in the hyperbolization theorem. An affirmative answer to Question 1.0.1 would therefore be a direct analogue of the hyperbolization theorem for 3-manifolds, but in the much broader context of all infinite discrete groups. This would provide perhaps the best vindication yet of Gromov's original definition of a hyperbolic group, and indeed would provide an alternative definition, since the converse holds for all hyperbolic groups.

On the other hand, a counterexample to Question 1.0.1 would be of great interest; it would be a group which fails to be hyperbolic for a completely new and unforeseen reason. Moreover, there are known ways to construct a possible counterexample, given the existence of other pathological hyperbolic groups. Suppose, for example, that there exists a hyperbolic group G, with a subgroup H that is *malnormal* (it intersects all its distinct conjugates trivially) and *distorted* (roughly speaking, the corresponding subgraph of the Cayley graph of G fails to be convex). Then, if H is of type F_∞, the *double* of G along H would be a counterexample to Question 1.0.1 (see Theorem 4.4.3 for a precise statement). This motivates the following question.

Question 1.0.2. Does there exist a hyperbolic group G with a finitely generated subgroup H which is malnormal and distorted in G?

We shall discuss Question 1.0.2 further in Section 4.4, but it provides a good justification for an effort to understand not only the geometry of negatively curved groups, but also the geometry of their *subgroups*. Question 1.0.2 is known to have a negative answer only for a few special cases—for example, finitely generated subgroups of free groups are all *quasiconvex* (that is, undistorted), and so the question is trivially answered in this case. Subgroups of free groups are, in fact, very well understood from a geometrical perspective, since they correspond to *immersions* (locally injective homomorphisms) between graphs, and it is reasonable to ask if this straightforward characterisation extends to a more general class of spaces. In recent years, a class of spaces has emerged that could provide the answer. These are the *special cube complexes*, defined by Haglund and Wise [HW08], and a good candidate for the higher dimensional generalisation of graphs. A special cube complex is a piecewise Euclidean cube complex whose *hyperplanes* avoid various pathologies (see Chapter 3 for more details), and a group is called *special* if it is the fundamental group of a special cube complex. Special, and virtually special, groups have a number of interesting properties, particularly from the point of view of *subgroup separability*, which is a property inherited from graphs. Their most prominent recent application has been in the proof of the Virtual Haken and Virtual Fibering theorems for 3-manifolds. After geometrization, these were two of the biggest remaining open questions in 3-dimensional topology, and both were recently resolved by Agol [Ago13] building on work of Wise [Wis12a, Wis12b], Kahn–Markovic [KM12] and others (see Section 3.8.2 for more details).

The structure theory of special cube complexes, and the variety of tools available for working with them, was central to this work, and much recent attention has been given to finding special cube complexes associated to a wide variety of groups. An answer to Question 1.0.2 in the special cube complex case would therefore be a very reasonable objective, and our work in Chapter 4 has this in mind.

As mentioned above, a positive solution to Question 1.0.1 would provide an additional justification for the role of hyperbolicity as the predominant notion of negative curvature in group theory. However, hyperbolicity is not the only option. It is a *coarse* notion—it is not sensitive to the geometry of the group on small scales—and one common theme within geometric group theory is to understand the relationship between this and more local notions of negative curvature. One such notion is the CAT(k) condition. A space is called CAT(k) if geodesic triangles in the space are "thinner" than those in the simply connected surface of constant sectional curvature k (for example, hyperbolic space $\mathbb{H}^2$ for $k = -1$); and a group is called CAT(k) if it possesses a geometric action (that is, a properly discontinuous, cocompact action by isometries) on a CAT(k) space. For example, the universal cover of a special cube complex is a CAT(0) space (indeed, a CAT(0) cube complex), and so virtually (compact) special groups are CAT(0). One can in fact say more in this setting. If the CAT(0) cube complex does not contain an isometrically embedded copy of two-dimensional Euclidean space, then the group is actually hyperbolic (see Section 2.4.3). This is another example of the genericity of hyperbolic groups, this time in the context of groups acting on cube complexes. However, the following question is still unanswered.

Question 1.0.3. Does there exist a hyperbolic group which is not CAT(0), or not CAT(-1)?

The opposite implications are well understood: every CAT(-1) group, and every CAT(0) group where the space in question has no flat planes, is hyperbolic. That Question 1.0.3 remains open is somewhat surprising—it means that, as far as is known, the "coarseness" which is central to the definition of a hyperbolic group may be superfluous; an equivalent definition may be that it is a group exhibiting a geometric action on a CAT(-1) space, or a CAT(0) space without flat planes. Again, either a positive or negative answer to this question in general would be of great interest.

The three questions above lie at the very heart of geometric group theory, and could form a research agenda for many decades. Our contribution in this thesis is, firstly, to suggest a method for approaching Question 1.0.2 in the setting of special cube complexes, and secondly, to partially answer Question 1.0.3 by extending the class of hyperbolic groups known to be CAT(-1).

The thesis is organised as follows. Chapter 2 summarises the introductory material which forms the mathematical background to the remaining chapters. We begin by recalling the notions of a graph of groups and graph of spaces, then discuss various notions of negative curvature for groups and metric spaces, particularly the CAT(k) condition and δ-hyperbolicity. Of central importance are two families of metric complex—non-positively curved cube complexes, and negatively curved simplicial complexes—so we give an account of the crucial properties of each family, with a focus on the interplay between group-theoretic and geometric concepts.

In Chapter 3, we focus on cube complexes. In particular, we outline the theory of *special* cube complexes (as originally described by Haglund and Wise [HW08]), and in Chapter 4 we describe a way to generalise the *folding* algorithm of Stallings to (a sub-class of) special cube complexes, exploiting the fact that they may be given a *hierarchy*; that is, decomposed as a graph of spaces with vertex spaces of lower complexity than the original space. Stallings' folding techniques provide a strong understanding of the geometry of subgroups of free groups, and our generalisation could provide similar insight into the geometry of subgroups of fundamental groups of special cube complexes. In particular, we give a proof that the folding algorithm terminates if and only if the subgroup is quasiconvex. Hence, analysing the geometry of the situation where the algorithm does *not* terminate gives a potential approach to answering Question 1.0.2 in the case of special cube complexes. We suggest some possible ways to do this in Section 4.4.

In Chapters 5 and 6, we shift our focus to Question 1.0.3. To this end, we leave cube complexes behind and look instead at 2-dimensional negatively curved simplicial complexes. As a warm-up, in Chapter 5 we exploit some of the flexibility of hyperbolic geometry to construct a negatively curved metric on the presentation complex of groups satisfying a *uniform* $C'(1/6)$ small cancellation condition, thus proving that such groups are CAT(-1) (a result first proved by Gromov). Finally in Chapter 6, we prove a combination theorem for negatively curved 2-complexes. This allows us once again to exploit hierarchical decompositions for certain families of groups, thus extending the class of groups known to be CAT(-1) to include hyperbolic *limit groups*, hyperbolic graphs of free groups with cyclic edge groups, and more generally hyperbolic groups whose JSJ components are 2-dimensionally CAT(-1). A substantial part of Chapter 6 formed the basis for a paper which has appeared in the Journal of the London Mathematical Society [Bro16].

We assume that the reader is familiar with elementary geometry (including hyperbolic geometry), topology and group theory.

Chapter 2

Preliminaries

In this chapter, we summarize the mathematical theory which gives the context for the remaining chapters. We begin with a brief account of graphs of groups and spaces, before covering in more detail some of the theory of CAT(k) spaces and hyperbolic groups. The material in Section 2.1 is covered more thoroughly in [Ser03] and [Hat02], and the remaining sections are covered comprehensively in the first two parts of [BH99].

2.1 Graphs of groups and graphs of spaces

Throughout this thesis, we will be interested in obtaining new groups and spaces by gluing together others. The best formalism for this is the algebraic notion of a *graph of groups*, and the analogous topological notion of a *graph of spaces*. We recall the essential definitions.

2.1.1 Graphs

We assume the reader is familiar with the topological notion of a graph. For convenience, we also recall Serre's definition of a graph, in which each edge e is replaced by a pair of edges corresponding to the two possible orientations of e. This simplifies notation later, as we can bypass the technicality of edges incident at a certain vertex having different orientations.

Definition 2.1.1. A *graph* consists of a set V of *vertices*, a set E of *edges*, and maps $\iota \colon E \to V$, $\tau \colon E \to V$, $\bar{\ } \colon E \to E$ satisfying $\bar{e} \neq e$, $\bar{\bar{e}} = e$ and $\tau(e) = \iota(\bar{e})$.

We may recover a topological graph from the above by taking a vertex for each $v \in V$, and then for each unordered pair $(e, \bar{e})$, we take a 1-cell whose ends are glued to $\iota(e)$ and $\tau(e)$. An *orientation* is a fixed choice of either e or $\bar{e}$ from each pair. We will sometimes write an edge as $e = uv$, meaning $\iota(e) = u$, $\tau(e) = v$.

An *immersion* of (topological) graphs is a combinatorial map (i.e. graph homomorphism) which is locally injective. Immersions of graphs are π_1-injective [Sta83]. A graph is *simplicial* if it has no loops or double edges, and a *core graph* for a graph is a subgraph with no vertices of degree 1, such that the inclusion map is a π_1-isomorphism.

2.1.2 Graphs of groups

A *graph of groups* B consists of a graph Γ_B together with the following: for each vertex v of Γ_B, a *vertex group* B_v; and for each edge e of Γ_B, an *edge group* B_e (where $B_{\bar{e}} = B_e$) and a pair of injective *edge maps* $\partial_e \colon B_e \hookrightarrow B_{\iota(e)}$, $\partial_{\bar{e}} \colon B_e \hookrightarrow B_{\tau(e)}$ (note that, using Serre's notation, we only need define one of the two maps). We refer to $\partial_e(B_e)$ as an *edge subgroup* of $B_{\iota(e)}$.

For a vertex v_0 of Γ_B, a *loop based at* v_0 in the graph of groups B is a sequence

$$c = (b_0, e_1, b_1, e_2, b_2, \ldots, b_{n-1}, e_n, b_n)$$

where $\tau(e_i) = \iota(e_{i+1}) = v_i$ for all i, $v_0 = v_n$, and $b_i \in B_{v_i}$. We may multiply two loops c, c' based at v_0 by the obvious rule

$$(b_0, e_1, \ldots, e_n, b_n) \cdot (b'_0, e'_1, \ldots, e'_m, b'_m) = (b_0, e_1, \ldots e_n, b_n b'_0, e'_1, \ldots, e'_m, b'_m).$$

We impose the following equivalence relation on the set of paths:

$$(\ldots e, \partial_{\bar{e}}(b), \bar{e}, \ldots) = (\ldots, \partial_e(b), \ldots).$$

This makes the set of closed paths based at v_0 into a group, called the *fundamental group of the graph of groups B based at v_0*, and written $\pi_1(B, v_0)$ (for more details see [Ser03]).

Definition 2.1.2. A loop $c = (b_0, e_1, b_1, e_2, b_2, \ldots, b_{n-1}, e_n, b_n)$ in a graph of groups B is called *reduced* if:

- If $n = 0$, then $b_0 \neq 1$.
- For every i, if $e_{i+1} = \bar{e}_i$ then $b_i \notin \partial_{\bar{e}}(B_e)$.

Theorem 2.1.3. *If c is a reduced loop, then $c \neq 1$ in $\pi_1(B, v_0)$. Moreover, the natural homomorphism $B_{v_0} \to \pi_1(B, v_0)$ is injective.*

The fundamental group is independent of the basepoint, and so the above theorem in fact implies that there are injective homomorphisms $B_v \hookrightarrow \pi_1(B, v_0)$ for all vertices v.

Definition 2.1.4. Let G be a group. If there exists a graph of groups B such that $G \cong \pi_1(B, v_0)$, then B is called a *graph of groups decomposition* or a *splitting* of G. If the edge groups of G are all of a given type, then B is said to be a *splitting over* subgroups of that type; for example, a splitting over free or cyclic subgroups.

It will be helpful for us to recall the following terminology. See Definition 2.5.7 for the definition of an almost malnormal subgroup.

Definition 2.1.5. A graph of groups B is called *thin* if the following two conditions hold:

1. For every vertex v of Γ_B, every edge subgroup of B_v is almost malnormal.
2. For every vertex v of Γ_B, any two conjugates of distinct edge subgroups of B_v have finite intersection.

Remark 2.1.6. Graphs of groups correspond to group actions on trees, and the study of this correspondence is the subject of *Bass–Serre theory*, for which we refer the reader to [Ser03]. If G acts on a tree T without inversions, then we may construct a *quotient graph of groups* whose fundamental group is G, whose underlying graph is the quotient T/G and whose vertex and edge groups correspond to stabilisers of the vertices and edges of T. Any subgroup H of G also acts on T, and the corresponding quotient graph of groups is called the *induced splitting* for H. An action of a group on a tree is said to be k*-acylindrical* if the stabiliser of any embedded length $k+1$ path is trivial. Wise defines an action of a group on a tree to be *thin* if the stabiliser of any length 2 path is finite [Wis02]. This is a slight weakening of 1-acylindrical, and is equivalent to Definition 2.1.5; indeed, edge stabilisers in the action on the Bass–Serre tree are precisely conjugates of edge groups, and so a stabiliser of a length 2 path would be an intersection between two such conjugates. It follows that the induced splitting for a subgroup of a thin graph of groups is also thin.

2.1.3 Graphs of spaces

Throughout, if not explicitly specified, we will assume all our spaces are CW complexes (they will usually be explicit polyhedral complexes).

Definition 2.1.7. A *graph of spaces* (X, Γ_X) consists of the following:

1. A graph Γ_X, called the *underlying graph*.

2. For each vertex v of Γ_X, a connected *vertex space* X_v.

3. For each edge pair $(e, \bar{e})$ of Γ_X, a connected *edge space* $X_e = X_{\bar{e}}$ and a pair $(\partial_e, \partial_{\bar{e}})$ of π_1-injective *attaching maps* from X_e to $X_{\iota(e)}$, $X_{\tau(e)}$ $\big(= X_{\iota(\bar{e})}\big)$ respectively.

Given the above data, we associate a space X, called the *total space*, as follows. Take a copy of X_v for each v, and a copy of $X_e \times [0,1]$ for each edge pair $(e, \bar{e})$. Now glue $X_e \times \{0\}$ to $X_{\iota(e)}$ using ∂_e, and glue $X_e \times \{1\}$ to $X_{\tau(e)}$ using $\partial_{\bar{e}}$. The edge space X_e is embedded inside X as $X_e \times \{\frac{1}{2}\}$. We will occasionally refer to $X_e \times [0,1]$ as an *edge cylinder*.

We will say that the total space X has a *graph of spaces decomposition*, or simply is a *graph of spaces*. There is an ambiguity here, since a given topological space may possess two different graph of spaces decompositions, so the decomposition in question will always be made clear.

To each graph of spaces, we may naturally associate a graph of groups with the same underlying graph, by replacing vertex and edge spaces with their fundamental groups and the attaching maps with the induced maps on fundamental groups (after choosing basepoints). The fundamental group of the graph of groups (as defined above) is then isomorphic to the fundamental group of the total space. When we discuss *vertex groups* and *edge groups* of a graph of spaces, we are referring to the vertex and edge groups of this corresponding graph of groups.

Given a graph of spaces X, there is a natural projection $\varphi_X \colon X \to \Gamma_X$ which maps each vertex space X_v to the vertex v, and each copy of $X_e \times [0,1]$ to the edge e by projection onto the second factor. A map $f \colon X \to X'$ of graphs of spaces is a map that respects the graph of spaces decomposition, in the sense that there is a graph homomorphism $\gamma \colon \Gamma_X \to \Gamma_{X'}$ such that $\gamma \circ \varphi_X = \varphi_{X'} \circ f$.

If X is a graph of spaces and $\hat{X} \to X$ is a covering map, then $\hat{X}$ inherits a graph of spaces structure. The vertex and edge spaces are the connected components of preimages of vertex and edge spaces of X, and the restriction of the covering map to a vertex (or edge) space of $\hat{X}$ is a covering map to the corresponding vertex (or edge) space of X. This induced structure as a graph of spaces corresponds to the induced splitting as a graph of groups of the subgroup $\pi_1(\hat{X})$.

There may be many graphs of spaces associated to the same graph of groups, but fortunately we have strong control over their homotopy type, thanks to the following lemma of Scott and Wall.

Lemma 2.1.8 ([SW79])**.** *Let X and Y be two graphs of spaces associated to the same graph of groups. Suppose all the vertex and edge spaces of X and Y are aspherical. Then X and Y are aspherical, and hence homotopy equivalent.*

In some accounts of the theory of graphs of groups and graphs of spaces (for example [Hat02]), a graph of groups is chosen first and then a graph of spaces is constructed by choosing a $K(G_v, 1)$ for each vertex group G_v, a $K(G_e, 1)$ for each edge group G_e, and then realising the attaching maps by π_1-injective maps between these spaces. The above lemma then essentially says that this construction is well-defined up to homotopy. In particular, we can safely replace vertex and edge spaces of a graph of spaces with homotopy equivalent spaces, and attaching maps with freely homotopic maps, without changing the homotopy type of the graph of spaces. We will make implicit use of this fact regularly in Chapters 4 and 6.

The theory of graphs of groups and graphs of spaces will be of central importance in all of the forthcoming chapters. In Chapter 3 we will discuss Wise's Hierarchy Theorem for special cube complexes (Theorems 3.8.2 and 3.8.6), and this will inform the work in Chapter 4 where we work in the more restricted setting of a graph of graphs. The work on negatively curved complexes in Chapter 6 is also designed to apply to groups with a certain type of graph of groups decomposition.

2.2 CAT(k) spaces

The CAT(k) criterion gives a method for describing the curvature of a metric space by comparing triangles in the space to triangles in a space of fixed constant curvature. The theory of CAT(k) spaces is well developed, and although we only summarise some of the essential details here, full details and proofs can be found in [BH99, Chapter II]. Our initial setting is a *geodesic metric space*, which is a metric space in which every pair of points is connected by a (not necessarily unique) geodesic.

2.2.1 The CAT(k) criterion

For any $k < 0$, we denote by M_k^n the space obtained by multiplying the metric on n-dimensional hyperbolic space $\mathbb{H}^n$ by $1/\sqrt{-k}$. For $k = 0$, M_k^n denotes Euclidean space $\mathbb{E}^n$, and for $k > 0$, M_k^n denotes the rescaled sphere obtained by multiplying the metric on the n-dimensional unit sphere $\mathbb{S}^n$ by $1/\sqrt{k}$.

A *geodesic triangle* $\Delta(a,b,c)$ in a geodesic metric space Y consists of three points a, b, c and a choice of three geodesics $[a,b]$, $[b,c]$, $[c,a]$. Given such a geodesic triangle and $k \leq 0$, there is a unique (up to isometry) triangle $\overline{\Delta}(\bar{a},\bar{b},\bar{c})$ in M_k^2 with the same edge lengths. For $k > 0$, the same holds provided $d(a,b) + d(b,c) + d(c,a) < 2\pi/\sqrt{k}$. The triangle $\overline{\Delta}$ is called a *comparison triangle* for Δ. For any point p on an edge (say, $[a,b]$) of Δ, there is a unique point $\bar{p}$ on the corresponding edge $[\bar{a},\bar{b}]$ of $\overline{\Delta}$ such that $d(a,p) = d(\bar{a},\bar{p})$, and $\bar{p}$ is called the *comparison point* for p.

Definition 2.2.1. For $k \leq 0$, a geodesic metric space Y is called CAT(k) if any geodesic triangle in Y is *thinner* than the comparison triangle in M_k^2, in the sense that $d(p,q) \leq d(\bar{p},\bar{q})$ for any two points p and q in Δ (see Figure 2.1). For $k > 0$, we require this only for triangles of perimeter less than $2\pi/\sqrt{k}$.

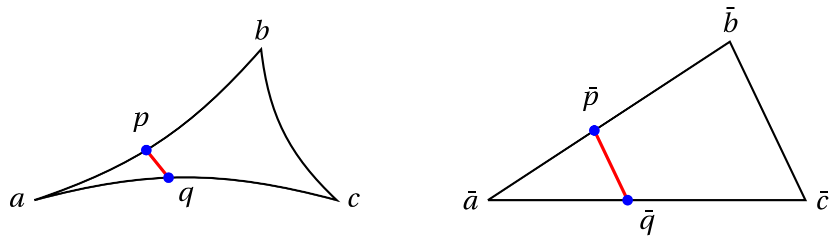

Figure 2.1: The right triangle is a comparison triangle for the left, and the left hand triangle is thinner than the right, as $d(p,q) \leq d(\bar{p},\bar{q})$ for any choice of points p and q. Note that equality is permitted.

Definition 2.2.2. A geodesic metric space is said to be *locally* CAT(k) if every point has a neighbourhood which is CAT(k). A space which is locally CAT(0) is called *non-positively curved*, and a space which is locally CAT(k) for some $k < 0$ is called *negatively curved*.

We summarize some basic properties of CAT(k) spaces for $k \leq 0$ below (there are analogous properties when $k > 0$, provided we restrict to balls of diameter less than $\pi/\sqrt{k}$).

Theorem 2.2.3 (Properties of CAT(k) spaces)**.** *Let X be a* CAT(k) *space, for $k \leq 0$. Then*

- *X is uniquely geodesic (each pair of points is connected by precisely one geodesic).*
- *Every local geodesic in X is a geodesic.*
- *X is contractible.*

Example 2.2.4. [BH99, Theorem II.1A.6] A smooth Riemannian manifold is locally CAT(k) if and only if has sectional curvature $\leq k$. In particular, M_k^n is CAT(k') for all $k' \geq k$.

Remark 2.2.5. It follows directly from the above example that a space which is (locally) CAT(k) for some k is (locally) CAT(k') for all $k' \geq k$.

Example 2.2.6. A metric graph is locally CAT(k) for all $k \in \mathbb{R}$, and a tree is CAT(k) for all $k \in \mathbb{R}$.

Remark 2.2.7. If a space is (locally) CAT(k) for some $k < 0$, then by multiplying the metric by $\sqrt{-k}$ we obtain a (locally) CAT(-1) space. That is, up to rescaling the metric (in particular, up to homeomorphism), a space which is (locally) CAT(k) for some $k < 0$ is (locally) CAT(k) for all $k < 0$. The same is true for $k > 0$.

2.2.2 Local isometries

When working with locally CAT(k) spaces, it is useful to have a well-behaved family of maps between such spaces. This is provided by the following definition.

Definition 2.2.8. Let X and Y be metric spaces. A map $f\colon X \to Y$ is called a *local isometry* if for every $x \in X$, there is a neighbourhood of x on which f restricts to an isometric embedding. A subspace Z of Y is called *locally convex* if the inclusion map $Z \hookrightarrow Y$ is a local isometry.

Remark 2.2.9. If Y is locally CAT(k), and there is a space X equipped with a local isometry to Y, then X is also locally CAT(k).

2.2.3 The Cartan-Hadamard Theorem

Already, we have both local and global notions of negative (and non-positive) curvature. These are related by the following fundamental theorem:

Theorem 2.2.10 (The Cartan-Hadamard Theorem [BH99, Theorem II.4.1])**.** *Let X be a geodesic metric space which is locally* CAT(k)*, for $k \leq 0$. Then the universal cover $\widetilde{X}$ is* CAT(k)*. In particular, a space which is locally* CAT(k) *and simply connected is* CAT(k).

By considering geodesics in X and Y, we see the following. This will be of particular interest when we define quasiconvex subgroups (see Theorem 2.5.6), especially when applied to cube complexes (see Section 3.5).

Corollary 2.2.11. *If Y is locally* $\mathrm{CAT}(k)$*, for $k \leq 0$, and X is equipped with a local isometry $f\colon X \to Y$, then any lift $\tilde{f}\colon \widetilde{X} \to \widetilde{Y}$ is an isometric embedding. Hence X is locally* $\mathrm{CAT}(k)$*, and moreover the induced map on fundamental groups $f_*: \pi_1(X,x) \to \pi_1(Y, f(x))$ is an injection.*

It follows from the above that, for $k \leq 0$, any local isometry into a $\mathrm{CAT}(k)$ space is in fact an isometric embedding, and (together with Theorem 2.2.3) a locally convex subspace of a $\mathrm{CAT}(k)$ space is convex.

2.2.4 M_k–complexes

We saw in Example 2.2.4 that the spaces M_k^n are $\mathrm{CAT}(k)$. A useful way to build more examples of (locally) $\mathrm{CAT}(k)$ spaces is to take polyhedra from M_k^n and glue them together in a compatible way. We will be predominantly interested in the cases where the polyhedra are negatively curved metric simplices or regular Euclidean cubes, and so we will not worry about the (rather technical) general definition of an M_k^n-polyhedron; the concerned reader should consult [BH99, Chapter I.7].

Definition 2.2.12. An *M_k-polyhedral complex* is a CW complex whose cells are M_k^n-polyhedra, attached by isometries of their faces.

Remark 2.2.13. As defined above, an M_k-polyhedral complex may fail to be a complete geodesic space. Consider, for example, the metric graph with vertices $\mathbb{N}$ and an edge of length $1/2^n$ connecting n to $n+1$ for each n; this is an M_k^1 complex for any k, and is isometric to the interval $[0,1)$ which is not complete. There are various assumptions which one may impose on an M_k-polyhedral complex X to rule out similar pathological behaviour. It is a theorem of Bridson that, under the assumption that X has finitely many isometry types of cells, X is a complete geodesic space, although it may not be locally compact or proper. See [BH99, Chapter I.7] for more details. In this thesis, all the complexes we consider will be finite dimensional, complete, and locally compact, in which case it follows from the Hopf–Rinow theorem that they are proper geodesic spaces.

We refer to spaces which are isometric to an M_{-1}, M_0 or M_1-polyhedral complex as *piecewise hyperbolic*, *piecewise Euclidean* or *piecewise spherical* respectively.

Example 2.2.14. Recall that an M_k^n*-simplex* is the convex hull of $n+1$ points in general position in M_k^n (if $k > 0$, these points are required to lie within an open ball of radius $\pi/\sqrt{k}$). These $n+1$ points are the vertices of the simplex, and its faces are precisely the lower dimensional simplices obtained by taking the convex hull of a subset of the vertices. An M_k-polyhedral complex whose cells are all M_k^n-simplices (for various n) is called an M_k*-simplicial complex*.

Remark 2.2.15. In general, by a *simplicial complex*, we will mean a CW complex whose cells are simplices, and where the attaching maps are induced by isometries between faces. This is an abuse of language, because the underlying combinatorial complex may not in fact be simplicial; in general, it is a Δ-complex in the sense of Hatcher (see [Hat02]). In particular, we assume neither that the attaching maps are injective, nor that the intersection between any two simplices is either empty or a single face. If a simplicial complex *does* satisfy these additional conditions, we say that it is *simple*. Any M_k-simplicial complex can be subdivided to make it simple, and so our definition will not cause any geometric problems. However, we will generally assume that vertex links are simple when considering cube complexes (see Section 2.3).

Example 2.2.16. A *cube complex* is an M_k-polyhedral complex all of whose cells are Euclidean cubes $[0,1]^n$ of various dimensions. (We sometimes identify the cubes with $[-1,1]^n$ rather than $[0,1]^n$).

We will investigate cube complexes in more detail in Section 2.3; for now, we restrict our attention to M_k-simplicial complexes.

2.2.5 The link condition in simplicial complexes

To study the geometry of an M_k-simplicial complex, it is very often necessary to focus on the local geometry around a vertex of the complex. For this, we need the notion of the *link* of a vertex, which is a description of the space of possible *directions* at that vertex. It is easiest to define in a purely combinatorial way. Note that we do not, a priori, assume links are simple (see Remark 2.2.15).

Definition 2.2.17 (Link of a vertex in an M_k-simplicial complex)**.** Let v be a vertex in an M_k-simplicial complex K. The *link* of v in K, written $\mathrm{Lk}(v,K)$, is a simplicial complex with a vertex e for each edge e of $K^{(1)}$ with $\iota(e) = v$, and an r-simplex with vertices $e_0,\ldots,e_r$ whenever $e_0,\ldots,e_r$ form the corner of an $r+1$-simplex at v.

Theorem 2.2.18. *If v is a vertex of an M_k-simplicial complex K, then* $\mathrm{Lk}(v,K)$ *possesses a piecewise spherical metric (that is, it is an M_1^n-simplicial complex.)*

Idea of proof. Individual simplices of $\mathrm{Lk}(v,K)$ are in correspondence with $\mathrm{Lk}(v,S)$, where S is a simplex of K containing v. We can think of $\mathrm{Lk}(v,S)$ as the intersection with S of a small sphere in M_k^n centred at v, and it is thus endowed with a spherical metric. It follows that the metrics on adjacent simplices of $\mathrm{Lk}(v,K)$ are compatible, and so there is a well-defined, piecewise spherical metric on the whole of $\mathrm{Lk}(v,K)$. See [BH99, Chapter I.7] for full details. □

Individual points of $\mathrm{Lk}(v,S)$ correspond to initial segments of geodesics through v, and with this metric, the distance between two such points is equal to the angle between the corresponding geodesics. Thus, the metric on links gives us a well-defined notion of angle between geodesics through vertices in K; in particular, it gives us a notion of angle between edges incident at K. Note that the angle between two geodesics issuing from a vertex may be greater than π, and in the case where the link is disconnected, it may take the value ∞. For more details, see [BH99].

Remark 2.2.19. A simplicial path γ in K is locally geodesic if and only if, at each vertex x along the path, the angle between the incoming and outgoing edges of γ at x is at least π. We call this the angle *subtended* by γ at x. A closed simplicial path which is locally geodesic is called a *closed geodesic.*

The following crucial result, observed by Gromov [Gro87] and proved by Bridson [Bri91], is what makes M_k-polyhedral complexes so important in the study of $\mathrm{CAT}(k)$ spaces.

Theorem 2.2.20 (The link condition for M_k-simplicial complexes)**.** *Let K be an M_k-simplicial complex with finitely many isometry types of simplices. Then K is locally* $\mathrm{CAT}(k)$ *if and only if for each vertex $v \in K$,* $\mathrm{Lk}(v,K)$ *is* $\mathrm{CAT}(1)$.

Proof. See [BH99, Chapter II.5]. □

Remark 2.2.21. In the case where K is 2-dimensional, the links of vertices are metric graphs, which are CAT(1) precisely when they do not contain any essential loops of length less than 2π. This makes it particularly easy to check the link condition in the 2-dimensional case.

Remark 2.2.22. Given a CAT(1) space X, the space obtained by connecting two points $x, y \in X$ with an arc of length l is CAT(1) if and only if $d(x,y) + l \geq 2\pi$. This is because no new non-degenerate triangles of perimeter $< 2\pi$ are introduced. This is a useful fact when checking that links remain CAT(1)—and hence, spaces remain locally CAT(k) for some k—after gluing operations.

For $k < 0$, we will sometimes refer to an M_k-simplicial complex which is locally CAT(k) as a *negatively curved complex (of curvature k)*.

2.3 Cube complexes

We now turn our attention to the case where all of the cells in an M_k-polyhedral complex are Euclidean cubes. We shall see (Theorem 2.3.8) that this will enable us to reduce the link condition to a purely combinatorial statement, a theme which will be investigated more deeply in Chapter 3. We begin by defining the link of a vertex in this setting, which is no different from the M_k-simplicial complex case. Again, the definition makes sense even if the link is not simple, although this will be an important additional condition shortly.

Definition 2.3.1 (Link of a vertex in a cube complex.)**.** Let v be a vertex in an cube complex X. The *link* of v in X, written $\mathrm{Lk}(v, X)$, is a simplicial complex with a vertex e for each edge e of $K^{(1)}$ with $\iota(e) = v$, and an r-simplex with vertices $e_0, \ldots, e_r$ whenever $e_0, \ldots, e_r$ form the corner of an $r+1$-cube at v. See Figure 2.2.

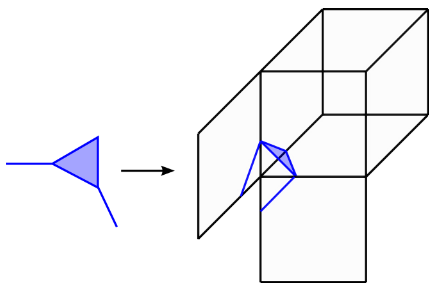

Figure 2.2: The link of a vertex in a cube complex.

Just as for a simplicial complex, we also obtain a piecewise spherical metric on the link; the proof is the same (see Theorem 2.2.18). In the cube complex case, however, we can give much more information about the geometry of the link.

Definition 2.3.2. A piecewise spherical simplicial complex L is called *all-right* if every edge of L has length $\pi/2$.

Remark 2.3.3. If L is an all-right simplicial complex, then links of vertices of L are also all-right (with their piecewise spherical metrics as defined in Theorem 2.2.18). This is because a spherical triangle with all sides of length $\pi/2$ has all three vertex angles equal to $\pi/2$.

Since every face angle in a cube is equal to $\pi/2$, we obtain the following theorem.

Theorem 2.3.4. *Let v be a vertex in a cube complex X. Then* $\mathrm{Lk}(v, X)$ *is an all-right piecewise spherical simplicial complex.*

Theorem 2.3.4 implies that to use the link condition in cube complexes, it is enough to understand when all-right piecewise spherical simplicial complexes are CAT(1). This understanding is provided by Theorem 2.3.6 below.

Definition 2.3.5. A simplicial complex is called *flag* if, for $r \geq 2$, every set of $r+1$ vertices which is pairwise connected by edges spans an r-simplex.

Theorem 2.3.6. *A simple, all-right simplicial complex is* CAT(1) *if and only if it is flag.*

The link condition for M_k-simplicial complexes (Theorem 2.2.20) carries over to cube complexes in the obvious way. We first make the following definition.

Definition 2.3.7. A cube complex X is called *simple* if the links of vertices are simple in the sense of Remark 2.2.15.

Applying Theorems 2.3.4 and 2.3.6, we therefore obtain:

Theorem 2.3.8 (The link condition for cube complexes)**.** *Let X be a finite-dimensional simple cube complex. Then X is non-positively curved if and only if for each vertex $v \in X$,* $\mathrm{Lk}(v, X)$ *is a flag complex.*

It follows in particular that any covering space of a non-positively curved cube complex is non-positively curved, and the universal cover of any non-positively curved cube complex is a CAT(0) cube complex. This is because the links of vertices in the covering space are isomorphic to the links of their projections.

Remark 2.3.9. From now on, we will always assume that cube complexes are simple unless we specify otherwise. This is a harmless assumption in the non-positively curved case.

Example 2.3.10. Every connected orientable surface without boundary of genus ≥ 1 has the structure of a non-positively curved cube complex. For the torus, this is the usual cube complex structure with a single square. For higher genus surfaces, it is the cube complex structure induced by the covering map to the genus 2 surface, which is non-positively curved as shown in Figure 2.3.

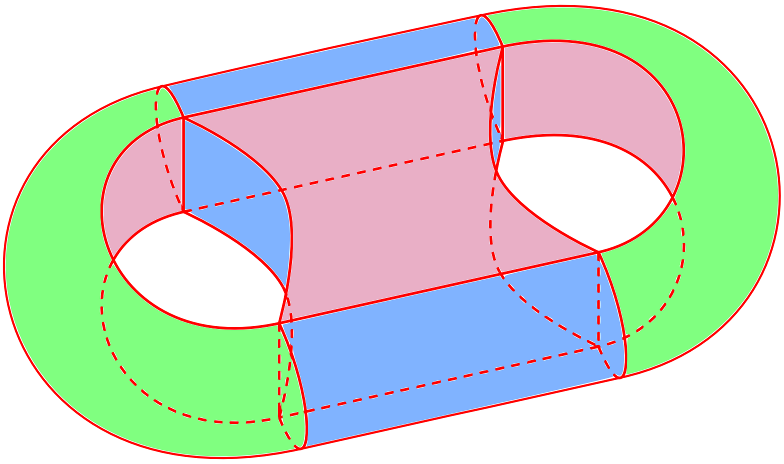

Figure 2.3: A non-positively curved square complex structure on the genus 2 surface. Five squares meet at each vertex, so the links are all pentagons, which are flag.

Thus, we have described non-positive curvature for cube complexes in purely combinatorial terms. Indeed, Theorem 2.3.8 is often used as a quick definition for non-positively curved cube complexes. In view of Theorem 2.2.10, the corresponding *global* notion is as follows.

Definition 2.3.11. A cube complex X is called CAT(0) if it is non-positively curved and simply connected.

2.3.1 Local isometries of cube complexes

When working with cube complexes, we will usually assume that maps are *combinatorial*; that is, they map cubes to cubes of the same dimension. The following theorem gives a combinatorial criterion for such a map to be a local isometry.

Theorem 2.3.12. *Let X and Y be cube complexes. A combinatorial map $f: X \to Y$ is a local isometry if and only if, for every vertex v of X, the induced map $\mathrm{Lk}(v) \to \mathrm{Lk}(f(v))$ is injective and maps $\mathrm{Lk}(v)$ onto a full subcomplex of $\mathrm{Lk}(f(v))$ (that is, if $f(a_0), \ldots, f(a_n)$ are the vertices of an n-simplex in $\mathrm{Lk}(f(v))$, then $a_0, \ldots, a_n$ are the vertices of an n-simplex in $\mathrm{Lk}(v)$).*

Note that when X and Y are non-positively curved, their links are flag, and this guarantees the second condition for $n \geq 2$. Hence, in this case, it is enough to require that if $f(a)$ and $f(b)$ are adjacent in $\mathrm{Lk}(f(v))$ then a and b are adjacent in $\mathrm{Lk}(v)$.

2.4 Negative curvature and group theory

As discussed in Chapter 1, understanding non-positive and negative curvature in terms of group theory is the central theme in this thesis. One key tool which allows us to translate between geometry and group theory is given by Definition 2.4.4. We begin by recalling some fundamental properties of group actions on metric spaces.

Definition 2.4.1. A metric space X is called *proper* if every closed ball is compact.

Definition 2.4.2. Let G be a group acting by isometries on a metric space X. The action is called *cocompact* if the orbit space X/G is compact. It is called *properly discontinuous* if every point $x \in X$ has a neighbourhood U such that $\{g \in G \mid gU \cap U \neq \emptyset\}$ is finite.

Remark 2.4.3. The definition given above of a properly discontinuous group action is not standard—indeed, our definition is sometimes called *wandering* [TL97], and it is inspired by the action of the fundamental group of a compact space on its universal cover. The more common definition of properly discontinuous is that for any compact subset $K \subset X$, there are only finitely many $g \in G$ such that $gK \cap K \neq \emptyset$. This latter definition is in fact weaker for actions on metric spaces, but the definitions coincide for actions of (discrete) groups on proper metric spaces, which is a sufficiently general setting for the work in this thesis.

2.4.1 Geometric actions and the Švarc–Milnor lemma

The following notion allows us to translate geometric properties between groups and spaces:

Definition 2.4.4. We say a finitely generated group acts on a metric space *geometrically* if it acts by isometries, properly discontinuously and cocompactly.

We can now make the following definition, which we restrict to complexes to avoid ambiguity later.

Definition 2.4.5. A group is called CAT(k) if it acts geometrically on a CAT(k) complex. A group is called *freely* CAT(k) if it acts freely and geometrically on a CAT(k) complex.

In light of Remark 2.2.7, Definition 2.4.5 gives us essentially one notion of negative curvature for a group, CAT(-1), as well as a notion of non-positive curvature, CAT(0). However, both of these notions rely on the existence of a particular metric space equipped with a specified geometric action of the group. Indeed, a CAT(-1) group, say, typically possesses many geometric actions on spaces which are not CAT(-1). It would seem preferable to have a notion which is *intrinsic* to the group. For this, we first need to recall the following notion.

Definition 2.4.6. A *quasi-isometric embedding* between metric spaces X and Y is a map $f\colon X \to Y$ for which there exist fixed constants $C \geq 0$, $K > 1$ such that:

- for all $x_1, x_2 \in X$:

$$\frac{1}{K} d(x_1, x_2) - C \leq d(f(x_1), f(x_2)) \leq K d(x_1, x_2) + C$$

A quasi-isometric embedding is called a *quasi-isometry* if it additionally satisfies:

- for every $y \in Y$, there is an $x \in X$ such that $d(f(x), y) \leq A$.

That is, f is an isometry "up to a fixed additive and multiplicative error".

It is straightforward to show that the existence of a quasi-isometry between spaces X and Y is an equivalence relation. The following fundamental result then allows us to discuss geometric properties of a group without specifying a particular space on which the group acts.

Theorem 2.4.7 (The Švarc–Milnor Lemma)**.** *Suppose G is a group acting geometrically on a proper geodesic space X. Then G is finitely generated by a set S, and the Cayley graph of G with respect to S, equipped with the path metric, is quasi-isometric to X. The quasi-isometry is given by $g \mapsto g \cdot x$ for any fixed $x \in X$.*

2.4.2 Hyperbolic spaces and groups

Definition 2.4.8. For some $\delta > 0$, a metric space X is called *δ-hyperbolic* if it satisfies the *δ-slim triangles condition*: namely, that each side of any geodesic triangle in X is contained in the δ-neighbourhood of the other two sides. See Figure 2.4. A space is called *hyperbolic* if it is δ-hyperbolic for some δ.

Definition 2.4.9. A *δ-hyperbolic group* is a group whose Cayley graph is a δ-hyperbolic metric space. A group is called *word hyperbolic* (or simply *hyperbolic*) if it is δ-hyperbolic for some δ.

Theorem 2.4.10. *If there is a quasi-isometric embedding $f\colon X \to Y$, and Y is δ-hyperbolic, then there exists $\delta' > 0$ such that X is δ'-hyperbolic. Hence, hyperbolicity of metric spaces is invariant under quasi-isometry.*

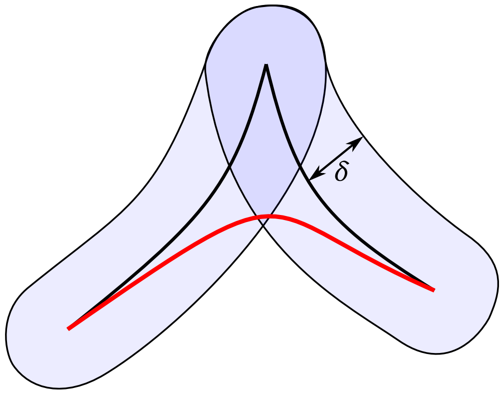

Figure 2.4: X satisfies the δ-slim triangles condition if, for every geodesic triangle in X, each side is contained in the δ-neighbourhood of the other two sides.

Remark 2.4.11. It follows from Theorems 2.4.10 and 2.4.7 that if G is a hyperbolic group, then any proper geodesic space equipped with a geometric action of G is δ-hyperbolic for some δ. In particular, the Cayley graph of G is a hyperbolic metric space whatever finite generating set we choose. In this sense, hyperbolicity is a property intrinsic to the group; we do not need to specify a particular action on a particular metric space in order to define it.

2.4.3 Relationships between the notions of negative curvature

Hyperbolicity is the most commonly studied notion of negative curvature in group theory, largely thanks to Remark 2.4.11. However, as we discussed in Chapter 1, the interplay between hyperbolicity and the CAT(-1) and CAT(0) notions is still an area of active research. We summarize here what is already known.

The strongest of the three notions is CAT(-1). Firstly, it follows from Remark 2.2.5 that any CAT(-1) group is CAT(0). Moreover, we have the following theorem.

Theorem 2.4.12. *Let G be a* CAT(-1) *group. Then G is word hyperbolic.*

Proof. It is enough to show that any CAT(-1) metric space X is δ-hyperbolic. Let $\Delta(a,b,c)$ be a geodesic triangle in X. Choose a side, say $[a,b]$, of X. Fix a point $p \in [a,b]$. For any $q \in [a,c] \cup [b,c]$, by Definition 2.2.1, $d(p,q) < d(\bar{p},\bar{q})$ where $\bar{p},\bar{q}$ are the comparison points on the comparison triangle $\overline{\Delta}$ in $\mathbb{H}^2$. Now, it is an exercise (see [BH99, Chapter III.H.1]) to see that $\mathbb{H}^2$ is δ-hyperbolic, and hence we can choose q such that $d(p,q) < \delta$. This completes the proof. $\square$

On the other hand, the following is an open question.

Question 2.4.13. Let G be hyperbolic. Is G CAT(-1)?

As mentioned earlier, we will answer this question for some classes of groups in Chapters 5 and 6.

To understand the interplay between hyperbolicity and the CAT(0) condition, recall first the following basic fact about hyperbolic groups.

Theorem 2.4.14 ([BH99, Corollary III.Γ.3.10])**.** *Let G be hyperbolic. Then G contains no subgroup isomorphic to $\mathbb{Z} \times \mathbb{Z}$.*

Clearly, $\mathbb{Z} \times \mathbb{Z}$ is a CAT(0) group; it acts geometrically on the Euclidean plane M_0^2. Hence, there exist CAT(0) groups which are not hyperbolic. On the other hand, we have the following theorem.

Theorem 2.4.15 (The Flat Plane Theorem [BH99, Theorem III.H.1.5])**.** *Let G be a group acting geometrically on a* CAT(0) *space X. Then G is hyperbolic if and only if X does not contain an isometrically embedded Euclidean plane.*

One would like to make the hypothesis on the group G, rather than the space X. In this case, we have the following.

Theorem 2.4.16 (The Flat Torus Theorem [BH99, Theorem II.7.1])**.** *Let G be a group containing a $\mathbb{Z} \times \mathbb{Z}$ subgroup, which acts geometrically on a* CAT(0) *space X. Then the $\mathbb{Z} \times \mathbb{Z}$ subgroup stabilises an isometrically embedded Euclidean plane in X.*

On the other hand, the converse to Theorem 2.4.16 is unknown; even in the cube complex case: it is *not* known that every group which acts geometrically on a CAT(0) space containing an embedded flat plane must in fact contain a $\mathbb{Z} \times \mathbb{Z}$ subgroup. In other words, given a CAT(0) group G and corresponding CAT(0) space X, ruling out flat planes in X is enough to ensure δ-hyperbolicity of G, but ruling out $\mathbb{Z} \times \mathbb{Z}$ subgroups in G is not known to be enough. If this were known, it would provide an affirmative answer to Question 1.0.1 in the setting of groups acting geometrically on a CAT(0) space (since the quotient would give a classifying space). For more details, see [McC09] or [BH99].

The situation is even less clear if we are provided with a hyperbolic group to begin with. A weaker form of Question 2.4.13, the following is also an open question.

Question 2.4.17. Let G be hyperbolic. Is G CAT(0)?

It is worth pointing out that the groups we consider in Chapter 6 are already known to be CAT(0), and so we do not make any progress with Question 2.4.17 in this thesis. Nonetheless, this is also a very active area of research, particularly in the setting of cube complexes; many classes of hyperbolic groups have recently been shown to act geometrically on a CAT(0) cube complex, a process which has become known as *cubulation*. Indeed, it turns out that a group acting on a CAT(0) cube complex is of much more interest than a general CAT(0) space, and the reasons for this come down to the theory of *special* cube complexes. As we mentioned in Chapter 1, special cube complexes have a variety of interesting properties, particularly regarding *subgroup separability*, as well as a powerful structure theory which allows us to arrange them in a hierarchy. We will discuss them in detail in Chapter 3.

2.4.4 Torsion and the fundamental group

The definition of a geometric action (Definition 2.4.4) is motivated by the action of the fundamental group of a compact space X on its universal cover. In this case, the action is not only geometric but also *free*—that is, it has no fixed points. The following fact simplifies matters in the case of torsion-free groups.

Lemma 2.4.18. *If a group G is torsion free, then any geometric action on a space X is a free action.*

Proof. For any $x \in X$, the stabiliser of x under the action is a subgroup of G. Since the action is properly discontinuous and by isometries, this stabiliser is finite. Since G is torsion-free, it must therefore be trivial. □

The converse to Lemma 2.4.18 is not true in general, as not every group possessing a free geometric action on a space is torsion-free. However, we can obtain a partial converse in a strong enough context. Recall that a $K(G,1)$ *complex* for a group G is a CW complex whose fundamental group is G and whose universal cover is contractible. Now recall the following fact (see [Hat02]):

Theorem 2.4.19. *If a group G has a finite $K(G,1)$ complex, then it is torsion free.*

Hence, we have the following partial converse to Lemma 2.4.18.

Lemma 2.4.20. *Suppose a group G acts freely and geometrically on a contractible, finite dimensional CW complex. Then G is torsion free. In particular (by Theorem 2.2.3) any group which* *acts freely and geometrically on a* $\mathrm{CAT}(k)$ *complex for* $k \leq 0$*, for example the fundamental group of a compact* $\mathrm{CAT}(k)$ *complex, is torsion free.*

2.4.5 Dimensions of groups

We make the following explicit for the avoidance of ambiguity:

Definition 2.4.21. The *geometric dimension* of a group G is the minimum dimension of a $K(G,1)$. The $\mathrm{CAT}(-1)$ *dimension* of G is the minimum dimension of a compact $\mathrm{CAT}(-1)$ complex which is a $K(G,1)$.

Remark 2.4.22. The requirement that the $K(G,1)$ be compact in the definition of $\mathrm{CAT}(-1)$ dimension is to ensure that a group cannot have a defined $\mathrm{CAT}(-1)$ dimension unless it is $\mathrm{CAT}(-1)$. Indeed, a version of Rips' construction can be used to build groups which have a non-compact $\mathrm{CAT}(-1)$ $K(G,1)$, but which are not finitely presented (see [BH99, II.5]). In particular, they are not hyperbolic, and so cannot be $\mathrm{CAT}(-1)$.

Remark 2.4.23. The $\mathrm{CAT}(-1)$ dimension of a group is at least the geometric dimension, but they may not be equal; Brady and Crisp [BC07] give examples of groups with geometric dimension 2 but $\mathrm{CAT}(-1)$ dimension 3.

Recall that any group with finite geometric dimension is torsion-free (Theorem 2.4.19). In particular, a group with finite $\mathrm{CAT}(-1)$ dimension is in fact freely $\mathrm{CAT}(-1)$, and a group has $\mathrm{CAT}(-1)$ dimension 2 if and only if it is the fundamental group of a compact negatively curved simplicial 2-complex, the main objects of study in Chapters 5 and 6.

2.5 Subgroups of negatively curved groups

Now that we have various notions on negative curvature in group theory, it is natural to wonder about the behaviour of subgroups with respect to these notions. A good candidate for a "well-behaved" family of subgroups is described in the following section.

2.5.1 Quasiconvexity

Definition 2.5.1. A subspace Y of a geodesic metric space X is called *quasiconvex* if there exists $k > 0$ such that any geodesic in X whose endpoints lie in Y is contained in the k-neighbourhood of Y. A subgroup H of a finitely generated group G with generating set S is called *quasiconvex with respect to S* if the corresponding subspace of the Cayley graph $\Gamma_S(G)$ is a quasiconvex subspace.

In general, quasiconvexity of a subgroup depends on the generating set—for example, the subgroup of $\mathbb{Z} \times \mathbb{Z}$ generated by $(1,1)$ is quasiconvex with respect to the generating set $\{(1,0), (1,1)\}$ but not with respect to the generating set $\{(1,0), (0,1)\}$. However, if the ambient group is hyperbolic, quasiconvexity is a well-defined notion. To be precise, the following holds. See [BH99, III.Γ.3] for details.

Theorem 2.5.2. *Let G be a hyperbolic group, and H a subgroup of G. The following are equivalent:*

1. *There exists a finite generating set S of G such that H is quasiconvex with respect to S.*
2. *For every generating set S of G, H is quasiconvex with respect to S.*
3. *There exists a proper geodesic space X, a geometric action of G on X, and a point $x \in X$, such that the orbit $H \cdot x$ is a quasiconvex subspace of X.*
4. *For any proper geodesic space X with a geometric action of G, and any point $x \in X$, the orbit $H \cdot x$ is a quasiconvex subspace of X.*
5. *For any finite generating set S of G, the subspace of $\Gamma_S(G)$ corresponding to H is quasi-isometrically embedded.*

If any of the above hold, we say H is a quasiconvex subgroup of G.

Sketch proof. By Theorem 2.4.7, Cayley graphs of G and geodesic spaces equipped with geometric actions of G are all quasi-isometric. Quasi-isometries of δ-hyperbolic spaces preserve quasiconvexity of subspaces, since they take geodesics to quasi-geodesics, and in δ-hyperbolic spaces, quasi-geodesics are uniformly close to geodesics (see [BH99, III.H.1.7]). Hence, we see the equivalence of 1—4. For 5, see [BH99, III.Γ.3.5]). □

Definition 2.5.3. A subgroup H of a hyperbolic group G is called *distorted* if it is not quasiconvex.

Theorem 2.5.4. *If G is a hyperbolic group and $H < G$ is a quasiconvex subgroup, then H is hyperbolic.*

Remark 2.5.5. If G is a hyperbolic group, $G' < G$ is a finite index subgroup and $H < G'$ is quasiconvex in G', then H is quasiconvex in G. This also follows from the characterisation of quasiconvex subgroups as those which are quasi-isometrically embedded in the Cayley graph.

The Cartan–Hadamard Theorem (see Section 2.2.3) hints at a way to construct quasiconvex subgroups in the context of non-positively curved spaces. Indeed, Corollary 2.2.11 provides a CAT(0) space $\widetilde{Y}$, a convex subspace $\tilde{f}(\widetilde{X}) \subseteq \widetilde{Y}$, and a geometric action of a group $\pi_1(Y)$ with a subgroup $\pi_1(X)$ which is cocompact and acts invariantly on the subcomplex. It follows that the orbit under $\pi_1(X)$ of some choice of basepoint in $\tilde{f}(\widetilde{X})$ is quasiconvex, and hence that $\pi_1(X)$ is a quasiconvex subgroup. In summary, we have the following extension of Corollary 2.2.11 in the hyperbolic case.

Theorem 2.5.6. *Let $f\colon X \to Y$ be a local isometry between compact geodesic spaces X and Y, where Y is non-positively curved and $\pi_1(Y)$ is hyperbolic. Then:*

- *X is non-positively curved.*
- *The induced map f_* on fundamental groups is injective.*
- *$f_*(\pi_1(X)) \subset \pi_1(Y)$ is a quasiconvex subgroup.*

Proof. Consider the induced map $\widetilde{X} \to \widetilde{Y}$. By Corollary 2.2.11, this is an isometric embedding, and it follows that $\widetilde{X}$ is CAT(0) and hence X is non-positively curved. The second conclusion is given explicitly in Corollary 2.2.11.

For the third conclusion, consider the action of $\pi_1(Y)$ on the universal cover $\widetilde{Y}$. This action is geometric because it is a covering space action, and by Corollary 2.2.11, the subspace $\bar{f}(\widetilde{X}) \to \widetilde{Y}$ is a convex subspace. Moreover, since Y is compact, the induced action of $f_*(\pi_1(X))$ on $\bar{f}(\widetilde{X})$ is cocompact. Therefore, it is quasi-isometric to the orbit $f_*(\pi_1(X)) \cdot y$ for some choice of basepoint y, and by Theorem 2.5.2, $f_*(\pi_1(X))$ is a quasiconvex subgroup. □

We will discuss the importance of quasiconvex subgroups in the cube complex setting in Section 3.5.

2.5.2 Malnormality and finite width

The definition of quasiconvexity is geometrically inspired. The following is a related, but algebraic, property of subgroups.

Definition 2.5.7. A subgroup $H \leq G$ is called *malnormal* if $g^{-1}Hg \cap H = \{1\}$ for any $g \notin H$, and *almost malnormal* if $g^{-1}Hg \cap H$ is finite for any $g \notin H$.

Note that in a torsion-free group, almost malnormal and malnormal are equivalent.

Remark 2.5.8. Malnormality is defined as an algebraic condition, but in a sufficiently strong geometric context, it is closely linked to geometric properties. For example, if Σ is a surface in a 3-manifold M, then a failure of malnormality of $\pi_1(\Sigma)$ in $\pi_1(M)$ corresponds to the existence of an annulus C with a map $C \to M$ taking the boundary components of C to essential closed curves in Σ, but such that the image of C is not homotopic into Σ. If Σ is the boundary of M, then malnormality of $\pi_1(\Sigma)$ in $\pi_1(M)$ is indeed equivalent to the non-existence of such an annulus, and M is called *acylindrical.*

Definition 2.5.7 can be generalised as follows, as studied by Gitik et al. in [GMRS98].

Definition 2.5.9. Two conjugates $g_1^{-1}Hg_1$, $g_2^{-1}Hg_2$ of a subgroup $H \leq G$ are called *distinct conjugates* if $Hg_1 \neq Hg_2$. The *width* of H is the maximal number of distinct conjugates of H whose pairwise intersections are infinite. If H is finite, the width is defined to be zero.

Remark 2.5.10. Note that the definition is a slight abuse of notation, since "distinct conjugates" may actually be the same subgroup of G. Note also that almost malnormal subgroups are precisely those with width 1. It is not hard to see that infinite index normal subgroups have infinite width, while any finite index subgroup has finite width.

The relationship between finite width and quasiconvexity is given by the following theorem, which is the main result of [GMRS98].

Theorem 2.5.11. *Quasiconvex subgroups of hyperbolic groups have finite width.*

The converse to this theorem is an open question (for finitely generated subgroups), of which Question 1.0.2 is a special case. We will discuss this further in Section 4.4.

Chapter 3

Special cube complexes

We introduced cube complexes in the previous chapter, as an example of a setting where non-positive curvature can be reduced to a purely combinatorial notion. In this chapter, we will describe a far-reaching application of this combinatorial structure, in an account of some of the theory developed in [HW08] and [Wis12a]. A good starting point is the well-understood case of 1-dimensional CAT(0) cube complexes: trees.

Let Γ be a group acting on a tree. If Γ acts freely (that is, only the identity element fixes any point), then Γ is in fact a free group. More generally, if Γ is allowed to fix points, then it can be given a graph of groups structure, whose vertex and edge groups correspond to the vertex and edge stabilisers of the action (this is the starting point of Bass–Serre theory; see [Ser03]).

To generalise this to actions on higher dimensional CAT(0) cube complexes, one might reasonably start by looking for some examples of groups which possess actions on CAT(0) cube complexes which are particularly easy to describe; indeed, we could look for groups which arise as fundamental groups of certain "easy" non-positively curved cube complexes. To this end, in Section 3.2 we will introduce *right-angled Artin groups* as the fundamental groups of a particular class of non-positively curved cube complexes called *Salvetti complexes*. Next, to widen our range of examples, we apply Corollary 2.2.11 and look for cube complexes admitting a local isometry to a Salvetti complex. This can be arranged by imposing restrictions on the *hyperplanes*, and inspires the definition of a special cube complex (Section 3.3). We thereby obtain an important foundational theorem: namely, that special cube complexes are precisely those whose fundamental groups are subgroups of right-angled Artin groups (Theorem 3.3.15).

In Section 3.4, we will discuss the group-theoretic property of *subgroup separability*, and describe how this property is motivated by geometric ideas; and in Sections 3.5 and 3.6 we will discuss the quasiconvex subgroups of fundamental groups of special cube complexes, proving that they are separable, and finding some explicit examples by looking again at hyperplanes. Moreover, we will be able to give sufficient conditions for a cube complex to be special in terms of the separability of certain subgroups.

In Section 3.8, we will indicate a few ways in which the theory of special cube complexes has been developed further, including the landmark applications to 3-manifold topology. An important development is the existence of a *hierarchy* for virtually special groups (see Corollary 3.8.4), and this provides a potential inductive framework for solving many problems in topology and group theory. We pick up on this idea further in Chapter 4.

The seminal paper which introduced special cube complexes was [HW08], and the account we give here (up to Section 3.6) is based on that paper. This followed earlier work of Daniel Wise in which he proved subgroup separability properties for a variety of classes of groups (for example [Wis00] and [Wis02]). Subgroup separability was also studied earlier by Scott (see [Sco78]) and earlier still by Marshall Hall (see [Hal49]). A good reference for further details on the combinatorial geometry of cube complexes is [Hag08]. Much of Wise's work is summarised very readably in [Wis12a], which is the primary reference for the hierarchy material outlined in Section 3.8, and more detail is given in [Wis12b]. A good reference for the applications to 3-manifolds is [AFW15].

3.1 Cube complexes and their hyperplanes

Let X be a cube complex. In this chapter, we identify each n-cube with $[-1,1]^n$, and we refer to 0-cells, 1-cells and 2-cells as *vertices*, *edges* and *squares* respectively. A cube complex whose cells have dimension at most 2 is called a *square complex*.

Each n-cube of X has n *midcubes*, given by setting one of the n coordinates to 0. A midcube in an n-cube is isometric to an $n-1$-cube. The inclusion of one cube of X as the face of another induces an inclusion of the corresponding midcubes. Taking the disjoint union of midcubes and gluing faces according to these induced inclusion maps, we obtain a family of cube complexes H, each equipped with a natural map $\varphi_H \colon H \to X$ sending cubes of H to midcubes of X. We say H is a *hyperplane* of X (see Figure 3.1 and Section 3.3 for more details).

An equivalent way to describe hyperplanes is as follows. Call two edges of X *elementarily parallel* if they are opposite sides of a square of X. This generates an equivalence relation on the edges of X, called *parallelism*. There is then an exact correspondence between hyperplanes and parallelism classes of edges; the set of midcubes of edges in a parallelism class is precisely the vertex set of a hyperplane. We say the hyperplane is *dual* to an edge in the corresponding parallelism class.

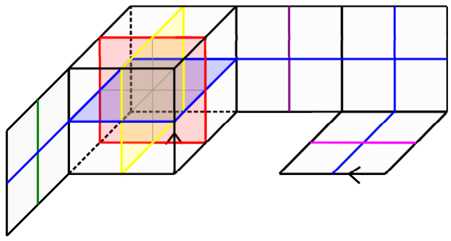

Figure 3.1: A cube complex, with the six hyperplanes shown. The two edges marked by the arrow are identified. The largest parallelism class of edges corresponds to the blue hyperplane, and contains all seven vertical edges and the three rightmost horizontal edges.

3.1.1 Sageev's dual cube complex construction

The following construction of Sageev [Sag95] provides an additional reason to work with cube complexes, rather than any other type of combinatorial structure. It has been applied recently to many different groups.

Definition 3.1.1. A *wallspace* $(X, \mathscr{W})$ is a space X with a set of bipartitions $\mathscr{W} = \{\{W_i, V_i\} \mid X = W_i \cup V_i \ \forall i \in \mathscr{I}\}$ (called *wall partitions*) of the space into *halfspaces* W_i and V_i, such that any two points in X are both in the same halfspace for all but finitely many wall partitions.

Theorem 3.1.2. *Let G be a group acting on a wallspace $(X, \mathscr{W})$, where the action permutes the wall partitions. Then there exists a* CAT(0) *cube complex C, called the* dual cube complex *to $(X, \mathscr{W})$, with an action of G on C.*

Wallspaces are highly ubiquitous in geometry, with wall partitions arising from (for example) incompressible immersed surfaces in a 3-manifold, and in group theory, where (for example) the subgroup $C \leq A *_C B$ gives rise to a wall partition. By choosing "sufficiently many" walls, the action on the dual cube complex can be controlled. For more details, see Section 3.8 and [Wis12a]. Note, however, that the cube complexes arising from Sageev's construction are typically very high-dimensional. Thus, if one wishes to obtain optimal information about the dimension of groups, it may be preferable to study different combinatorial stuctures, as we show in Chapters 5 and 6.

3.2 Right-angled Artin groups

The following class of groups is constructed by associating a simple presentation to a simplicial graph.

Definition 3.2.1. Let Γ be a graph with vertices $V(\Gamma)$. The *right-angled Artin group* corresponding to Γ is

$$A_\Gamma = \langle V(\Gamma) \,\big|\, \{ vwv^{-1}w^{-1} \mid vw \in E(\Gamma) \} \rangle.$$

Free and free abelian groups are both examples of right-angled Artin groups, corresponding respectively to graphs with no edges and complete graphs (see Figure 3.2).

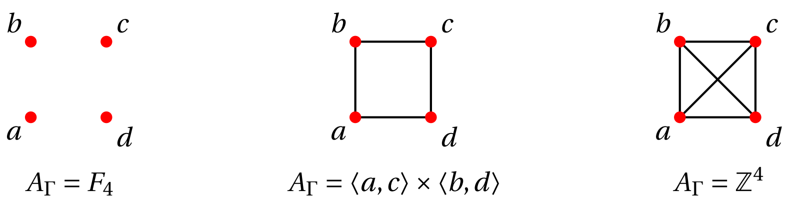

Figure 3.2: Examples of right-angled Artin groups

Remark 3.2.2. The simplicity of this construction ensures that right-angled Artin groups have many appealing properties. For example, if Γ is a *finite* graph then A_Γ is linear (see [HW99])—indeed, any finitely generated right-angled Artin group can be embedded in $\mathrm{SL}_m(\mathbb{Z})$ for some m. Therefore, a group is linear if it can be realised as a subgroup of a finitely generated right-angled Artin group. As we will see later, this is a much less restrictive condition than it first appears (see Theorem 3.3.15).

3.2.1 Salvetti complexes

The reason right-angled Artin groups are interesting when studying cube complexes is that we can construct a cube complex with any given right-angled Artin group as its fundamental group. We describe the construction for the case where Γ is finite, but the infinite case is similar.

Construction 3.2.3. As before, let Γ be a graph with vertices $v_1, \ldots, v_n$. The Salvetti complex for the group A_Γ is constructed via a subcomplex S' of the n-dimensional unit cube spanned by $x_1, \ldots, x_n$. The r-cube spanned by $x_{i_1}, \ldots, x_{i_r}$ is included in S' if and only if the vertices $v_{i_1}, \ldots, v_{i_r}$ of Γ are pairwise connected. In the resulting cube complex S', the link of the origin is the unique flag complex with 1-skeleton Γ. Identifying opposite sides of the n-cube, S' becomes a subcomplex $S(\Gamma)$ of the n-dimensional torus. This is defined to be the *Salvetti complex* of A_Γ.

The Salvetti complex is built from a bouquet of n loops, and contains a 2-torus spanned by a pair of loops if and only if the corresponding vertices are joined by an edge in Γ. Indeed, this fully describes the 2-skeleton of $S(\Gamma)$, and so $\pi_1(S(\Gamma)) = A_\Gamma$. Furthermore, we can show that $S(\Gamma)$ is non-positively curved.

Proposition 3.2.4. *$S(\Gamma)$ is non-positively curved.*

By Theorem 2.3.8, it is enough to show that the link L of the single vertex of $S(\Gamma)$ is a flag complex. L has vertex set $\{v_1^\pm, \ldots, v_n^\pm\}$. Consider an r-cube in S' (from Construction 3.2.3) spanned by $x_{i_1}, \ldots, x_{i_r}$. This contributes 2^r $r-1$-simplices to L, which are the links of the 2^r vertices of the r-cube in S'. So, the vertex set of each simplex corresponds to a choice of superscripts in $\{v_{i_1}^\pm, \ldots, v_{i_r}^\pm\}$, and moreover any such choice appears as the vertex set of a simplex. In particular, since the subcomplex obtained by taking all superscripts as $+$ is flag (as in Construction 3.2.3), the whole complex L is flag.

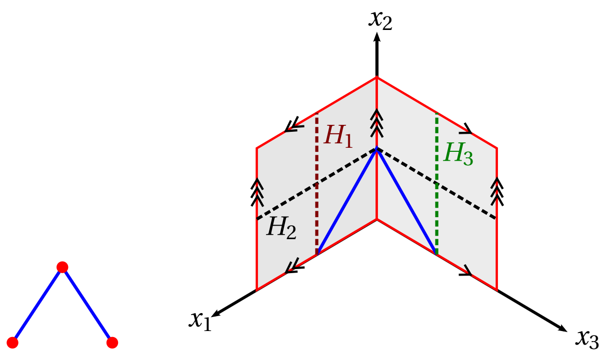

Figure 3.3: Example of a graph Γ and corresponding Salvetti complex $S(\Gamma)$. Side identifications are indicated by arrows. The three hyperplanes are labelled H_1, H_2 and H_3.

By Corollary 2.2.11, we know that if we can find a cube complex X which admits a local isometry to the Salvetti complex $S(\Gamma)$, then $\pi_1(X)$ will be isomorphic to a subgroup of the

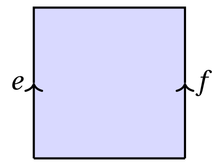

Figure 3.4: Edges e and f have a consistent orientation.

right-angled Artin group A_Γ. As we will see in the next section, special cube complexes are designed precisely to ensure that this local isometry exists.

3.3 Special cube complexes

In order to give the definition of a special cube complex, we must first look in more detail at the behaviour of hyperplanes. Recall that hyperplanes correspond to parallelism classes of edges. We may choose an orientation on any edge e of X. If e is elementarily parallel to f, it is opposite f in a square of X, and e and f may be given a *consistent* orientation (see Figure 3.4).

Definition 3.3.1. If every edge in a parallelism class can be oriented consistently, the corresponding hyperplane is called *two-sided.* Otherwise, it is called *one-sided.* A one-sided hyperplane is shown in the leftmost picture of Figure 3.6.

There is also a more sophisticated way to describe two-sidedness.

Definition 3.3.2. For a hyperplane H in a cube complex X, the *cubical neighbourhood* of H, denoted U_H, is the union of cubes in X intersecting H (that is, the union of cubes whose midcubes comprise H).

Definition 3.3.3. Let M be a midcube of a cube C of X. There is an obvious retraction map $r\colon C \to M$, which we may think of as an interval bundle $N_M \to M$. Since the retraction map commutes with the inclusion of one cube as the face of another, we may piece together these bundles for all midcubes in a given hyperplane H. The resulting bundle is the *normal bundle* N_H, which restricts to N_M over each cube M of H, and is equipped with a map $\chi_H\colon N_H \to X$ with image U_H. See Figure 3.5 for an example.

A hyperplane H is two-sided if and only if N_H is isometric to $H \times [-1,1]$. In this case, the boundary ∂N_H has two components, $\partial N_H^+ = H \times \{1\}$ and $\partial N_H^- = H \times \{-1\}$, each isometric to H, and we identify H with the 0-section $H \times \{0\}$ of the bundle (see Figure 3.5). Otherwise,

N_H is isometric to a twisted interval bundle, ∂N_H has one component, and the hyperplane is one-sided.

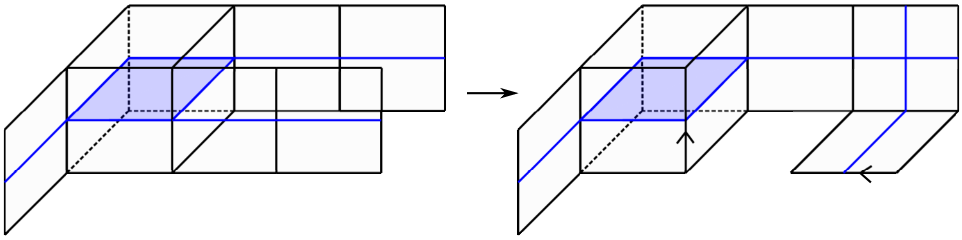

Figure 3.5: The map χ_H from the normal (product) bundle N_H over a hyperplane H into a cube complex X. The restriction of this map to the 0-section (shown in blue) is precisely the map $\varphi_H\colon H \to X$.

From now on, we will assume that all edges dual to a given two-sided hyperplane are consistently oriented.

Definition 3.3.4 (Hyperplane pathologies)**.** Consider two-sided hyperplanes H_1 and H_2 in X with dual edges e_1 and e_2, respectively, which are both incident at some vertex v of X. If e_1 and e_2 form the corner of a square at v, then we say H_1 and H_2 *intersect*. If e_1 and e_2 do not form the corner of a square at v, we say H_1 and H_2 *osculate* at v. If $H_1 = H_2$, then we use the terms *self-intersect* and *self-osculate* accordingly. If H self-osculates and the corresponding edges are both oriented towards (or both away from) v, we say the self-osculation is *direct*, otherwise it is *indirect*. Finally, a pair of distinct hyperplanes is said to *interosculate* if they both intersect and osculate. See Figure 3.6.

Remark 3.3.5. Intersection of two hyperplanes H_1, H_2 is equivalent to non-injectivity of the map $H_1 \cup H_2 \to X$, and osculation of two hyperplanes implies that the map $\partial N_{H_1} \cup \partial N_{H_2} \to X$ is non-injective. Direct self-osculation of H implies non-injectivity of either $\partial N_H^+ \to X$ or $\partial N_H^- \to X$, and indirect self-osculation implies that there is an intersection of the images of these two maps.

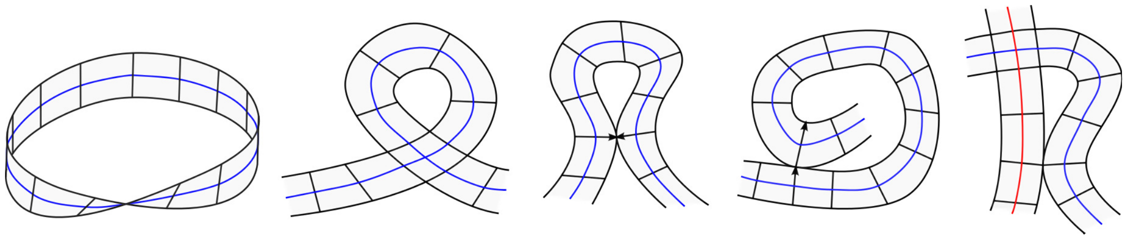

Figure 3.6: Hyperplane pathologies. From left to right: one-sidedness, self-intersection, direct self-osculation, indirect self-osculation and interosculation.

Definition 3.3.6. A *special cube complex* X is a non-positively curved cube complex such that:

1. All hyperplanes are two-sided.
2. No hyperplane self-intersects.
3. No hyperplane directly self-osculates.
4. No pair of hyperplanes interosculate.

That is, special cube complexes are those non-positively curved cube complexes which display none of the behaviour shown in Figure 3.6 (except perhaps indirect self-osculation).

Remark 3.3.7. Haglund and Wise do not insist that special cube complexes are non-positively curved. In fact, any simple cube complex satisfying the given hyperplane conditions can be made non-positively curved without altering the 2-skeleton. See [HW08].

We start with two easy lemmas.

Lemma 3.3.8. *Salvetti complexes are special.*

Proof. We know from Section 3.2.1 that Salvetti complexes are non-positively curved. Each hyperplane is dual to a unique oriented edge, with both endpoints at the single vertex v. This guarantees 2-sidedness, and rules out self-intersection and direct self-osculation (although every hyperplane indirectly self-osculates). For interosculation, observe that if a pair of hyperplanes H_1, H_2 intersect with corresponding dual edges e_1, e_2, then the reversed edges $\bar{e}_1$, $\bar{e}_2$ also span a square at v, hence osculation is impossible for hyperplanes which intersect. □

Lemma 3.3.9. *Covering spaces of special cube complexes are special.*

Proof. Let X be special and $p\colon \hat{X} \to X$ be a covering map. By definition p is a local isometry, so there is a natural induced cube complex structure on X, which is non-positively curved by Theorem 2.2.10. Hyperplanes in $\hat{X}$ project to hyperplanes in X—more specifically, if e and f are parallel edges in $\hat{X}$, then $p(e)$ and $p(f)$ are parallel edges in X. It is straightforward to see that any of the forbidden hyperplane pathologies in $\hat{X}$ would project to the same pathology in X. Thus no pathology can occur in $\hat{X}$, and so it is special. □

Definition 3.3.10. A group is called *(compact) special* if it is the fundamental group of a (compact) special cube complex.

Example 3.3.11. Any orientable surface of genus ≥ 1 is homeomorphic to a special cube complex. This is clearly true for the torus (as it is a Salvetti complex for the right angled Artin group $\mathbb{Z}^2$). All higher genus surfaces cover the genus 2 surface, so by Lemma 3.3.9 it is enough to show that this is special, and it is easy to check by hand that the cube complex structure illustrated in Figure 2.3 is special. (This is not optimal; there is a retraction onto a square complex with only four squares.)

Example 3.3.12. The non-orientable surface of Euler characteristic -1 is not homeomorphic to a special cube complex. This can be shown by an Euler characteristic argument, considering the possible cases for the degrees of vertices in the cube complex structure.

Example 3.3.13. Any CAT(0) cube complex is special. The proof of this uses the fact that hyperplane neighbourhoods are convex subcomplexes, which follows from Lemma 3.6.1 together with Corollary 2.2.11. For full details, see [HW08].

Definition 3.3.14. A property is said to *virtually* hold in a group if it holds in a finite-index subgroup, and similarly in a space if it holds in a finite-sheeted covering space.

In particular, we will often refer to *virtually special* cube complexes, which are cube complexes with a finite-sheeted special covering space.

By the above, we know that right-angled Artin groups are special, and subgroups of special groups are special. We are now ready to state the main theorem of this section, the proof of which will motivate each of the conditions in Definition 3.3.6.

Theorem 3.3.15. *A group G is the fundamental group of a special cube complex if and only if G is a subgroup of a right-angled Artin group.*

Proof. The reverse implication follows from the remarks above, since if G is a subgroup of a right-angled Artin group, it is the fundamental group of a covering space of a Salvetti complex, which is a special by Lemmas 3.3.8 and 3.3.9.

Let X be a special cube complex with fundamental group G. We will construct a local isometry from X to the Salvetti complex of a right-angled Artin group, and this will prove the forward implication by Corollary 2.2.11.

Define $\Gamma(X)$ to be the graph with vertices for each hyperplane of X and an edge between two vertices when the hyperplanes they represent intersect in X. Since no hyperplanes self-intersect, this will be a simplicial graph.

Consider the Salvetti complex $S = S(\Gamma(X))$. We construct a combinatorial map $X \to S$. First, map all the vertices in X to the single vertex in S. Next, we must map edges. Since all the hyperplanes of X are two-sided, we may fix a consistent orientation on all edges of S. Suppose some edge e in X is dual to a hyperplane H_i. Then H_i corresponds to a vertex v_i of $\Gamma(X)$, which corresponds in turn to the oriented edge x_i in the Salvetti complex S. We map e to x_i preserving the orientation.

To map the squares of X to S, we observe that if two edges e_i and e_j in X span a square C then their dual hyperplanes H_i and H_j intersect in C, and hence the vertices v_i and v_j are joined by an edge in $\Gamma(X)$, so there is a corresponding square C' in S spanned by x_i and x_j. We can therefore extend the map to C, mapping it to C'.

Now, since X is non-positively curved, the links of all its vertices are flag, and the same is true for S. Hence, once we have mapped the edges and faces of X into S, the image of the 2-skeleton of any k-cube in X will be the 2-skeleton of a k-cube in S. Therefore, we can extend the map across k-cubes, obtaining a map Φ_X defined on all of X.

To verify that Φ_X is a local isometry, by Theorem 2.3.12 it is enough to check that Φ_X is injective on links, and that whenever the image of Φ_X on a given link contains two adjacent vertices, their preimages were adjacent in the original link. So firstly, suppose two vertices a, $b \in \mathrm{Lk}(v)$ are mapped to the same vertex in $\mathrm{Lk}(\Phi_X(v))$. By construction of the map Φ_X, this means the edges in X corresponding to a and b are parallel (that is, they have the same dual hyperplane), and moreover, are both oriented away from or both towards v. Therefore, the corresponding hyperplane self-intersects or directly self-osculates, which is a contradiction.

Finally, suppose that a, $b \in \mathrm{Lk}(v)$ are vertices whose images in $\mathrm{Lk}(\Phi_X(v))$ are joined by an edge. By construction of Γ, this means the hyperplanes dual to a and b intersect. If a and b are not joined by an edge in $\mathrm{Lk}(v)$, then they osculate at v, so they interosculate, which is forbidden. So a and b are joined by an edge, as required. □

By Remark 3.2.2, the fact that virtually linear groups are linear, and the fact that finitely generated linear groups are residually finite [Mal40], we obtain the following.

Corollary 3.3.16. *Virtually compact special groups are linear. In particular, they are residually finite.*

3.4 Subgroup separability

The proof of Theorem 3.3.15 provides a neat justification for each of the conditions in the definition of a special cube complex. However, it was not the original motivation for the construction. This was to enable a generalisation of a standard technique to prove subgroup separability. In this section we will discuss this technique and its consequences for special groups. Later, we will be able to use subgroup separability to give a criterion for determining virtual specialness.

Definition 3.4.1. A subset H of a group G is called *separable* if, for any $x \in G - H$, there exists a finite index subgroup $G' \leq G$ containing H but not containing x.

Sometimes it will be convenient to express this in the following way:

Definition 3.4.2. The *profinite topology* on a group G is the topology whose closed basis is given by cosets of finite index subgroups of G.

Then a subset is separable if and only if it is closed in the profinite topology. Note that this applies to sub*sets* of G, rather than just subgroups.

Remark 3.4.3. Given any group homomorphism $\varphi\colon G \to H$, if $H' \leq H$ is a finite-index subgroup, then the preimage $\varphi^{-1}(H')$ must have finite index in G. Therefore, group homomorphisms are continuous in the profinite topology. Similarly, if $g \in G$, then gG' is a coset of a finite index subgroup only if G' is a finite index subgroup, so left multiplication is continuous, as is right multiplication. This makes it easy to find continuous maps in the profinite topology.

Definition 3.4.4. A group G is called *residually finite* if the trivial subgroup is separable, and *subgroup separable* if all its finitely generated subgroups are separable.

Proposition 3.4.5. *Let $H \leq G' \leq G$ where G' is finite index in G. Then any subgroup of H which is closed in G' is also closed in H and in G. Moreover, if G' is subgroup separable, then so are H and G.*

Proof. It is straightforward to check this using the profinite topology. For an alternative proof, see (for example) [Sco78]. □

Example 3.4.6. Finite groups are subgroup separable, because every subgroup is finite index and hence closed. The infinite cyclic group is also subgroup separable, because the only subgroup not of finite index is the trivial subgroup, which is the intersection of all finite index subgroups and hence closed. Hence, virtually infinite cyclic groups are subgroup separable by Proposition 3.4.5.

Example 3.4.7. Finitely generated abelian groups are subgroup separable. To see this, first note that they are all residually finite, by their classification. Now let H be a subgroup of the finitely generated abelian group G. Since G is abelian, H is normal, and the quotient G/H is finitely generated abelian; hence H is the preimage under the quotient map of the closed set $\{1\} \subset G/H$, and is therefore closed in G.

In the following lemma, H is said to be a *virtual retract* of G.

Lemma 3.4.8 (The virtual retract criterion)**.** *Let G be residually finite, $G' \leq G$ a finite index subgroup, $H \leq G' \leq G$, and let $\rho\colon G' \to H$ be a retraction (i.e. a homomorphism satisfying $\rho(h) = h$ for all $h \in H$). Then H is separable in both G' and G.*

Proof. Let $f\colon G' \to G'$ be the map sending g to $g^{-1}\rho(g)$. Observe that f is continuous in the profinite topology (see Remark 3.4.3). Applying Proposition 3.4.5 to the trivial subgroup, $\{1\}$ is closed in G' and so $f^{-1}(\{1\}) = H$ is closed in G', and hence in G. □

3.4.1 Topological motivation

Our interest in subgroup separability comes from topology. As we will see later, proving separability for a subgroup H of a group $G = \pi_1(X)$ often relies on taking some *immersion* (that is, locally injective map) $Y \to X$ representing H, and factoring it as the composition of an inclusion and a finite-sheeted covering map. In other words, we will lift an immersion to an embedding in a finite covering space, and deduce separability. In fact, this can be thought of as a characterisation of separability, as noticed by Peter Scott in [Sco78]:

Theorem 3.4.9. *Let X be any cell complex, $G = \pi_1(X)$, and $X_H \to X$ a covering map corresponding to a finitely generated subgroup $H < G$. Then H is separable if any only if, for any compact subcomplex $C \subset X_H$, the covering map $X_H \to X$ factors through a finite-sheeted cover $X' \to X$ such that C projects homeomorphically into X'.*

Therefore, if we have an immersion $Y \to X$, and we know the corresponding subgroup of $\pi_1(X)$ is separable, we will be able to lift this immersion to an embedding in a finite cover. This tool is used time and again in the later development of the theory of special cube complexes (see Theorem 3.6.6, Section 3.8, and [Wis12a] for more examples).

3.4.2 Free groups

One of the first results concerning subgroup separability was provided by Marshall Hall in [Hal49]:

Theorem 3.4.10. *Free groups are subgroup separable.*

Proof. Let F be a free group, with a finitely generated subgroup H. Let B be a bouquet of circles such that $F = \pi_1(B)$, and consider the covering map $X_H \to B$ corresponding to H. In general, X_H is an infinite graph with $\pi_1(X_H) = H$. However, since H is finitely generated, there is a finite, connected subgraph C of X_H with $\pi_1(C) = H$.

The covering map $X_H \to B$ restricts to an immersion $\varphi\colon C \to B$. Let us colour B with a different colour for each edge, and also fix an orientation for each edge. Pulling back under φ, we obtain a colouring and orientation of C.

Consider the subgraph C_1 of C containing every vertex and all edges of some fixed colour (red, say). Every vertex of C_1 has degree at most two, so the connected components are directed cycles, paths or isolated vertices. For each isolated vertex v of C_1, add in a red loop at v. For each component of C_1 which is a non-closed path P, add in to C a red edge e connecting the endpoints, oriented to make $P \cup e$ a directed circle. After doing this for each colour, the resulting graph C' has an incoming and outgoing edge of every colour at every vertex, and so it is a covering map. Note that there is also a naturally defined (non-combinatorial) retraction map $r\colon C' \to C$ given by mapping each new edge continuously onto the monochromatic path or vertex to which we added it. Moreover $G' = \pi_1(C')$ is a finite index subgroup of F (since C' is a finite cover).

The retraction r induces a retraction $G' \to H$. Thus H is separable in G by Lemma 3.4.8 (and the fact that free groups are residually finite—see [Sta83]). □

Remark 3.4.11. This proof is a model for future proofs of subgroup separability, via the stronger property of a subgroup being a virtual retract. We began by taking a covering space X_H corresponding to H. We then found a finite cover C', with a retraction realising H as a virtual retract. To do this, we found a finite subcomplex C of X_H, with an immersion to B given by restricting the covering map. We then *completed* C to a covering space C' (by adding extra edges). C contained all the nontrivial topology of X_H, and will be referred to as a *compact core* of X_H. Scott used a similar technique to prove the following result in [Sco78].

Theorem 3.4.12. *Surface groups are subgroup separable.*

Special cube complexes were originally constructed in order to generalise this technique to a more sophisticated combinatorial setting (see Section 3.4.3 and [Wis12a]). In particular, given a restricted covering map (or any other local isometry) between special cube complexes, we will be able to factor this map as the composition of an inclusion and a finite-sheeted covering map, where the inclusion will have a retraction as a left inverse.

Remark 3.4.13. Hall's original proof of Theorem 3.4.10 does not use the virtual retraction criterion for subgroup separability; instead, it keeps track of a representative loop for a chosen element $x \in F - H$, enlarging C to contain a non-closed path mapping to this loop, which then does not represent an element of $\pi_1(C')$. One advantage of the virtual retraction criterion is to remove the need to keep track of such an element. A disadvantage is that it assumes residual finiteness: indeed, one way to prove residual finiteness for free groups is to repeat the above proof in the particular case where C is a subdivided line mapping to a fixed element of F.

Note also that in Scott's reformulation (Theorem 3.4.9), one condition was for the map $X_H \to X$ to factor through the finite cover C'. This indeed holds in our proof of Theorem 3.4.10, since $H < G'$, but we did not need to use it explicitly in the proof.

3.4.3 Canonical completion and retraction

The first step in studying the separability properties of special cube complexes is to find a method for "completing" a map between special cube complexes to a covering map. In the one dimensional case, where the image complex was a bouquet of circles, this was achieved by adding extra edges. We mimic this in Construction 3.4.16 below. To simplify matters slightly, we use the following terminology: a cube complex B is called *fully special* if it is special and, in addition, hyperplanes do not indirectly self-osculate. That is, there is at most one edge in each parallelism class incident at each vertex of B. This is not a significant restriction—see Remark 3.5.13.

We will require the following categorical notion:

Definition 3.4.14. Given objects X, Y and Z, and morphisms $f\colon X \to Z$ and $g\colon Y \to Z$, the *fibre product* $X \times_Z Y$ is defined to be $\{(x, y) \in X \times Y \mid f(x) = g(y)\}$, equipped with the projections p_1 and p_2 to X and Y. The following diagram commutes.

$$\begin{tikzcd}
X \times_Z Y \arrow[r, "p_2"] \arrow[d, "p_1"'] & Y \arrow[d, "g"] \\
X \arrow[r, "f"'] & Z
\end{tikzcd}$$

Remark 3.4.15. When X and Y are cube complexes, there is an inherited cube complex structure on the fibre product. In the case of graphs, this was used extensively by Stallings (see [Sta83]), who noticed that if X, Y and Z are graphs and f and g are immersions, then one connected component of the fibre product represents the subgroup $\pi_1(X) \cap \pi_1(Y) \leq \pi_1(Z)$. This leads to a quick proof that the intersection of finitely generated free subgroups of a free group is finitely generated.

Construction 3.4.16. Let $f\colon A \to B$ be a local isometry from a compact cube complex A to a fully special compact cube complex B. Then we wish to construct a special cube complex $C = C(A,B)$, equipped with a finite-sheeted covering map $p\colon C \to B$, an inclusion $j\colon A \hookrightarrow C$ and a retraction $r\colon C \to A$ as in the following diagram[1]:

$$\begin{tikzcd}
& C = C(A,B) \arrow[dl, bend left, "r"] \arrow[dr, "p"] & \\
A \arrow[ur, hook, bend left, "j"] \arrow[rr, "f"] & & B
\end{tikzcd}$$

First, let B^* be the Salvetti complex $S(\Gamma(B))$, with $\Phi_B\colon B \to B^*$ defined as in the proof of Theorem 3.3.15. We will begin by constructing a cube complex $C^* = C(A, B^*)$. The 1-skeleton of B^* is a bouquet of circles $x_1, \ldots, x_n$, so as in the proof of Theorem 3.4.10, there is a graph C', a covering map p^*, an inclusion j^* and a retraction r^* as in the diagram below. We take C' as the 1-skeleton of C^*. Note that the vertices of C' are in one-to-one correspondence with the vertices of A. We regard the edges of C' as coloured by colours $x_1, \ldots, x_n$ (according to their image in B^*).

$$\begin{tikzcd}
& C' \arrow[dl, bend left, "r^*"] \arrow[dr, "p^*"] & \\
A^{(1)} \arrow[ur, hook, bend left, "j^*"] \arrow[rr, "\Phi_B \circ f"] & & B^{*(1)}
\end{tikzcd}$$

Since B is fully special, A has at most one edge at any vertex mapping to any given parallelism class of B, and hence the subgraph $A^{(1)}$ of C' has at most one edge of any colour incident at any vertex. Therefore, referring back to the construction in the proof of Theorem

[1]This diagram does *not* commute, because r is defined on the whole of C, but the other maps do commute. The rest of the diagrams in the construction behave similarly.

3.4.10, the single-coloured components of C' are all either loops or double edges; in the latter case, exactly one of these edges is in $A^{(1)}$.

Consider a square S of B^*, with boundary $x_i x_j x_i^{-1} x_j^{-1}$. This loop lifts as a path π at each vertex of C'. By definition, the hyperplanes H_i, H_j corresponding to x_i, x_j intersect in B, and so they cannot osculate in A. Therefore, if some vertex $v \in A$ has an incident edge of each colour x_i, x_j, these edges span a square, and in C', the four edges become double edges. This means that if two consecutive edges of π are non-loops, π is in fact a closed path. If not, then both lifts in π of at least one of x_i and x_j are loops. Again, this means π is a closed path. Hence, we can attach a square to π, mapping under p^* to S. All squares of A will appear in this way, so the inclusion j^* extends to $A^{(2)}$. Each square takes one of three forms depending on whether 0, 2 or 4 of its edges are loops, and the retraction r^* sends such a square onto a square, edge or vertex of A respectively. In this way, we extend $C' = C^{*(1)}$ to a square complex $C^{*(2)}$ satisfying:

$$\begin{tikzcd}
& C^{*(2)} \arrow[dl, bend left, "r^*"] \arrow[dr, "p^*"] & \\
A^{(2)} \arrow[ur, bend left, "j^*"] \arrow[rr, "\Phi_B \circ f"] & & B^{*(2)}
\end{tikzcd}$$

Using the fact that squares in $C^{*(2)}$ take one of three forms, it is easy to verify that we may always glue in higher-dimensional cubes to obtain a non-positively curved complex C^*, and that the maps extend across these cubes (for full details, see [HW08]).

$$\begin{tikzcd}
& C^* = C(A, B^*) \arrow[dl, bend left, "r^*"] \arrow[dr, "p^*"] & \\
A \arrow[ur, bend left, "j^*"] \arrow[rr, "\Phi_B \circ f"] & & B^*
\end{tikzcd}$$

We complete the construction using the fibre product. $C(A,B)$ is defined to be $B \times_{B^*} C(A,B^*)$. This is summarised in the following diagram (as before, all the maps commute except for the retractions).

$$\begin{tikzcd}
B \times_{B^*} C(A,B^*) \arrow[dr, equal] \arrow[drrr, dashed, bend left, "p_2"] \arrow[dddr, dashed, bend right, "p_1"'] & & & \\
& C(A,B) \arrow[rr, "\text{collapse } B"] \arrow[dd, "p"'] \arrow[dr, "r"] & & C(A,B^*) \arrow[dl, dashed, bend left, "r^*"] \arrow[dd, "p^*"] \\
& & A \arrow[ul, bend right, "j"'] \arrow[ur, dashed, bend left, "j^*"] \arrow[dl, "f"] \arrow[dr, "\Phi_B \circ f"] & \\
& B \arrow[rr, "\Phi_B"'] & & B^*
\end{tikzcd}$$

Note that the vertices of $C(A,B)$ are in correspondence with $B^{(0)} \times A^{(0)}$: the map labelled "collapse B" restricts to vertices as a projection onto the second factor, and p as a projection onto the first factor, as in Definition 3.4.14. The inclusion j is the map that takes an edge vw of A to the edge $(f(v),v)(f(w),w)$ of $B \times_{B^*} C(A,B^*) = C(A,B)$. The bold diagram in the centre contains the desired canonical completion and retraction.

Definition 3.4.17. The map j in the construction above is called the *canonical completion* of the map f. Sometimes we will also refer to the complex $C(A,B)$ by this name. The map r is called the *canonical retraction*.

3.5 Quasiconvexity in cube complexes

We have shown that a local isometry $f: A \to B$ between (fully) special cube complexes can be factored as a composition of an inclusion (the canonical completion) and a covering map. To deduce that $f_*(\pi_1(A))$ is a separable subgroup of $\pi_1(B)$, we need this covering map to be finite sheeted, and if B is compact this requires A to be compact. In the graph case this is easy: for any finitely generated subgroup $H \leq \pi_1(B)$, we can take a core graph of the covering space corresponding to H, so that the canonical completion is a finite graph and hence the covering space is finite sheeted.

In the cube complex case, it is not enough to merely take a compact core; the map we are to complete must be a map of cube complexes, so at the very least, the core must be a subcomplex. Moreover, since the canonical completion is only defined for local isometries, we require the core to be locally convex (otherwise, the restriction of the cover fails to be a local isometry). Rephrasing this (via Corollary 2.2.11) in terms of the universal cover $\widetilde{B}$, we wish to find a subcomplex Y of $\widetilde{B}$ which is *convex*, and invariant and cocompact under the action of H on $\widetilde{B}$. Then, H will be the fundamental group of the compact cube complex $A = Y/H$, and we can apply the canonical completion to show that H is separable.

This motivates the following definition.

Definition 3.5.1. Let G act geometrically on a CAT(0) cube complex X. A subgroup H of G is called *combinatorially convex cocompact* (or *CCC*) if there is a convex subcomplex Y on which H acts cocompactly.

However, there is a problem with this approach. The property of being CCC depends on the choice of complex $\widetilde{X}$; a priori, a subgroup $H \leq G$ might be CCC with respect to one action of G but not another. The following example illustrates this.

Example 3.5.2. Consider $\mathbb{Z} \times \mathbb{Z}$ acting in the standard way on $\mathbb{R}^2$ (the standard action being given by $(a,b) \cdot (x,y) = (x+a, b+y)$), with a cube complex structure given by unit square tiles and vertices at $\mathbb{Z} \times \mathbb{Z}$. The diagonal subgroup D is not CCC, since any invariant subcomplex must contain the diagonal line $y = x$, and no subcomplex containing this line is convex, except for $\mathbb{R}^2$ itself, upon which the action of D is not cocompact.

However, if we change the action to, say, $(a,b) \cdot (x,y) = (x+a-b, y+a+b)$, then D acts cocompactly on the y axis. Hence, with respect to this action, D is CCC.

It is not a coincidence that this is reminiscent of the ambiguity in defining a quasiconvex subgroup that we saw in Section 2.5.1, and indeed, restricting to hyperbolic groups again resolves the ambiguity.

Theorem 3.5.3. *If G is a hyperbolic group acting geometrically on a CAT(0) cube complex, then the quasiconvex subgroups of G coincide precisely with the CCC subgroups. In particular, the property of being CCC is independent of the choice of action on a* CAT(0) *cube complex.*

The full proof of the above theorem relies on the work of Haglund [Hag08], and we will give some more details in Section 3.5.1. In fact, we have already seen one direction of this theorem. Since a CCC subgroup corresponds to an isometrically embedded CAT(0) cube complex, applying Theorem 2.5.6 gives another proof that CCC subgroups are quasiconvex. However, the other direction involves a strengthening from the coarse notion of quasiconvexity to the full convexity in the definition of CCC, and this is where Haglund's work is required.

To summarise the importance of quasiconvex subgroups in hyperbolic cube complexes, we give the following restatement of Theorems 3.5.3 and 2.5.6.

Theorem 3.5.4. *Let X be a compact, non-positively curved cube complex with hyperbolic fundamental group G. Then a subgroup H of G is quasiconvex if and only if there exists a compact cube complex Y equipped with a local isometry $f\colon Y \to X$, such that f induces the inclusion $H \hookrightarrow G$ of fundamental groups.*

3.5.1 Convex subcomplexes of CAT(0) cube complexes

In this section, we recall from Haglund's paper [Hag08] some results concerning convex subcomplexes of CAT(0) cube complexes. These will play an important role in Chapter 4, but their main relevance to the current chapter is to give some insight into the proof of Theorem 3.5.3. Given a CAT(0) cube complex $\widetilde{X}$, the main idea behind proving the results in this section is to disregard the intrinsic metric on the cube complex coming from the Euclidean metrics on each cube (as in Section 2.2.4), and instead to consider the combinatorial distance function on $\widetilde{X}^{(0)}$ given by length of edge paths in $\widetilde{X}$.

First, recall that since CAT(0) cube complexes are special (Example 3.3.13), hyperplanes are 2-sided, and hence removing an open hyperplane neighbourhood $\mathring{U}_H$ divides $\widetilde{X}$ into two disjoint complexes, each equipped with a natural inclusion map to $\widetilde{X}$.

Definition 3.5.5. If H is a hyperplane in a CAT(0) cube complex $\widetilde{X}$, the two connected components of $\widetilde{X} - \mathring{U}_H$ are called the *halfspaces* corresponding to H.

It is straightforward to show that halfspaces, as well as hyperplane neighbourhoods, are convex subcomplexes (see Lemma 3.6.1 and [Hag08]).

Theorem 3.5.6 ([Hag08, Proposition 2.17])**.** *If C is a convex subcomplex of a* CAT(0) *cube complex $\widetilde{X}$, then C coincides with the intersection of halfspaces containing C.*

Haglund also introduces the following coarser notion of convexity:

Definition 3.5.7. Let $\widetilde{X}$ be a CAT(0) cube complex. A *combinatorial geodesic* between two vertices x and y is a path in the 1-skeleton $\widetilde{X}^{(1)}$ of minimal length. A subcomplex $C \subset \widetilde{X}$ is called *combinatorially quasiconvex* if there exists $K \geq 0$ such that, for any two vertices $x, y \in C$, any combinatorial geodesic connecting x and y is contained within the K-neighbourhood of C. If G is a group acting geometrically on $\widetilde{X}$, then a subgroup $H \subset G$ is called *combinatorially quasiconvex* if there exists a point $x \in \widetilde{X}$ such that the orbit $H \cdot x$ is combinatorially quasiconvex.

Remark 3.5.8. Combinatorial quasiconvexity is independent of the choice of basepoint, and moreover it is equivalent to ordinary quasiconvexity if G is a hyperbolic group (equivalently, if $\widetilde{X}$ is a hyperbolic space), since quasi-geodesics are close to geodesics (see [BH99, III.H.1.7]).

Lemma 3.5.9 ([Hag08, Lemma 2.25])**.** *Let Y and C be two subcomplexes of a* CAT(0) *cube complex $\widetilde{X}$, such that Y is combinatorially quasiconvex and C is contained within the R-neighbourhood of Y for some R. Then C is combinatorially quasiconvex.*

Moreover, the following result gives a strong relationship between combinatorial quasiconvexity and full convexity. The *convex hull* of a subset C is the intersection of all halfspaces containing C; by Theorem 3.5.6 this is equivalent to the usual notion of convex hull.

Theorem 3.5.10 ([Hag08, Theorem 2.28])**.** *Let $\widetilde{X}$ be a uniformly locally finite* CAT(0) *cube complex, and let C be a combinatorially quasiconvex subcomplex. Then there exists R such that the convex hull of C is contained within the R-neighbourhood of C.*

In the context of groups acting on CAT(0) cube complexes, we have the following.

Corollary 3.5.11 ([Hag08, Corollary 2.29])**.** *Let $\widetilde{X}$ be a* CAT(0) *cube complex, let G be a group acting geometrically on $\widetilde{X}$, and let $H \subset G$ be a subgroup. Then H is combinatorially quasiconvex if and only if H is CCC.*

Together with Remark 3.5.8, Corollary 3.5.11 implies Theorem 3.5.3.

3.5.2 Separability of quasiconvex subgroups

We are now ready to state the following theorem.

Theorem 3.5.12. *Let X be a compact, fully special cube complex. Let $G = \pi_1(X)$. Then every combinatorially convex cocompact (with respect to the universal covering action) subgroup H of G is separable.*

Proof. Since H is CCC, we may find a convex subcomplex Y of $\widetilde{X}$ on which H acts cocompactly. Let A be the compact cube complex Y/H. There is a well-defined combinatorial map $f\colon A \to X$, mapping $\pi_1(A)$ isomorphically onto H. Since Y is convex, it follows that f is a local isometry. We may therefore apply the canonical completion to f, as in the following diagram:

$$\begin{tikzcd}
Y \arrow[rr, hook] \arrow[dd, "q"'] & & \widetilde{X} \arrow[dd] \\
& C(A,X) \arrow[dl, bend left, "r"] \arrow[dr, "p"] & \\
Y/H = A \arrow[ur, bend left, "j"] \arrow[rr, "f"] & & X
\end{tikzcd}$$

Since G is residually finite (Corollary 3.3.16), p is a finite sheeted cover and r is a retract, we may apply Lemma 3.4.8 to $p_*(j_*(H)) \le p_*(\pi_1(C(A,X))) \le G$ to see that $H = \pi_1(A)$ is separable in G, as required. $\square$

Remark 3.5.13. The restriction to fully special complexes does not restrict the class of groups: taking the first subdivision (where we divide each edge in two, each square into four and so on) of any special cube complex gives one which is fully special (and has simplicial 1-skeleton). Moreover, by Proposition 3.4.5, it is sufficient to prove subgroup separability by showing that it holds in a finite index subgroup, and hence we obtain:

Corollary 3.5.14. *Let X be a virtually compact special cube complex with fundamental group G. Then every combinatorially convex cocompact (with respect to the universal covering action) subgroup of G is separable.*

Finally we can use Theorem 3.5.3 to deduce:

Corollary 3.5.15. *Let X be a virtually compact special cube complex with δ-hyperbolic fundamental group G. Then every quasiconvex subgroup of G is separable.*

Haglund and Wise [HW08] also prove a converse to this:

Theorem 3.5.16. *Let X be a compact non-positively curved cube complex with δ-hyperbolic fundamental group G. If every quasiconvex subgroup of G is separable, then X is virtually special.*

The proof is postponed to the following section.

3.6 Hyperplane subgroups

We have shown that quasiconvex (or CCC) subgroups of virtually compact special groups are separable, but for this to be useful, we would like to be able to apply it to some explicit subgroups. Recall that when $G = \pi_1(X)$ acts on $\widetilde{X}$, a subgroup is CCC if it preserves a convex subcomplex Z. We have already introduced one type of subcomplex: the regular neighbourhood of a hyperplane. Lemma 3.6.1 implies that, in the universal cover, this subcomplex will be convex, and hence hyperplanes will provide explicit examples of CCC subgroups. It is also an important ingredient in the proof of the results in Section 3.5.1.

Lemma 3.6.1. *Let H be a hyperplane in a non-positively curved cube complex X. The map χ_H is a local isometry.*

Proof. In the two-sided case, there is a map $\chi_H\colon N_H = H \times [-1,1] \to X$, whose image is U_H. Each vertex v of $H \times [-1,1]$ has exactly one edge e_v dual to H (that is, the fibre over a vertex of H). Pick two non-adjacent vertices a, b of the link of some vertex $v \in H \times [-1,1]$. Neither corresponds to the edge e_v, since this edge bounds a square with all the other edges at v. The vertices $\chi_H(a), \chi_H(b) \in \mathrm{Lk}(\chi_H(v))$ are both adjacent to the vertex $s \in \mathrm{Lk}(\chi_H(v))$ corresponding to $\chi_H(e_v)$. Hence if $\chi_H(a)$ and $\chi_H(b)$ are adjacent, then s, $\chi_H(a)$ and $\chi_H(b)$ form a 3-cycle. Since X is non-positively curved, the corresponding three squares bound a 3-cube, which is in U_H since one of its edges is $\chi_H(e_v)$. But this cube must be the image of a cube in $H \times [-1,1]$ spanned by e_v and the edges corresponding to a and b, which is a contradiction as a and b are not adjacent. Hence, $\chi_H(a)$ and $\chi_H(b)$ are not adjacent, and χ_H is a local isometry by Theorem 2.3.12. The one-sided case is identical (since the twisted interval bundle is locally isometric to $H \times [-1,1]$). □

Remark 3.6.2. The hypothesis that χ_H is injective is equivalent to assuming the hyperplane does not self-intersect or self-osculate, and so it will always be satisfied in the (fully) special case. Note, however, that we did not need this assumption to show that χ_H is a local isometry. This is a subtle distinction: all hyperplane neighbourhoods U_H are the image of a local isometry $\chi_H \colon N_H \to U_H$, but this is not always the same map as the inclusion $U_H \hookrightarrow X$, and so not all hyperplane neighbourhoods are locally convex as subcomplexes.

The bundle N_H over H has the same fundamental group as H, and so by Corollary 2.2.11, χ_H induces an injection $\pi_1(H) \to \pi_1(X)$. This holds for any hyperplane, though when χ_H is non-injective, the subgroup will not necessarily be the same as the one induced by the inclusion $U_H \hookrightarrow X$.

Definition 3.6.3. For each hyperplane H, the subgroup $\chi_{H*}(\pi_1(H)) \leq \pi_1(X)$ is called the *hyperplane subgroup* of H.

Lemma 3.6.4. *Hyperplane subgroups of compact non-positively curved cube complexes are CCC.*

Proof. Let H be a hyperplane in a compact non-positively curved cube complex X. The map $\chi_H \colon N_H \to X$ is a local isometry by Lemma 3.6, and so by Corollary 2.2.11 it lifts to an isometry onto a convex subcomplex of $\widetilde{X}$, upon which the hyperplane subgroup acts cocompactly (since the quotient is the compact cube complex U_H). □

Applying Corollary 3.5.14, we obtain:

Corollary 3.6.5. *Hyperplane subgroups of compact virtually special cube complexes are separable.*

As in the previous section, there is a converse to this:

Theorem 3.6.6. *Let X be a compact non-positively curved cube complex. Suppose all the hyperplane subgroups $\pi_1(H)$ are separable, and all the hyperplane double cosets $\pi_1(H)g\pi_1(K)$ are separable for intersecting pairs H, K of hyperplanes. Then X is virtually special.*

Idea of proof. We use the characterisation of separability from Theorem 3.4.9. For each hyperplane H of X, we can find an index 2 cover in which H is two-sided. Then we may find a further finite cover (say X^H) in which χ_H lifts to an embedding, which rules out self-intersection and self-osculation. Since X is compact there are only finitely many hyperplanes, so we can find a finite cover factoring through X^H for every hyperplane H.

To rule out interosculation uses the double coset separability hypothesis, which allows us to find a finite cover in which pairs of interosculating hyperplanes are lifted to pairs of hyperplanes which cross, but do not osculate. For more details on this step, see [HW08]. □

The following theorem is due to Minasyan [Min06].

Theorem 3.6.7. *If G is a hyperbolic group in which all quasiconvex subgroups are separable, then all double cosets of quasiconvex subgroups are separable.*

Proof of Theorem 3.5.16. X is a compact non-positively curved cube complex with δ-hyperbolic fundamental group G, in which every quasiconvex subgroup of G is separable. Hyperplane subgroups are quasiconvex by Lemma 3.6.4, and Theorem 3.5.3. Hence by Theorem 3.6.7, all hyperplane double cosets are separable. Then by Theorem 3.6.6, X is virtually special. □

3.7 VH complexes

Before the definition of special cube complexes, Wise had studied the following class of square complexes (see [Wis02]).

Definition 3.7.1. A *VH complex* is a square complex with two distinct families of edges: *vertical* edges and *horizontal* edges, such that the boundary of each square alternates between horizontal and vertical edges.

It follows immediately from Definition 3.7.1 that VH complexes also have two families of hyperplanes: *vertical* hyperplanes which are dual only to horizontal edges, and *horizontal* hyperplanes which are dual only to vertical edges. After either taking a double cover or subdividing (for example, in the case of a Möbius strip [Wis06]), all hyperplanes are two-sided.

Such a VH complex is called *non-singular*, and it has a natural structure as a graph of spaces. Vertex spaces correspond to the connected components of the vertical 1-skeleton, and edge spaces correspond to vertical hyperplanes. Both edge and vertex spaces are graphs. If the VH complex is simple, then it is non-positively curved (by Remark 2.2.21), and the attaching maps in the graph of spaces are immersions. Such a graph of spaces is also called a *graph of graphs*, and these are the main objects of study in Chapter 4 (see Definition 4.2.1).

Definition 3.7.2. A simple, non-singular VH complex is called *thin* if the corresponding graph of spaces is thin in the sense of Definition 2.1.5. It is called *clean* if the attaching maps are embeddings.

The main result of [Wis02] is:

Theorem 3.7.3. *Let X be a compact, thin VH complex. Then X has a finite sheeted cover which is clean.*

Note that hyperplanes of a non-singular VH complex cannot self-intersect, and cleanness rules out direct self osculation (at least for vertical hyperplanes). This suggests the following, which was proved in [HW08].

Proposition 3.7.4. *A compact, virtually clean VH complex X is virtually special.*

Combining Theorem 3.7.3 and Proposition 3.7.4, we obtain:

Theorem 3.7.5. *Any compact, thin VH complex is virtually special.*

Remark 3.7.6. In a sense, Theorem 3.7.3 is a separability theorem; it says that the immersed edge spaces in X can all be lifted to embeddings in a finite sheeted cover. The separability theorems for special cube complexes may be thought of as a higher dimensional generalisation, and indeed this was the original motivation for the definition of a special cube complex (as mentioned earlier). Indeed, by Theorem 3.7.5, Corollary 3.6.5 is an exact generalisation of Theorem 3.7.3.

3.8 Further developments

We have described some of the foundational results concerning special cube complexes, as originally developed in [HW08]. However, it is beyond the scope of this thesis to explain in detail those applications which have been of the most recent interest. In this final section, we will give a brief introduction to two important, but more difficult, aspects of the theory: Wise's Hierarchy Theorems for special cube complexes, and the applications to 3-manifold topology.

3.8.1 Hierarchies

A good reason to look for embedded, codimension 1 subcomplexes of any complex is that we can cut the complex along such a subcomplex, simplifying it, and elucidating its structure. Via the Seifert–van Kampen theorem, we can also elucidate the structure of the fundamental group. Hyperplanes in a special cube complex are particularly well suited to this goal.

If X is a compact fully special cube complex with a hyperplane H, then removing the open cubical neighbourhood $\mathring{U}_H$ gives a cube complex with either one component Y, or two components Y and Z. By the Seifert–van Kampen theorem, $\pi_1(X)$ splits as either an HNN extension $\pi_1(Y)*_C$ or an amalgamated product $\pi_1(Y) *_C \pi_1(Z)$, where $C = \pi_1(H)$ is the hyperplane subgroup corresponding to H. Since the total number of cells decreases at each stage, we may continue until all components are simply connected (for example, they may be individual edges or cubes).

Definition 3.8.1. The class $\mathcal{QH}$ of groups with a *quasiconvex hierarchy* is defined to be the smallest class such that:

1. $1 \in \mathcal{QH}$
2. If $G = A *_C B$, C is a quasiconvex subgroup of G, and $A, B \in \mathcal{QH}$, then $G \in \mathcal{QH}$.
3. If $G = A*_C$, C is a quasiconvex subgroup of G, and $A \in \mathcal{QH}$, then $G \in \mathcal{QH}$

Theorem 3.8.2. *If X is a compact special cube complex and $G = \pi_1(X)$ is δ-hyperbolic, then $G \in \mathcal{QH}$.*

Proof. By Remark 3.5.13, we can assume without loss of generality that X is fully special. For any hyperplane H, we need to show that $X - \mathring{U}_H$ is fully special. The hyperplane pathologies are easily ruled out, because they do not occur in X. To see that $X - \mathring{U}_H$ is non-positively curved, note that when removing the open neighbourhood $\mathring{U}_H$, for each vertex v in the image of ∂N_H, there is precisely one vertex removed from $\mathrm{Lk}(v)$ (as X is fully special), and every simplex which is removed has v as a vertex. Thus $\mathrm{Lk}(v)$ remains flag, as required. Hence, by the Seifert–van Kampen theorem, we can decompose $G = \pi_1(X)$ as $G = A *_C B$ or $G = A*_C$, where $C = \pi_1(H)$. Quasiconvexity of $C < G$ follows from the fact that hyperplane subgroups are quasiconvex in the hyperbolic case (Lemma 3.6.4 and Theorem 3.5.3). The result then follows by induction on the number of cubes. □

In fact, we may strengthen this theorem using malnormality (see Section 2.5.2). We need the following lemma.

Lemma 3.8.3. *Let X be a compact fully special cube complex with hyperbolic fundamental group G. Then there is a finite cover $\hat{X}$ in which all hyperplane subgroups are malnormal.*

Proof. Let $H = \pi_1(Y)$ be a hyperplane subgroup of G. Then $H < G$ is quasiconvex (Lemma 3.6.4 and Theorem 3.5.3), and hence by Theorem 2.5.11, it has finite width. Let $g_1^{-1}Hg_1, \ldots, g_k^{-1}Hg_k$ be a maximal set of distinct conjugates for H. We can separate H from each g_i in a finite index subgroup, and hence (by taking the intersection of these), we can separate H from the set of all g_i in a finite index subgroup H'. As H' contains H and not g_i, it has empty intersection with the coset Hg_i, and hence no element of H' conjugates H to a subgroup with infinite intersection with H. That is, H is an almost malnormal subgroup of H'. Since G is torsion-free by Lemma 2.4.20, H is in fact a malnormal subgroup of H'.

For each hyperplane subgroup $H_i = \pi_1(Y_i)$, form the finite-index subgroup H_i' as above, and consider the intersection K' of all H_i'. This is finite index. Moreover, since a finitely generated group contains finitely many subgroups of a given index, the intersection of K' with all its conjugates is a finite index normal subgroup $K \leq G$.

Now consider the finite sheeted regular covering space $X^K \to X$ corresponding to K. This factors through the cover $X^{H_i'} \to X$ corresponding to H_i' for every i. Let $D_i = \pi_1(Z_i)$ be a hyperplane subgroup of K, where Z_i projects to Y_i. Then, by normality of K, D_i is conjugate (in G) to $K_i = H_i \cap K$. If K_i nontrivially intersects a conjugate $k^{-1}K_ik$ for $k \in K$, then $k^{-1}H_ik$ nontrivially intersects H_i. As $K < H'$, this contradicts malnormality of H_i in H_i' unless $k \in K_i$. Therefore, K_i is a malnormal subgroup of K, and since D_i is conjugate to K_i, D_i is also malnormal in K. Therefore, $\hat{X} = X^K$ is our required finite covering space. □

We say a group is in $\mathcal{MQH}$ (or has a *malnormal quasiconvex hierarchy*) if it is in $\mathcal{QH}$ with the additional hypothesis that the edge groups C are malnormal. Then, by combining Theorem 3.8.2 with Lemma 3.8.3, we immediately obtain:

Corollary 3.8.4. *Let X be a virtually compact special cube complex with δ-hyperbolic fundamental group G. Then G has a finite index subgroup with a malnormal quasiconvex hierarchy.*

Corollary 3.8.4 is very useful, because it gives us a strong basis for inductive proofs: if we can prove a property holds for the trivial group, and that it is stable under taking malnormal quasiconvex amalgams and HNN extensions, then it will virtually hold for any virtually special group. In Section 4.4.1, we will describe a problem which we hope can be tackled using this approach. However, Theorem 3.8.2 and Corollary 3.8.4 are far from the pinnacle of achievement in the theory of special cube complexes. Remarkably, both results have a converse. In particular, we have:

Theorem 3.8.5. *Let G be a δ-hyperbolic group with an (almost) malnormal quasiconvex hierarchy. Then G is virtually the fundamental group of a compact special cube complex.*

Theorem 3.8.6. *Let G be a δ-hyperbolic group with a quasiconvex hierarchy. Then G is virtually the fundamental group of a compact special cube complex.*

Theorem 3.8.5 requires two principal ingredients: firstly, we must find a cube complex structure on the space given by amalgamating two virtually special cube complexes along a subspace corresponding to a malnormal, quasiconvex subgroup (see [HW15]); and secondly, we must ensure that this cube complex structure can itself be taken to be virtually special (see [HW12]). Finally, Theorem 3.8.6 is deduced from Theorem 3.8.5 using a deep theorem called the Malnormal Special Quotient Theorem, described in [Wis12a] and [Wis12b].

3.8.2 3-manifolds

One of the most significant applications of special cube complexes has been towards the theory of three-dimensional manifolds. As we discussed in Chapter 1, much of the research into 3-manifolds over the last few decades has been guided by an article of Thurston [Thu82], in which he posed a list of twenty-four questions, including the (now proven) Geometrization Conjecture. As a consequence of this, combined with the JSJ decomposition theorem, any prime 3-manifold is known to be built from pieces which are either Seifert-fibred or hyperbolic. Seifert-fibred manifolds are circle bundles over 2-dimensional orbifolds, and their structure is reasonably well understood. Hence, after Perelman's proof of geometrization, the remaining task was to understand hyperbolic 3-manifolds—those whose universal cover is $\mathbb{H}^3$.

Two in particular of Thurston's questions remained unanswered in early 2012, and these both concerned hyperbolic manifolds. Recall that an embedded surface Σ in a 3-manifold

M is called *incompressible* if it is not a disk or a 2-sphere, and it is π_1-injective. A *Haken* 3-manifold is one which contains an incompressible surface. This is a useful notion because, as we saw with cube complexes, it enables us to cut along the surface and decompose the manifold into smaller pieces, thus obtaining a *Haken hierarchy* for M. A 3-manifold is called *fibred* if it is a surface bundle over a circle. Thurston's questions are the following:

Question 3.8.7. Is every closed hyperbolic 3-manifold virtually Haken?

Question 3.8.8. Is every closed hyperbolic 3-manifold virtually fibred?

The motivation for Wise's work towards Theorem 3.8.6 was to settle Question 3.8.8 in the case of Haken manifolds. In particular, if M is Haken, its Haken hierarchy is quasiconvex, and so it is virtually compact special by Theorem 3.8.6. By a theorem of Agol (see [Ago08]), it is then virtually fibred. However, the true culmination of the theory has been to apply these ideas to non-Haken manifolds, which do not automatically possess a suitable hierarchy. In [KM12], Kahn and Markovic showed that any hyperbolic 3-manifold M contains a suitably large set of immersed, quasiconvex surfaces, and Bergeron and Wise [BW12] were able to apply Sageev's construction (see Section 3.1.1) to find a geometric action of $G = \pi_1(M)$ on a CAT(0) cube complex. Agol then proved the following remarkable theorem (see [Ago13]), significantly strengthening Theorem 3.5.16:

Theorem 3.8.9. *Let G be a δ-hyperbolic group acting geometrically on a CAT(0) cube complex. Then G is virtually compact special.*

In particular, $G = \pi_1(M)$ is virtually compact special for any hyperbolic 3-manifold, and hence M is virtually fibred by Agol's previous theorem [Ago08]. To see that M is virtually Haken, we may argue as follows. Take a special cube complex X corresponding to a finite index subgroup of G, and take the finite cover M' of M corresponding to the same subgroup. The group isomorphism between $\pi_1(M')$ and $\pi_1(X)$ can be realised by a continuous map $f\colon M' \to X$, since the spaces are both aspherical; moreover, this map can be homotoped so that it is transverse to the hyperplanes of X. Now, the preimage of any codimension 1 hyperplane H of X is an embedded surface $f^{-1}(H) = \Sigma_H$ in M'. If Σ_H is compressible, we may repeatedly homotope f and decrease the genus of Σ_H; this terminates either with an incompressible surface, or a collection of spheres (which we can remove by a further homotopy of f because M' is aspherical). In the former case, M' is Haken, as required. In the latter case,

$f(M')$ does not intersect H, and so we can cut X along H and repeat the argument with another hyperplane; since X is compact, we will eventually reach a hyperplane whose preimage is an incompressible surface, or else M' is trivial. Therefore, M' is Haken in this case too, and M is virtually Haken. Indeed, we could now deduce from Wise's answer to Question 3.8.8 in the Haken case that M is also virtually fibred, as an alternative to using Agol's result [Ago08] that virtually compact special cube complexes are virtually fibred. Both of Thurston's questions are thus answered in the affirmative.

Chapter 4

Folding cube complexes

We have already seen that an essential ingredient in the study of negatively curved groups is to understand their quasiconvex subgroups. As well as hyperbolic groups in general (see Theorem 2.5.4), we have seen this in particular in the case of hyperbolic cube complexes. Theorem 3.5.4 summarises exactly why they are so important: a local isometry from a compact cube complex to a non-positively curved cube complex with hyperbolic fundamental group corresponds exactly to the inclusion of a quasiconvex subgroup.

The main purpose of this chapter is to devise a *folding algorithm* that allows us to explicitly construct a map of complexes corresponding to the inclusion of some subgroup, and hence check whether or not it is quasiconvex. That is, given a non-positively curved cube complex X, with $G = \pi_1(X)$ hyperbolic and $H \subset G$ quasiconvex, we can construct a cube complex Y equipped with a local isometry $Y \to X$ which induces the inclusion $H \hookrightarrow G$ on fundamental group. The main theorem describing this algorithm is Theorem 4.3.9. Understanding the geometry of the complex Y could therefore be used to show that certain families of subgroups are quasiconvex, and as we discussed in Chapter 1, an important open question (Question 1.0.2) asks whether this holds for malnormal subgroups. This problem was our original motivation for describing the algorithm, and we discuss a potential application to it in Section 4.4.

Folding was originally introduced by Stallings in the context of graphs, and we will begin by describing the procedure in this case. Our main folding algorithm will then be described in terms of a graph of groups, inspired by the graph of groups description of special cube complexes given in Wise's Hierarchy Theorem (Theorem 3.8.2). There have been other accounts of folding for a graph of groups [KWM05, Dun98, BF91], but we exploit the explicit geometric structure of cube complexes to give a more geometric formulation than previous authors.

We will focus on the two dimensional case for technical simplicity, but the higher dimensional generalisation is conceptually straightforward, and we shall discuss it in Section 4.5. More recently, Beeker and Lazarovich [BL16] have also described a folding algorithm for general CAT(0) cube complexes, which takes place on the wallspace level (see Section 3.1.1). Our algorithm can be thought of as a more explicit version of theirs in the two dimensional case.

4.1 Folding for graphs

The fundamental tool used throughout this chapter is that of *folding* graphs. First made explicit by Stallings [Sta83], folding gives a way to factor any map between graphs as a product of elementary maps known as *folds*, and an immersion.

Let X be a graph, and let e_1 and e_2 be distinct edges such that $e_1 \neq \bar{e}_2$ and $\iota(e_1) = \iota(e_2) = v$. Consider the quotient of X obtained by identifying the edges e_1 and e_2, and the vertices $\tau(e_1)$ and $\tau(e_2)$. This gives a graph X' with one less edge than X, and one less vertex in the case where $\tau(e_1) \neq \tau(e_2)$. The map $X \to X'$ is called a *fold*. Sometimes we say X' is obtained by *folding together* e_1 and e_2.

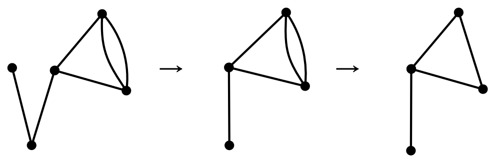

Figure 4.1: Two folds on a graph. Only the first is a homotopy equivalence.

Remark 4.1.1. In the notation above, a fold $X \to X'$ is a homotopy equivalence if and only if $\tau(e_1) \neq \tau(e_2)$.

The main theorem obtained by Stallings [Sta83] is:

Theorem 4.1.2. *Any homomorphism of finite graphs factors as a product of finitely many folds and an immersion.*

Given a homomorphism $\varphi\colon Y \to X$ between graphs, it is not difficult to see how one constructs a sequence of folds as in Theorem 4.1.2. If φ is an immersion, there is nothing to prove, so assume that φ fails to be injective on the link of some $v \in Y$. This means there are two edges e_1 and e_2 at v both mapping to the same edge e of X, and hence φ factors through the graph Y' obtained by folding together e_1 and e_2. Since Y' has one fewer edge than Y, we can repeat this process only finitely many times.

Definition 4.1.3. Suppose we have two graphs A, B, with immersions $\alpha\colon A \to X$, $\beta\colon B \to X$, and chosen basepoints $a \in A$, $b \in B$ mapping to the same point x in X. The wedge sum $W_0 = A \vee_{a=b} B$ is naturally equipped with a map ψ_0 to X (restricting to α and β on A and B respectively), and we may apply Theorem 4.1.2 to obtain a graph Z equipped with an immersion ψ to X.

We say Z is obtained by *folding together A and B at a and b, with respect to the maps α and β*. There are natural maps (induced by the inclusion to the wedge product) $i_A\colon A \to Z$ and $i_B\colon B \to Z$ such that $\psi \circ i_A = \alpha$, $\psi \circ i_B = \beta$. These maps are necessarily immersions but are not in general injective.

Lemma 4.1.4. *Suppose Z is obtained by folding together graphs A and B, where B is a tree. Then there exists a homotopy retraction $Z \twoheadrightarrow A$.*

Proof. The natural map $i_A\colon A \to Z$ is an immersion, hence π_1-injective. However, it is equal to the composition of a π_1-isomorphism $A \to W_0$ (in the notation above) with a π_1-surjection $W_0 \to Z$. Hence i_A is a π_1-isomorphism. It follows that i_A is injective, and that Z retracts onto $i_A(A) = A$. □

The following lemma will be useful in Section 4.3.

Lemma 4.1.5. *Let A, B, C and X be graphs with immersions α, β, γ between them as in the diagram below. Let $a \in A$ and $b \in B$ satisfy $\alpha(a) = \beta(b)$. Let Z be obtained by folding together A and B at a and b with respect to the maps $\gamma \circ \alpha$ and $\gamma \circ \beta$:*

$$\begin{tikzcd}
A \arrow[d, "i_A"'] \arrow[dr, "\alpha"] \arrow[drr, bend left, "\gamma\circ\alpha"] & & \\
Z \arrow[r, red, "\varphi"'] \arrow[rr, bend left, "\psi"] & C \arrow[r, "\gamma"'] & X \\
B \arrow[u, "i_B"] \arrow[ur, "\beta"'] \arrow[urr, bend right, "\gamma\circ\beta"'] & &
\end{tikzcd}$$

Then there exists an immersion $\varphi\colon Z \to C$, as shown, such that the diagram commutes. Hence, by uniqueness of Stallings folding, Z is the graph obtained by folding together A and B at a and b with respect to the maps α and β.

Proof. Recall Z is obtained by applying folding moves to the wedge product $W_0 = A \vee_{a=b} B$, to make the map $\psi_0 \colon W_0 \to X$ restricting to $\gamma \circ \alpha$ on A and $\gamma \circ \beta$ on B an immersion. Since $\alpha(a) = \beta(b)$, there is a well-defined map $\varphi_0 \colon W_0 \to C$ defined by α on A and β on B. The diagram below commutes for $j = 0$ (we omit the indices on the maps i_A, i_B). Now let W_j be the graph obtained after j elementary folds on W_0, and let ψ_j be the corresponding map to Z. Assume for induction that $\varphi_j \colon W_j \to C$ is defined and makes the diagram below commute:

$$\begin{tikzcd}
A \arrow[d, "i_A"'] \arrow[dr, "\alpha"] \arrow[drr, bend left, "\gamma\circ\alpha"] & & \\
W_j \arrow[r, "\varphi_j"'] \arrow[rr, bend left, "\psi_j"] & C \arrow[r, "\gamma"'] & X \\
B \arrow[u, "i_B"] \arrow[ur, "\beta"'] \arrow[urr, bend right, "\gamma\circ\beta"'] & &
\end{tikzcd}$$

If $W_j \neq Z$ then ψ_j is not an immersion, so there are two coincident edges e, e' mapping under ψ_j to the same edge in X. Since γ is an immersion and the diagram commutes, this happens if and only if they map under φ_j to the same edge in C. Hence φ_{j+1} is defined on the graph obtained by folding together e and e'. We see that the process of folding together A and B with respect to the maps α and β is identical to the process with respect to the maps $\gamma \circ \alpha$ and $\gamma \circ \beta$; in particular, ψ_j is an immersion iff φ_j is an immersion. This completes the proof. $\square$

Remark 4.1.6. Categorically, Lemma 4.1.5 says precisely that Z is the pushout of the diagram

$$\begin{tikzcd}
Z & B \arrow[l, "i_B"'] \\
A \arrow[u, "i_A"] & * \arrow[l, hook'] \arrow[u, hook]
\end{tikzcd}$$

in the category of graphs equipped with immersions to X, where $*$ maps to $a \in A$ and $b \in B$, as described in [Sta83]. It can also be formed in the category of graphs as the pushout of $A \leftarrow P \to B$, where P is the connected component of the pullback of $A \to X \leftarrow B$ which contains (a, b).

Remark 4.1.7. We will usually apply Lemma 4.1.5 in the following context. Suppose there is another graph C' such that the given immersions from A and B to X factor through C' as well as C; that is, there exist $\alpha' \colon A \to C'$, $\beta' \colon B \to C'$, $\gamma' \colon C' \to X$ such that $\gamma \circ \alpha = \gamma' \circ \alpha'$, $\gamma \circ \beta = \gamma' \circ \beta'$.

Then the graph obtained by folding together A and B with respect to the maps α, β is the same as the graph obtained by folding together A and B with respect to the maps α', β'.

4.2 Graphs of graphs

Definition 4.2.1. A graph of spaces in which all the vertex and edge spaces are graphs, and the attaching maps are combinatorial immersions of graphs, is called a *graph of graphs*.

As we discussed in Section 3.7, graphs of graphs are precisely simple non-singular VH complexes. The vertical edges are edges of vertex spaces of Y, and the horizontal edges are edges of the form $\{a\} \times I$ contained in the mapped-in copy of $Y_e \times I$ for some edge space Y_e. We denote such an edge $\{a\} \times I$ by $[a]$. The boundary of each square alternates between horizontal and vertical edges.

Remark 4.2.2. A priori, it is not clear that any graph of free groups can be realised by a graph of graphs, since choosing arbitrary graphs for the vertex spaces, it may not be possible to choose edge spaces equipped with immersions to the vertex spaces on either side.

In light of the fact that thinness of a compact graph of graphs is enough to imply virtual specialness (Theorem 3.7.5), we will assume throughout this chapter that graphs of graphs are thin. Moreover, thinness can be assumed in the hierarchy for special cube complexes (Theorem 3.8.2), and so it is a reasonable assumption for the higher dimensional case too (see Section 4.5).

Since our algorithm will be designed to determine quasiconvexity, and this is stable under taking the intersection with a finite index subgroup, we may also exploit Proposition 3.7.4 and assume that our graphs of graphs are clean (see Remark 4.3.22). The geometric interpretation of these two properties is contained in the following lemma.

Lemma 4.2.3. *Let X be a clean, thin graph of graphs. Then the intersection of the images of two edge spaces in any vertex space X_v is simply connected.*

Proof. Suppose not. Then there is an embedded loop in X_v in the image of two edge spaces. This corresponds (up to conjugacy) to an element of $\pi_1(X_v)$ in the intersection of two edge subgroups, which is infinite order as $\pi_1(X_v)$ is free. This contradicts thinness. □

The following technical lemma will be important later. Note that if A is a graph and B is a subgraph, $A - B$ denotes the smallest subgraph of A containing only those edges which are not edges of B.

Lemma 4.2.4. *Let A be a connected graph, and let $k > 0$. Suppose there is a family of connected subgraphs $B_1, \dots, B_n$ such that $\bigcup B_i = A$, and for each $i \neq j$, the intersection $B_i \cap B_j$ is a forest, each of whose components has diameter at most k. Furthermore, assume that there exists a core graph C for A which intersects every B_i. For each i, denote by B_i^{unique} the (possibly disconnected or empty) subgraph $B_i - \bigcup_{j \neq i} B_j$ obtained by deleting from B_i the intersections with each other B_j. Then each subgraph B_i is contained in the k-neighbourhood of $C \cup B_i^{unique}$.*

Proof. $A - C$ is a forest. Denote its components by $T_1, \dots, T_m$. In each T_j, there is a single vertex v_j which intersects C. It is enough to show that for each j, $T_j - N_k(v_j)$ intersects exactly one B_i.

Since B_i is connected and C is a core graph, $B_i \cap T_j$ is connected for every i and j. Therefore, if B_i and B_k intersect inside T_j, this intersection is connected and hence has diameter at most k. Since C intersects both B_i and B_k, and these are connected, it follows that $v_j \in B_i \cap B_k \cap T_j$. Therefore, any point of T_j outside the k-neighbourhood of v_j is contained in at most one (and hence, exactly one) B_i. □

4.2.1 Lollipops

The following simple class of subgraphs will be used in our folding algorithm.

Definition 4.2.5. Let Γ be a graph with a vertex a. A *lollipop of length n based at a* is an embedding $\lambda \colon \ell \to \Gamma$ of a graph ℓ consisting of a path of length k (where $0 \leq k < \frac{n}{2}$) from a vertex α to a vertex α', together with a cycle of length $n - 2k$ through α', such that $\lambda(\alpha) = a$. See Figure 4.2. We often refer to the graph ℓ, or its image $\lambda(\ell)$, as a lollipop.

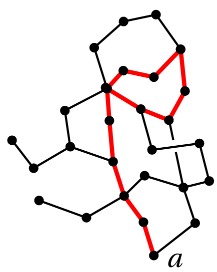

Figure 4.2: A lollipop of length 17 based at a.

It will be crucial in our applications that lollipops are *embedded* into graphs. However, this is not a serious limitation, since any non-embedded lollipop may easily be replaced with a union of embedded ones. Indeed, any immersed based loop in a graph is homotopic to a concatenation of embedded lollipops with the same basepoint. In particular, homotopy classes of embedded lollipops generate the fundamental group of a connected graph. The union of generating lollipops is the core of the graph, and if we then add in a tree consisting of a path from the basepoint to each leaf, the union is the whole graph.

4.3 Folding graphs of graphs

In order to develop a notion of folding in the setting of graphs of graphs, we must define some elementary folds. As in the graph case, these are all designed to correct a local failure of local isometry for a morphism of graphs of graphs. There are at least two notions for such a morphism, and one such notion is given in Definition 4.3.1. We will, in fact, use a more restrictive notion when we develop the folding algorithm; see Remark 4.3.3.

Definition 4.3.1. Let Y and X be graphs of graphs. A *morphism* is a continuous map $f\colon Y \to X$ which respects the graph of spaces structure and the graph structure of vertex and edge spaces. That is, it induces graph homomorphisms $\underline{f}\colon \Gamma_Y \to \Gamma_X$, $f_v\colon Y_v \to X_{\underline{v}}$, and $f_e\colon Y_e \to X_{\underline{e}}$, where we write $\underline{v}$, $\underline{e}$ for $\underline{f}(v)$, $\underline{f}(e)$. A morphism is called a *local isometry* if it is a local isometry of cube complexes.

Remark 4.3.2. There are four possible ways that a morphism f, as defined above, can fail to be a local isometry.

- **Case 1)** f identifies two vertical edges. That is, for some vertex space Y_u, f_u identifies two coincident edges of Y_u.
- **Case 2)** f identifies two coincident horizontal edges $[a]$, $[a']$ where a and a' are in the same edge space of Y.
- **Case 3)** f identifies two coincident horizontal edges $[a]$, $[a']$ where a and a' are in different edge spaces of Y.
- **Case 4)** For some vertex s of Y, f fails to map $\mathrm{Link}(s)$ to a full subcomplex of $\mathrm{Link}(f(s))$. This means that there is a pair of coincident edges in Y which do not span a square in Y but whose images span a square in X.

Remark 4.3.3. The four cases in Remark 4.3.2 are reminiscent of the corresponding situation for graphs, where a map fails to be an immersion if and only if it identifies two coincident edges at some vertex. Accordingly, we may try to define a *fold* for graphs of graphs to be the simplest possible map which "corrects" one of the four cases. For example, to correct Case 4 we may glue in the corresponding square to Y (considered as a VH complex), and to correct one of the other three cases we simply perform the corresponding fold on the 1-skeleton. However, this approach is problematic, since after performing a basic fold on a vertex space to correct an occurrence of Case 1, the attaching maps of incident edge spaces may no longer be immersions, and hence the complex is no longer a graph of graphs. In order to remain within this category, and hence retain the link with the other notions of folding that have been developed for graphs of groups [BF91, Dun98, KWM05], we will insist throughout that morphisms restrict to immersions on vertex spaces.

Definition 4.3.4. A morphism between graphs of graphs is called a *local immersion* if it restricts to an immersion on vertex spaces.

It is easy to verify that a local immersion between graphs of graphs also restricts to an immersion on edge spaces. A morphism is a local immersion if and only if Case 1 of Remark 4.3.2 does not occur, and a local immersion is a local isometry if and only if none of Case 2, Case 3 or Case 4 occurs. Being *locally injective* is stronger than being a local immersion: a locally injective map can only fail to be a local isometry due to Case 4.

Definition 4.3.5. Given a local immersion $f\colon Y \to X$ of graphs of graphs, a morphism $\varphi\colon Y \to Y'$ of graphs of graphs is called a *fold (with respect to f)* if:

- φ is π_1-surjective.
- There is a local immersion $f'\colon Y' \to X$ satisfying $f' \circ \varphi = f$.
- φ induces a surjection from the set of hyperplanes of Y to the set of hyperplanes of Y'.
- φ is not the identity map.

It follows from the second condition of Definition 4.3.5 that a fold φ is itself a local immersion. Combining the first two conditions in Definition 4.3.5, we see that the images of the induced maps $f_*\colon \pi_1(Y) \to \pi_1(X)$ and $f'_*\colon \pi_1(Y') \to \pi_1(X)$ are equal. The third condition will be needed to understand when a sequence of folds terminates. In practice, we will define and use just two more specific types of fold, each designed to correct a failure of the map f to be a local isometry (see Section 4.3.2.

4.3.1 Outline of the folding algorithm

In this section, we give a brief outline of the folding algorithm which will be formalised in Theorem 4.3.9 and its proof.

Let H be a subgroup of $\pi_1(X)$. To begin, we take a set S of n generators for H, and let $f\colon B \to X$ be a corresponding map of a bouquet of n circles. Since B is a graph, we may apply Stallings folding to obtain a graph Y^0 equipped with an immersion f to the 1-skeleton of X.

Since X is a graph of graphs, there is an induced graph of graphs structure on Y^0. Edge spaces are points (so the corresponding subspaces of Y^0 are edges of the form point $\times$ I), and vertex spaces are graphs which immerse under f to vertex spaces of X. In particular, Y^0 is clean and f is locally injective. It can fail to be a local isometry only due to Case 4 above.

In general, if Y is a graph of graphs equipped with a morphism to X that is locally injective but not a local isometry, this means there is a pair of coincident edges—an edge a in Y_u and a horizontal edge $[b]$ corresponding to $b \in Y_e$—which is mapped by f to the corner of a square in X. To correct this, we must perform a folding move which glues in the corresponding square in Y. However, such a move may introduce additional instances of Case 4 at some adjacent edge to a. To control this, we would like to attach squares to some chosen subgraph of Y_u, which may just consist of the edge a, or may be larger.

The method we shall adopt is to select a loop (in general, a lollipop) l containing a and mapping under f into the edge subspace $\partial_{\underline{e}}(X_{\underline{e}}) \subset X_{\underline{u}}$. This loop intersects $\partial_e(Y_e)$, but is not contained in it.

Attaching a square to a, and then continuing to attach squares where necessary to the other edges in the loop, has the effect of adding an entire strip of new squares to the complex, along that section of l which was outside $\partial_e(Y_e)$. The opposite side of this strip of new squares must be attached to the opposite vertex space (that is, $Y_{\tau(e)}$), and this corresponds to folding together a copy of the loop l and the opposite vertex space, with respect to the appropriate maps to X. Similarly, the new edge space is that obtained by folding together a copy of the loop l with the old edge space. A typical effect would be to simply add an arc to the vertex space, however, this may introduce instances of Case 1, so in general, the new vertex space may have undergone further folding. By using Definition 4.1.3, we circumvent the need to deal with this directly. However, Case 2, Case 3 and Case 4 may still occur.

To correct Case 2 observe that, for any path connecting the pair of offending vertices in the corresponding edge space, the image of this path under the attaching map is a loop

l. Thus, pulling lollipops across this edge space will eventually be sufficient to remove this instance of Case 2.

To correct Case 3, we must identify the two offending edge spaces into a single edge space, by folding them together. Since this fold induces a Stallings fold on the underlying graph, it will only be needed finitely many times.

Our folding algorithm will work as follows. Firstly, we check if Case 3 occurs, and if so, we perform the above fold to correct it. Otherwise, we pull across a suitable loop to correct a failure of local isometry due to Case 2 or Case 4. Each fold will be defined using the folding together (i.e. pushout) construction on vertex spaces (Definition 4.1.3), which prevents new occurrences of Case 1 (and hence keeps us in the category of graphs of graphs equipped with local immersions). Since Case 3 occurs only finitely many times, all folds will eventually consist of pulling loops across edge spaces.

Having verified that the above moves are indeed folds, it will follow that the map induced on fundamental group by f still has image H. We will be able to use finite presentability of H to deduce that eventually, in the finitely presented case, we arrive at a complex whose fundamental group is H (see Section 4.3.5). By a local finiteness argument, we will show that the folding process corrects all failures of local isometry in arbitrarily large subcomplexes, and hence we can construct the direct limit complex Y_∞ which maps to X via a local isometry (after possibly pulling across some trees); see Section 4.3.4. We will also show that the algorithm terminates (that is, Y_∞ is compact) if and only if H is quasiconvex. The proof of this part will make use of Haglund's characterisation of convex subcomplexes in CAT(0) cube complexes (see Section 3.5.1), together with the fact that folding cannot introduce new hyperplanes.

We shall now describe each folding move more precisely.

4.3.2 The folding moves

Let $f\colon Y \to X$ be a local immersion of graphs of graphs, where X is clean.

Basic underlying graph fold

Suppose f fails to be a local isometry due to Case 3. Let d and e be edges of Γ_Y with $\iota(d) = \iota(e) = u$, such that f identifies a horizontal edge $[y_d]$ in $Y_d \times I$ with a coincident horizontal edge $[y_e]$ in $Y_e \times I$. Then the graph of graphs Y' is obtained as follows.

Let Y'_e be the graph obtained by folding together Y_d and Y_e at y_d and y_e, with respect to the maps ∂_d and ∂_e. The resulting immersion $\partial'_e \colon Y'_e \to Y_u$ is the new attaching map, and an application of Lemma 4.1.5 (in the spirit of Remark 4.1.7) gives us the immersion $f'_e \colon Y'_e \to X_{\underline{e}}$.

Let Y'_v be the graph obtained by identifying $\partial_{\bar{d}}(y_d) \in Y_{\tau(d)} = Y_w$ and $\partial_{\bar{e}}(y_e) \in Y_{\tau(e)} = Y_v$, and then applying Stallings folds to make the induced map to $X_{\underline{v}}$ an immersion (in the case where $v \neq w$ this is the same as the graph obtained by folding together Y_v and Y_w with respect to the two maps to $X_{\underline{v}}$). The resulting immersion to $X_{\underline{v}}$ is f'_v. Note that the natural maps $i_w \colon Y_w \to Y'_v$ and $i_v \colon Y_v \to Y'_v$ are immersions.

We may arrange all these maps into a commutative diagram as follows.

$$\begin{tikzcd}
 & Y_d \arrow[rr, "\partial_{\bar{d}}"] \arrow[dl, "\partial_d"'] \arrow[dr, "i_{Y_d}"] \arrow[dddr, "f_d"] & & Y_w \arrow[dr, "i_w"] \arrow[dddr, "f_w"] & \\
Y_u \arrow[ddd, "f_u"'] & & Y'_e \arrow[ll, "\partial'_e"'] \arrow[rr, "\partial'_{\bar{e}}"] \arrow[dd, "f'_e"] & & Y'_v \arrow[dd, "f'_v"] \\
 & Y_e \arrow[ul, "\partial_e"] \arrow[ur, "i_{Y_e}"] \arrow[rr, "\partial_{\bar{e}}"] \arrow[dr, "f_e"'] & & Y_v \arrow[ur, "i_v"] \arrow[dr, "f_v"'] & \\
X_{\underline{u}} & & X_{\underline{e}} \arrow[ll, hook, "\partial_{\underline{e}}"] \arrow[rr, hook, "\partial_{\underline{\bar{e}}}"'] & & X_{\underline{v}}
\end{tikzcd}$$

A further application of Lemma 4.1.5, using the right hand half of the diagram, gives us the attaching map $\partial'_{\bar{e}}$.

Lemma 4.3.6. *The map φ above is a fold.*

Proof. The map φ is clearly surjective (in particular, induces a surjection on the set of hyperplanes), and the map $f' \colon Y \to X$ satisfies $f' \circ \varphi = f$ by commutativity of the diagram. Hence, it only remains to show that φ is π_1-surjective, and for this it suffices to show that any element of a vertex group of Y' has a preimage in $\pi_1(Y)$. This is clearly true for every vertex group except $H'_v = \pi_1(Y'_v)$, so suppose $b \in H'_v$. Since Y'_v is obtained by folding together Y_w and Y_v, b has a preimage in $H_w * H_v$ (where H_* denotes $\pi_1(Y_*)$). Let $(p_1, q_1, \ldots, p_r, q_r)$ be a normal form in $H_w * H_v$ for this preimage (so $p_i \in H_w$, $q_i \in H_v$, and only p_1 and q_r may equal the identity). Now consider the (not necessarily reduced) loop in the graph of groups for Y given by $(p_1, \bar{d}, 1, e, q_1, \bar{e}, 1, d, p_2, \ldots, p_r, \bar{d}, 1, e, q_r)$. The corresponding element of $\pi_1(Y)$ is clearly a preimage for b, which completes the proof. □

Note that after performing this fold, the map f' is still a local immersion of graphs of graphs (since f'_v is an immersion by construction).

Pulling a subgraph across an edge space

To correct failures of local isometry due to Case 2 or Case 4, we may select a subgraph of a vertex space of Y which maps into the edge subspace of the corresponding vertex space of X, and add it into the appropriate edge space of Y. This move modifies both that edge space and the vertex space on the other end of that edge space, which is why we refer to it as "pulling the subgraph across the edge space".

To be precise, we consider pullbacks of edge spaces inside vertex spaces. As before, let $f\colon Y \to X$ be a local immersion of graphs of graphs, where X is clean.

For each edge space Y_e attached to some vertex space Y_u, consider the pullback:

$$\begin{tikzcd}
Y_u \arrow[d, "f_u"] & P \arrow[l, hook'] \arrow[d] \\
X_{\underline{u}} & X_{\underline{e}} \arrow[l, hook', "\partial_{\underline{e}}"]
\end{tikzcd}$$

By the existence of the map f and the connectedness of the edge space Y_e, we know that $\partial_e(Y_e)$ must be contained in some connected component of P. Denote this connected component by P_e. Note that P_e embeds in Y_u, but we do not need to assume that Y_e does.

Let Y_e, Y_u and P_e be as above, and let $i\colon Q \to Y_u$ be the inclusion of a subgraph of Y_u which factors through P_e. Assume that the intersection $i(Q) \cap \partial_e(Y_e)$ is non-empty. Thus, there are vertices $\alpha \in Q$ and $y \in Y_e$ such that $i(\alpha) = \partial_e(Y_e)$. The following diagram commutes:

$$\begin{tikzcd}
 & & Y_e \arrow[dll, "\partial_e"'] \arrow[d] \arrow[dddd, bend left, "f_e"] \\
Y_u \arrow[ddd, "f_u"'] & & P_e \arrow[ll, hook'] \arrow[ddd] \\
 & Q \arrow[ul, hook', "i"] \arrow[ur, hook] \arrow[ddr, "f_Q"] & \\
 & & \\
X_{\underline{u}} & & X_{\underline{e}} \arrow[ll, hook', "\partial_{\underline{e}}"]
\end{tikzcd}$$

The fold $\varphi\colon Y \to Y'$ which we now define is called *pulling* Q *across* Y_e. Y' is the graph of graphs obtained as follows:

- Let Y'_e be the graph obtained by folding together Y_e and Q at y and α, with respect to the maps ∂_e and i. The resulting immersion $Y'_e \to Y_u$ is the new attaching map ∂'_e. By commutativity of the diagram, the immersions from Y_e and Q to $X_{\underline{u}}$ which factor through Y_u also factor through $X_{\underline{e}}$. Therefore, applying Lemma 4.1.5 guarantees existence of the immersion f'_e.

$$\begin{tikzcd}
& Y_e \arrow[dl, "\partial_e"'] \arrow[dr, "i_{Y_e}"] \arrow[dddr, "f_e"] & \\
Y_u \arrow[dd, "f_u"'] & & Y'_e \arrow[ll, "\partial'_e"'] \arrow[dd, "f'_e"] \\
& Q \arrow[ul, "i"] \arrow[ur, "i_Q"] \arrow[dr, "f_Q"'] & \\
X_{\underline{u}} & & X_{\underline{e}} \arrow[ll, hook', "\partial_{\underline{e}}"]
\end{tikzcd}$$

- Let Y'_v (where $v = \tau(e)$) be the graph obtained by folding together Y_v and Q at $z = \partial_{\bar{e}}(y)$ and α, with respect to the maps $f_v\colon Y_v \to X_{\underline{v}}$ and $\partial_{\underline{\bar{e}}} \circ f_Q\colon Q \to X_{\underline{v}}$. The resulting immersion to $X_{\underline{v}}$ is f'_v, and the following diagram commutes:

$$\begin{tikzcd}
& Y_e \arrow[dl, "\partial_e"'] \arrow[dr, "i_{Y_e}"] \arrow[dddr, "f_e"] \arrow[rr, "\partial_{\bar{e}}"] & & Y_v \arrow[dr, "i_{Y_v}"] \arrow[dddr, "f_v"] & \\
Y_u \arrow[dd, "f_u"'] & & Y'_e \arrow[ll, "\partial'_e"'] \arrow[dd, "f'_e"] \arrow[rr, "\partial'_{\bar{e}}"] & & Y'_v \arrow[dd, "f'_v"] \\
& Q \arrow[ul, "i"] \arrow[ur, "i_Q"] \arrow[urrr, "i'_Q"] \arrow[dr, "f_Q"'] & & & \\
X_{\underline{u}} & & X_{\underline{e}} \arrow[ll, hook', "\partial_{\underline{e}}"] \arrow[rr, hook, "\partial_{\underline{\bar{e}}}"'] & & X_{\underline{v}}
\end{tikzcd}$$

See Figure 4.3 for a pictorial version of this diagram which illustrates the folding move.

- The attaching map $\partial'_{\bar{e}}$ is defined by two more applications of Lemma 4.1.5, since the immersions from Y_e and Q to $X_{\underline{v}}$ which factor through $X_{\underline{e}}$ also factor through Y'_v.
- The folding map $\varphi\colon Y \to Y'$ is induced by the maps $i_{Y_e}\colon Y_e \to Y'_e$ and $i_{Y_v}\colon Y_v \to Y'_v$, and the identity maps on all other vertex and edge spaces.

Lemma 4.3.7. *The map φ defined above is a fold.*

Proof. To see that φ induces a surjection on the set of hyperplanes, recall that vertical hyperplanes of Y' are in correspondence with edge spaces, and hence clearly no additional such hyperplanes are introduced by φ. Horizontal hyperplanes correspond to parallelism classes of vertical edges. The only such edges in Y' not necessarily in the image of φ are those edges in the copy of Q which is folded together with Y_v, and by construction, each such edge is elementarily parallel in Y' to the corresponding edge in $i(Q)$. Hence, each hyperplane of Y' contains the image under φ of some hyperplane of Y, as required.

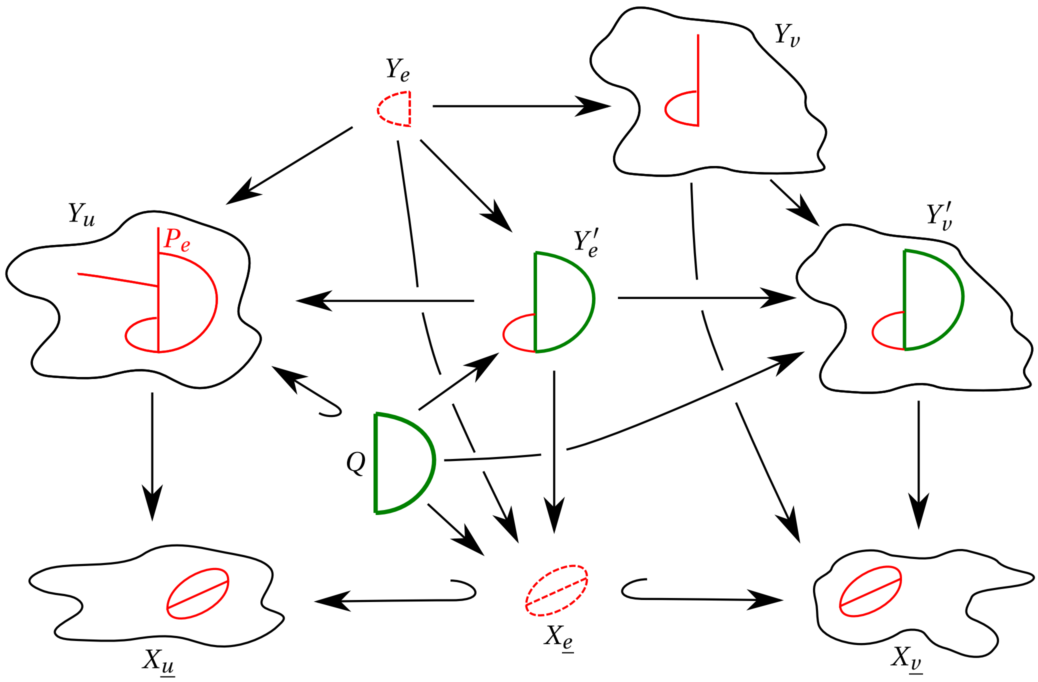

Figure 4.3: Pulling a subgraph across an edge space. We leave it to the reader to specify suitable immersions; for example, the immersion from P_e to $X_{\underline{e}}$. See the commutative diagram above for the names of all the maps shown.

The second condition in Definition 4.3.5 is satisfied by construction, so it remains to show that φ is also π_1-surjective. As before, we argue using the definition of the fundamental group of the graph of groups. Let φ_* be the induced map on fundamental groups.

Let b be an element of the vertex group H'_v in the graph of groups corresponding to Y'. Since Y'_v is obtained by folding together Y_v and Q, the map $\varphi_{*v}\colon H_v \to H'_v$ factors as the inclusion $H_v \hookrightarrow H_v * A$ and a surjection $H_v * A \to H'_v$, where $A = \pi_1(Q)$. Suppose $(h_1, a_1, \ldots, h_r, a_r)$ is a normal form for a preimage of b in $H_v * F$ (where $h_i \in H_v$, $f_i \in f_v$, and only h_1 and a_r are allowed to equal the identity). We may also consider the a_i as elements of H_u, via the inclusion map $Q \hookrightarrow Y_u$. So, consider the loop $(h_1, \bar{e}, a_1, e, \ldots, h_r, \bar{e}, a_r)$ in the graph of groups for Y. This may not be reduced, however, its image under φ in Y' is equivalent to $(h_1, a_1, \ldots, h_r, a_r)$, and hence to b. It follows that φ_* is π_1-surjective. □

Note that, after pulling a subgraph across an edge space as above, the map $f'\colon Y' \to X$ still restricts to immersions on vertex spaces (because f'_v is an immersion), and so, as before, we remain in the category of graphs of graphs equipped with local immersions to X.

Remark 4.3.8. We will apply the above construction in the case that Q is an embedded lollipop. However, in most accounts of folding graphs of groups (such as [KWM05]), the entire pullback P_e is pulled across. This is a natural thing to do insofar as it corrects more occurrences of Case 4, however, *any* choice of subgraph will do provided the subgraphs chosen exhaust P_e as the folding process continues. We will appeal to local finiteness to see that pulling across lollipops (followed possibly by trees) is sufficient.

4.3.3 The folding algorithm

We are now ready to state the main theorem of this chapter.

Theorem 4.3.9. *Let X be a clean, thin, compact graph of graphs, with $G = \pi_1(X)$, and let $H = \langle S \rangle < G$ be a finitely generated subgroup. Then there is an algorithm that generates a (finite or infinite) sequence $Y^0, Y^1, \ldots$ of compact graphs of graphs, equipped with morphisms $f^i\colon Y^i \to X$, such that:*

1. *Each map f^i is a local immersion.*
2. *The induced maps on fundamental groups f^i_* all have image H.*
3. *Y^0 is a topological graph.*
4. *For each i, there is a fold $\varphi^i\colon Y^i \to Y^{i+1}$ (such that $f^{i+1} \circ \varphi^i = f^i$).*
5. *The algorithm terminates (i.e. the sequence is finite $Y^0, \ldots, Y^n$) if and only if H is quasiconvex. In this case, f^n_* is an isomorphism onto H.*
6. *The direct limit Y^∞ of $Y^0 \xrightarrow{\varphi^0} Y^1 \xrightarrow{\varphi^1} \ldots$ exists, and the corresponding map $f^\infty\colon Y^\infty \to X$ induces an isomorphism f^∞_* onto H.*
7. *If H is finitely presented, then there exists m such that, for all $i \geq 0$, the fold φ^{m+i} is a homotopy equivalence and the map f^{m+i}_* is an isomorphism onto H.*

Remark 4.3.10. It follows from the Bestvina–Feighn combination theorem [BF92] that every compact, thin graph of graphs has hyperbolic fundamental group, and so G above is automatically hyperbolic.

Proof. The remainder of this section is dedicated to the proof of Theorem 4.3.9. We shall describe the algorithm inductively.

Base case (initial setup of the algorithm)

Each generator $s \in S$ can be realised as a combinatorial immersed loop in the 1-skeleton $X^{(1)}$, based at x. Let B be a finite bouquet of circles, with basepoint b, equipped with a combinatorial map f^B to X that realises S. We may apply Stallings folding (Theorem 4.1.2) to obtain a graph Y^0 mapping by an immersion f^0 into X. The graph of graphs structure on X induces a graph of graphs structure on Y^0. Each edge space of Y^0 is a point y_e equipped with two maps ∂_e, $\partial_{\bar{e}}$ into the vertex spaces on either side. Call the two images z_e, $z_{\bar{e}}$ respectively—we may regard these as "basepoints" of the two edge subspaces.

Inductive step

Let $f^i\colon Y^i \to X$ be a local immersion of graphs of graphs. Assume f^i is not a local isometry, so that at least one of Case 2, Case 3 or Case 4 holds. Note that $f^0\colon Y^0 \to X$ satisfies these hypothesis (unless f^0 is a local isometry, in which case the folding sequence ends with Y^0).

We apply a fold from Y^i to Y^{i+1} as follows.

If Y^i satisfies Case 3, then we may perform a basic underlying graph fold φ^i, to obtain a graph of graphs Y^{i+1} with a map f^{i+1} to X. This map still restricts to an immersion on vertex spaces. Note that since this fold reduces the number of edges of the underlying graph, it will only ever be applicable finitely many times.

Otherwise, Y^i satisfies either Case 2 or Case 4. Now, let $\mathscr{L}$ be the set of all lollipops satisfying the following conditions:

- $\lambda\colon \ell \to Y^i_u$ is an embedded lollipop in a vertex space Y^i_u of Y^i.
- For some e with $\iota(e) = u$, λ is based at some $z \in \partial_e(Y_e)$ and factors through P^i_e.
- $\lambda\colon \ell \to Y^i_u$ does not lift to $\partial_e\colon Y^i_e \to Y^i_u$ (that is, $\lambda(\ell)$ is not already the image of a lollipop in Y^i_e, though it may be the image of a path).

For each lollipop in $\mathscr{L}$ we may define the *complexity* to be $d_{Y^i_u}(z_e, z) + n$, where n is the length of the lollipop.

If f satisfies Case 2, there is an edge space Y^i_e of Y^i with distinct vertices a, b such that the horizontal edges $[a]$ and $[b]$ have a common endpoint $y \in Y^i_u$ (where $u = \iota(e)$), and such that $f(a) = f(b)$ in $X_{\underline{e}}$. Then for γ an immersed path in Y^i_e connecting a and b, $\partial_e \circ \gamma$ is an immersed loop in Y^i_u based at y, which does not lift as a loop to $Y_e)i$. It follows that $\mathscr{L}$ is nonempty.

In the case that f satisfies Case 4 but not Case 2, it is possible that $\mathscr{L}$ is empty. In this case, we **skip ahead to the final stage of the algorithm**—see below. Otherwise, $\mathscr{L}$ is nonempty. Let $\lambda\colon \ell \to X_u$ be a lollipop in $\mathscr{L}$ of minimal complexity, and perform the fold φ^i pulling this lollipop across the edge space Y_e^i, to obtain a graph of spaces Y^{i+1} with a map f^{i+1} to X. Again, f^{i+1} restricts to an immersion on vertex spaces.

Continuing this process indefinitely, we either obtain an infinite sequence of complexes as described in the statement of the theorem, or we reach a position where $\mathscr{L}$ is empty. In this case, we move to the final stage below.

Final stage

Suppose now that we have reached a complex Y^i, with a map f^i to X which is locally injective, and that $\mathscr{L}$ is empty. If f^i is a local isometry, then we stop with Y^i as the final graph of graphs in the sequence. Otherwise, f^i fails to be a local isometry due to Case 4. Below we write Y for Y^i.

Choose a vertex around which f fails to be a local isometry; say the vertex y in a vertex space Y_u. Then there is horizontal edge $[a]$ (where $a \in Y_e$) and a vertical edge b at y, which do not span a square in Y, but whose images under f span a square in X. Since $\mathscr{L}$ is empty and $\partial_e(Y_e)$ is connected, removing b from P_e separates P_e into two components, one of which is a tree T', and the other contains $\partial_e(Y_e)$, as in the following diagram. Now pull the tree $T = B \cup T' \subset P_e \subset Y_u$ across Y_e, to obtain a graph of graphs Y'.

$$\begin{tikzcd}
 & Y_e \arrow[rr, hook, "\partial_{\bar e}"] \arrow[dl, hook', "\partial_e"'] \arrow[dr, "i_{Y_e}"] \arrow[dddr, "f_e"] & & Y_v \arrow[dr, "i_{Y_v}"] \arrow[dddr, "f_v"] & \\
Y_u \arrow[ddd, "f_u"'] & & Y'_e \arrow[ll, hook', "\partial'_e"'] \arrow[rr, hook, "\partial'_{\bar e}"] \arrow[dd, "f'_e"] & & Y'_v \arrow[dd, "f'_v"] \\
 & T \arrow[ul, "\lambda"] \arrow[ur, hook, "i_T"] \arrow[urrr, hook, "i'_T"] \arrow[dr, "f_T"'] & & & \\
X_{\underline{u}} & & X_{\underline{e}} \arrow[ll, hook', "\partial_{\underline{e}}"] \arrow[rr, hook, "\partial_{\bar{\underline{e}}}"'] & & X_{\underline{v}}
\end{tikzcd}$$

Y_v' is obtained from Y_v by folding it together with a tree T. Since T is a tree, Lemma 4.1.4 applies, so we may consider Y_v as a subgraph of Y_v', and every immersed loop in Y_v' is in fact contained in Y_v. Therefore, Y_v is a core graph for Y_v'. It follows that $\mathscr{L}$ remains empty in Y', and $f'\colon Y' \to X$ may still only fail to be a local isometry due to Case 4.

So, we keep applying the above move, pulling trees across edge spaces to correct failures of local isometry due to Case 4, to obtain a sequence of graphs of graphs $Y^i = Y, Y^{i+1} = Y', Y^{i+2}, Y^{i+3}$ and so on, with folds $\varphi_j \colon Y^j \to Y^{j+1}$ between them. We claim that we may only perform the move finitely many times, and hence the algorithm terminates.

Consider some sequence of vertex spaces $Y^i_u, Y^{i+1}_u, \ldots$. By the above argument, Y^i_u embeds as a core graph in all Y^j_u for $j > i$. Since X is thin and f induces an immersion on vertex spaces, we may apply Lemma 4.2.3 to deduce that the family of connected components of pullbacks of edge spaces $\{P_e \mid \iota(e) = u\}$ in each Y^j_u forms a family of subgraphs satisfying the conditions of Lemma 4.2.4, where the core graph C is the identical image of Y^i_u, and k is uniform. Fix some $j > i$ and denote $Y_u = Y^j_u$ for simplicity.

Now, suppose T is a tree in Y_u to be pulled across some edge space Y_e. This means that T is contained within P_e but is edge-disjoint from the subgraph $\partial_e(Y_e)$. In the notation of Lemma 4.2.4, $P_e^{\text{unique}} - C$ must be inside $\partial_e(Y_e)$. To see this, note that it is edge-disjoint from C and thus consists entirely of edges which have been pulled across in the final stage of the algorithm; therefore, every edge in $P_e^{\text{unique}} - C$ is in the image of some edge space, and this edge space must be that corresponding to e by the definition of P_e^{unique}.

It follows that T is edge-disjoint from $P_e^{\text{unique}} - C$, but contained within P_e. Thus, by Lemma 4.2.4, it is inside the k-neighbourhood of C. By local finiteness, there are therefore only finitely many such trees T which are ever available to pull across edge spaces incident to Y_u, and similarly for all other vertex spaces. Therefore, only finitely many such moves may be applied.

It follows that, after pulling across finitely many trees, we arrive at a graph of spaces Y^n (for some $n > i$), where there are no trees left to pull across edge spaces: in particular, Case 4 no longer applies, and so the map f^n is a local isometry. This finishes the algorithm in the case where $\mathscr{L}$ is eventually empty.

4.3.4 Termination, stabilisation and the limit complex

The algorithm described above satisfies the first four conclusions of Theorem 4.3.9, by construction and the definition of a fold. One direction of conclusion 5 follows directly from Theorem 3.5.4; namely, if the algorithm terminates, then H is quasiconvex. For the reverse direction, if H is quasiconvex then Theorem 3.5.4 guarantees that there *exists* a cube complex Y equipped with a local isometry to X inducing $H \subset G$; we must demonstrate that our algorithm will, at some point, arrive at this complex. To show this, we will first describe the direct limit construction given as conclusion 6 in Theorem 4.3.9. If the algorithm terminates, then the direct limit obviously exists, so assume from now on that it does not terminate.

To define the direct limit of $Y^0 \xrightarrow{\varphi^0} Y^1 \xrightarrow{\varphi^1} \dots$ we need to show that the algorithm *locally* stabilises, in the sense that for any fixed subspace, the folding maps eventually induce isometries. This is made precise for vertex spaces in Lemma 4.3.11.

In any graph Γ with a vertex x, we let $B_R(x)$ denote the induced subgraph on vertices of distance at most R from x.

Lemma 4.3.11. *Let Γ^i denote the underlying graph of Y^i, and let k be such that the Γ^i are the same for all $i \geq k$ (i.e. the underlying graph has stabilised). Choose a vertex v of Γ^k. Consider the vertex spaces Y_v^i for $i \geq k$ and the maps $\varphi_v^i \colon Y_v^i \to Y_v^{i+1}$ obtained by restricting the folding maps φ^i. Choose a vertex y_v^k of Y_v^k, and denote by y_v^i its image under the folding maps in each Y_v^i. Then for any $R > 0$, the ball of radius R around y_v^k eventually stabilises in the following sense: there exists N such that, for all $j \geq 1$, the map $\varphi_v^{N+j-1} \circ \cdots \circ \varphi_v^N$ from Y_v^N to Y_v^{N+j} restricts to a graph isomorphism from $B_R\left(y_v^N\right)$ to $B_R\left(y_v^{N+j}\right)$.*

Proof. Y_v^i immerses into $X_{\underline{v}}$, which is locally finite, and so a labelling of the edges of $X_{\underline{v}}$ induces a finite labelling of the edges of Y_v^i. Thus $B_R\left(y_v^i\right)$ is a uniformly locally finite graph of fixed radius, and so it is uniquely determined by a uniformly finite set of labelled loops based at y_v^i (for example, the standard generating set with respect to a spanning tree), together with a finite set of labelled paths (corresponding to the degree 1 vertices). The folding process preserves labelled loops and paths. For fixed R, it follows that for large enough N, the set of labelled based loops and paths in $B_R\left(y_v^i\right)$ is the same for every $i \geq N$, and since they immerse, the graphs are therefore isomorphic by Theorem 4.1.2. □

Remark 4.3.12. We may extend the above lemma to the whole graph of spaces Y^i (for some $i > k$) as follows. Again, choose vertices y_v^i in each vertex space Y_v^i, such that $\varphi^i(y_v^i) = y_v^{i+1}$. Now, define Y_R^i to be the following subspace of Y^i. The vertex spaces of Y_R^i are the subspaces $B_R\left(y_v^i\right)$ of Y_v^i. For the edge spaces, let $Y_{e,R}^i$ be the intersection $\partial_e^{-1}\left(B_R\left(y_{\iota(e)}^i\right)\right) \cap \partial_{\bar{e}}^{-1}\left(B_R\left(y_{\tau(e)}^i\right)\right)$ (that is, we take the edge spaces as large as possible with the given vertex spaces). By the same argument as before (applied to all edge spaces and vertex spaces), we may find $N(R)$ such that the folding maps φ^i restrict to the identity map $Y_R^i \to Y_R^{i+1}$ for $i > N(R)$, and we may define Y_R to be the complex Y_R^i for all $i > N(R)$.

There is a natural inclusion $Y_R \subset Y_{R+1}$. Let Y^∞ be the direct limit of $Y_R \subset Y_{R+1} \subset \cdots$. This is a graph of spaces with the underlying graph Γ^k, and is equipped with a natural map to X. If the folding algorithm terminates, the terminal complex is compact and hence equal to Y_R for all large enough R, and indeed equal to Y^∞. We may think of Y^∞ as the "limit" of the folding algorithm, even when it does not terminate.

Proposition 4.3.13. *Y^∞ is clean.*

Proof. If Y^∞ is not clean, then Case 2 of Remark 4.3.2 applies in Y^∞, and so we may pick R such that Case 2 applies in $Y^{N(R)}$ and all Y^i for $i > N(R)$. This means there is a lollipop of uniformly bounded complexity (say, complexity K) in $\mathscr{L}$ at stage i for each $i > N(R)$. Since we choose lollipops of minimal complexity in $\mathscr{L}$ to pull across, this means that for all $i > N(R)$, Y^i is obtained from Y^{i-1} by pulling a lollipop of complexity $\leq K$ across an edge space. In particular, the vertex spaces may never extend beyond a distance K of the union of the basepoints z_e. Therefore Y^∞ is bounded, and local finiteness implies that Y^∞ is compact and the algorithm terminates. But this only happens if $\mathscr{L}$ is eventually empty, which contradicts our assumption. □

The above proof formalises the notion that "all lollipops are eventually pulled across", which is also fundamental to the next proposition. The proof is, in essence, an infinite version of the final stage of the algorithm in the case where it terminates (Section 4.3.3).

Proposition 4.3.14. *The natural map $f^\infty \colon Y^\infty \to X$ induces the inclusion $H \subset G$ of fundamental groups.*

Proof. If f^∞ is a local isometry, the proposition follows from Corollary 2.2.11. However, f may not yet be a local isometry; although every lollipop has been pulled across, there may still be trees in $P_e - \partial_e(Y_e)$, just as in the final stage of the algorithm in the terminating case (see Section 4.3.3). We may therefore apply the same move as before, pulling trees across edge spaces to correct each failure of local isometry, each time pulling across the trees closest to the basepoints of the vertex spaces. It follows from Lemma 4.1.4 that we obtain a nested sequence $Y^\infty = Y_0$, $Y' = Y_1$, Y_2, $Y_3 \ldots$ of graphs of graphs in this way, with folds $\varphi_i \colon Y_i \to Y_{i+1}$, and since the vertex and edge spaces are only altered by folding with trees, the vertex and edge groups remain the same and Lemma 2.1.8 implies that the Y_i are all homotopy equivalent. Then, the direct limit $\hat{Y}$ of $Y_0 \subset Y_1 \subset \ldots$ is a graph of graphs homotopy equivalent to Y^∞ (in fact, it deformation retracts to Y^∞) equipped with a natural map $\hat{f}$ to X which is a local isometry (by local finiteness and the fact that we pulled across trees in order of their distance from the basepoints). The result follows. □

We are now ready to complete the proof of Conclusion 5 of Theorem 4.3.9. We work in the universal cover, and we must establish some more notation to proceed.

Recall that we began by folding a bouquet of circles B, with basepoint b, to obtain a graph Y^0, equipped with a map $f^0 \colon Y^0 \to X$. Denote by b^i the image (under the sequence of folding maps) of b in Y^i. By the definition of a fold, every b^i has the same image (under f^i) in X; call this x. We also fix a basepoint $\tilde{x}$ of $\widetilde{X}$, such that $p(\tilde{x}) = x$ where $p \colon \widetilde{X} \to X$ is the universal covering map.

Let $q^i \colon \widetilde{Y^i} \to Y^i$ the universal cover of Y^i, and choose basepoints $\widetilde{b^i} \in \widetilde{Y^i}$ projecting to b^i. By the lifting criterion, the maps $\varphi^i \circ q^i \colon \widetilde{Y^i} \to Y^{i+1}$ each lift to maps $\widetilde{\varphi^i} \colon \widetilde{Y^i} \to \widetilde{Y^{i+1}}$ between the universal covers, and moreover we may choose these maps to send $\widetilde{b^i}$ to $\widetilde{b^{i+1}}$. Likewise, each map $f^i \circ q^i \colon \widetilde{Y^i} \to X$ lifts to a map $\widetilde{f^i} \colon \widetilde{Y^i} \to \widetilde{X}$, which sends $\widetilde{b^i}$ to $\tilde{x}$. Having fixed these maps, we hereafter abuse notation and refer to all of the b^i as b, and to all of the $\widetilde{b^i}$ as $\tilde{b}$.

Since Y^∞ is the direct limit of the Y^i, there is a natural map $\widetilde{f^\infty} \colon \widetilde{Y^\infty} \to \widetilde{X}$ which also fits into the commutative diagram, and likewise a natural map $\tilde{\hat{f}} \colon \tilde{\hat{Y}} \to \widetilde{X}$ (where $\hat{Y}$ is the complex constructed in the proof of Proposition 4.3.14).

In summary, the following diagram commutes.

$$\begin{tikzcd}
 & & & \widetilde{X} \arrow[ddd, "p"] & & & \\
\cdots \arrow[r] & \widetilde{Y^i} \arrow[r, "\widetilde{\varphi^i}"] \arrow[d, "q^i"] \arrow[urr, "\widetilde{f^i}"] & \widetilde{Y^{i+1}} \arrow[r] \arrow[d, "q^{i+1}"] \arrow[ur, "\widetilde{f^{i+1}}"] & \cdots \arrow[r] & \widetilde{Y^\infty} \arrow[r] \arrow[d, "q^\infty"] \arrow[ul, "\widetilde{f^\infty}"'] & \cdots \arrow[r] & \tilde{\hat{Y}} \arrow[d, "\hat{q}"] \arrow[ulll, "\tilde{\hat{f}}"'] \\
\cdots \arrow[r] & Y^i \arrow[r, "\varphi^i"] \arrow[drr, "f^i"'] & Y^{i+1} \arrow[r] \arrow[dr, "f^{i+1}"] & \cdots \arrow[r] & Y^\infty \arrow[r] \arrow[dl, "f^\infty"'] & \cdots \arrow[r] & \hat{Y} \arrow[dlll, "\hat{f}"] \\
 & & & X & & &
\end{tikzcd}$$

Consider the map $\widetilde{f^i}$. This is a locally injective combinatorial map, but it may not be injective if f^i is not a local isometry. Denote the image $\widetilde{f^i}\left(\widetilde{Y^i}\right) \subset \widetilde{X}$ by Z^i. Our first claim is the following.

Lemma 4.3.15. *If H is quasiconvex, Z^0 is combinatorially quasiconvex. Moreover, it is contained within a bounded neighbourhood of the H-orbit $H \cdot \tilde{x}$.*

Proof. Since Y^0 is a compact graph with basepoint b, $\widetilde{Y^0}$ is contained within a bounded neighbourhood of the set $\left(q^0\right)^{-1}(b)$ of lifts of b. We claim that, in fact, $\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right)$ is the H-orbit $H\cdot\widetilde{x}$. Then, by Lemma 3.5.9, it is enough to show that $\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right)$ is combinatorially quasiconvex. The proof is elementary covering space theory, but we outline the argument for convenience.

First, we show $H\cdot\widetilde{x}\subset\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right)$. Choose some point $z\in H\cdot\widetilde{x}$. There is a unique element $h\in H$ such that $h\cdot\widetilde{x}=z$. Recall that Y^0 is obtained by folding from B, a bouquet of circles labelled by the generators in a generating set S for H. Hence, we may choose a closed path in Y^0 which maps under f^0 to a representative loop for h in X. This path lifts to $\widetilde{Y^0}$, with one endpoint lifting to the selected basepoint $\widetilde{b}$, and the other lifting to some $w\in\left(q^0\right)^{-1}(b)$. By commutativity of the diagram above (with $i=0$), we must have $\widetilde{f^0}(w)=z$, as required.

Next, we show $\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right)\subset H\cdot\widetilde{x}$. Every point $w\in\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right.$ is the endpoint of the lift to $\widetilde{b}$ of some closed path in Y^0 based at b. This path maps under f^0 to a representative loop in X for some $h\in H$, and hence lifts to $\widetilde{X}$ as a closed path from $\widetilde{x}$ to some z in $H\cdot\widetilde{x}$. In particular, commutativity of the diagram implies that $\widetilde{f^0}(w)=z$, and hence $\widetilde{f^0}\left(\left(q^0\right)^{-1}(b)\right)=H\cdot\widetilde{x}$, as required.

Finally, by hyperbolicity of G (Remark 4.3.10) and Remark 3.5.8, quasiconvexity of H implies that $H\cdot\widetilde{x}$ is combinatorially quasiconvex. This completes the proof of the lemma. □

By Lemma 4.3.15 and Theorem 3.5.10, the combinatorial convex hull Z of Z^0 is contained within a bounded neighbourhood of Z^0. In particular, since Z^0 is contained within a bounded neighbourhood of an H orbit, the action of H on Z is cocompact. Recall that Z is the intersection of a family of halfspaces. To complete the proof that the algorithm terminates in the quasiconvex case, we first show that every image Z^i is contained within Z. This is the content of the following lemma.

Lemma 4.3.16. *If $\mathfrak{h}$ is a halfspace of $\widetilde{X}$, and $Z^i\subset\mathfrak{h}$, then $Z^{i+1}\subset\mathfrak{h}$.*

Proof. By construction (indeed, by the definition of a fold) φ^i is surjective on the set of hyperplanes of Y^i. It follows that the induced folding map $\widetilde{\varphi_i}\colon\widetilde{Y^i}\to\widetilde{Y^{i+1}}$ is surjective on the set of hyperplanes of $\widetilde{Y^i}$. Now let $\mathscr{H}$ be the hyperplane corresponding to the halfspace $\mathfrak{h}$, and denote by $\mathfrak{h}^c$ the other halfspace corresponding to $\mathscr{H}$. If $Z^i\subset\mathfrak{h}$, then in particular $\widetilde{x}\in\mathfrak{h}$, and so by construction of the map $\widetilde{f^{i+1}}$, we have that $Z^{i+1}\cap\mathfrak{h}\neq\varnothing$. Hence, if $Z^{i+1}\not\subseteq\mathfrak{h}$, then since Z^{i+1} is connected, it must intersect the hyperplane $\mathscr{H}$.

Since $\widetilde{\varphi_i}$ is surjective on the set of hyperplanes, there must be a hyperplane $\mathscr{K}$ in $\widetilde{Y^i}$ such that $\widetilde{f^{i+1}} \circ \widetilde{\varphi_i}(\mathscr{K}) \subset \mathscr{H}$. But, by commutativity of the diagram above, we have $\widetilde{f^i} = \widetilde{f^{i+1}} \circ \widetilde{\varphi_i}$, and hence $\widetilde{f^i}(\mathscr{K}) \subset \mathscr{H}$. In particular, Z^i intersects $\mathscr{H}$, which contradicts our assumption. □

By Lemma 4.3.16 and induction, the map $\widetilde{f^\infty}$ has image contained in Z. Moreover, applying the same argument to the folds used to build $\hat{Y}$ from Y^∞, the map $\tilde{\hat{f}}$ also has image contained in Z. However, the latter map is the lift to the universal cover of a local isometry of cube complexes, which is an isometric embedding by Corollary 2.2.11. Since Z is cocompact under the action of H, it follows that $\tilde{\hat{Y}}$ is cocompact under the universal covering action, and hence $\hat{Y}$ is compact. This is a contradiction, which completes the proof that the folding algorithm terminates in the case that H is quasiconvex.

Conclusion 6 of Theorem 4.3.9 follows from Remark 4.3.12 and Proposition 4.3.14. It remains only to prove Conclusion 7, and this is dealt with in the following section.

Remark 4.3.17. There is a slight inconsistency in our notation between the terminating and non-terminating cases. In the terminating case, the limit Y^∞ is equipped with a local isometry to X, because both loops and trees are pulled across as part of the algorithm. In the non-terminating case, only loops have been pulled across in the definition of Y^∞, and so it does *not* have a local isometry to X, even though it is a deformation retract of a space $\hat{Y}$ which does. It would be straightforward to rectify this by pulling across trees at the same time as loops in the main stage of the algorithm, but we have chosen not to do this in order to keep the folding process as canonical as possible.

4.3.5 Folding and presentations

We have described how to construct, step by step, a possibly infinite graph of graphs Y^∞ corresponding to any finitely generated subgroup of a graph of graphs X, where we constructed Y^∞ as a direct limit of spaces Y^i. We would now like to associate a group presentation $\langle S_i \mid R_i \rangle$ to each of the complexes Y^i. Recall from Section 4.3.3 that we began with a generating set S for H, and Y_0 was made by folding the corresponding bouquet of circles. Since all the folding maps are π_1-surjective, we may therefore use $S_i = S$ for all $i > 0$, and $R_0 = \varnothing$. Then either $R_{i+1} = R_i$ (in the case that the fold φ^i is a homotopy equivalence), or R_i is obtained from R_{i-1} by adding a finite number of relations (each elementary graph fold is either a homotopy equivalence, or it kills one loop, and our folds are finite combinations of these). Note that, if $R = \bigcup_{i\geq 1} R_i$, then $\pi_1(Y^\infty) = H = \langle S \mid R \rangle$.

We recall the following theorem:

Theorem 4.3.18. *[Mil04, Theorem 2.10] Let H be a finitely presentable group, with a presentation $\langle S \mid R\rangle$, where S is finite and R is infinite. Then there is a finite subset R' of R such that $H = \langle S \mid R'\rangle$.*

In our setting, this implies that if H is finitely presented, then for some i_0, $\pi_1(Y^{i_0}) = H$. In particular, all subsequent folding moves are π_1-isomorphisms, and hence homotopy equivalences by Whitehead's theorem (asphericity of each Y^i follows from non-positive curvature). This proves Conclusion 7, and hence completes the proof of Theorem 4.3.9. □

We close this section with some remarks about our algorithm, and how it relates to other accounts in the literature.

Corollary 4.3.19. *If X is a compact, thin VH complex, $G = \pi_1(X)$, and $H < G$ intersects each vertical hyperplane subgroup of G in a finitely generated group, then H is quasiconvex.*

Proof. The intersections between H and vertical hyperplane subgroups of G are precisely the edge groups of Y^∞. If they are all finitely generated, then it is clear that the folding algorithm terminates, so the result follows from Theorem 4.3.9. □

Remark 4.3.20. Corollary 4.3.19 is a special case of the main result of Beeker and Lazarovich's paper [BL16, Theorem 1.2], which is an alternative description of folding for cube complexes. Our corollary is weaker in that it is restricted to VH complexes, but it is worth noting that we do not need to assume finitely generated intersection of H with *every* hyperplane, but just with the vertical hyperplanes.

Remark 4.3.21. As well as the above account for cube complexes, there have also been multiple other generalisations of Stallings folding to graphs of groups. Bestvina and Feighn's account [BF91] is built using equivariant Stallings folding on the Bass–Serre tree, as is Dunwoody's generalisation [Dun98]. Our account is more closely related to that of Kapovich, Weidmann and Miasnikov [KWM05], which iteratively constructs generating sets for the vertex groups of the graph of groups being folded.

Remark 4.3.22. We have described the folding algorithm in the case where X is clean. However, if X is thin but not clean, then by Theorem 3.7.3 there is a finite sheeted cover X' which is clean. So, suppose G is the fundamental group of a thin graph of graphs, let $H < G$ be a subgroup, and let G' be the finite index subgroup corresponding to the clean cover. Since a subgroup is quasiconvex if and only if its intersection with some finite index subgroup is quasiconvex, we may apply our algorithm to $H \cap G' < G'$ to obtain a map of cube complexes $Y' \to X'$ inducing $H \cap G' \hookrightarrow G'$, and hence determine quasiconvexity of $H < G$. This does not automatically give us a local isometry between cube complexes $Y \to X$ realising $H < G$, however we can construct such a cube complex Y from the diagram $Y' \to X' \to X$. This is similar to the construction of the intermediate covering space between $H \cap G'$ and G representing H.

4.4 Folding for subgroups of finite width

4.4.1 Motivation: quasiconvexity of finite width subgroups

An important ingredient in the proof that virtually special groups possess a malnormal quasiconvex hierarchy was the theorem of Gitik–Mitra–Rips–Sageev (Theorem 2.5.11) that quasiconvex subgroups of hyperbolic groups have finite width. A natural question is whether the converse holds; that is:

Question 4.4.1. Are all finitely generated finite width subgroups of hyperbolic groups quasiconvex?

Remark 4.4.2. The finite generation assumption is necessary, because quasiconvex subgroups of hyperbolic groups are hyperbolic (Theorem 2.5.4); in particular, they are finitely generated.

This is a slightly more general version of Question 1.0.2, and as we discussed in Chapter 1, a good motivation for studying it is that either a positive or a negative answer would be of interest. If positive, this would give a purely algebraic definition of quasiconvexity in the δ-hyperbolic setting, compared to the existing notions which are heavily dependent on additional geometric structure. In the context of Chapter 3 it would, in particular, unite the twin assumptions of malnormality and quasiconvexity which are required in the definition of a malnormal quasiconvex hierarchy (see Section 3.8.1).

Although the question seems ambitious in the general setting of hyperbolic groups, there is reason to believe it might have a positive answer in many cases. For free groups and surface groups, the answer is obvious as *every* subgroup is quasiconvex. Less obviously, it also holds for hyperbolic 3-manifold groups (by the Tameness Theorem proved by Agol [Ago04] and independently by Calegari–Gabai [CG06]). For example, in the case of a 3-manifold M with a surface boundary Σ, we mentioned in Remark 2.5.8 that malnormality of $\pi_1(\Sigma)$ in $\pi_1(M)$ is equivalent to M being acylindrical. In this case, it is a theorem of Thurston that if M is also irreducible and atoroidal, it possesses a hyperbolic metric in which Σ is totally geodesic (for an account, see [Bon02]). This is stronger than quasiconvexity; indeed, it implies that the lift of Σ to $\widetilde{M}$ is convex. The close relationship between the theory of special cube complexes and 3-manifolds, which we described in Section 3.8.2, then gives a heuristic reason to believe that the question might have a positive answer in the special cube complex case.

Mitra's work provides another step towards an affirmative answer to Question 4.4.1. The main theorem of [Mit04] is that if a hyperbolic group G splits over a subgroup H, and the attaching maps are quasi-isometric embeddings, then H is of finite width in G if and only if it is quasiconvex in G. A slightly weaker statement was given, with an alternative proof, in [Kap01]. In the final section of [Mit04], Mitra also provides suggestions as to how to approach Question 4.4.1 in more generality.

On the other hand, a negative answer to Question 4.4.1 would also be a very interesting result. One reason is that it would provide a new example of a non-quasiconvex subgroup of a hyperbolic group; such subgroups do exist, but most examples arise as the subgroup K in a short exact sequence:

$$1 \to K \to \Gamma \to Q \to 1,$$

where Γ is hyperbolic and Q is infinite; in particular, K is actually an infinite-index normal subgroup, very far from having finite width (indeed, by Theorem 2.5.11, this is exactly how they are known to be non-quasiconvex).

One source of such examples is due to Rips [Rip82], where Γ is a $C'(1/6)$ *small cancellation* group (see Chapter 5 or [LS77, Chapter V] for an introduction to small cancellation theory), and Q is any finitely presented group. Another example arises when Γ is the fundamental group of a certain type of fibred 3-manifold: take $\Sigma \times [0,1]$ for some hyperbolic surface Σ, and identify the ends $\Sigma \times \{0\}$ and $\Sigma \times \{1\}$ by a pseudo-Anosov homeomorphism of Σ. The resulting manifold is fibred as a Σ bundle over S^1, and the fibration gives rise to a short exact sequence as above, with $K = \pi_1(\Sigma)$ and $Q = \mathbb{Z}$.

It is also possible to construct a short exact sequence as above where Γ is the fundamental group of a compact thin VH complex, as Haglund and Wise showed in [HW08]. This gives at least one example where our folding algorithm would not terminate, although performing the algorithm explicitly would be rather difficult given the complexity of the example.

If a negative answer to Question 4.4.1 could be found in the case where the subgroup was not only finite width but malnormal, then this would lead to another very interesting example of a non-hyperbolic group—in fact, it is close to a counterexample to the first fundamental open question we mentioned, Question 1.0.1. This follows from Theorem 4.4.3 below, of which we give a proof for convenience. Recall that a group is said to be of *type* F_n if it has a classifying space with a finite n-skeleton; *type* F_∞ if it has a classifying space with finite n skeleton for all n, and *type* F if it has a finite classifying space. Hyperbolic groups are of type F_∞, and of type F if they are torsion-free.

Theorem 4.4.3. *Suppose H is a type F_n malnormal subgroup of a hyperbolic group Γ, but H is not quasiconvex in Γ. Then the double $D = \Gamma *_H \Gamma$ is a non-hyperbolic group of type F_{n+1} with no Baumslag–Solitar subgroups.*

The proof depends on the following lemma of Bass–Serre theory (see Remark 2.1.6).

Lemma 4.4.4. *Suppose a Baumslag–Solitar group $BS(m,n) = \langle a,t \mid t^{-1}a^m t = a^n \rangle$ acts 1-acylindrically on a tree without edge inversions. Then it stablises a vertex.*

Proof. Isometries of a tree without edge-inversions are either hyperbolic (translate along a unique axis) or elliptic (fix a vertex). It is elementary to check that if $s^{-1}ps = q$, and q is hyperbolic, then p is hyperbolic, the translation lengths of q and p are equal, and if L is the axis of q then sL is the axis of p. So assume first that a is hyperbolic. Then a^m and a^n must have the same translation length, which implies that $m = \pm n$, and moreover t must fix the axis of a. Hence the whole group acts on L, and therefore admits a map to D_∞, whose kernel (which is infinite) fixes L pointwise. This contradicts acylindricity of the action.

The other case is that a is elliptic, so fixes a vertex v. If t also fixes v, then the proof is complete, so assume v and tv are distinct vertices. Now, a^{mn} fixes v, but since $a^{mn} = ta^{n^2}t^{-1}$, we have $a^{mn}tv = ta^{n^2}v = tv$, so a^{mn} also fixes tv and hence the line from v to tv. This contradicts 1-acylindricity unless v and tv are adjacent. In this latter case, since the action is without inversions, $d(v,t^2v) = 2$, and a similar argument with $a^{m^2n^2}$ shows that this also contradicts 1-acylindricity. This completes the proof of the lemma. □

Proof of Theorem 4.4.3. Firstly, we build a classifying space for D. Take classifying spaces for Γ with finitely many cells in each dimension, and a classifying space for H with finite n-skeleton. Now we may build a graph of spaces for $\Gamma *_H \Gamma$ using these as the vertex and edge spaces respectively. The $n+1$ cells of this space are of two types; either $n+1$ cells in the vertex spaces, or n-cells in the edge space crossed with an interval. Thus, the total space has finite $n+1$-skeleton, and by Lemma 2.1.8 it is aspherical and hence a classifying space for D.

Now, if D has a Baumslag–Solitar subgroup, by Lemma 4.4.4 it is conjugate into a vertex group, since malnormality of H in Γ guarantees 1-acylindricity (see Remark 2.1.6). However, hyperbolic groups cannot contain Baumslag–Solitar subgroups, so this is a contradiction.

Finally, we show that D cannot be hyperbolic. Suppose it is, and let X be the Cayley graph of D with the generating set given by the union of two copies of a finite generating set of Γ. This has a subgraph X_H corresponding to the subgroup H, and two subgraphs X_1, X_2 corresponding to each vertex subgroup Γ. Note that $X_1 \cap X_2 = X_H$. Let δ be the hyperbolicity constant for X.

Since X_H is not quasiconvex in Γ, for any R we may choose two points x, y in X_H such that, for $i = 1,2$, any choice of geodesic γ_i in X_i joining x and y is not contained in the R-neighbourhood of X_H. These two paths are also geodesics in X, by uniqueness of normal forms (or by the fact that Γ is a retract of D). Hence, they form a geodesic bigon.

Let z_1 be a point on γ_1 such that $d(z_1, X_H) > R$, and let z_2 be the closest point on γ_2 to z_1. It follows immediately from δ-thinness that the two sides of a geodesic bigon in any δ-hyperbolic metric space are of distance at most δ, and hence $d(z_1, z_2) \le \delta$.

However, by the Švarc–Milnor Lemma (Theorem 2.4.7), X is quasi-isometric to the universal cover of the classifying space for D. Therefore, any path joining z_1 to z_2 must pass within some uniform distance r of X_H, and hence has length at least $R-r$. Thus, for $R > \delta + r$, we obtain a contradiction. This completes the proof of the theorem. □

If we could find a distorted, malnormal subgroup of a hyperbolic group which was type F_∞, then Theorem 4.4.3 would provide a counterexample to Question 1.0.1. Indeed, if Γ were compact special, then H and Γ would both be type F, and D would be a counterexample to the weaker claim that any type F group without Baumslag–Solitar subgroups is hyperbolic.

Even in the case where the subgroup is just finitely generated and Γ is not necessarily torsion-free, D would be a finitely presented non-hyperbolic group without Baumslag–Solitar subgroups, and very few such examples are known (they are due to Brady, and arise as finitely presented subgroups of hyperbolic groups which are not of type F_3). If such a subgroup could be found that were finitely presented, then D would be the first known example of a type F_3 non-hyperbolic group without Baumslag–Solitar subgroups.

Note that if H is malnormal and quasiconvex, then $D = \Gamma *_H \Gamma$ is known to be hyperbolic by the Bestvina–Feighn combination theorem [BF92].

4.4.2 Asymptotic injectivity radius

Our original motivation for describing the folding algorithm was to understand the geometry of finite width subgroups, with a view to answering Question 4.4.1. The following lemma shows that finite width does have geometric consequences.

Lemma 4.4.5. *Let Γ be a finite graph, let $G = \pi_1(\Gamma)$, and let $H < G$ be a subgroup of finite width k. Let $p\colon \Gamma' \to \Gamma$ be an immersion realising the inclusion $H < G$. Then Γ' contains only finitely many distinct embedded combinatorial loops of each length.*

Proof. Suppose there are infinitely many embedded loops in Γ' of length n. Since Γ is finite, we may choose a set $l_1, \dots, l_{k+1}$ of $k+1$ embedded loops of length n in Γ', all mapping under p to the same loop l in Γ. Choose a vertex $w \in l$, and for each i let y_i be the lift of w to l_i. Denote $y = y_1$, let p_1 be the constant path at y_1, and for each $i \geq 2$ let p_i be a path in Γ' from y to y_i. We take the fundamental groups of Γ and Γ' with respect to w and y respectively.

If there is a pair i, j such that p_i and p_j represent elements in the same right coset of H in G, then $p_i p_j^{-1}$ represents an element of H, and so it is a loop in Γ'. This is a contradiction because the points y_i and y_j are distinct. Thus, the element $[l]$ of $\pi_1(\Gamma)$ represented by l is an infinite order element in the intersection of any pair of conjugates of H by the elements of G represented by the p_i. This contradicts the fact that H has width k. □

The following corollary is immediate, and gives a geometric consequence of finite width in the case where Γ' is an infinite graph.

Corollary 4.4.6. *Let Γ and Γ' be as in Lemma 4.4.5, and fix basepoints $x \in \Gamma$, $x' \in \Gamma'$. Then the injectivity radius of Γ' is asymptotically infinite, in the sense that for any n, there exists R such that $\Gamma' - B(x', R)$ contains no embedded loops of length $< n$.*

Lemma 4.4.5 and Corollary 4.4.6 both generalise in the obvious way to a local isometry between cube complexes, however one might hope to apply them in the graph case by looking at vertex spaces, as in the following subsection.

4.4.3 Potential approach: asymptotic geometry of the vertex spaces

Lemma 4.4.7. *Let G be the fundamental group of a graph of groups and let $H < G$ be a subgroup of finite width. Then the vertex groups of H are finite width subgroups of the vertex groups of G.*

Proof. Let H_v be a vertex group of (the induced graph of groups for) H, and G_v the corresponding vertex group of G. These are subgroups of H and G respectively, and $H_v = H \cap G_v$ (up to conjugacy). It is a general fact that if $A < B$ and A has width k in B, then for any other subgroup $C < B$ we have that $A \cap C$ has width $\leq k$ in C. Indeed, if $A \cap C$ has width $> k$ in C, we can choose $k+1$ elements $p_1, \dots, p_{k+1}$ in C, such that for every pair $i \neq j$, $p_i p_j^{-1} \notin A$ but $p_i^{-1} A p_i \cap p_j^{-1} A p_j$ is infinite. This contradicts the fact that A has width k in B. □

This suggests one of the possible approaches to understanding finite width subgroups of graphs of graphs: studying the geometry of graphs with asymptotically infinite injectivity radius. In the setting of our folding algorithm, we may bring more to bear on this problem. Using the thinness of X, and the fact that each infinite vertex space of Y^∞ is comprised of a small initial section (the image of Y^0) combined with finitely many connected edge spaces, we obtain the following (we omit a rigourous proof, but the properties all follow quickly from the folding algorithm).

Proposition 4.4.8. *Consider the setting of Theorem 4.3.9 in the case where the folding algorithm does not terminate. Assume H has finite width in G, and choose a vertex space Y_v of $Y = Y^\infty$. Then Y_v is a graph satisfying the following conditions.*

- *Y_v is uniformly locally finite.*
- *Y_v is a union of subgraphs $Y_v = C \cup Y_{e_1} \cup \cdots \cup Y_{e_n}$ where C_v is finite, and the pairwise intersections of different Y_{e_i} are forests whose components have uniformly bounded size.*
- *The fundamental group $H_v = \pi_1(Y_v)$ is generated by the subgroups corresponding to C_v and the Y_{e_i}.*
- *Y_v (and each Y_{e_i}) has asymptotically infinite injectivity radius in the sense of Corollary 4.4.6.*

In light of Proposition 4.4.8, it is helpful to consider the different subgraphs Y_{e_i} to correspond to *colours* on the graph Y_v; this induces a multiple colouring on the edges of Y_v (possibly excluding some edges of C_v). One may then consider the representation of some long, monochromatic loop l in Y_v as a product of generators from the generating set $C_v \cup Y_{e_1} \cup \cdots \cup Y_{e_n}$ as a disc diagram with boundary l, and with each 2-cell corresponding to a loop which is either monochromatic or contained within C_v. By choosing l to be a loop outside of some large neighbourhood of the basepoint in Y_v, we may arrange for an arbitrarily strong small cancellation constant in a large neighbourhood of the boundary loop in this disc diagram. One could hope to derive a contradiction directly from this.

Alternatively, one could try to find information about the distortion of $H < G$ as follows. First, in each vertex space Y_u, choose a generating set of the subgraph C_u. Now, for the given product decomposition of l, we have a certain number of generators from C_v and a certain number of monochromatic loops, each of which has its own product decomposition in some adjacent vertex space. Using the small cancellation of each vertex space, one might be able to find a function linking the length of the loop l to the number of generators from the full collection of subgraphs C_u (from all vertex spaces) required to express it as an element of H. This function would be closely related to the distortion of H in G. For this approach to work, it seems necessary to have a better understanding of the injectivity radius growth in vertex spaces, and we hope that future work will make progress here.

Remark 4.4.9. There are only finitely many embedded loops of each length in the complex X of Theorem 4.3.9, and so we can always use separability to take a finite sheeted covering space with arbitrarily high injectivity radius, regardless of finite width. Such a cover would typically have many more edge spaces than X. This implies that any approach to Question 4.4.1 using the small cancellation style ideas outlined above would certainly have to make use of the fact that the small cancellation constant was *asymptotically* very small, while the number of edge spaces remains the same.

4.4.4 Other potential approaches

We close this section with a few further ideas for how one could use the folding algorithm to approach Question 4.4.1.

Remark 4.4.10. It is not difficult to see that, considering the Y_i as cube complexes, the folding process does not add any hyperplanes but rather enlarges existing hyperplanes. It follows that if the algorithm does not terminate, then Y^∞ is an infinite cube complex with finitely many hyperplanes. In particular, there are finitely many *horizontal* hyperplanes, as well as the vertical hyperplanes which correspond to edge spaces. It seems likely that an approach to Question 4.4.1 would have to make good use of the fact that both these sets were finite. It is not obvious whether or not such a special cube complex exists with asymptotically infinite injectivity radius (see Corollary 4.4.6). If one could construct such a cube complex, it would be a finite width, non-convex cocompact subgroup of a finitely generated right angled Artin group.

Remark 4.4.11. In the finitely presented case, all folds eventually consist of pulling across loops, and all folds are eventually homotopy equivalences. However, there are still two types of such fold. *Injective folds* induce an injection $Y^i \to Y^{i+1}$, and *surjective folds* induce a surjection on vertex spaces. This depends on whether the loop ℓ being pulled across from one vertex space, Y_u to another, Y_v, lifts as a *path* to $P_{\bar{e}} \subset Y_v$. In this case, Y_v does not embed in the new vertex space Y'_v, but rather surjects onto it. If one could eventually rule out surjective folds—either by imposing hypotheses or by changing the order with which loops are chosen—then the associated generating sets of the vertex groups would be easier, and this would simplify the analysis of the product decompositions we discussed in the previous section.

4.5 Higher dimensional folding

Theorem 4.3.9 only applies to graphs of graphs—a subclass of compact special cube complexes. In order to generalise it to a folding algorithm which works for higher dimensional special cube complexes, we may exploit Corollary 3.8.4 and use induction on the level in the malnormal quasiconvex hierarchy (graphs of graphs providing the first non-trivial case). In this section, we outline how such an induction might work in the case where the subgroup is known to be quasiconvex. We move freely between the algebraic notion of a malnormal quasiconvex amalgam (or HNN extension) and the corresponding geometric notion of a compact special cube complex with a specified hyperplane.

Let G be a virtually compact special hyperbolic group, and let G' be the finite index subgroup of G in $\mathcal{MQH}$. Let h denote the *height* of G' with respect to $\mathcal{MQH}$—that is, G' is an amalgam $A *_C B$ or HNN extension $A*_C$, where A and B have height $h-1$, and C is malnormal and quasiconvex in G'. We proceed by induction on h.

Suppose, for induction, that if G is a group admitting a malnormal quasiconvex hierarchy of height $n-1$, with corresponding cube complex X, and if $Y \to X$ is a suitably defined morphism, then there is a sequence of folds (in a suitable sense) taking Y to Y', where $Y' \to X$ is a local isometry and $\pi_1(Y')$ is quasiconvex if and only if Y' is compact if and only if the folding sequence is finite. If the case $n = 0$ corresponds to Stallings folding for free groups, then Theorem 4.3.9 provides the case $n = 1$.

Now to prove the above for height n, we would like to mimic the proof of Theorem 4.3.9. So, we begin with a special cube complex X with fundamental group $G = A *_C B$, and a quasiconvex subgroup $H < G$. We aim to construct a sequence $Y_0, Y_1 \cdots$ of cube complexes (with induced graph of spaces decompositions) as in the statement of Theorem 4.3.9. To begin, we define our elementary folds.

Just as in Remark 4.3.2, there are different ways in which a morphism $f\colon Y \to X$ can fail to be a local isometry. If f fails to be a local isometry on a vertex space, then our inductive hypothesis can be applied to fold that vertex space. If f is a local isometry on all vertex spaces but still fails to be a local isometry globally, then it may be non-injective on some link. This can be corrected by identifying vertices in an edge space and a vertex space (or a pair of edge/vertex spaces) and folding using the inductive hypothesis. The final possibility is that f fails to map onto a full subcomplex somewhere, and as before, this is corrected by pulling (some subcomplex of) the appropriate component of a pullback across an edge space, just as before. Again, the inductive hypothesis will allow us to build the new vertex and edge spaces by folding.

Having defined these folds, the algorithm should proceed just as before—we begin with a folded bouquet of circles, and apply the above folds to correct each failure of local isometry. Local finiteness arguments should still ensure the existence of the limit complex. The principal subtlety is the need to deal with the fact that, even if H is quasiconvex, the vertex groups of Y^i may not be quasiconvex subgroups of the vertex groups of X at each stage i (since we lose the free group-specific fact that every finitely generated subgroup is quasiconvex). Thus, we either have to fold in such a way that this quasiconvexity *can* be ensured (which may be impossible), or we have to fold “diagonally”, by carrying out a certain amount of folding on a vertex space, enough to stabilise a large enough subcomplex that we can pull around a chosen loop to an adjacent vertex space, but without waiting for the folding algorithm on the vertex space to terminate. Provided that this can be achieved, writing down the folding algorithm precisely in higher dimensions should be doable.

Chapter 5

Negatively curved metrics on small cancellation groups

In Section 2.4.3, we discussed one of the most fundamental open questions in geometric group theory: to what extent do the various definitions of negative curvature for groups coincide? For the remainder of the thesis, we leave cube complexes behind, and return to one of the other notions of negative curvature, namely the CAT(-1) condition. In Chapter 6, we will fully settle the ambiguity between the notions of negative curvature in the case of *limit groups*, by proving that they are hyperbolic if and only if they are CAT(-1) (for a stronger statement, see Theorem 6.3.10). The geometric objects we work with are negatively curved simplicial 2-complexes. In the current chapter, to get a feel for these objects, we use them to prove Theorem 5.1.4 below. A presentation is called *uniformly* $C'(1/6)$ if *pieces* (overlaps between relators) are all shorter than a sixth of the length of the shortest relator. Our theorem is then:

Theorem (Theorem 5.1.4)**.** *Let G be a group with a uniformly $C'(1/6)$ presentation. Then G acts geometrically on a 2-dimensional* CAT(-1) *space.*

Although the uniform $C'(1/6)$ condition we use is in general stronger than the standard $C'(1/6)$ condition, it still holds for an important class of $C'(1/6)$ groups; namely, random groups in the density model at density $< 1/12$. We therefore have the following corollary:

Corollary (Corollary 5.1.6)**.** *Random groups in the density model, for density $d < 1/12$, act geometrically on a 2-dimensional* CAT(-1) *space.*

Wise showed in [Wis04] that $C'(1/6)$ groups are CAT(0), and hence so are random groups at density $< 1/12$. Ollivier and Wise then improved this to density $< 1/6$ [OW11]. However, since both results use cubulation, the CAT(0) spaces obtained are of high dimension, and so Corollary 5.1.6 represents an improvement in dimension as well as curvature.

Our proof consists of two steps. Firstly, we find an explicit hyperbolic 2-complex structure on the universal cover of a presentation complex for our group, using singular hyperbolic polygons. The resulting complex has some points of local positive curvature, but we are able to perform folds to remove this positive curvature and obtain our desired complex.

Remark 5.0.1. After completing this proof, we became aware that this argument was known to Gromov [Gro01], and a more general version of it (in the context of small cancellation over graphs of groups) is described in [Mar13]. The latter paper deals with a CAT(0) metric, but the author points out that the argument also works in the CAT(-1) case. We would like to thank Alexandre Martin and Anthony Genevois for bringing these two papers to our attention.

5.1 Preliminaries

Recall from Section 2.2.4 the definition of an M_k-simplicial complex. In this chapter we refer to an M_k-simplicial complex which is locally CAT(k), where $k < 0$, as a *negatively curved simplicial complex (of curvature k)*, or just a *negatively curved complex*. Both in this chapter and in Chapter 6, we will assume that all complexes are locally finite (see Remark 6.1.7), and will be concerned only with 2-dimensional complexes (see Remark 6.1.2).

5.1.1 Small cancellation conditions

A good reference for classical small cancellation theory is [LS77], and we refer the reader there for full details. We only state here what is necessary for us to give our main theorem.

Definition 5.1.1. Let $R = \{r_1, \dots, r_n\}$ be a set of cyclically reduced words on an alphabet $S \sqcup S^{-1}$, closed under taking cyclic permutations and inverses. A *piece* in R is a word w which appears as an initial segment of at least two elements of R.

Definition 5.1.2. Let $\mathscr{P} = \langle S \mid R \rangle$ be a presentation for a group G. Without loss of generality, assume R is closed under taking cyclic permutations and inverses. We say $\mathscr{P}$ is $C'(1/6)$ if every piece in R has length strictly less than $1/6$ of the length of any relator in which it appears. Now let g be the minimal length of any relator in R. We say $\mathscr{P}$ is *uniformly* $C'(1/6)$ if every piece in R has length strictly less than $g/6$, and moreover no element of R is a proper power.

Remark 5.1.3. Groups which are $C'(1/6)$ are torsion-free if and only if no relator is a proper power, and proper powers are forbidden by our uniform $C'(1/6)$ condition; hence, any uniformly $C'(1/6)$ group is torsion-free. This is necessary for our argument, since we produce a *free* action on a CAT(−1) space, and all groups possessing such an action are torsion-free. However, we do not know whether the uniform small cancellation condition can be relaxed to the standard $C'(1/6)$ condition in the torsion-free case.

Our main theorem is the following.

Theorem 5.1.4. *Let G be a group with a uniformly $C'(1/6)$ presentation, $\langle S \mid R \rangle$. Then G has* CAT(−1) *dimension 2.*

Remark 5.1.5. Random groups in the density model, for density $< 1/12$, satisfy the ordinary (non-uniform) $C'(1/6)$ condition [Gro93]. Since they have all relations of equal length, they satisfy the uniform $C'(1/6)$ condition too. This provides an assurance that the uniform condition is not too much of a restriction; indeed we obtain the following immediate corollary of Theorem 5.1.4.

Corollary 5.1.6. *Random groups at density* $< 1/12$ *have* CAT(−1) *dimension 2.*

5.1.2 Geometry of regular polygons

The first part of our argument relies on choosing suitable metrics on the 2-cells in the presentation complex. These metrics are based on small regular hyperbolic polygons, however since sufficiently small hyperbolic polygons closely resemble Euclidean polygons, we will argue in the Euclidean case for technical simplicity. Proposition 5.1.11 makes the conversion to a hyperbolic metric explicit.

We first establish some terminology about such polygons.

Definition 5.1.7. Let P be a regular (hyperbolic or Euclidean polygon). A *diagonal* is a geodesic connecting two (possibly consecutive) vertices of P. A *segment* of P is the smaller of the two pieces obtained by cutting P along a diagonal (in the case where the diagonal is an edge of P, the segment is also this single edge). The diagonal is said to *subtend* the corresponding segment. The *length* of the segment is the number of edges of P it contains, and

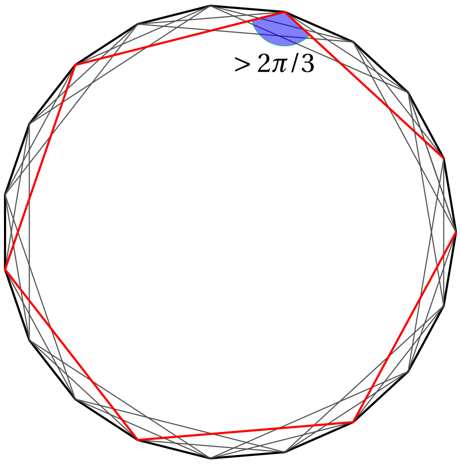

Figure 5.1: The extremal angle between diagonals of length ≤ 3 in a regular 19-gon. No internal angle between diagonals in the picture is less than the highlighted angle.

the *length* of the diagonal is the length of the corresponding segment. The *radius* of a regular polygon is the distance from the centre to any boundary vertex.

Definition 5.1.8. Let P be an n-gon, and let d_1, d_2 be two diagonals of length less than $n/2$. Suppose d_1 and d_2 intersect at a point p. Precisely one of the connected components of $P - (d_1 \cup d_2)$ contains the centre of P, and the angle at p which is inside this component is called the *internal angle* between d_1 and d_2.

Lemma 5.1.9. *Consider two intersecting diagonals in a regular Euclidean $n+1$-gon, each of length at most $\lfloor n/6 \rfloor$. The minimal internal angle between such diagonals is $> 2\pi/3$.*

Proof. Clearly, the case realising the minimal internal angle is where the two diagonals share an endpoint and are of the maximal permitted length. To see this, take any other intersection of diagonals d_1, d_2, where d_i has endpoints v_i and w_i. Without loss of generality, the clockwise order of the endpoints around the boundary of the polygon is v_1, v_2, w_1, w_2. If d_1 is not of maximal length, then increase its length by keeping w_1 fixed and moving v_1 anticlockwise. This clearly decreases the internal angle between d_1 and d_2. Similarly, increase the length of d_2 by fixing v_2 and moving w_2 clockwise; this also decreases the angle. Finally, fix d_1 and rotate d_2 by moving both v_2 and w_2 clockwise, until v_2 coincides with w_1. This process also decreases angle, and we have arrived at the extremal case.

Since six consecutive maximal length diagonals fail to complete a hexagon, the internal angle between each pair must be $> 2\pi/3$. See Figure 5.1. $\square$

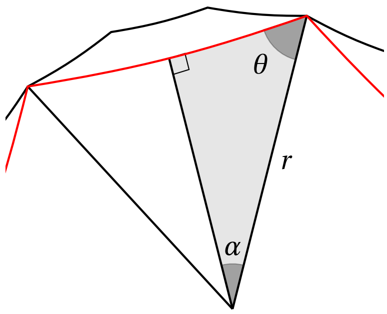

Figure 5.2: A right-angled triangle in a regular hyperbolic $6n+1$-gon (for the case $n=3$).

Remark 5.1.10. Lemma 5.1.9 also applies to sufficiently small regular hyperbolic polygons, since the metric differs from a Euclidean metric by an arbitrarily small amount. The following proposition makes this more precise:

Proposition 5.1.11. *For each integer $n > 6$, there exists $r > 0$ such that Lemma 5.1.9 also holds for regular hyperbolic n-gons of radius $< r$.*

Proof. It is enough to show that, for any $n \geq 1$ and $s \geq 1$, we may pick r small enough that the angle between two diagonals of length n, with a common endpoint, in a regular hyperbolic $(6n+s)$-gon of radius r, is $> 2\pi/3$ (that is, the picture in Figure 5.1 still applies). Clearly, the extremal case is $s = 1$, so it suffices to consider this case. Take a right-angled triangle with one vertex at the centre o of such a polygon, one vertex at the endpoint of a diagonal of length n, and one vertex at the midpoint of this diagonal. Denote by α and θ the two non-right angles in this triangle, as in Figure 5.2. Our goal is to calculate the range of values of r such that $\theta > \pi/3$.

Since this triangle is obtained by bisecting the isosceles triangle with angle $2\pi \times \frac{n}{6n+1}$, we see that $\alpha = \frac{n\pi}{6n+1}$. It then follows from the second hyperbolic cosine rule that

$$\cosh(r) = \cot(\theta)\cot\left(\frac{n\pi}{6n+1}\right).$$

We would like to find r such that $\theta > \pi/3$; equivalently, $\cot(\theta) < 1/\sqrt{3}$. Hence

$$\frac{\cosh(r)}{\cot(\frac{n\pi}{6n+1})} < 1/\sqrt{3}$$
$$r < \cosh^{-1}\left(\frac{1}{\sqrt{3}}\cot\left(\frac{n\pi}{6n+1}\right)\right).$$

□

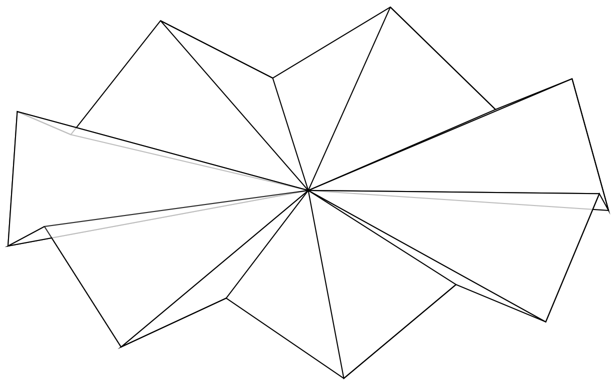

Figure 5.3: A singular 16-gon. The 2-simplices are all isometric and isosceles, with angle (say) $\pi/6$ at the centre and radius r, so this is a picture of $D(16,12,r)$.

Denote the right hand side of the above inequality by $r_{\max}(n)$. This decreases roughly logarithmically with n. For illustration: $r_{\max}(1) \approx 0.62$, $r_{\max}(10) \approx 0.20$, $r_{\max}(100) \approx 0.06$ and $r_{\max}(1000) \approx 0.02$.

5.1.3 Singular polygons

A regular hyperbolic or Euclidean n-gon, with radius r, can be regarded as a simplicial complex with a vertex o at the centre, and n isometric isosceles 2-simplices with two sides of length r either side of an angle of $2\pi/n$, identified in a cycle around o. Denote such an isosceles 2-simplex by $T(n,r)$.

For any integer m, we may obtain a singular 2-complex structure on a disc by identifying m copies of $T(n,r)$ in the analogous manner. For $m < n$, the central vertex o has local positive curvature: that is, its link is a loop of length $< 2\pi$ (see Remark 2.2.21). For $m = n$, this is the usual regular n-gon, and for $m > n$, the central vertex o has local negative curvature: the link has length $> 2\pi$. We denote this singular disc by $D(m,n,r)$. See Figure 5.3.

Remark 5.1.12. For any $k < n/2$, we may define a segment of length k in $D(m,n,r)$ in exactly the same way as for the regular n-gon $D(n,n,r)$. The isometry type of such a segment depends on n, r, and the underlying metric (i.e. hyperbolic or Euclidean), but it does not depend on m. Moreover, Lemma 5.1.9 still holds. See Figure 5.4.

5.2 Proof of the main theorem

Proof of Theorem 5.1.4. By Definition 5.1.2 no relator from R is a proper power. Also assume, without loss of generality, that the relators are distinct up to cyclic permutation and inverses

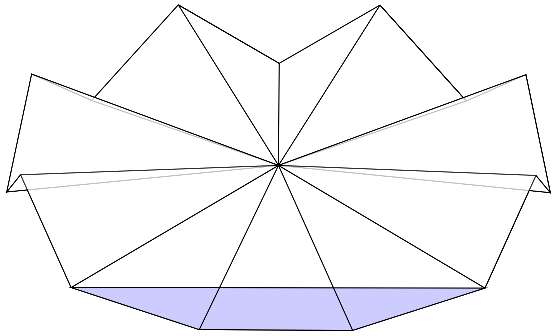

Figure 5.4: A segment in a singular $D(m,n,r)$ is isometric to a segment of the same length in $D(n,n,r)$.

(if not, we can simply delete relators until this is the case). The small cancellation condition now reduces to the intuitive condition that, in the disjoint set of labelled cycles corresponding to the relators, the maximum length of any labelled path which appears at least twice is $< g/6$, where g is the minimum length of any relator. For the rest of the proof, we refer to these labelled cycles (with either orientation) as relators.

Denote by X the presentation complex corresponding to $\langle S \mid R \rangle$. This consists of a bouquet of circles B, labelled by S, and for each relator r_i a disc D_i whose boundary is attached to the path labelled by r_i. We equip X with the following metric. Applying Proposition 5.1.11, we may choose r such that Lemma 5.1.9 holds for the regular hyperbolic g-gon of radius r. For each $i \geq 1$, let $g_i = |r_i|$, and metrize each disc projecting to D_i by a singular hyperbolic g_i-gon $D(g_i, g, r)$. Since the boundary of $D(g_i, g, r)$ is a g_i-cycle with all edges of length some constant λ, the metrics on each disc induce a well-defined metric on X. Moreover, by equipping each disc with the simplicial complex structure depicted in Figure 5.3, we obtain a simplicial complex structure on X. Denote by Y the universal cover of X with the induced metric, and by Z the preimage in Y of the bouquet B.

A piece from the presentation corresponds to a maximal path in Z which is contained in the boundary of two distinct discs (either two distinct lifts of the same D_i, or lifts of two different D_i). We refer to such paths also as “pieces”. Each such piece subtends a segment of each of these two discs, and these two segments are isometric by construction. Therefore, we may subdivide to make the segments into simplicial subcomplexes, and pass to a quotient complex of Y in which the two segments are identified. We will refer to such an identification as a “fold”. Note also that, after such an identification, the images of the two affected boundary loops both still bound well-defined, isometric discs; in particular, we can safely apply the same operation again to other segments whether or not they intersect the pair already folded. Of course, we may need to subdivide Y again if we wish to retain a simplicial complex structure.

So, let $\overline{Y}$ denote the quotient space obtained by folding the corresponding pair of segments for every piece in Y. Denote by $\overline{Z}$ the image of Z under this map. We claim that $\overline{Y}$ is a CAT(-1) space with a geometric action of G.

Firstly, although we identify infinitely many pieces, Y is a locally finite complex and there is an upper bound on the length of pieces. Thus there are only finitely many pieces containing each vertex of Z, and so $\overline{Y}$ is a well defined, locally finite, piecewise hyperbolic simplicial 2-complex. It is therefore sufficient to check the link condition on vertices.

There are two types of vertex to consider: those which are images of vertices in Z, and those which are not; i.e. images of points in the interior of discs of Y. So first, let v' be a vertex in the image of Z, and let v be the vertex of Z mapping to v'. This is unique, since vertices in Z are never identified under $Y \to \overline{Y}$.

Topologically, $\mathrm{Lk}_Y(v)$ is a graph with two vertices $s^{\pm}$ for each $s \in S$, and an edge $s^- t^+$ whenever st appears as a subword in any relator. These edges all have the same length, equal to the interior angles in the discs. There are additional vertices corresponding to the simplicial subdivision of the discs, but these have degree 2 and we ignore them.

The quotient $Y \to \overline{Y}$ induces a map $\mathrm{Lk}_Y(v) \to \mathrm{Lk}_{\overline{Y}}(v')$, and we can describe this map very precisely. For each pair st which appears as a subword of some piece (or whose inverse appears), there are at least two edges connecting s^- and t^+; one edge for each appearance of st in a relator. In $\mathrm{Lk}_{\overline{Y}}(v')$, all of these edges are folded (in the genuine Stallings sense) to a single edge.

The other operation performed by the map $\mathrm{Lk}_Y(v) \to \mathrm{Lk}_{\overline{Y}}(v')$ corresponds to the case where v is the initial or final vertex of a piece in Z. Suppose there is a piece p ending with s, and let $t_1, \dots, t_k$ be the set of all generators which occur in the set of relators immediately following the piece p. Then, for each i, there is an edge in $\mathrm{Lk}_Y(v)$ from s^- to t_i^+, and the map $\mathrm{Lk}_Y(v) \to \mathrm{Lk}_{\overline{Y}}(v')$ is a metric fold which identifies a short initial subpath of each of these edges. This second operation is a homotopy equivalence; in particular, any unbased loop in $\mathrm{Lk}_{\overline{Y}}(v')$ has a preimage (up to homotopy equivalence) in the graph obtained by identifying multiple edges in $\mathrm{Lk}_Y(v)$. Lemma 5.1.9 implies that, even if initial and final subpaths of some edge of $\mathrm{Lk}_Y(v)$ have been folded in this way, the length of the central path is still $> 2\pi/3$.

We may therefore express the map from $\mathrm{Lk}_Y(v) \to \mathrm{Lk}_{\overline{Y}}(v')$ as a composition of two folding maps, one folding multiple edges, and one folding short initial segments of edges, as shown

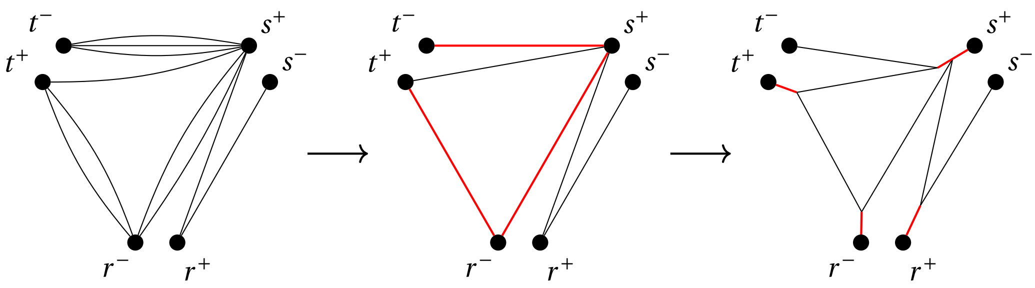

Figure 5.5: A planar example of the map on links of vertices of the first type. Thick red paths show regions where the map just applied was non-injective.

in Figure 5.5. The intermediate graph contains no bigons, and since the second map is a homotopy equivalence, all loops in the final graph must contain at least three unidentified central subpaths of edges, and thus have length $> 2\pi$ as required.

We now address the second type of vertex in $\overline{Y}$. Let $v' \in \overline{Y}$ be a vertex which is the image under $Y \to \overline{Y}$ of a point in $Y - Z$. If v' is the image of one of the singular points in the centre of a disc, then its link is a circle of length $\geq 2\pi$. This is because no segments intersect this point, so its link is unaltered by the map $Y \to \overline{Y}$. Hence we may assume v' is not of this type.

Let $v_1, \ldots, v_k$ be the set of all points in Y mapping to v_i; note that this is indeed finite by local finiteness of Y. For each i, let D_i be the disc in Y containing v_i in its interior. The link $L = \mathrm{Lk}_{\overline{Y}}(v')$ is a quotient of a disjoint union of k round circles $C_1 \sqcup \cdots \sqcup C_k$. We abuse notation and refer to the image of each C_i in L as C_i.

The quotient map $\varphi \colon C_1 \sqcup \cdots \sqcup C_k \to L$ is induced by the process of folding segments in Y. If such a segment contains one of the vertices v_i, then v_i is contained either in the diagonal bounding the segment or in its interior. The map is therefore a multiple composition of two possible operations: identifications of subarcs of length π between the C_i, or identifications of complete circles C_i. The latter operation does not affect L and so we can assume that it does not occur. We now show that L cannot contain any closed geodesics of length $< 2\pi$. This is trivial in the case $k = 1$, so assume $k \geq 2$.

The map φ consists of repeatedly identifying length π subarcs of different C_i. Now, each v_i is the point of intersection of a number of diagonals bounding segments in D_i. The intersection of all these segments is a polygonal region in D_i bounded by two of these diagonals, and it follows from Lemma 5.1.9 that the subarc α_i of C_i corresponding to this intersection has length at least $2\pi/3$ (see Figure 5.6).

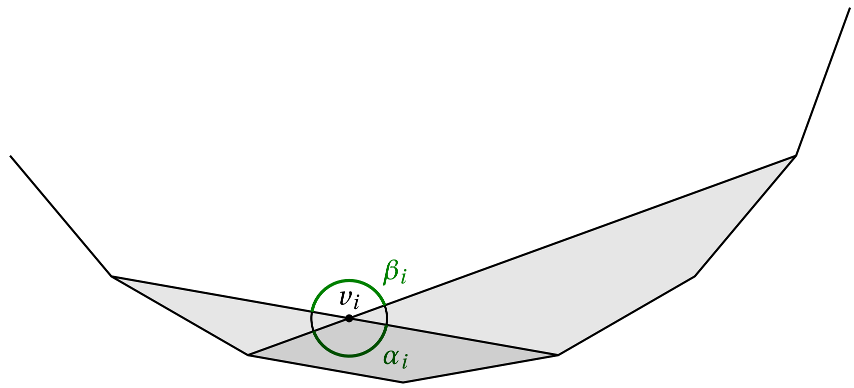

Figure 5.6: The link of the vertex v_i in the disk D_i. Lemma 5.1.9 implies that both arcs α_i and β_i have length $> 2\pi/3$.

Each identification of arcs under φ is induced by an isometry between segments in discs D_i and D_j. It follows that φ isometrically identifies the arcs α_i for all i. Refer to the image of these arcs in L as α. In particular, the set $C_1 \cap \cdots \cap C_n \subset L$ is nonempty.

Also, for each i, there is an open arc in L which is contained in the image of only C_i, and not C_j for $j \neq i$. Call this arc β_i (see Figure 5.6), using the same name to refer to the image in L or the subarc in C_i.

Consider the subgraph $L_0 = L - \bigcup_i \beta_i$ obtained by deleting every β_i from L. This is equal to the subgraph $\bigcup_{i \neq j} \left(C_i \cap C_j\right)$, and each $C_i \cap C_j$ is a connected arc containing α. Moreover, intersections of three or more C_i are also connected arcs containing α. It follows that L_0 is simply connected.

Now let ℓ be a geodesic loop in L. Since L_0 is simply connected, ℓ contains at least one β_i. If ℓ contains only one β_i (say, β_1), then $C_1 - l_1$ is a geodesic path of length $2\pi - |\beta_i|$ between the endpoints of β_1 in L_0, and so must coincide with ℓ by simple connectedness of L_0; hence ℓ coincides with C_1 and has length 2π. If ℓ contains at least three β_i, then it has length $> 2\pi$ since each β_i has length $> 2\pi/3$.

The remaining case is that ℓ contains precisely two β_i; say β_1 and β_2. Since L_0 is a tree, each component of $\ell - \{\beta_1 \cup \beta_2\}$ is contained in $C_1 \cup C_2$, and hence $\ell \subset C_1 \cup C_2$. Now $C_1 \cup C_2 \subset L$ is obtained from the circles C_1 and C_2 by identifying arcs of length π, all containing the common arc α. If only a single arc of length π is identified, then the 1-complex $C_1 \cup C_2$ clearly contains no loops of length $< 2\pi$, in particular ℓ has length at least 2π. Otherwise, C_1 and C_2 intersect in L an arc longer than π. In this situation, there must be two intersecting segments in D_1 which are identified respectively with two overlapping segments in D_2. Therefore, there

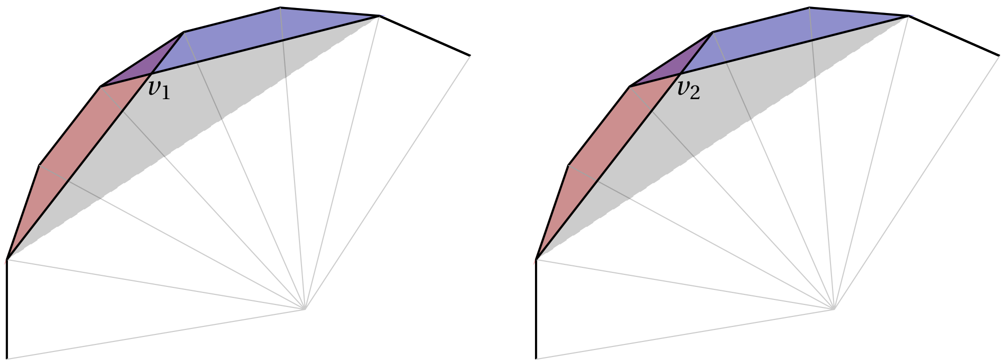

Figure 5.7: If the red and blue pairs of segments are identified, then the union of the corresponding pieces is also a piece, and so the grey segments are also identified.

is a larger piece whose corresponding pair of segments was not identified (see Figure 5.7). This contradicts the construction of the map $Y \to \overline{Y}$.

Hence, $\overline{Y}$ is CAT(-1). Since the action of G takes pieces to pieces, the invariance of the metric under the action of G is clear. Since Y was the universal cover of X, it follows that the action of G on $\overline{Y}$ is also a universal covering action; in particular, it is geometric. This completes the proof.

□

Remark 5.2.1. Proposition 5.1.11 enables us to compute an approximate volume for the negatively curved complex constructed when we prove Theorem 5.1.4. The radius of each disc used in the construction is approximately $r_{\max}(g/6)$, and the area of $D(g_i, g, r)$ is approximately $\pi r^2 g_i / g$ (approximating a flat polygon as a Euclidean disc). Hence, the area of the metrized presentation complex, before any folding is carried out, is approximately equal to

$$A \approx \frac{\Sigma_i g_i}{g} \times \pi r_{\max}(g/6)^2,$$

where

$$r_{\max}(n) = \cosh^{-1}\left(\frac{1}{\sqrt{3}}\cot\left(\frac{n\pi}{6n+1}\right)\right).$$

Of course, this is a slight overestimate, and decreases when folding is applied in a way that depends precisely on the pieces of the presentation. We expect that it is asymptotically accurate for large g and large small cancellation constant.

Chapter 6

A gluing theorem for negatively curved complexes

In Chapter 5, we described an explicit negatively curved metric on certain small cancellation groups using negatively curved 2-complexes. The aim was to provide a partial answer to Question 2.4.13, which asked whether every hyperbolic group was CAT(-1). In this chapter, we combine this with the setting of a thin graph of spaces (which was central to Chapter 4), in order to answer Question 2.4.13 in a different context. The main class of groups we are interested in are *limit groups*, introduced by Sela [Sel01] and widely studied due to their usefulness in understanding homomorphisms from a finitely generated group to a free group (for more details, see Section 6.3 and the references therein). Limit groups were shown in [AB06] to be CAT(0), and we improve this to the following theorem.

Theorem 6.3.1. Let G be a limit group. Then G is CAT(-1) if and only if G is hyperbolic.

To prove this, continuing to work with negatively curved 2-complexes, we devise a gluing theorem (Theorem 6.2.1), in the spirit of the gluing theorems presented in [BH99, II.11]. Essentially, we would like to take two negatively curved 2-complexes, glue them together along a tube, and find a hyperbolic metric on the resulting complex. Imposing a hyperbolic metric on the tube is problematic, since hyperbolic annuli have a non-geodesic boundary component. To get around this issue, we modify the metric on the two complexes, concentrating negative curvature at vertices, and giving "room" to glue in the tube. Care is needed to ensure that the pieces we glue on do not themselves combine to give positive curvature, and this is dealt with in Lemma 6.1.18.

The most closely related theorem to ours in the literature is Bestvina and Feighn's gluing theorem [BF92] for δ-hyperbolic spaces, and thus for hyperbolic groups. Our theorem is complementary to theirs; our hypotheses are stronger, but so is our conclusion.

We present the proof of the gluing theorem in Section 6.2, and then in Section 6.3 we introduce limit groups. After giving some background and listing some useful properties, we will exploit the rich structure theory of limit groups to allow us to apply our gluing theorem, and hence to prove Theorem 6.3.1. In Section 6.4 we provide two more applications of the gluing theorem, showing that it can be applied to the JSJ decomposition of a torsion-free hyperbolic group (Theorem 6.4.4), and deducing that hyperbolic graphs of free groups with cyclic edge groups are CAT(-1) (Theorem 6.4.8). A consequence of the application to the JSJ decomposition, together with the Strong Accessibility Theorem of Louder–Touikan [LTar] (see also [DP01]), is that we may reduce the question of whether a hyperbolic group is two-dimensionally CAT(-1) to its rigid subgroups (those which do not admit a non-trivial free or cyclic splitting).

Remark 6.0.2. This chapter has appeared as the main part of the paper [Bro16]. We would like to thank the editors, and the anonymous referee for several helpful comments which have improved the exposition.

6.1 Preliminaries

6.1.1 Metrizing graphs of spaces

Our definition of a graph of spaces (Section 2.1.3) is purely topological, and a priori, a graph of spaces X is not equipped with a metric. However, if we have a metric on each vertex and edge space, then we may metrize the cylinders $X_e \times [0,1]$ using the product metric (with the standard metric on $[0,1]$), and then the quotient pseudometric on X will be a true metric if the attaching maps are suitably nice. A sufficient condition is that they be local isometries, in the sense of Definition 2.2.8.

However, even in the case that the attaching maps are not local isometries—or even when metrics on the edge spaces are not specified—a graph of spaces may possess other metrics. Theorem 6.2.1 defines a metric on a graph of spaces which does not come from a product metric on the edge space cylinders.

6.1.2 Thinness

As we saw in Chapters 3 and 4, malnormality of edge groups in vertex groups (or *thinness*) is a common assumption made on graphs of groups in certain contexts. In the gluing theorem presented in this chapter, we will make this assumption on a subset of the vertex groups of the graph of groups. This is a stronger condition than the “annuli flare” condition used by Bestvina and Feighn in their gluing theorem for δ-hyperbolic spaces [BF92], which could be thought of as a coarse analogue of our theorem. As discussed in Section 4.4, the geometric consequences of malnormality for a subgroup are not fully understood (in particular, the circumstances under which it implies quasiconvexity), and it may be possible to replace it with a weaker assumption in the future. See Remark 6.5.3 for more details.

The technique at the heart of our proof of Theorem 6.2.1 is to replace simplices in a negatively curved complex by *comparison simplices* taken from a rescaled hyperbolic space; replacing hyperbolic simplices with simplices from, say, $M^2_{1/2}$. This is outlined in the following section.

6.1.3 Excess angle and comparison simplices

Remark 6.1.1. Let $k \le 0$ and let S be an M^2_k-simplex with 1-skeleton $S^{(1)}$. Then for any $k' \le 0$, there exists a $M^2_{k'}$-simplex S' with 1-skeleton $S'^{(1)}$ isometric to $S^{(1)}$; indeed, $S'^{(1)}$ is a comparison triangle for $S^{(1)}$. Each angle of S' is strictly larger than the corresponding angle of S if $k < k' \le 0$, and strictly smaller if $k' < k \le 0$. We say S' is a *comparison simplex for S of curvature k'* (or a *comparison $M^2_{k'}$-simplex* for S).

Remark 6.1.2. Remark 6.1.1 does not have a direct analogue for higher dimensional simplices—although we may define a comparison simplex in dimension n, it might not exist if the difference between k' and k is too great. Indeed, given any $-1 < k \le 0$, it is easy to construct a hyperbolic 3-simplex S for which there is no k-comparison simplex; we can even construct one with any 2-simplex as its base (this is a generalisation of a construction mentioned in [CD95]). Take a hyperbolic plane P. Let Δ be a hyperbolic triangle in P with vertices x, y and z. Denote by c the incentre of Δ, and the inradius by r. Now, embed P isometrically in $\mathbb{H}^3$, and take a point p in $\mathbb{H}^3 - P$ some small distance ϵ from c. Consider the 3-simplex S with vertices x, y, z, p. One face of S is the simplex bounded by Δ, and call the other three faces X, Y and Z according to the vertex opposite them in S. In Z, note that $d\big([x,y],p\big) < r+\epsilon$ (the 2-simplex Z has *bounded height*), and it follows that $d(x,p)+d(y,p) < d(x,y)+2r+2\epsilon$. It

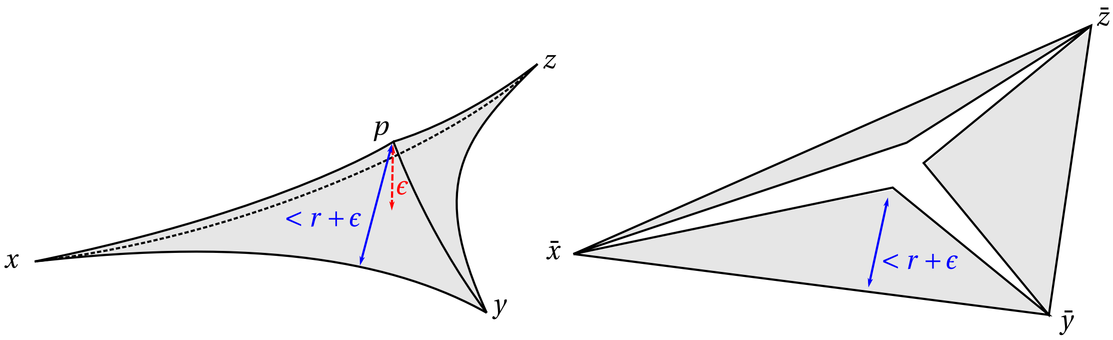

Figure 6.1: The hyperbolic 3-simplex on the left has no Euclidean comparison simplex.

follows from two applications of the reverse triangle inequality that any comparison simplex $\overline{Z}$ must also have height $d\left([\bar{x}, \bar{y}], \bar{p}\right)$ bounded by $r + \epsilon$. The same is true for X and Y. Now, let us try to construct a comparison simplex $\bar{S}$ for S with curvature k. Begin with a k-comparison simplex $\overline{\Delta}$ for Δ. Its inradius $\bar{r}$ is greater than r. Set $\epsilon = \bar{r} - r$. Now, in the purported comparison simplex $\overline{S}$, the point $\bar{p}$ is of distance $< r + \epsilon$ from each edge $[\bar{x}, \bar{y}]$, $[\bar{y}, \bar{z}]$, $[\bar{z}, \bar{x}]$. It follows that the inradius $\bar{r}$ of $\overline{\Delta}$ is less than $r + \epsilon$; a contradiction. See Figure 6.1.

Definition 6.1.3. If K is a negatively curved simplicial 2-complex, then we may form a combinatorially isomorphic complex K' by replacing each M_k^2-simplex S with the comparison simplex S' of curvature k', where $k < k' < 0$. Then K' is an $M_{k'}$-simplicial 2-complex, which we call the *comparison $M_{k'}$-complex* for K. Note that gluing maps between faces of simplices remain isometries, and so this is a well-defined $M_{k'}$-simplicial complex.

There is a natural combinatorial isomorphism $' \colon K \to K'$, which can be realised by a homeomorphism. We also use the same notation for the induced combinatorial isomorphism $' \colon \mathrm{Lk}(v, K) \to \mathrm{Lk}(v', K')$ for each vertex $v \in K^{(0)}$.

Lemma 6.1.4. *Let K and K' be as in Definition 6.1.3. Then K' is locally* $\mathrm{CAT}(k')$*; in particular, it is a negatively curved 2-complex.*

Proof. By Remark 6.1.1, angles at vertices in K' are larger than in K. Since links in K and K' are metric graphs, it follows that the map $' \colon \mathrm{Lk}(v, K) \to \mathrm{Lk}(v', K')$ strictly increases the distance between pairs of points. Remark 2.2.21 then implies that links in K' are $\mathrm{CAT}(1)$, and hence K' is locally $\mathrm{CAT}(k')$. □

We quantify this as follows.

Definition 6.1.5. For a negatively curved simplicial 2-complex K and comparison complex K', the *excess angle* δ is defined by:

$$\delta(K', K) = \inf\{\theta' - \theta\}$$

where θ ranges over all vertex angles of 2-simplices $S \subset K$ and θ' is the corresponding angle in the comparison simplex $S' \subset K'$. Equivalently:

$$\delta(K', K) = \inf\{d_{\mathrm{Lk}(v',K')}(a', b') - d_{\mathrm{Lk}(v,K)}(a, b) \mid a, b \in \mathrm{Lk}(v, K)^{(0)}, v \in K^{(0)}\}$$

If K is finite (or more generally, if the set $\mathrm{Shapes}(K)$ of isometry types of simplices is finite), then $\delta(K', K) > 0$.

Remark 6.1.6. Suppose $\gamma = (e_1, \ldots, e_n)$ is a locally geodesic simplicial path in K, with $\iota(e_i) = v_{i-1}$, $\tau(e_i) = v_i$ for $i = 1, \ldots, n$. Then $\gamma' = (e'_1, \ldots, e'_n)$ is a local geodesic in K', and for each $i = 1, \ldots, n-1$ the angle subtended by γ at v'_i is at least $\pi + 2\delta(K', K)$; that is:

$$d_{\mathrm{Lk}(v'_i,K')}(e'_i, e'_{i+1}) \geq \pi + 2\delta(K', K).$$

In the case where γ is a closed local geodesic, the same also holds for the angle at $v'_0 = v'_n$ between e'_n and e'_1 and so γ' is still a closed geodesic.

Remark 6.1.7. We will typically use the excess angle only in the spirit of Remark 6.1.6 above. In this setting, we could relax the requirement that K or $\mathrm{Shapes}(K)$ is finite, provided we insist that K is *locally* finite. This is because we only need to consider the finitely many angles around the finite subset of K consisting of simplices which intersect γ.

Remark 6.1.8. The process of replacing hyperbolic simplices by simplices of different curvature is equivalent to rescaling the 1-skeleton by a constant factor and keeping the curvature fixed, by definition of the rescaled spaces M^n_k. We have chosen the former approach here, though the latter may have advantages in future; for example, calculating explicit volumes for the complexes constructed.

Remark 6.1.9. In this chapter, we will rarely specify the exact value of the curvature for a simplicial 2-complex, and indeed we will not always ensure that all the simplices in such a complex have equal curvature. This is not a problem provided that the curvatures are all negative and bounded above, since letting $\bar{k}$ be the upper bound (i.e. the curvature of smallest absolute value which appears among the different simplices) and replacing every simplex with its $\bar{k}$-comparison simplex, we obtain a simplicial 2-complex with isometric 1-skeleton, greater or equal face angles, and all simplices of equal curvature.

6.1.4 Hyperbolic triangles, quadrilaterals and annuli

Remark 6.1.10. For any angle $0 < \theta < \pi$, any $a, b > 0$, and any $k \leq 0$, there exists a 2-simplex in M_k^2 with angle θ between two sides of length a, b. For brevity, we will sometimes refer to such a simplex as an (a, θ, b)-*fin*.

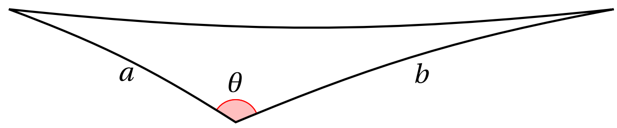

Figure 6.2: An (a, θ, b)-fin.

Definition 6.1.11. A *Lambert quadrilateral* is a quadrilateral in M_k^2 (for $k < 0$) with three angles equal to $\pi/2$, as illustrated in Figure 6.3. The bottom edge, called the *base*, has length a, and the top edge, called the *summit*, has length c. The single angle θ not equal to $\pi/2$ is called the *summit angle*.

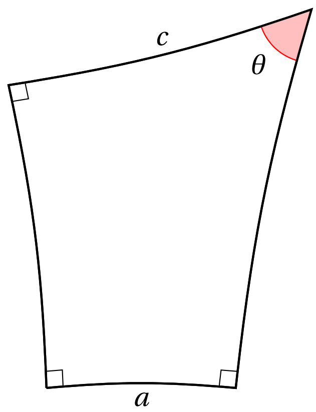

Figure 6.3: A Lambert quadrilateral

Lemma 6.1.12. *In a Lambert quadrilateral with $k = -1$,*

$$\sin\theta = \frac{\cosh a}{\cosh c},$$

where labels are as in Figure 6.3.

Proof. We will give a short proof using the hyperbolic sine and cosine rules. Many such identities can also be found in [Bea83, Chapter 7], proved using the definition of the hyperbolic metric on the upper half plane.

Divide the quadrilateral along diagonals, and label as follows:

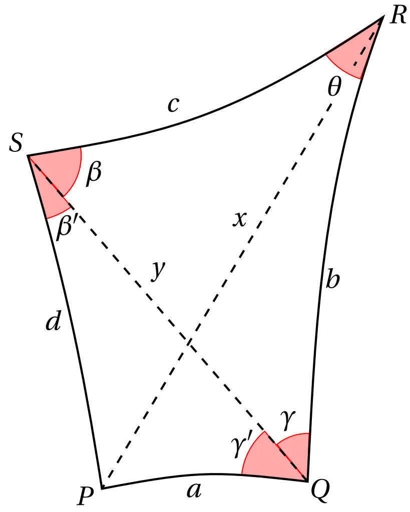

Note that $\beta + \beta' = \pi/2 = \gamma + \gamma'$, and hence $\cos\gamma = \sin\gamma'$, and $\cos\beta' = \sin\beta$. Firstly, applying the (hyperbolic) Pythagorean theorem in triangles PQR and PSR gives

$$\cosh a \cosh b = \cosh x = \cosh c \cosh d. \tag{6.1}$$

The cosine rule in triangle QRS gives

$$\begin{aligned}\cosh c &= \cosh b \cosh y - \sinh b \sinh y \cos\gamma \\ &= \cosh b \cosh y - \sinh b \sinh y \sin\gamma'.\end{aligned} \tag{6.2}$$

The sine rule in the right-angled triangle PQS gives

$$\sinh y \sin\gamma' = \sinh d,$$

and substituting this into (6.2) gives

$$\cosh c = \cosh b \cosh y - \sinh b \sinh d. \tag{6.3}$$

Similarly applying the cosine rule to triangle PQS and substituting for $\cos\beta' = \sin\beta$ using the sine rule in triangle QRS, we obtain

$$\cosh a = \cosh d \cosh y - \sinh b \sinh d \sin\theta. \tag{6.4}$$

Now, from (6.1), we have

$$\cosh d = \frac{\cosh b \cosh a}{\cosh c},$$

and substituting this into (6.4) gives

$$\begin{aligned} & \cosh a = \frac{\cosh b \cosh a \cosh y}{\cosh c} - \sinh b \sinh d \sin\theta. \\ \implies \quad & \cosh c = \cosh b \cosh y - \sinh b \sinh d \sin\theta \left(\frac{\cosh c}{\cosh a}\right). \qquad (6.5) \end{aligned}$$

Finally, equating (6.3) and (6.5) we obtain

$$\begin{aligned} & \cosh b \cosh y - \sinh b \sinh d = \cosh b \cosh y - \sinh b \sinh d \sin\theta \left(\frac{\cosh c}{\cosh a}\right) \\ \implies \quad & 1 = \sin\theta \left(\frac{\cosh c}{\cosh a}\right) \\ \implies \quad & \sin\theta = \frac{\cosh a}{\cosh c}, \end{aligned}$$

as required. □

Lemma 6.1.13. *A Lambert quadrilateral in* $\mathbb{H} = M^2_{-1}$ *with summit length* c *and summit angle* θ *exists if and only if* $\theta_c < \theta < \pi/2$, *where*

$$\theta_c = \sin^{-1}\left(\frac{1}{\cosh c}\right).$$

Proof. Recall that for any two non-crossing geodesics in $\mathbb{H}$, there is a unique geodesic perpendicular to both. Now, construct a geodesic segment C of length c in $\mathbb{H}$, with a perpendicular geodesic B at one end, and a perpendicular geodesic B' at the other end. Clearly B' and B do not cross, but no (non-degenerate) Lambert quadrilateral exists with summit C, since the unique geodesic perpendicular to B and B' is C. Now continuously decrease the angle θ between C and B'. By convexity of the metric, the unique geodesic perpendicular to B and B' will form a Lambert quadrilateral on the same side of C as the angle θ, until θ reaches the value θ_c at which B and B' form an ideal hyperbolic triangle with C. This can be calculated (for example) by setting $a = 0$ in the formula from Lemma 6.1.12:

$$\theta_c = \sin^{-1}\left(\frac{1}{\cosh c}\right),$$

as required. □

Note that, given such c and θ, the quadrilateral is then uniquely determined. Note also that the fact that the minimal distance between B and B' increases continuously from zero as θ increases from θ_c follows directly from the formula in Lemma 6.1.12. For an alternative proof of Lemma 6.1.14 below, we could show first that this distance is a continuous function of the distance between the points B and B' in the boundary circle of the hyperbolic plane (which is straightforward to verify), and then perturb a right-angled ideal triangle to obtain the required Lambert quadrilateral.

Lemma 6.1.14. *For any* $0 < \theta < \pi$*, and for any* a*,* c *such that* $c > a > 0$*, there exists* $k < 0$ *and a Lambert quadrilateral in* M^2_k *with base length* a*, summit length* c *and summit angle greater than* $\theta/2$*.*

Proof. First, let

$$\bar{c} = \cosh^{-1}\left(\frac{1}{\sin\left(\frac{\theta}{2}\right)}\right),$$

so that $\frac{\theta}{2} = \theta_{\bar{c}}$. For any $\bar{a} < \bar{c}$, we have

$$1 > \frac{\cosh \bar{a}}{\cosh \bar{c}} > \frac{1}{\cosh \bar{c}} = \sin\left(\frac{\theta}{2}\right)$$

and hence

$$\frac{\pi}{2} > \sin^{-1}\left(\frac{\cosh \bar{a}}{\cosh \bar{c}}\right) > \frac{\theta}{2}.$$

Applying Lemma 6.1.13, we see that a Lambert quadrilateral exists in M^2_{-1} with base length $\bar{a}$, summit length $\bar{c}$ and summit angle $\sin^{-1}\left(\frac{\cosh \bar{a}}{\cosh \bar{c}}\right) > \theta/2$. By multiplying the metric by a factor $c/\bar{c}$, we obtain a Lambert quadrilateral in M^2_k, where $k = -(\bar{c}/c)^2$, satisfying the required conditions. □

Lemma 6.1.15. *For any* $0 < \theta < \pi$*, and any* A*,* C *such that* $C > A > 0$*, there exists* $k < 0$ *and a locally* CAT(k) *annulus with one locally geodesic boundary component of length* A*, and one* *boundary component of length C which is locally geodesic everywhere except for one point where it subtends an angle greater than θ.*

Proof. Apply Lemma 6.1.14 with the same θ, $a = A/2$ and $c = C/2$. Now take two copies of the resulting Lambert quadrilateral, and glue together two pairs of sides to obtain the required annulus (see Figure 6.4). □

Remark 6.1.16. We have not been able to find any previous reference to the existence of a rectangle as in Lemma 6.1.14 (and hence an annulus as in Lemma 6.1.15), and it is worth pointing out that the existence of such a rectangle runs counter to the common intuition that, in the hyperbolic case, small rectangles are "approximately Euclidean". In fact, as in the proof of Lemma 6.1.14, we can find a hyperbolic rectangle with arbitrarily low area (since the area is equal to the difference between 2π and the sum of the interior angles), and yet arbitrarily high ratio between the lengths of a pair of opposite sides.

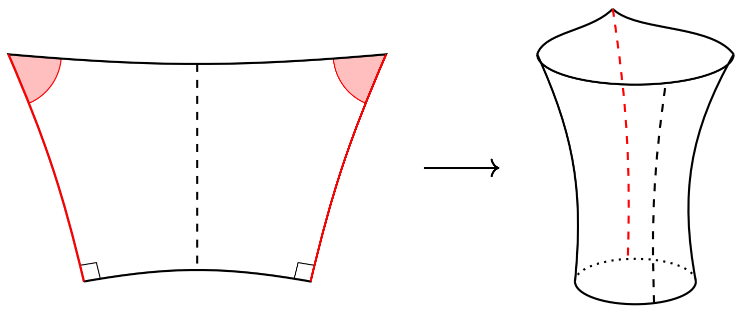

Figure 6.4: Gluing two Lambert quadrilaterals to obtain an annulus. Identify the two sides of the left picture. The bottom edge of the annulus has length A, the top edge has length C, and the angle in the top edge is $> \theta$.

6.1.5 Transversality

As discussed at the beginning of the chapter, we would like to build new negatively curved complexes by gluing fins and annuli to existing ones. To avoid introducing positive curvature, for example by identifying a pair of adjacent sides in two fins, we would like to glue along paths that intersect transversely. The following lemma ensures that this can always be arranged.

Definition 6.1.17. Let $\gamma = (e_1, \ldots, e_n)$ and $\gamma' = \left(e'_1, \ldots, e'_{n'}\right)$ be two closed geodesics in K. We say γ and γ' *intersect transversely* if $e_i \neq e'_j$ and $\bar{e}_i \neq e'_j$ for all i and j (that is, the two loops do not share an edge). Similarly, we say γ *intersects itself transversely* if $e_i \neq e_j$ and $\bar{e}_i \neq e_j$ for all $i \neq j$.

Lemma 6.1.18. *Let K be a negatively curved simplicial 2-complex of curvature k. Let $\{\gamma^{(1)}, \dots, \gamma^{(m)}\}$ be a collection of simplicial closed geodesics such that the corresponding cyclic subgroups of $\pi_1(K)$ form a malnormal family. Then there is a 2-complex $\bar{K}$ such that:*

- *$\bar{K}$ is a negatively curved simplicial complex of curvature $\bar{k}$, where $k \le \bar{k} < 0$;*
- *there is an inclusion $i\colon K \to \bar{K}$ and a deformation retraction $r\colon \bar{K} \to K$;*
- *there are closed geodesics $\bar{\gamma}^{(1)}, \dots, \bar{\gamma}^{(m)}$ in $\bar{K}$, such that $r\left(\bar{\gamma}^{(i)}\right) = \gamma^{(i)}$, which intersect themselves and each other transversely.*

Proof. The complex $\bar{K}$ will be obtained from K by gluing on fins to shorten the intersection between the closed geodesics $\gamma^{(i)}$. We will ensure that there is "room" to glue these fins by repeatedly taking the comparison complex and obtaining an excess angle at the vertices.

Let $\gamma^{(i)} = \left(e_1^{(i)}, \dots, e_{n(i)}^{(i)}\right)$. Let I be the number of repetitions (ignoring orientation) in the list

$$\left\{e_j^{(i)} \;\middle|\; i = 1, \dots, m,\ j = 1, \dots, n(i)\right\};$$

that is, I counts $d-1$ for each edge of K that occurs d times in the union of the $\gamma^{(i)}$. Thus, I counts the number of failures of transversality of the set of geodesics; if $I = 0$, then the $\gamma^{(i)}$ intersect themselves and each other transversely, and the conclusions of the lemma hold already. To prove the lemma, we will show that if $I > 0$, it is always possible to find a complex satisfying all the conclusions of the lemma *other* than transverse intersection of the geodesics, but with I reduced by 1 compared to K. The required result follows from this by induction.

So, suppose $I > 0$. The first stage is to replace K by a comparison complex K' of curvature (say) $k/2$. Since the isomorphism $'\colon K \to K'$ is a homeomorphism and preserves the $\gamma^{(i)}$, we may safely proceed with K' instead of K. Let δ be the excess angle $\delta(K, K')$; if Shapes(K) is not finite, then we may ensure $\delta > 0$ by taking the infimum in Definition 6.1.5 only over the finitely many angles at vertices contained in the $\gamma^{(i)}$ (see Remark 6.1.7).

Since $I > 0$, we may assume that there are two closed geodesics $\gamma = \{e_1, \dots, e_r\}$, $\gamma' = \{e'_1, \dots, e'_s\}$, obtained by relabelling and possibly reversing the $\gamma^{(i)}$, such that $e_1 = e'_1$ with orientation (note that γ and γ' may be distinct relabellings or reversals of the same $\gamma^{(i)}$). Without loss of generality, $r \le s$.

Now, suppose that, for all $z \in \mathbb{Z}$, $e_{z \bmod r} = e'_{z \bmod s}$. In the case that γ and γ' are relabellings of different $\gamma^{(i)}$, this implies that the loops are both powers of a common loop, contradicting the malnormal family assumption. If γ and γ' are relabellings without reversal of the same $\gamma^{(i)}$, then this again implies that $\gamma^{(i)}$ is a proper power of a loop, which contradicts the assumption. The only remaining case is that γ and γ' are relabellings of the same $\gamma^{(i)}$, but one of them is reversed; that is, e'_1 $(= e_1) = \bar{e}_t$ for some $t > 1$. It follows that $e_2 = \bar{e}_{t-1}$, $e_3 = \bar{e}_{t-2}$ and so on—hence, γ must contain either an edge $e = \bar{e}$ (which is forbidden by definition) or an adjacent pair $e, \bar{e}$, which contradicts the fact that the loops are local geodesics.

It follows that there is some $j < s$ such that $e_j = e'_j$ but $e_{j+1} \neq e'_{j+1}$. Let $a = |e_j|$ and $b = \left|e'_{j+1}\right|$. Now, form a complex K^+ from K' by gluing an $(a, \pi-\delta, b)$-fin, also of curvature $k/2$, along e_j and e'_{j+1} (see Figure 6.5). For any closed geodesic ℓ in K' containing $\left(e'_j, e'_{j+1}\right)$, the loop ℓ^+ in K^+ obtained from ℓ by replacing $\left(e'_j, e'_{j+1}\right)$ with the new edge e' (along the top of the fin) is a local geodesic in K^+, and there is a clear deformation retraction to K' sending ℓ^+ to ℓ. So, consider K^+ together with the set of closed geodesics obtained from $\{\gamma^{(1)}, \ldots, \gamma^{(m)}\}$ by replacing any occurrence of $\left(e'_j, e'_{j+1}\right)$ with e'. This satisfies the second conclusion of the lemma, and the first half of the final conclusion; moreover, the value of I is strictly lower than for the original set of loops and the complex K (because we have removed at least one edge in the intersection between γ and γ'). It remains only to check that K^+ is negatively curved.

K^+ is an $M^2_{k/2}$-simplicial complex by construction, so we may check that it is negatively curved using the link condition. The only three links to check are those at v_{j-1}, v'_{j+1} and v_j, since all the others are unchanged from K'. To obtain $\mathrm{Lk}\left(v_{j-1}, K^+\right)$ and $\mathrm{Lk}\left(v'_{j+1}, K^+\right)$, we have simply glued a leaf to the corresponding links in K', and to obtain $\mathrm{Lk}\left(v_j, K^+\right)$, we have (in light of Remark 6.1.6), glued an edge of length $\pi - \delta$ between two vertices of distance $\geq \pi + 2\delta$ in $\mathrm{Lk}\left(v_j, K'\right)$. Neither of these can introduce a failure of the CAT(1) condition (by Remark 2.2.22) and so K^+ is negatively curved. This completes the proof. □

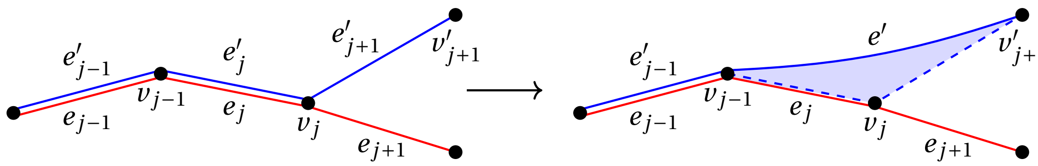

Figure 6.5: Gluing on a fin

Remark 6.1.19. It is not difficult to see that we may improve the above lemma to make the geodesics not only transverse, but disjoint. Indeed, if two geodesics intersect transversely at a vertex v, we can make them locally disjoint by gluing a fin along one of the two geodesics at v. Alternatively, it can be seen as an application of Theorem 6.2.1 in the next section, in the special case where the underlying graph is a star with vertex spaces K for the central vertex, circles for the leaves, and circles for each edge space attached to each of the loops $\gamma^{(i)}$.

6.2 The gluing theorem

We are now ready to state our main gluing theorem. As explained in the subsequent remarks, the statement below is not the strongest available; however, it allows for a straightforward application to the limit group situation (see Section 6.3 for more details).

Theorem 6.2.1. *Let X be a graph of spaces with finite underlying graph Γ, such that:*

1. *Vertex spaces are one of three types:*

   type P *single points,*

   type N *simplicial circles, or*

   type M *connected negatively curved simplicial 2-complexes that are neither points nor circles.*

2. *Edge spaces are either circles or points.*
3. *Each circular edge space connects a type N vertex space to a type M vertex space.*
4. *The images of attaching maps of circular edge spaces are (simplicial) closed geodesics.*
5. *For each type M vertex group in the corresponding graph of groups, the family of subgroups corresponding to incident edge groups is a malnormal family.*

Then X is homotopy equivalent to a negatively curved simplicial 2-complex $\bar{X}$. Moreover, if X has compact vertex spaces, then $\bar{X}$ is compact.

Remark 6.2.2. We believe that the assumption in condition 1 that type N vertex spaces are circles can be relaxed, although this is not necessary for our applications; see Remark 6.5.1. Condition 3 in fact requires only that two type N vertex spaces are not connected by a circular edge space, since if a circular edge space connected two type M vertex spaces we could then subdivide the corresponding edge of Γ and insert a circular vertex space in between. Condition 4 is always achievable; indeed, if K is a locally CAT(-1) space, then every conjugacy class in $\pi_1(K)$ is represented by a unique closed geodesic (see [BH99]). If K is a simplicial complex, then by subdivision we may assume this closed geodesic is simplicial. Conjugacy classes in

$\pi_1(K)$ correspond to free homotopy classes of unbased loops in K, and hence each attaching map is homotopic to an isometry to a simplicial closed geodesic.

Remark 6.2.3. Condition 5 can also be unravelled a little. As mentioned above, in each vertex space, the loops corresponding to circular edge spaces determine conjugacy classes in the fundamental group. As in Lemma 6.1.18, the malnormal family assumption then says that, whichever representatives of these conjugacy classes are chosen when defining the graph of groups, the corresponding edge subgroups have trivial intersection. The other requirement is that the subgroups are individually malnormal; which, for cyclic subgroups of a torsion-free CAT(-1) group, is equivalent to requiring that they are maximal cyclic—that is, not generated by a proper power. This is because torsion-free CAT(-1) groups, being hyperbolic, cannot contain any Baumslag–Solitar subgroup.

Proof of Theorem 6.2.1. To begin, we consider X as a topological graph of spaces, with specified metric on the type M vertex spaces only. We will describe how to metrize the rest of the vertex and edge spaces in the proof.

Consider a type M vertex space X_v of X. Note that X_v together with the set $\left\{\partial_e(X_e) \mid e \in E(\Gamma),\ \iota(e) = v,\ X_e \simeq \mathbb{S}^1\right\}$ of images of incident circular edge spaces satisfies the conditions of Lemma 6.1.18. We may therefore find a negatively curved 2-complex $\bar{X}_v$, equipped with a deformation retraction $r_v\colon \bar{X}_v \to X_v$, inclusion $i_v\colon X_v \to \bar{X}_v$, and a transverse set of loops $\left\{\gamma_e \mid e \in E(\Gamma),\ \iota(e) = v,\ X_e \simeq \mathbb{S}^1\right\}$ such that $r_v(\gamma_e) = \partial_e(X_e)$. Without loss of generality (by taking a comparison complex if necessary) we may assume that $\bar{X}_v$ has an excess angle $\delta_v > 0$ around each loop γ_e (as in Remark 6.1.6) so that the angle subtended through each vertex by each loop γ_e is at least $\pi + 2\delta_v$.

Now consider a type N vertex space X_w. For each circular incident edge space X_e with $\tau(e) = w$, the attaching map $\partial_{\bar{e}}$ is a d to 1 covering map for some $d = d(e)$. Choose a positive number a_w satisfying

$$a_w < \min\left\{\frac{l(\gamma_e)}{d(e)}\ \middle|\ e \in E(\Gamma),\ \tau(e) = w\right\}.$$

We will now describe how to modify the complex X into the complex $\bar{X}$. By Lemma 2.1.8, the two complexes will be homotopy equivalent.

1. Replace each type M vertex space X_v with $\bar{X}_v$ as above, then replace each attaching map $i_v \circ \partial_e$ with a homotopic map to the closed geodesic γ_e.

2. Metrize each type N vertex space X_w to have length a_w.

3. For each circular edge space X_e, by condition 3, we may assume $\iota(e) = v$ and $\tau(e) = w$ where X_w is type N. After the above, the two ends of $X_e \times [0,1]$ are identified with closed geodesics of length $l\left(\gamma_e\right)$ and $a_w d(e)$. Since $a_w d(e) < l\left(\gamma_e\right)$, there exists by Lemma 6.1.15 a negatively curved annulus A_e with one geodesic boundary component of length $a_w d(e)$, and one boundary component of length $l\left(\gamma_e\right)$ subtending an angle greater than $\pi - \delta_v$ at a vertex and geodesic elsewhere. Therefore, we may replace $X_e \times [0,1]$ with (a copy of) A_e such that the metrics on A_e, X_v and X_w are all compatible (ensuring we position the vertex of A_e at some vertex of X_v).

4. Triangulate A_e to ensure it is simplicial.

5. For each type P vertex space X_v, take a point $\bar{X}_v$.

6. For each edge e such that X_e is a point, attach a line between $\bar{X}_{\iota(e)}$ and $\bar{X}_{\tau(e)}$. Note that the endpoints of this line (that is, the attaching maps), as well as its length, are irrelevant from the perspective of homotopy, since the vertex spaces are path connected.

Topologically, to obtain $\bar{X}$ from X we have simply replaced edge cylinders and vertex spaces with homotopy equivalent spaces, and attaching maps with homotopic maps. It follows from Lemma 2.1.8 that $\bar{X}$ and X are homotopy equivalent. Note that $\bar{X}$ is still, topologically, a graph of spaces, and it supports a global metric, but this is no longer the standard metric on a graph of spaces as described in Section 6.1.1.

We must now check that $\bar{X}$ is a negatively curved simplicial 2-complex. By construction, it is a simplicial 2-complex, all of whose simplices are negatively curved. It is sufficient, therefore, to check the link condition on vertices (possibly after an application of Remark 6.1.9). We already have that each vertex space $\bar{X}_v$, and each annulus A_e, is negatively curved, so it suffices to check the link condition for vertices at which the annuli A_e are glued to vertex spaces.

For vertices in type N vertex spaces, links consist of two vertices connected by several (possibly subdivided) arcs, of length π. These satisfy the link condition, since all circles in the space have length at least 2π.

Now let $\bar{X}_v$ be a type M vertex space, and let x be a vertex of $\bar{X}_v$ contained in at least one γ_e. To obtain $\mathrm{Lk}(x, \bar{X})$ from $\mathrm{Lk}(x, \bar{X}_v)$, we glue on a number of arcs of length $> \pi - \delta_v$. We may prove by induction that the resulting space remains CAT(1).

Firstly, recall from above that the geodesics γ_e subtend angles of at least $\pi + 2\delta_v$ at each vertex. In particular, the first arc glued on to $\mathrm{Lk}(x, \bar{X}_v)$ connects two points of distance at least $\pi + 2\delta_v$, and is itself of length $> \pi - \delta_v$; Remark 2.2.22 then says that the space remains CAT(1).

For subsequent arcs, let us assume for induction that after gluing on $r-1$ arcs, the space is still CAT(1). Now suppose the rth arc is to be glued between points a and b. By Lemma 6.1.18), a and b are distinct from any previous points to which arcs have been glued, and hence any path between a and b which is not contained in $\mathrm{Lk}(x, \bar{X}_v)$ must begin and end with an edge from $\mathrm{Lk}(x, \bar{X}_v)$. All such edges have length at least δ_v, and hence any such path has length at least $2\delta_v + \pi - \delta_v = \pi + \delta_v$. On the other hand, any path connecting a and b which does lie in $\mathrm{Lk}(x, \bar{X}_v)$ is of length at least $\pi + 2\delta_v$ as noted before. Thus Remark 2.2.22 applies again, and so the space again remains CAT(1).

Therefore, the link condition holds for $\bar{X}$, and so $\bar{X}$ is a negatively curved simplicial complex. Moreover, it is clear by construction that if all type M vertex spaces of X are compact, then $\bar{X}$ is compact. This completes the proof. □

6.3 Limit groups

The motivating application for Theorem 6.2.1 was to prove the following theorem.

Theorem 6.3.1. *Let G be a limit group. Then G is* CAT(-1) *if and only if G is δ-hyperbolic.*

This is a simplified version of Theorem 6.3.10, which will follow quickly from one of the defining characterisations of a limit group (see Theorem 6.3.6), but before we launch into this we will give some context as to the relevance and usefulness of limit groups. To this end, we begin by stating the simplest definition.

Definition 6.3.2. A *limit group* (also *finitely generated fully residually free group*) is a finitely generated group G such that, for any finite subset $S \subset G$, there is a homomorphism $h\colon G \to \mathbb{F}$ to a non-abelian free group $\mathbb{F}$ such that h is injective on S.

Limit groups were first introduced in the study of equations over free groups; their importance lies in the fact that, in the study of the sets of solutions of these equations, limit groups correspond to irreducible varieties. The original definition (motivating the name) was given by Sela [Sel01], and is quite different (although equivalent) to Definition 6.3.2. We refer the interested reader to [Sel01], alongside [BF09] and [Wil09], for more information and references.

We will use several facts concerning limit groups without proof, and we state these below for convenience. Properties 1, 3 and 5 are easy to see from the definition, and proofs of all the others can be found in [Sel01] or [BF09] except where otherwise indicated.

Theorem 6.3.3 (Properties of limit groups)**.** *Let G be a limit group. Then:*

1. *G is torsion-free.*
2. [Sel01, Corollary 4.4] *G is finitely presented.*
3. *Every finitely generated subgroup of G is a limit group.*
4. [Sel01, Corollary 4.4] *Every abelian subgroup of G is finitely generated.*
5. [Sel01, Lemma 1.4] *Every non-trivial abelian subgroup of G is contained in a unique maximal abelian subgroup.*
6. [Sel01, Lemma 1.4] *Every maximal abelian subgroup of G is malnormal.*
7. [Sel01, Lemma 2.1] *If $G = A *_C B$ for C abelian, then any non-cyclic abelian subgroup $M \subset G$ is conjugate into A or B.*
8. [AB06] *G is* CAT(0).

Much of the work that has been done on limit groups, including the proof of 8 above, depends upon a powerful structure theory. For full details of this (often expressed in terms of *constructible limit groups*), the reader is referred again to [Sel01] and subsequent papers in that series, as well as the expositions [BF09, Wil09, CG05] (see also [KM98a] and subsequent papers). A particularly elegant result that comes out of this theory is Theorem 6.3.6 below, proved in [KM98b] and [CG05]. To state it, we must first introduce the following notion. Our overview follows that given in [Wil08].

Definition 6.3.4. Let G' be a group, let $g \in G'$, and let $C(g)$ denote the centralizer of g. Let $n \geq 1$. Then the group $G' *_{C(g)} \left(C(g) \times \mathbb{Z}^n\right)$ is called a *centralizer extension of G' by $C(g)$*.

Definition 6.3.5. A group is said to be an *iterated centralizer extension* if it is either a finitely generated free group, or can be obtained from one by taking repeated centralizer extensions. The class of iterated centralizer extensions is denoted ICE; sometimes we will refer to a group in ICE as an ICE *group.* An ICE group obtained by taking a free group and then taking a centralizer extension n times is said to have *height n*.

Theorem 6.3.6 ([KM98b, CG05])**.** *Limit groups coincide with finitely generated subgroups of* ICE *groups.*

Remark 6.3.7. [see also [Wil08, Remark 1.14]] If G is an ICE group (and hence a limit group by Theorem 6.3.6) which is not free, then it follows from property 5 of Theorem 6.3.3 that centralizers in G are abelian. Since G is in ICE, there is an ICE group G' such that $G = G' *_{C(g)} \left(C(g) \times \mathbb{Z}^n\right)$. If $C(g)$ is non-cyclic, then by property 7, $C(g)$ is conjugate into one of the two components in the iterated centralizer decomposition of G'; by induction, it follows that $C(g)$ is conjugate into some previously attached $C(g') \times \mathbb{Z}^m$. At this stage, we could therefore have attached $C(g') \times \mathbb{Z}^{m+n}$ instead. It follows that, when building limit groups, we may assume all centralizer extensions are by (infinite) cyclic centralizers.

A consequence of Theorem 6.3.6 is that there is a natural graph of spaces decomposition for any ICE group G. If G is free, it is a single vertex space: a compact graph with fundamental group G. Otherwise, $G = G' *_{\langle g \rangle} \left(\langle g \rangle \times \mathbb{Z}^n\right)$, where we can assume (by induction) that $G' = \pi_1(Y')$ for some graph of spaces Y'. Then G has a graph of spaces decomposition Y with underlying graph an edge, one vertex space $M = Y'$, one vertex space N which is an $n+1$-torus $\mathbb{T}^{n+1}$, and edge space a circle A. The attaching maps send A to a closed curve $\gamma \subset M$ representing g in $\pi_1(M) = G'$, and to a coordinate circle of $N = \mathbb{T}^{n+1}$. We may assume without loss of generality that M and N are simplicial, the attaching maps are combinatorial local isometries, and their images are simplicial closed geodesics.

Remark 6.3.8. The above graph of spaces decomposition for an ICE group G induces a graph of spaces decomposition for any subgroup $H < G$. Therefore, any limit group H has a graph of spaces decomposition X with edge spaces either circles or points, and vertex spaces covering spaces of either M or N. Moreover, since all limit groups are finitely generated, we can assume that X has finite underlying graph.

We are now ready to prove the following consequence of Theorem 6.2.1.

Theorem 6.3.9. *Let H be a limit group which does not contain any subgroup isomorphic to $\mathbb{Z}^2$, and let X be the graph of spaces induced by an embedding of H into an* ICE *group. Then there exists a compact negatively curved simplicial 2-complex $\bar{X}$ which is homotopy equivalent to X.*

Proof. The proof uses Theorem 6.3.6, along with Theorem 6.2.1 and induction on height.

Clearly the result holds if H embeds into an ICE group of height 0. So suppose that H embeds into an ICE group $G = \pi_1(M *_A N)$ of height n, and assume that the result is proved for all limit groups that do not contain $\mathbb{Z}^2$ and that embed in ICE groups of height $\leq n-1$. The graph of spaces X for H has two types of vertex space: those mapping to M and those mapping to N. Call these type M and type N respectively.

Edge spaces of X are either lines or circles. Since H is finitely generated, this implies that the vertex groups are all finitely generated. Since H does not contain $\mathbb{Z}^2$, each type M vertex group is therefore a limit group that does not contain $\mathbb{Z}^2$, and each type N vertex space is homotopy equivalent to either a point or a circle. By the inductive hypothesis, each type M vertex space is therefore homotopy equivalent to a compact negatively curved simplicial 2-complex.

By Lemma 2.1.8, we may replace vertex and edge spaces X_v, X_e of X with homotopy equivalent spaces X'_v, X'_e, to obtain a graph of spaces X' with the same homotopy type as X. We do this as follows:

1. For each edge space X_e which is a line, let X'_e be a point.

2. For each edge space X_e which is a circle, let X'_e be a circle.

3. For each type N vertex space X_v which is homotopy equivalent to a point, let X'_v be a point.

4. For each type N vertex space X_v which is homotopy equivalent to a circle, let X'_v be a circle.

5. For each type M vertex space X_v, apply the inductive hypothesis to find a homotopy equivalent compact negatively curved simplicial 2-complex X'_v.

As in the proof of Theorem 6.2.1, the attaching maps are defined by composing the attaching maps in X with the homotopy equivalences applied to the vertex spaces, followed by a further homotopy to ensure that the images of attaching maps of circular edge spaces are closed geodesics. That is, once we have fixed an edge and vertex space, we choose as our attaching map a local isometry which represents the corresponding attaching map in the graph of groups. As before, for those edge spaces which are points, this can be any map.

At this stage, X' is a compact graph of spaces satisfying conditions 1 to 4 of Theorem 6.2.1. To show that it also satisfies condition 5, note that the set of non-trivial incident edge subgroups in a type M vertex group of X' is a set of cyclic subgroups of a limit group. We claim that these cyclic subgroups are maximal cyclic. By Remark 6.3.7 this is true in the ICE group

G, and so it follows from the definition of the induced splitting that edge subgroups are also maximal cyclic in type M vertex groups of X'. Since H does not contain $\mathbb{Z}^2$, they are therefore maximal abelian, and so malnormality of individual edge subgroups follows from property 6 of Theorem 6.3.3.

It remains to show that if a and b generate two distinct edge subgroups of a type M vertex group $H_v = \pi_1\left(X'_v\right)$ of X', then conjugates of $\langle a\rangle \subset H_v$ and conjugates of $\langle b\rangle \subset H_v$ have trivial intersection in H_v. Since the ICE group G has only one edge group (which we may take without loss of generality to be $\langle a\rangle$) it follows from the definition of the induced splitting that a and b are conjugate as elements of G, say $g^{-1}ag = b$, and that the conjugating element g is contained in the vertex group G_v ($= \pi_1(M)$) of G, but not in the subgroup H. If a and b are also conjugate in H, say $h^{-1}ah = b$, then a commutes with gh^{-1} in G. Since (by definition of a centralizer extension) $C_{G_v}(a) = \langle a\rangle$, it follows that $gh^{-1} \in \langle a\rangle$; in particular, $g \in H$, which is a contradiction. Therefore, a and b cannot be conjugate in H. Likewise, if a and b^{-1} are conjugate in H, say $h^{-1}ah = b^{-1}$, then it follows that a commutes with $\left(gh^{-1}\right)^2$ in G, and hence $\left(gh^{-1}\right)^2 \in \langle a\rangle$. Since $\langle a\rangle$ is maximal cyclic, we again see that $gh^{-1} \in \langle a\rangle$, and obtain a contradiction as before.

Now, suppose $t \in H_v$ satisfies $t^{-1}a^p t = b^q$. Since H_v does not contain $\mathbb{Z}^2$, property 5 implies that the two cyclic subgroups generated by $t^{-1}at$ and b must coincide; but since the edge groups are maximal cyclic, neither $t^{-1}at$ nor b is a proper power. Thus $t^{-1}at = b^{\pm 1}$, which is a contradiction according to the previous paragraph. It follows that conjugates of $\langle a\rangle \subset H_v$ and conjugates of $\langle b\rangle \subset H_v$ have trivial intersection as required.

It follows that X' satisfies all the conditions of Theorem 6.2.1, and so we can apply it to find the complex $\bar{X}$ as required. □

In summary, we have proved the following theorem:

Theorem 6.3.10. *Let G be a limit group. Then the following are equivalent:*

1. *G is hyperbolic.*

2. *G is* $\mathrm{CAT}(-1)$.

3. *G has* $\mathrm{CAT}(-1)$ *dimension 2.*

4. *G does not contain* $\mathbb{Z}^2$.

Proof. 3 $\implies$ 2 follows from our definition of CAT(-1) dimension (Remark 2.4.22). 1 $\implies$ 4 and 2 $\implies$ 1 are true for any group. 4 $\implies$ 3 is precisely Theorem 6.3.9. □

Note in particular that this provides an alternative proof of the fact due to Sela [Sel01, Corollary 4.4] that a limit group G is hyperbolic if and only if every abelian subgroup is cyclic. Sela's proof also uses a combination theorem, namely that of Bestvina and Feighn [BF92].

Remark 6.3.11. In combination with the fact that limit groups are CAT(0) (item 8 of Theorem 6.3.3), Theorem 6.3.10 provides a complete answer in the limit group case to the questions posed in Section 2.4.3 about the relationships between the different notions of negative curvature. That is, it shows that being CAT(0) without $\mathbb{Z}^2$ subgroups, CAT(-1), and hyperbolic are all equivalent notions of negative curvature for limit groups.

6.4 Further applications of the gluing theorem

There are several contexts in which cyclic splittings of groups are of interest, and our gluing theorem therefore has the potential to shift the question of whether such groups are CAT(-1) to their vertex groups under a cyclic splitting. With this in mind, we can give two more consequences of Theorem 6.2.1. The first consequence concerns JSJ decompositions, for which we will need to recall some technical background. The second consequence concerns graphs of free groups with cyclic edge groups, and this will follow quickly from the JSJ material.

6.4.1 JSJ decompositions of torsion-free hyperbolic groups

JSJ decompositions were originally invented to study toroidal decompositions of 3-manifolds [JS79, Joh79], and can be thought of as the second stage in the decomposition of a 3-manifold—the stage after cutting along essential spheres. Analogous notions for groups have been studied by Bowditch [Bow98] (in the case of hyperbolic groups) and Rips–Sela [RS97] (in the case of general finitely presented groups), among many other generalisations. We will give only the details essential for our argument, and for these we follow [Bow98].

Dunwoody's Accessibility Theorem [Dun85] shows that any hyperbolic group can be decomposed as a graph of groups whose vertex groups are either finite or one-ended, and whose edge groups are finite. In the torsion-free case, this reduces to a decomposition with trivial edge groups and one-ended, hyperbolic vertex groups (this is also called the Grushko decomposition). Analogously to the 3-manifold setting, the JSJ decomposition then describes how to decompose these components further, and in the torsion-free case, this is a decomposition along infinite cyclic subgroups.

The following statement is a special case of [Bow98, Theorem 0.1]:

Theorem 6.4.1 (JSJ decomposition for torsion-free hyperbolic groups)**.** *Let* Γ *be a torsion-free, one-ended hyperbolic group. Then* Γ *is the fundamental group of a well-defined, finite, canonical graph of groups with infinite cyclic edge groups and vertex groups of three types:*

type S *fundamental groups of non-elementary surfaces, whose incident edge groups correspond precisely to the subgroups generated by the boundary components;*

type N *infinite cyclic groups, and*

type M *non-elementary hyperbolic groups not of type S or type N.*

These three types are mutually exclusive, and no two of the same type are adjacent.

In the above statement, a cyclic splitting of Γ is called *canonical* if it has a common refinement with any other cyclic splitting of Γ, where a *refinement* of a splitting is obtained by taking further splittings of vertex groups so that the images of attaching maps in the original splitting are still contained in vertex groups. In this sense, a canonical splitting contains information about any cyclic splitting of the group. The splitting in the JSJ decomposition is well-defined because it is the deepest canonical splitting possible; it has no further canonical refinement.

Remark 6.4.2. In Bowditch's original statement, it is taken as a condition that Γ is not a co-compact Fuchsian group, which, in the torsion-free case, is the same thing as a closed hyperbolic surface group. We do not need to rule out this case, but note that its JSJ decomposition consists of just one vertex group, of type S, with no incident edge groups. Indeed, a closed surface group cannot have a non-trivial canonical splitting—each cyclic splitting corresponds to a simple closed curve on the surface, and given any such curve, we can always choose another simple closed curve which cannot be homotoped disjoint from it. Then the two splittings defined by these two curves can never admit a common refinement.

We would like to apply our gluing theorem to the JSJ decomposition of a hyperbolic group. For our theorem to apply, we will need to assume (as before) that the type M vertex spaces have CAT(-1) dimension 2; however, no further assumptions are required due to the following fact.

Lemma 6.4.3. *For each type M or type S vertex group in the JSJ decomposition of a hyperbolic group Γ, the images of incident edge groups form a malnormal family of subgroups.*

Proof. In the type S case, this follows from Theorem 6.4.1—in particular, the subgroups concerned are the subgroups generated by the boundary components, and these always form a malnormal family. In the type M case, it follows from the proof of the fact that the action of Γ on the Bass–Serre tree corresponding to the JSJ decomposition is 2-acylindrical (see [GL11]). □

The next proposition then follows quickly from Theorem 6.2.1 and Lemma 6.4.3.

Proposition 6.4.4. *Let Γ be a torsion-free, one-ended hyperbolic group. Suppose that all the type M vertex groups in the JSJ decomposition of Γ are fundamental groups of (compact) negatively curved 2-complexes. Then Γ is also the fundamental group of a (compact) negatively curved 2-complex.*

Proof. We construct a graph of spaces corresponding to the JSJ decomposition of Γ. By definition of the JSJ decomposition, we may choose a circle for each edge space. By assumption, we can choose (compact) negatively curved 2-complexes for the type M vertex spaces, and we can clearly also choose (compact) negatively curved 2-complexes for the type S vertex spaces. We choose a circle for each type N vertex space. If a type M and type S vertex space are adjacent, we insert an additional type N vertex space in between, so that each type M or S vertex space is then only adjacent to type N vertex spaces. Thus the first three conditions of Theorem 6.2.1 hold (where type S vertex spaces are included in the type M vertex spaces of Theorem 6.2.1). Remark 6.2.2 implies that we may assume that the images of attaching maps are then simplicial closed geodesics, hence condition 4 also holds, and Lemma 6.4.3 gives condition 5. Hence Theorem 6.2.1 applies, and the result follows. □

In the JSJ decomposition, we do not necessarily know any more about the type M vertex groups of Γ than we know about Γ itself. In particular, they may themselves have a non-trivial JSJ decomposition, or even a free decomposition. However, we may appeal to the following theorem ([LTar], see also [DP01]).

Theorem 6.4.5 (Strong Accessibility Theorem)**.** *Let Γ be a torsion-free hyperbolic group. Consider the hierarchy obtained by taking either the free (if freely decomposable) or JSJ decomposition of Γ, and then taking a free or JSJ decomposition of the resulting vertex groups, and so on. Then this hierarchy is finite.*

Definition 6.4.6. A torsion-free hyperbolic group is called *rigid* if it does not have a non-trivial free or cyclic splitting.

The Strong Accessibility Theorem says that if we continue decomposing vertex spaces using free products or JSJ decompositions, we must eventually terminate at a decomposition whose vertex groups are rigid (note that vertex groups are always hyperbolic—this is clear for free decompositions, and in the JSJ case is given by Theorem 6.4.1). In the context of our gluing theorem, it implies the following proposition.

Proposition 6.4.7. *A torsion-free hyperbolic group* Γ *has* CAT(-1) *dimension 2 if each rigid subgroup of* Γ *has* CAT(-1) *dimension 2.*

6.4.2 Graphs of free groups with cyclic edge groups

In the above subsection, we saw how the gluing theorem can be applied to JSJ decompositions of hyperbolic groups. However, we could only conclude that a group was CAT(-1) if the vertex groups in the JSJ decomposition were CAT(-1) of dimension 2, which is a strong requirement. Here we describe a context where this requirement is met.

In [HW10], Hsu and Wise show that a group G which splits as a finite graph of finitely generated free groups with cyclic edge groups is CAT(0) if and only if it contains no non-Euclidean Baumslag–Solitar subgroups. Their method uses Sageev's construction, and as such gives little control over the CAT(0) dimension. If G is also hyperbolic (so that it contains no Baumslag–Solitar subgroups at all), then we may improve their result in two ways: firstly, showing G is in fact CAT(-1), and secondly, showing that the CAT(-1) dimension (and hence the CAT(0) dimension) is equal to 2.

Theorem 6.4.8. *Let G be a hyperbolic group which splits as a finite graph of finitely generated free groups with cyclic edge groups. Then G has* CAT(-1) *dimension 2.*

To prove this, it is tempting to try to apply Theorem 6.2.1 directly to the graph of spaces corresponding to the given graph of free groups. However, it may not be possible to ensure that incident edge groups form malnormal families in vertex groups. We circumvent this difficulty by appealing to the JSJ machinery of the previous subsection.

Proof of Theorem 6.4.8. We would like to apply Proposition 6.4.7, and so we need to check that each rigid subgroup of G has CAT(-1) dimension 2. So, let H be a rigid subgroup of G, and consider the splitting of H induced by the decomposition of G as a graph of free groups with cyclic edge groups. Since H is rigid, this induced splitting must consist of a single vertex group, so H is a subgroup of a vertex group of the original splitting. Hence H is free (indeed, since H cannot split freely, it is trivial), and so certainly has CAT(−1) dimension 2, as required. □

6.5 Remarks

Remark 6.5.1. Theorem 6.2.1 was designed for the limit groups case. However, Theorem 6.5.2 below is a slightly more general theorem which follows from the same argument. The essential difference is to allow *any* negatively curved 2-complex as a type N vertex space. In the proof of Theorem 6.2.1, we used the fact that the type N vertex spaces were circles twice: firstly, we gave them a specific length which was sufficiently small for there to exist some negatively curved annulus A_e; and secondly, we explicitly described the structure of the links in $\bar{X}$ of vertices in type N vertex spaces to verify the link condition. To prove Theorem 6.5.2, we may rescale the type N vertex spaces so that all the closed geodesics γ_e are sufficiently short for there to exist suitable annuli A_e. The proof then proceeds as before, and we may appeal to Remark 2.2.22 to verify the link condition for vertices in type N vertex spaces in the resulting complex. We do not know whether this generalisation has any applications beyond those of Theorem 6.2.1.

Theorem 6.5.2. *Let $\mathbb{X}$ be a graph of groups with underlying graph Γ. Suppose:*

1. *Γ is bipartite, with corresponding vertex sets M and N;*
2. *vertex groups of X are fundamental groups of negatively curved 2-complexes;*
3. *edge groups of X are infinite cyclic or trivial, and*
4. *for vertex groups corresponding to vertices in M, the incident edge groups form a malnormal family of subgroups.*

Then $\mathbb{X}$ has a corresponding graph of spaces homotopy equivalent to a negatively curved 2-complex.

Remark 6.5.3. We expect that the malnormal family assumption we make in our main theorem can be relaxed. Indeed, it is used only to show that it is possible to make the corresponding set of closed geodesics transverse (Lemma 6.1.18), which in turn is used only to ensure that there is a safe place to attach the corner of a hyperbolic annulus (see the proof of Theorem 6.2.1). It is easy to design a hyperbolic annulus with *reflex* angle of, say, $\pi+\delta_1$ in the geodesic side, and an angle in the other side of $\pi - \delta_2$, provided $\delta_2 > \delta_1$. Using such annuli for edge cylinders, we may attach multiple annuli along the same closed geodesic in a vertex space. It may be that such a technique allows us to replace the malnormal family assumption with a weaker assumption, such as k-acylindricity for larger k, or the "annuli flare condition" used in [BF92].

Remark 6.5.4. Remark 6.1.1 does not directly generalise to n-dimensional simplices, as we showed in Remark 6.1.2. Moreover, even when comparison simplices *can* be found (which can be ensured by choosing k' sufficiently close to k), the dihedral angles may *decrease* when passing to a comparison simplex of lesser negative curvature (i.e. $k < k' < 0$). This means that Remark 6.1.6 may not hold in dimensions > 2, and consequently our proof of Theorem 6.2.1 is valid only in two dimensions. To prove a higher dimensional version, one would need to generalise "excess angle" to dimensions > 2 using the notion of a complex with *extra large* links. An M_k-complex is said to have extra large links if the systole of the link of each vertex is strictly greater than 2π, and moreover the links themselves are complexes with extra large links. It has the useful feature that it is stable under small perturbations of the metric (see [Mou88] or [Dav08]). However, it is not clear whether it is always possible to preserve this property under gluings, or even subdivision (for example, to make some loop simplicial). See [CD95, CDM97, Riv05] for more information.

Remark 6.5.5. We have indicated the applicability of our method to families of hyperbolic groups built hierarchically, namely limit groups, and graphs of free groups with cyclic edge groups (via the JSJ decomposition). There are many other families of hyperbolic groups built in a similarly hierarchical way—for example, hyperbolic special groups (see Section 3.8.1)—and these may lend themselves to investigation using similar techniques. It does not seem unreasonable to seek a more general CAT(-1) gluing theorem, valid for larger classes of graphs of spaces—for example, the case where the edge groups are not cyclic; where the vertex spaces are of dimension > 2 (as discussed above), or where the groups are allowed to have torsion; however, it appears that any such generalisation requires an understanding of the higher dimensional case, and is thus subject to the difficulties described above.

Bibliography

[AB06] Emina Alibegović and Mladen Bestvina. Limit groups are CAT(0). *J. London Math. Soc.*, 74(1):259–272, 2006.

[AFW15] Matthias Aschenbrenner, Stefan Friedl, and Henry Wilton. *3-manifold groups*. EMS Series of Lectures in Mathematics. European Mathematical Society (EMS), Zürich, 2015.

[Ago04] Ian Agol. Tameness of hyperbolic 3–manifolds. arXiv: math.GT/0405568, 2004.

[Ago08] Ian Agol. Criteria for virtual fibering. *J. Topol.*, 1(2):269–284, 2008.

[Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[BC07] Noel Brady and J Crisp. CAT(0) and CAT(-1) dimensions of torsion free hyperbolic groups. *Commentarii mathematici helvetici*, 82(1):61–85, 2007.

[Bea83] Alan F. Beardon. *The geometry of discrete groups*, volume 91 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1983.

[BF91] Mladen Bestvina and Mark Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.

[BF92] Mladen Bestvina and Mark Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

[BF09] Mladen Bestvina and Mark Feighn. Notes on Sela's work: limit groups and Makanin-Razborov diagrams. In *Geometric and cohomological methods in group theory*, volume 358 of *London Math. Soc. Lecture Note Ser.*, pages 1–29. Cambridge Univ. Press, Cambridge, 2009.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BL16] Benjamin Beeker and Nir Lazarovich. Stallings folds for CAT(0) cube complexes and quasiconvex subgroups. arXiv:1605.02758, 2016.

[Bon02] Francis Bonahon. Geometric structures on 3-manifolds. In *Handbook of geometric topology*, pages 93–164. North-Holland, Amsterdam, 2002.

[Bow98] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.

[Bri91] Martin R. Bridson. Geodesics and curvature in metric simplicial complexes. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 373–463. World Sci. Publ., River Edge, NJ, 1991.

[Bro16] Samuel Brown. A gluing theorem for negatively curved complexes. *Journal of the London Mathematical Society*, 93(3):741–762, 2016.

[BW12] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[CD95] Ruth Charney and Michael Davis. The polar dual of a convex polyhedral set in hyperbolic space. *The Michigan Mathematical Journal*, 42(3):479–510, 1995.

[CDM97] Ruth Charney, Michael W. Davis, and Gabor Moussong. Nonpositively curved, piecewise Euclidean structures on hyperbolic manifolds. *The Michigan Mathematical Journal*, 44(1):201–208, 1997.

[CG05] Christophe Champetier and Vincent Guirardel. Limit groups as limits of free groups. *Israel Journal of Mathematics*, 146(1):1–75, 2005.

[CG06] Danny Calegari and David Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 19(2):385–446, 2006.

[Dav08] Michael W. Davis. *The geometry and topology of Coxeter groups*, volume 32 of *London Mathematical Society Monographs Series*. Princeton University Press, Princeton, NJ, 2008.

[DP01] Thomas Delzant and Leonid Potyagailo. Accessibilité hiérarchique des groupes de présentation finie. *Topology*, 40(3):617–629, 2001.

[Dun85] Martin J Dunwoody. The accessibility of finitely presented groups. *Inventiones Mathematicae*, 81(3):449–457, 1985.

[Dun98] M. J. Dunwoody. Folding sequences. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 139–158 (electronic). Geom. Topol. Publ., Coventry, 1998.

[GL11] Vincent Guirardel and Gilbert Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

[GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[Gro87] Mikhail Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro93] Mikhail Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[Gro01] Mikhail Gromov. CAT(κ)-spaces: construction and concentration. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 280(Geom. i Topol. 7):100–140, 299–300, 2001.

[Hag08] Frédéric Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.

[Hal49] Marshall Hall, Jr. Coset representations in free groups. *Trans. Amer. Math. Soc.*, 67:421–432, 1949.

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

[HW99] Tim Hsu and Daniel T. Wise. On linear and residual properties of graph products. *Michigan Math. J.*, 46(2):251–259, 1999.

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[HW10] Tim Hsu and Daniel T. Wise. Cubulating graphs of free groups with cyclic edge groups. *Amer. J. Math.*, 132(5):1153–1188, 2010.

[HW12] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Ann. of Math. (2)*, 176(3):1427–1482, 2012.

[HW15] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. *Invent. Math.*, 199(2):293–331, 2015.

[Joh79] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.

[JS79] William H. Jaco and Peter B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Amer. Math. Soc.*, 21(220):viii+192, 1979.

[Kap01] Ilya Kapovich. The combination theorem and quasiconvexity. *Internat. J. Algebra Comput.*, 11(2):185–216, 2001.

[KM98a] Olga Kharlampovich and Alexei Myasnikov. Irreducible affine varieties over a free group: I. Irreducibility of quadratic equations and Nullstellensatz. *Journal of Algebra*, 200(2):472–516, 1998.

[KM98b] Olga Kharlampovich and Alexei Myasnikov. Irreducible affine varieties over a free group: II. Systems in triangular quasi-quadratic form and description of residually free groups. *Journal of Algebra*, 200(2):517–570, 1998.

[KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[KWM05] Ilya Kapovich, Richard Weidmann, and Alexei Miasnikov. Foldings, graphs of groups and the membership problem. *Internat. J. Algebra Comput.*, 15(1):95–128, 2005.

[LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin-New York, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.

[LTar] Larsen Louder and Nicholas Touikan. Strong accessibility for finitely presented groups. *Geometry and Topology*, to appear.

[Mal40] Anatoly Malcev. On isomorphic matrix representations of infinite groups. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):405–422, 1940.

[Mar13] Alexandre Martin. Complexes of groups and geometric small cancellation over graphs of groups. arXiv:1306.6847, 2013.

[McC09] Jon McCammond. Constructing non-positively curved spaces. In *Geometric and cohomological methods in group theory*, volume 358 of *London Math. Soc. Lecture Note Ser.*, pages 162–224. Cambridge Univ. Press, Cambridge, 2009.

[Mil04] Charles F. Miller III. Combinatorial group theory. Lecture notes, University of Melbourne, 2004.

[Min06] Ashot Minasyan. Separable subsets of GFERF negatively curved groups. *J. Algebra*, 304(2):1090–1100, 2006.

[Mit04] Mahan Mitra. Height in splittings of hyperbolic groups. *Proc. Indian Acad. Sci. Math. Sci.*, 114(1):39–54, 2004.

[Mou88] Gabor Moussong. *Hyperbolic Coxeter groups*. PhD thesis, The Ohio State University, 1988.

[OW11] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than 1/6. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.

[Per02] Grisha Perelman. The entropy formula for the Ricci flow and its geometric applications. arXiv:math/0211159, 2002.

[Rip82] Eliyahu Rips. Subgroups of small cancellation groups. *Bull. London Math. Soc.*, 14(1):45–47, 1982.

[Riv05] Igor Rivin. Extra-large metrics. arXiv:math/0509320, 2005.

[RS97] Eliyahu Rips and Zlil Sela. Cyclic splittings of finitely presented groups and the canonical JSJ decomposition. *Annals of Mathematics*, 146(1):53–109, 1997.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Sco78] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[Sel01] Zlil Sela. Diophantine geometry over groups. I. Makanin-Razborov diagrams. *Publ. Math. Inst. Hautes Études Sci.*, 93:31–105, 2001.

[Ser03] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[SW79] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.

[Thu82] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.

[Thu86] William P. Thurston. Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds. *Ann. of Math. (2)*, 124(2):203–246, 1986.

[TL97] William P. Thurston and Silvio Levy. *Three-dimensional geometry and topology*, volume 1. Princeton university press, 1997.

[Wil08] Henry Wilton. Hall's theorem for limit groups. *Geometric and Functional Analysis*, 18(1):271–303, 2008.

[Wil09] Henry Wilton. Solutions to Bestvina & Feighn's exercises on limit groups. In *Geometric and cohomological methods in group theory*, volume 358 of *London Math. Soc. Lecture Note Ser.*, pages 30–62. Cambridge Univ. Press, Cambridge, 2009.

[Wis00] Daniel T. Wise. Subgroup separability of graphs of free groups with cyclic edge groups. *Q. J. Math.*, 51(1):107–129, 2000.

[Wis02] Daniel T. Wise. The residual finiteness of negatively curved polygons of finite groups. *Invent. Math.*, 149(3):579–617, 2002.

[Wis04] Daniel T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14(1):150–214, 2004.

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

[Wis12a] Daniel T. Wise. *From Riches to Raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012.

[Wis12b] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. Preprint. http://www.math.mcgill.ca/wise/papers.html, 2012.