HIERARCHIES FOR RELATIVELY HYPERBOLIC VIRTUALLY COMPACT SPECIAL NON-POSITIVELY CURVED CUBE COMPLEXES

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Eduard Einstein

August 2018

© 2018 Eduard Einstein

ALL RIGHTS RESERVED

HIERARCHIES FOR RELATIVELY HYPERBOLIC VIRTUALLY COMPACT SPECIAL NON-POSITIVELY CURVED CUBE COMPLEXES

Eduard Einstein, Ph.D.

Cornell University 2018

Cube complexes and hierarchies of cube complexes have been studied extensively by Wise and feature prominently in Agol's proof of the Virtual Haken Conjecture for hyperbolic 3-manifolds. Among hyperbolic groups, Wise characterized hyperbolic virtually compact special groups as the hyperbolic groups that virtually admit a quasiconvex hierarchy terminating in finite groups. The main result of this thesis is that every relatively hyperbolic fundamental group of a virtually compact special non-positively curved cube complex virtually admits a quasiconvex hierarchy terminating in peripheral subgroups, answering a question due to Wise.

The proof of the main theorem roughly follows the outline of Agol, Groves and Manning's New Proof of Wise's Malnormal Special Quotient Theorem for hyperbolic groups, but instead uses relatively hyperbolic geometric tools to prove that Wise's double dot hierarchy construction yields a quasiconvex hierarchy. Group theoretic relatively hyperbolic Dehn filling and the main theorem are used in the final chapter to provide a new proof of a relatively hyperbolic analog of Wise's malnormal special quotient theorem.

BIOGRAPHICAL SKETCH

Eduard Einstein was born in Santa Monica, CA to Dr. K. Alice Chang and Dr. Thomas Einstein in November, 1987. He attended Harvard-Westlake School, graduating in 2006. He received a bachelor of arts in mathematics from Pomona College in Claremont, CA in 2010, a master of arts in mathematics from the University of California at Santa Barbara in 2012 and a Ph.D. in mathematics from Cornell University in Ithaca, NY in August 2018.

To Mom and Thomas, for always supporting my dreams

ACKNOWLEDGEMENTS

Jason Manning, my advisor, made invaluable contributions to this work. Without his precise explanations, breadth of knowledge, patience, generosity and obsessive attention to quantifiers, this thesis would not have been possible. My colleague Yen Duong explained the problem that would become the heart of this thesis, helped make mischief with me in Santa Barbara and has always been there in times of need. To my family, especially Alice, Thomas and Benita Campos-Ramirez, thanks always for your love and support. I would like to thank Lucien Clavier, Daniel Groves, Chris Hruska, Michael Hull and Dani Wise for useful conversations that contributed to developing the geometric ideas in this thesis. In addition to those mentioned above, there have been many mathematicians who have provided mentoring and supported my career; these include Chuck Akemann, Erica Flapan, Stephan Garcia, Ken Goodearl, Birge Huisgen-Zimmermann, Martin Kassabov, Tim Riley, Ghassan Sarkis, and Shahriar Shahriari. A special thanks as well to Kelly Delp, Maria Terrell, Joy Jones, Melissa Totman and Medina Price whose underappreciated administrative talents made my job a lot easier over the years. To my various contemporaries in Ithaca, including but not limited to Aaron Palmer, Andrew Marshall, Christine McMeekin, Elliot Cartee, Florian Frick, Jason Anema, Laura Escobar-Vega, Kristen Pueschel, Margarita Amchislavska, Mark Cerenzia, Romin Abdolazahdi and Ruth Fahey, thanks for making Ithaca fun and keeping the winters bearable. Thanks to Arielle Leitner and Thomas Ng for being dependable conference buddies, especially on antisocial occasions. Finally, I would like to thank the Big Red Barn, Kris Corda, the barn staff and Saranac pale ale for their instrumental roles in both helping and hindering the production of this thesis.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Background, History and Motivation

One of the main goals of cube complex theory is to use the geometry and combinatorial structure of cube complexes to better understand groups. The study of cubical groups has played an important role in recent developments in the study of hyperbolic 3-manifold groups; in particular, cubical groups are essential to Agol's proof of the virtual Haken conjecture in [1]. Scott [29] showed that **subgroup separability**, defined in Chapter 5.1, is closely related to determining when an immersion lifts to an embedding in a finite cover. In [21], Kahn and Markovic found that closed hyperbolic 3-manifolds have many immersed closed surfaces, which reduced proving the virtual Haken conjecture to proving that certain surface subgroups of 3-manifolds are separable. **Virtually special cube complexes**, developed by Wise and others, were used to show that these surface subgroups are separable. Specifically, in [15], Haglund and Wise showed that a hyperbolic group is the fundamental group of a compact virtually special cube complex if and only if every quasiconvex subgroup is separable.

A group that is the fundamental group of a virtually special cube complex acts properly and cocompactly on a **CAT(0) cube complex** with virtually special quotient. Such an action provides a natural link between the geometry of the group and the geometry of CAT(0) and **non-positively curved** cube complexes. Natural questions include: which groups act properly and cocompactly on CAT(0) cube complexes? When does an action on a CAT(0) cube complex have a virtually special quotient? or does the group also act on a different

CAT(0) cube complex with special quotient?

When a group G has many **"codimension-one subgroups,"** Sageev [26] showed there is a dual CAT(0) cube complex that G acts on. However, depending on the nature of the codimension-one subgroups, the action may fail to be proper and/or cocompact. A natural example of codimension-one subgroups are the Kahn-Markovic surface subgroups in the fundamental group of a closed hyperbolic 3-manifold group. In that case, the action on the dual CAT(0) cube complex turns out to be proper and cocompact (see [6]).

Even when a group G admits a proper and cocompact action on a CAT(0) cube complex, it is difficult to say whether G admits an action on a CAT(0) cube complex with virtually special quotient. In the case where G is a hyperbolic group, Agol [1] proved that if G acts properly and cocompactly on a CAT(0) cube complex, then it acts virtually specially. On the other hand, there are examples such as the fundamental group of complex X in [33] that are fundamental groups of non-positively curved square complexes but fail to have a virtually special action on a CAT(0) cube complex. Agol's theorem relies on ingredients of Wise's characterization of hyperbolic virtually compact special groups as those groups that admit a **virtual quasiconvex hierarchy**.

Definition 1.1.1 ([35, Definition 11.5])**.** *Let $\mathcal{QVH}$ be the smallest class of hyperbolic groups closed under the following operations.*

1. $\{1\} \in \mathcal{QVH}$.
2. *If $G = A *_C B$ and $A, B \in \mathcal{QVH}$ and C is finitely generated and quasi-isometrically embedded then $G \in \mathcal{QVH}$.*
3. *If $G = A*_C$, $A \in \mathcal{QVH}$ and C is finitely generated and quasi-isometrically embed-*

ded, then $G \in \mathcal{QVH}$.

4. *If $H \leq G$ with $|G : H| < \infty$ and $H \in \mathcal{QVH}$, then $G \in \mathcal{QVH}$.*

In other words, groups in $\mathcal{QVH}$ are hyperbolic groups that can be built from the trivial group by taking finite index subgroups or taking amalgamations and HNN extensions over quasiconvex subgroups.

Theorem 1.1.2 ([35, Theorem 13.3])**.** *Let G be a hyperbolic group. Then $G \in \mathcal{QVH}$ if and only if G is virtually compact special.*

In Bass-Serre theory (see e.g. [5] and [30]), a group G acting on a tree T without inversions gives a splitting of G into HNN extensions and amalgamations of groups isomorphic to vertex stabilizers of the tree over edge stabilizers of the tree. This splitting can be represented by a **graph of groups** that is a quotient of the action of T on G. A **hierarchy** is a "multilevel" generalization of a graph of groups that allows each vertex group at a level of the hierarchy to have a graph of groups splitting; see Chapter 3.2 for details. A **quasiconvex hierarchy** for a group G terminating in $\mathcal{P}$, a collection of subgroups of G, is an iterative method of building an isomorphism from G to a group constructed iteratively by taking HNN extensions and amalgamations starting with groups isomorphic to $\mathcal{P}$ over quasi-isometrically embedded subgroups of G. A quasiconvex hierarchy for a group G is somewhat weaker than the case where G is an element of $\mathcal{QVH}$ in that passing to finite index subgroups in between taking HNN extensions and amalgamations is no longer allowed. Graphs of groups and quasiconvex hierarchies will be defined and discussed in detail in Section 3.

The geometry of relatively hyperbolic groups is also a major theme of this thesis. Relatively hyperbolic groups were first introduced by Gromov in [12].

In [17], Hruska showed that many of the different notions of relative hyperbolicity and relative quasiconvexity for finitely generated groups including those developed by [7, 10, 11, 24] are equivalent. Chapter 7 uses yet another equivalent notion of relative hyperbolicity from [13] and relative quasiconvexity from [2, 22]. Roughly speaking, a finitely generated group G is hyperbolic relative to a finite collection of subgroups $\mathcal{P}$ if coning off each coset of every element of $\mathcal{P}$ produces a hyperbolic space and pairwise coarse intersections of cosets have uniformly bounded diameter. A typical example of a relatively hyperbolic group is the fundamental group of a finite volume hyperbolic 3-manifold with torus cusps; the group is hyperbolic relative to the $\mathbb{Z}^2$ cusp subgroups.

This thesis extensively uses relatively thin triangles to generalize geometric techniques from the hyperbolic setting to the relatively hyperbolic setting. Relatively thin triangles were originally developed by Drutu and Sapir [10] who used asymptotic cones to study to the geometry of the Cayley graph of a relatively hyperbolic group. Chapter 2.3 introduces many of the basic tools used to work with relatively thin triangles. In [4], Baker and Cooper developed a combination theorem for convex hyperbolic manifolds. Proposition 4.2.11, the main result of Chapter 4.2 is a relatively hyperbolic analog for a space with a geometric ation by a CAT(0) relatively hyperbolic group. The statement of Proposition 4.2.11 is styled after Theorem 4.5 in [3], but the proof is inspired by Lemma 2.2 in [20]. Relatively thin triangles are also used several times in Chapter 6 to adapt arguments from δ-hyperbolic spaces.

The geometry of CAT(0) spaces with a proper cocompact action of a relatively hyperbolic group by isometries is especially relevant. Hruska and Hruska-Kleiner [16, 18] studied the special case of **CAT(0) spaces with isolated**

flats. If a group G acts geometrically on a CAT(0) space with isolated flats, then G is hyperbolic relative to virtually abelian subgroups. Hruska and Kleiner showed that CAT(0) spaces with isolated flats have the **relative fellow traveler property**, a relatively hyperbolic analog of quasigeodesic stability in hyperbolic spaces. In fact, as shown in Chapter 6.1, the relative fellow traveler property holds in CAT(0) spaces with a geometric action by a relatively hyperbolic group with arbitrary peripheral subgroups where the peripheral subgroups have quasiconvex orbits.

The main goal of this thesis is to prove relatively hyperbolic analogs of important ingredients in the proof of Theorem 1.1.2. The first result answers [35, Question 16.31] posed by Wise:

Theorem A. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let G be a virtually compact special group. Then G has a finite index subgroup G_0 so that G_0 has a quasiconvex, malnormal and fully $\mathcal{P}$-elliptic hierarchy terminating in groups isomorphic to elements of $\mathcal{P}$.*

Virtually compact special groups are defined in Chapter 5. See Chapter 3 for details about quasiconvex, malnormal and fully $\mathcal{P}$-elliptic hierarchies. The proof of Theorem A largely follows the outline of [3, Theorem 2.11] (which only applies to hyperbolic results) except that intermediate steps relying on hyperbolic geometry are replaced by relatively hyperbolic geometric results.

Another key ingredient in the proof of Theorem 1.1.2 is Wise's malnormal special quotient theorem:

Theorem 1.1.3 (Wise, [35, Theorem 12.3])**.** *Let G be a hyperbolic and virtually special group with G hyperbolic relative to a collection of subgroups $\{P_1, \ldots, P_m\}$. Then there* *exist finite index subgroups $\dot{P}_i \leq P_i$ such that if $\overline{G} = G(N_1, \ldots, N_m)$ is any peripherally finite Dehn filling with $N_i \leq \dot{P}_i$, then $\overline{G}$ is hyperbolic and virtually special.*

The malnormal special quotient theorem together with virtually special amalgamation criteria from [20] are used to prove Theorem 1.1.2.

In [32], Thurston announced that generic **classical Dehn fillings** of finite volume hyperbolic three manifolds with torus boundary components produce hyperbolic three manifolds. A classical Dehn filling involves gluing solid tori by diffeomorphisms into the torus cusps to create a new manifold. Later work by Osin [25] and Groves-Manning [13] developed a group theoretic analog, **relatively hyperbolic Dehn filling**. In a group theoretic Dehn filling, the gluing of solid tori into boundary components is replaced by taking a quotient whose kernel is normally generated by subgroups of the peripheral subgroups. As in Agol, Groves and Manning's new proof of Theorem 1.1.3, Theorem A can be combined with a relatively hyperbolic group theoretic Dehn filling to give a new proof of [35, Lemma 16.13] in the cocompact case without the machinery of cubical small cancellation theory:

Theorem B ([35, Lemma 16.13])**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair with $\mathcal{P} = \{P_1, \ldots, P_m\}$. There exist subgroups $\{\dot{P}_i \triangleleft P_i\}$ where $\dot{P}_i$ is finite index in P_i such that if $\overline{G} = G(N_1, \ldots, N_m)$ is any peripherally finite filling with $N_i \triangleleft \dot{P}_i$, then $\overline{G}$ is hyperbolic and virtually special.*

The strategy of the proof is to show that the hierarchy structure implied by Theorem A is compatible with certain Dehn fillings so that the resulting quotient is hyperbolic and has an induced quasiconvex hierarchy structure so that the result follows from Theorem 1.1.2.

1.2 A Brief Outline

Chapter 2 contains an introduction to the geometry of CAT(0) and non-positively curved cube complexes and basic ideas in the geometry of CAT(0) spaces and the geometry of relatively hyperbolic groups. Chapter 3 defines and gives examples of graphs of groups and hierarchies. Chapter 3 also introduces different types of hierarchies such as quasiconvex hierarchies, malnormal hierarchies and fully $\mathcal{P}$-elliptic hierarchies for a relatively hyperbolic group pair $(G, \mathcal{P})$. Chapter 4 proves that certain subspaces of CAT(0) spaces acted upon geometrically by a relatively hyperbolic group pair are quasi-isometrically embedded. There are two parts: Chapter 4.1 where the main result, Proposition 4.1.5, shows that to prove these subspaces are quasi-geodesically embedded, it suffices to show that a paths of a specific type with favorable geometric features are quasi-isometrically embedded, and Chapter 4.2 which contains a combination lemma 4.2.2 that applies only to certain types of piecewise geodesics and a more general combination lemma, Proposition 4.2.11 that uses Proposition 4.1.5 and Lemma 4.2.2 to show these subspaces are quasi-isometrically embedded. Chapter 5 defines special cube complexes, discusses residual and separability properties of special cube complexes. Chapter 5.2 shows how to use the separability of the hyperplane subgroups to build finite covers where elevations of a hyperplane are far apart. Chapter 5.3 uses results of Sageev and Wise [28] to represent peripheral subgroups of a relatively hyperbolic compact special group G as immersed complexes in a NPC cube complex X with $\pi_1 X = G$. Chapter 6.1 contains a proof of the relative fellow traveling property for relatively hyperbolic groups with arbitrary peripheral subgroups. The remainder of Chapter 6 follows the outline of Section 5 of [3] and uses Wise's double dot hierarchy construction to prove Theorem A. Finally, Chapter 7 uses group theoretic Dehn Filling along with Theorem A to prove Theorem B.

CHAPTER 2

CUBE COMPLEXES, CAT(0) GEOMETRY AND RELATIVELY HYPERBOLIC GEOMETRY

2.1 Non-Positively Curved Cube complexes

Cube complexes are the basic objects studied in this thesis:

Definition 2.1.1. *An n-cube is a Euclidean cube of the form $[0,1]^n$. A **cube complex** is a union of Euclidean cubes glued by isometries along faces.*

Each vertex of X has an associated simplicial complex called the **link**:

Definition 2.1.2. *Let X be a cube complex and let v be a vertex of X. The **link** of v in X, link(v), is a simplicial complex defined as follows:*

1. *The vertices of link(v) are in one to one correspondence with the edges of X adjacent to v.*
2. *Any $k+1$ vertices of link(v) determine a k-simplex if their corresponding edges are edges of a common k-cube of X.*

*The link is a **flag simplicial complex** if and only if it is simplicial and every $k+1$-clique of vertices determine a k-simplex.*

Informally, the link of a vertex v in X is flag if every collection of $(k \geq 2)$-cubes with v as a vertex that outlines the corner of a $(k+1)$–cube is a collection of faces of a $(k+1)$–cube having v as a vertex.

The vertex links of a cube complex play an important role in determining its geometry. Here is the formal statement of the local non-positive curvature condition:

Definition 2.1.3. *A cube complex X is* ***non-positively curved (NPC)*** *if for all $v \in X^{(0)}$, link(v) is a flag simplicial complex.*

If X is non-positively curved and simply connected, then X is a ***CAT(0) cube complex****.*

The geometric consequences of the non-positive curvature will be revisited in Chapter 2.2. Note that if X is a non-positively curved cube complex, then X is completely determined by its 2-skeleton because every vertex link is flag simplicial.

Hyperplanes are an essential tool in the study of cube complexes:

Definition 2.1.4. *In an n-cube, two edges are* ***parallel*** *if they do not form a corner of a dimension-2-face. If e, f are edges, then say $e \sim f$ if and only if e is parallel to f. The relation $\sim$ defines an equivalence relation. Let E be an equivalence class of parallel edges, the convex hull of the midpoints of the edges of E is an $n-1$-cube called a* ***dual codimension-1-midcube****.*

The relation extends to an equivalence relation on the edges of the cube complex X as follows: two edges e, f are ***parallel in*** *X, $e \sim f$ if and only if there exists a sequence of edges $e_0 = e, e_1, \cdots, e_m = f$ such that e_i and e_{i-1} are parallel in an n-cube of X.*

Let P be a parallelism class of edges in X. The disjoint union of the codimension-1-midcubes dual to elements of P glued naturally along common faces is a ***hyperplane*** *H of X. The inclusion of each codimension$-1-$midcube into its respective n-cube also*

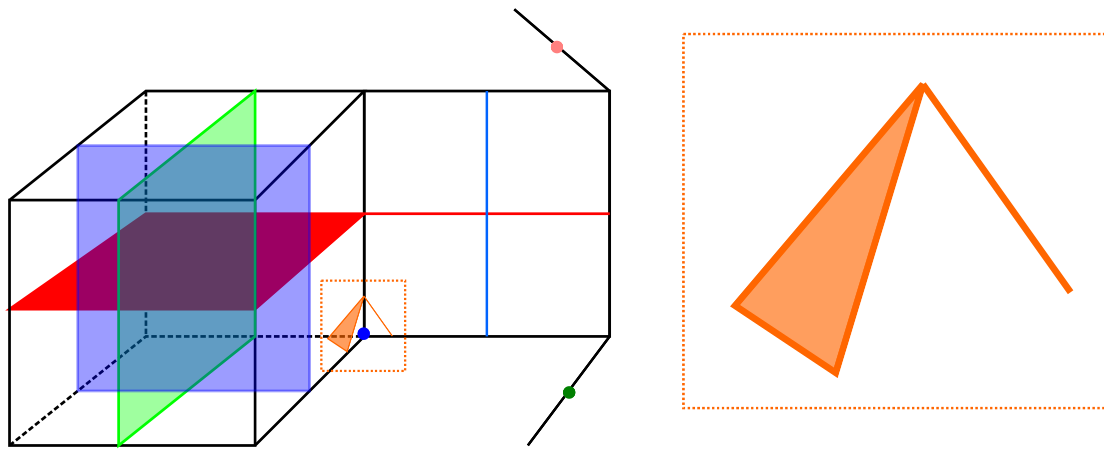

Figure 2.1: An example of a NPC cube complex (including a 3-cube) with its hyperplanes as well as the link of the blue vertex shown in orange and enlarged on the right.

induces a natural immersion $H \to X$.

See Figure 2.1 for an example of a NPC cube complex and the link of a vertex.

Cube complexes have special immersions called **local isometries** which respect the cubical structure:

Definition 2.1.5. *Let X, Y be non-positively curved cube complexes and let $f : X \to Y$ be a combinatorial map. Note that for each vertex $v \in X^{(0)}$, f induces a map $f_v : link(v) \to link(f(v))$. The map f is **full** if for all $v \in X^{(0)}$ and vertices v_1, v_2 of $link(v)$, if $f_v(v_1), f_v(v_2)$ are joined by an edge e in $link(f(v))$, then e is the image of an edge between v_1 and v_2 under f_v. The map f is a **local isometry** or **local isometric immersion** if f is full and is an immersion.*

A combinatorial map of NPC cube complexes $f : X \to Y$ is an immersion if and only if for all vertices $v \in X^{(0)}$, the induced map $f_v : \text{link}(v) \to \text{link}(f(v))$ is injective. Figure 2.2 shows an example of a local isometric immersion and Figure 2.3 gives a non-example.

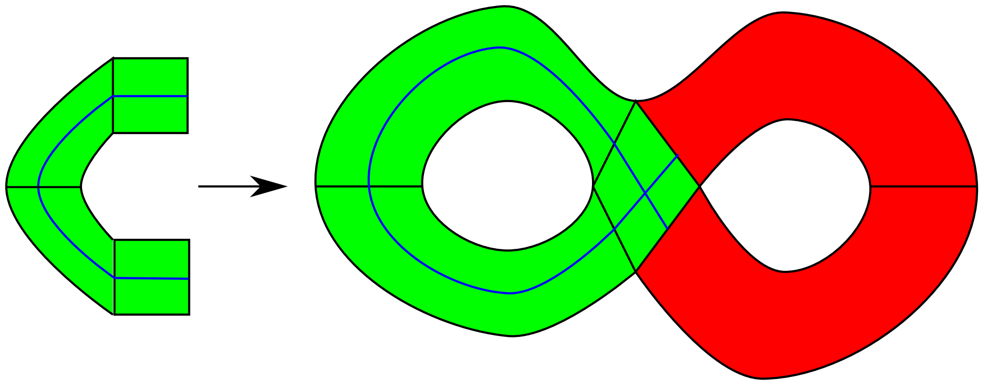

Figure 2.2: Mapping the four green squares by following the green hyperplane is an example of a local isometry which is an immersion rather than an embedding.

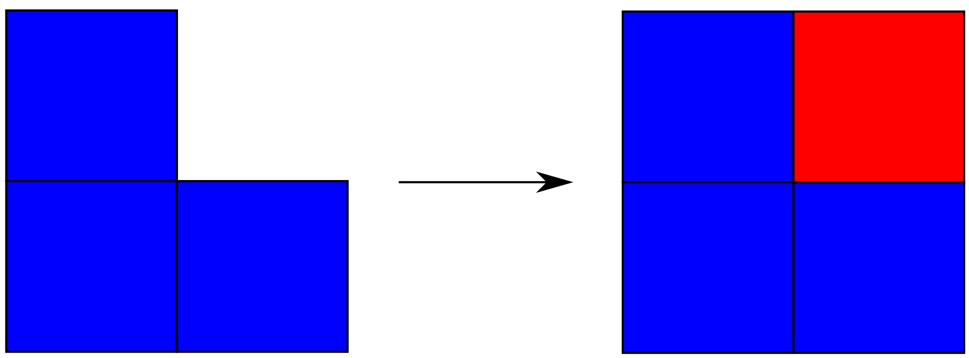

Figure 2.3: Mapping the three blue squares into the four squares on the right fails to be a local isometry because the induced map on links is not full at the central vertex of the four squares.

Local isometries also induce injections on fundamental groups:

Proposition 2.1.6 ([15] Lemma 2.11)**.** *Let X, Y be finite dimensional NPC cube complexes, and let $f : Y \to X$ be a local isometric immersion. Then the induced map $f_* : \pi_1 Y \to \pi_1 X$ is injective.*

2.2 CAT(0) geometry

This subchapter contains basic facts about CAT(0) geometry and explains that a CAT(0) cube complex has a CAT(0) metric. These details are explained thoroughly in Parts I and II of [9]. Roughly speaking, a CAT(0) metric space is one where the triangles are "thinner than Euclidean."

Definition 2.2.1. *Let (X,d) be a geodesic metric space and let $a,b,c \in X$ and let $\triangle abc$ be a geodesic triangle with vertices a,b,c. There exists a unique geodesic triangle $\triangle \bar{a}\bar{b}\bar{c}$ (up to isometry) in $\mathbb{R}^2$ together with a map $f : \triangle \bar{a}\bar{b}\bar{c} \to \triangle abc$ such that f is isometric on each side, $f(a) = \bar{a}$, $f(b) = \bar{b}$ and $f(c) = \bar{c}$. The triangle $\overline{\triangle}\bar{a}\bar{b}\bar{c}$ is called a **comparison triangle** and f is called a **comparison map**. In the following, the comparison triangle of $\triangle abc$ will be denoted by $\overline{\triangle}(a,b,c)$ and the images $\bar{a}, \bar{b}$ and $\bar{c}$ of a,b,c may be left implicit.*

*The metric space (X,d) is **CAT(0)** if for all triangles $\triangle abc$ with comparison triangles $\overline{\triangle}(a,b,c)$, comparison maps $f : \overline{\triangle}(a,b,c) \to \triangle abc$ and all pairs $x,y \in \triangle abc$, $\bar{d}(f(x),f(y)) \leq d(x,y)$ where $\bar{d}$ is the standard metric on $\mathbb{R}^2$.*

For the remainder of this section, let (X,d) to be a CAT(0) metric space. CAT(0) spaces also have a well-defined notion of angle.

Definition 2.2.2 ([9] Definition I.1.12)**.** *Let $c : [0,a] \to X$ and $c' : [0,a'] \to X$ be geodesics issuing from the point $c(0) = c'(0)$. Let $0 < t \leq a$ and $0 < t' \leq a'$ and consider the comparison triangle $\overline{\triangle}(c(0), c(t), c'(t'))$. Let $\overline{\angle}_{c(0)}(c(t), c'(t'))$ denote the Euclidean angle at $\overline{c'(0)}$ in the comparison triangle formed by the preimages of these three points. The angle $\overline{\angle}_{c(0)}(c(t), c'(t'))$ is called the **comparison angle** between*

$c(t), c'(t')$ *at* $c(0)$. *The* ***(Alexandrov) angle at*** $c(0)$ ***between*** c ***and*** c' *is given by:*

$$\angle_{c(0)}(c, c') := \limsup_{t,t' \to 0} \overline{\angle}_{c(0)}(c(t), c'(t')) = \lim_{\epsilon \to 0} \sup_{0<t,t'<\epsilon} \overline{\angle}_{c(0)}(c(t), c'(t'))$$

Whenever the limit $\lim_{t,t'\to 0} \overline{\angle}_{c(0)}(c(t), c'(t'))$ *exists, the angle is said to* ***exist in the strict sense****.*

If B *is a closed convex subspace of* X *with* $x \in B$ *and* γ *is a geodesic with an endpoint at* x*, then* $\angle_x^X(\gamma, B) := \inf_{\rho \in \Lambda} \angle_x^X(\gamma, \rho)$ *where* Λ *is the set of germs of geodesics in* B *based at* x*.*

In a CAT(0) space, the angle *always* exists in the strict sense:

Proposition 2.2.3 ([9] II.3.1)**.** *Let* c, c' *be geodesics issuing from the same point as in Definition 2.2.2. Then the comparison angle* $\overline{\angle}_{c(0)}(c(t), c'(t'))$ *is a non-decreasing function of both* t *and* $t' \geq 0$*, the angle exists in the strict sense and:*

$$\angle_{c(0)}(c, c') = \lim_{t,t' \to 0} \overline{\angle}(c(t), c'(t')).$$

CAT(0) spaces have unique geodesics, so given three points $a, b, c \in X$, it makes sense to define $\angle_b(a, c)$ as the angle created by the geodesics $[b, a]$ and $[b, c]$. Usually, the ambient space is left implicit, but sometimes, the notation $\angle_{c(0)}^X(c, c')$ will be used to denote that the angle is being taken in the metric on X. Angles will be revisited and used extensively in Chapter 4.1.

The Cartan-Hadamard theorem shows the non-positive curvature condition and the local Euclidean metric on each cube together induce a CAT(0) metric on the universal cover of a non-positively curved cube complex:

Theorem 2.2.4 (Cartan-Hadamard Theorem, see [9] Theorem II.4.1)**.** *Let* X *be a complete connected metric space.*

1. *If the metric on X is locally convex, then the induced length metric on the universal cover $\widetilde{X}$ is a convex uniquely geodesic metric and geodesic segments in $\widetilde{X}$ vary continuously with their endpoints.*
2. *If every point of X is contained in a ball on which the induced metric is CAT(0), then the induced length metric on $\widetilde{X}$ is CAT(0).*

The hyperplanes of a CAT(0) cube complex are embedded (see [34] Theorem 3.2), so they are also locally convex in the length metric. This yields the following important corollary:

Corollary 2.2.5. *Let H be a hyperplane of a CAT(0) cube complex $\widetilde{X}$. Then H is a convex subspace of $\widetilde{X}$ in the standard CAT(0) metric on $\widetilde{X}$.*

Hausdorff distance is a useful way of measuring distance between subsets:

Definition 2.2.6. *Let (X, d) be a metric space and let A, B be subsets of X. Then*

$$d_{Haus}(A, B) := \max\{\sup_{a \in A} \inf_{b \in B} d(a, b), \sup_{b \in B} \inf_{a \in A} d(a, b)\}$$

2.3 (Hyperbolic) and Relatively Hyperbolic Geometry

A key notion in geometric group theory is that of quasi-isometry:

Definition 2.3.1. *Let X and Y be metric spaces and let $f : X \to Y$ be a continuous map. Let $\lambda \geq 1$ and $\epsilon > 0$. If for all $x, y \in X$,*

$$\frac{1}{\lambda} d_X(x, y) - \epsilon \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y) + \epsilon,$$

*then f is called a (λ, ϵ)-**quasi-isometric embedding**.*

*If f is a quasi-isometric embedding, and for all $z \in Y$, there exists $x \in X$ so that $d_Y(z, f(x)) \leq \epsilon$, then f is called a (λ, ϵ)-**quasi-isometry**.*

The quasi-isometry constants will often be left implicit. If $f : X \to Y$ is a quasi-isometry, there exists a **quasi-inverse**, that is, a map $g : Y \to X$ and $\lambda' \geq 1$, $\epsilon' \geq 0$ and $k \geq 0$ such that g is a (λ', ϵ')-quasi-isometry, for all $x \in X$, $d(gf(x), x) \leq k$ and for all $z \in Y$, $d(fg(z), z) \leq k$. The map g can be constructed by mapping $z \in Y$ to x where $x \in X$ such that $d_Y(f(x), z) \leq \epsilon$.

For the remainder of this section, let X be a geodesic metric space.

Definition 2.3.2. *Let $a, b, c \in X$ and let $\triangle abc$ be a geodesic triangle. There exists a map $f : \triangle abc \to T(a, b, c)$ where $T(a, b, c)$ is a unique tripod (up to isometry) with center point x such that f is isometric on each side of the triangle and the three legs of the tripod are $[f(a), x]$, $[f(b), x]$ and $[f(c), x]$. The tripod $T(a, b, c)$ is called a **comparison tripod** for $\triangle abc$. The map f is the **comparison map**.*

The length of the tripod leg $[f(a), x]$ is equal to $(b \cdot c)_a := \frac{1}{2}(d(a, b) + d(a, c) - d(b, c))$, which is called the **Gromov product of b and c with respect to a**

A hyperbolic metric space is one where preimages of points in the comparison tripod have uniformly bounded diameter:

Definition 2.3.3. *Let X be a geodesic metric space. Let $\delta > 0$. Let $\triangle abc$ be a geodesic triangle in X with comparison tripod $T(a, b, c)$ and comparison map $f : \triangle abc \to T(a, b, c)$. If for all $t \in T(a, b, c)$, $\operatorname{diam} f^{-1}(t) < \delta$, then $\triangle abc$ is called δ-**thin**.*

*If every triangle in X is δ-thin, then X is called δ-**hyperbolic**. The space X is called **hyperbolic** if there exists some $\delta > 0$ such that X is δ-hyperbolic.*

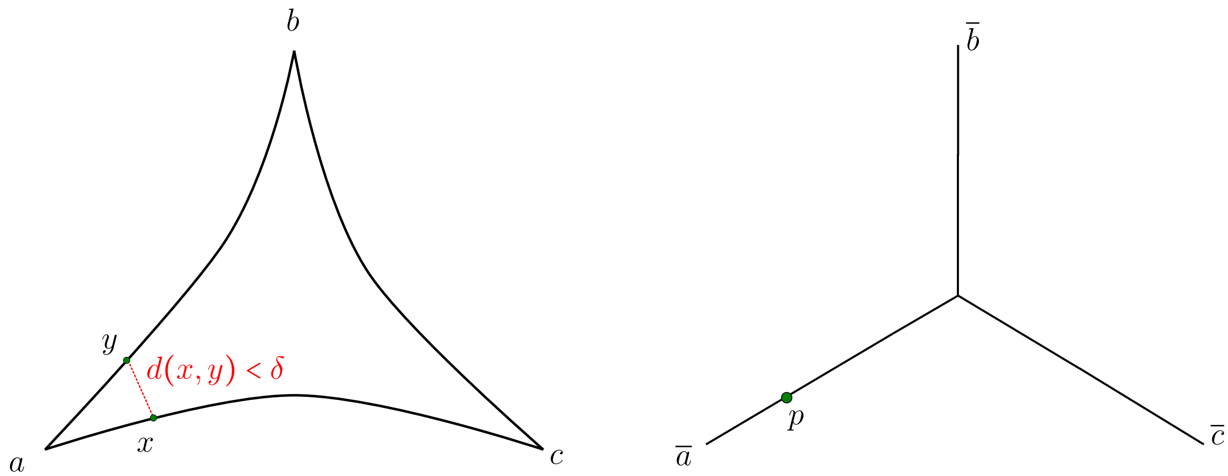

Figure 2.4: A δ-thin triangle and its comparison tripod with points x, y in the triangle that are preimages of a point p in the tripod under the comparison map.

Examples of δ-hyperbolic spaces include trees and the real hyperbolic spaces $\mathbb{H}^n$. Being hyperbolic is a quasi-isometry invariant.

Proposition 2.3.4 ([9] Theorem III.H.1.9)**.** *Let X and Y be geodesic metric spaces and let $f : X \to Y$ be a quasi-isometry, then X is hyperbolic if and only if Y is hyperbolic.*

Although hyperbolicity is a quasi-isometry invariant, the quasi-isometry may change the hyperbolicity constant. For example if T is a tree (so T is 0-hyperbolic), $T \times [0,1]$ is quasi-isometric to T, and is not 0–hyperbolic because T contains triangles which are not tripods.

Definition 2.3.5. *A finitely generated group G is **hyperbolic** if a Cayley graph for G is hyperbolic.*

By the preceding proposition, if G is hyperbolic, then every Cayley graph of G is hyperbolic. Further, by the Milnor-Švarc lemma, if X admits a proper cocompact action by a group G, X is hyperbolic if and only if G is.

An advantage of hyperbolic spaces is that quasigeodesics do not stray far from geodesics with the same endpoints.

Proposition 2.3.6 (*Quasigeodesic Stability*, see [9] Theorem III.H.1.7)**.** *Let X be a δ–hyperbolic metric space and let $\lambda \geq 1$ and $\epsilon > 0$. There exists $R = R(\delta, \lambda, \epsilon) \geq 0$ such that if α is a (λ, ϵ)-quasigeodesic in X and γ is a geodesic in X with the same endpoints as α, then $d_{Haus}(\alpha, \gamma) \leq R$.*

A major theme of this thesis is to take ideas from hyperbolic geometry and apply them to the more general situation where a **relatively hyperbolic group** (see Definition 2.3.8) acts properly and cocompactly on a CAT(0) space. Rather than having quasigeodesic stability, there will be a notion of "relative fellow-traveling."

Roughly speaking, a relatively hyperbolic group is a group together with a collection of "peripheral subgroups" such that coning off the cosets of the peripheral subgroups in the Cayley graph yields a hyperbolic graph where quasigeodesics sharing endpoints pass through the same peripheral coset cone points. The definition given here is originally due to Farb [11], but the version given here is taken from [17].

Let G be a group and let $\mathcal{P} := \{P_1, \ldots, P_m\}$ be a collection of finitely generated subgroups of $\mathcal{P}$. A set S is a **symmetrized relative generating set** for the pair $(G, \mathcal{P})$ if $S \cup P_1 \cup \cdots \cup P_m$ generate G and $S = S^{-1}$. Let $\Gamma = \Gamma(G, S)$ be the Cayley graph with respect to S (which may or may not be connected). The **coned-off Cayley Graph of** $(G, \mathcal{P})$ **with respect to** S is a graph $\hat{\Gamma}(G, \mathcal{P}, S)$ formed from Γ by adding a vertex called a **cone point** for every distinct coset of the form gP with $g \in G$ and $\mathcal{P}$ in P and for every vertex v in Γ corresponding to an element of gP, adding an edge of length $\frac{1}{2}$ connecting v to cone point for gP. When S is a relative generating set, $\hat{\Gamma}(G, \mathcal{P}, S)$ is connected.

An oriented path γ in $\hat{\Gamma}(G, \mathcal{P}, S)$ **penetrates the coset** gP if γ passes through the cone point for gP. If γ is a path which never passes through the same cone point twice, then γ is a **path without backtracking**. When γ passes through a cone point, the vertex immediately preceding that cone point is called an **entering vertex for** gP and the vertex immediately after that cone point is called an **exiting vertex for** gP.

Definition 2.3.7 ([17] Definition 3.5)**.** *Let S be a relative generating set for $(G, \mathcal{P})$. The triple $(G, \mathcal{P}, S)$ has **Bounded Coset Penetration** if and only if for all $\lambda \geq 1$, there exists $a = a(\lambda) > 0$ such that if γ and γ' are two $(\lambda, 0)$-quasigeodesics in $\hat{\Gamma}(G, \mathcal{P}, S)$ without backtracking whose initial endpoints are the same and whose terminal endpoints are at most distance 1 apart, then the following hold:*

1. *If γ penetrates gP but γ' does not penetrate gP, then the distance in $\Gamma(G, S)$ between the entering and exiting vertices of γ for gP is at most a.*
2. *If γ and γ' both penetrate gP, then the entering vertices for gP of both γ and γ' are distance at most a apart in $\Gamma(G, S)$ and similarly, the exiting vertices are at most a apart in $\Gamma(G, S)$.*

Definition 2.3.8 ([17] Definition 3.6)**.** *Let G be finitely generated relative to $\mathcal{P}$ with each $P \in \mathcal{P}$ finitely generated. The pair $(G, \mathcal{P})$ is a **relatively hyperbolic pair** if for some finite relative generating set S, the coned-off Cayley graph $\hat{\Gamma}(G, \mathcal{P}, S)$ is hyperbolic and $(G, \mathcal{P}, S)$ has bounded coset penetration.*

*The elements of $\mathcal{P}$ and their conjugates are called **peripheral subgroups** and the cosets $\{gP : g \in G,\ P \in \mathcal{P}\}$ are called peripheral cosets.*

When S is finitely generated, $(G, \mathcal{P})$ is a relatively hyperbolic pair if and only if Definition 2.3.8 holds for every finite relative generating set.

There are many notions of relatively hyperbolic, but most of the commonly used definitions are equivalent in the case that G is finitely generated and $\mathcal{P}$ has finitely many finitely generated elements (see [17] for details).

Here are two examples to help illustrate the preceding definition:

Example 2.3.9 ($\mathbb{Z}^2$ is not relatively hyperbolic relative to one of the $\mathbb{Z}$ factors.)**.** *Let*

$$G = \mathbb{Z}^2 = \langle (1,0), (0,1) \,|\, (1,0) + (0,1) - (1,0) - (0,1) = (0,0) \rangle,$$

let S be the specified generating set and let $\mathcal{P} = \{\langle (0,1) \rangle\}$ so that in the Cayley graph $\Gamma(G,S)$ the cosets of $\langle (0,1) \rangle$ correspond to the vertical lines. Let γ be the path connecting $(0,0)$ to $(1,0)$ in $\Gamma(G,S)$ and let γ_n be the path in $\hat{\Gamma}(G, \mathcal{P}, S)$ consisting of the following three segments:

1. *A path of length 1 from $(0,0)$ to $(0,n)$ via the cone point for $\langle (0,1) \rangle$,*
2. *a path of length 1 from $(0,n)$ to $(1,n)$ in $\Gamma(G,S)$, and*
3. *a path of length 1 from $(1,n)$ to $(1,0)$ passing through the cone point for $(1,0) + \langle (0,1) \rangle$.*

The path γ_n has length 3 and is therefore a $(3,0)$ quasigeodesic without backtracking, but γ_n penetrates $\langle (0,1) \rangle$ and has entering and exiting points that are distance n apart in $\Gamma(G,S)$.

Therefore, the $(\mathbb{Z}^2, \mathcal{P}, S)$ does not have bounded coset penetration.

Note that in the preceding example, the coned-off Cayley graph is a quasi-line which is hyperbolic, but that $(\mathbb{Z}^2, \mathcal{P})$ is not hyperbolic relative to $\langle(0,1)\rangle$.

Example 2.3.10. *Let $G = \langle a, b, c \,|\, bcb^{-1}c^{-1}\rangle$, let S be the specified generating set, and let $\mathcal{P} = \{P\}$ where $P = \langle b, c\rangle \cong \mathbb{Z}^2$. Observe G is isomorphic to $\pi_1(S^1 \vee T^2)$ where T^2 is a 2-torus, so the Cayley graph is quasi-isometric to a "tree of flats" consisting of flat 2-planes corresponding to lifts of the torus and length 1 edges connecting lattice points of flats corresponding to lifts of S^1. Coning off the cosets of P turns the flats into diameter $1 + \sqrt{2}$ spaces, so the coned off Cayley graph is a quasi-tree (and is therefore hyperbolic).*

That $(G, \mathcal{P}, S)$ has bounded coset penetration roughly follows from the fact that given any two flats F_1 and F_2, every geodesic path in $\hat{\Gamma}(G, \mathcal{P}, S)$ without backtracking between F_1 and F_2 passes through a sequence of flats $F_1 = E_1, E_2, \ldots, E_m = F_2$ in that order and the entry and exit points must be the same as the entry and exit points of γ.

In this case, $(G, \mathcal{P})$ is a relatively hyperbolic pair.

The reason why bounded coset penetration failed in the first example is that the cosets of $\langle b\rangle$ have k-neighborhoods which have infinite diameter intersection while in the second example, k-neighborhoods of flats corresponding to distinct cosets of P intersect in sets of diameter at most $2k - 2$. Following Example 2.3.9, k-neighborhoods of peripheral cosets have intersections with uniformly bounded diameter:

Proposition 2.3.11. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let $k \geq 0$. Let S be a finite generating set for G. For all $R \geq 0$, there exists a $M = M(R)$ such if gP, $g'P'$ is a pair of distinct peripheral cosets, then $\operatorname{diam} \mathcal{N}_R(gP) \cap \mathcal{N}_R(g'P') \leq M$ in the word metric on $\Gamma(G, S)$.*

Proof. Let $z, z' \in \mathcal{N}_R(gP) \cap \mathcal{N}_R(g'P')$ Let $x, y \in gP$ and $x', y' \in g'P'$ such that $d_S(z,x), d_S(z,y) \leq R$ and $d(z',x'), d(z',y') \leq R$. Therefore $d_S(x,x'), d_S(y,y') \leq 2R$.

If $x = x'$ and $y = y'$, let γ be the geodesic path of length 1 in $\hat{\Gamma}(G, \mathcal{P}, S)$ connecting x to y via the cone point for gP and let γ' be the geodesic path of length 1 connecting $x' = x$ to $y' = y$ via the cone point for $g'P'$. In this case γ and γ' are geodesics, γ penetrates only gP and γ' penetrates only $g'P'$.

The path γ is geodesic in $\hat{\Gamma}(G, \mathcal{P}, S)$ and penetrates only gP while γ' is also geodesic in $\hat{\Gamma}(G, \mathcal{P}, S)$ and penetrates only $g'P' \neq gP$. These paths have the same entry and exit points. Therefore, by bounded coset penetration, there exists $a > 0$ such that $d_S(x,y) = d_S(x',y') \leq a$.

Otherwise, WLOG, assume $x \neq x'$, so $d_{\hat{\Gamma}}(x,x') \geq 1$. In the coned-off Cayley graph $\hat{\Gamma}(G, \mathcal{P}, S)$, construct the path γ from x to x' from the following 3 paths:

1. the path of length 1 from x to y,
2. the path of length at most $2R$ from y to y', and
3. the path of length 1 from y' to x'.

Since $1 \leq d_{\hat{\Gamma}}(x,x') \leq 2R$, the path γ is a $(2R+2, 0)$-quasigeodesic in $\hat{\Gamma}(G, \mathcal{P}, S)$. Let γ' be the S-geodesic in $\Gamma(G,S)$ connecting x and x' that has length at most $2R$ and so is also a $(2R+2)$-quasigeodesic in $\hat{\Gamma}$ which does not penetrate any peripheral cosets. The points x, y are the entry and exit points of γ from gP. Similarly, x', y' are the entry and exit points of γ from $g'P'$. Since γ' does not penetrate any peripheral cosets, by bounded coset penetration, there exists $b \geq 0$ such that $d_S(x, y) \leq b$ and $d_S(x', y') \leq b$.

Let $M = \max\{a, b\} + 4R$. Then $d(z, z') \leq M$. □

The goal is to use the geometry of a relatively hyperbolic group G to study the geometry of spaces that G acts on geometrically.

Corollary 2.3.12. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair with finite generating set S and let G act properly and cocompactly on a geodesic space X by isometries. Let $x \in X$ be a base point. For all $R \geq 0$, there exists $M = M(R, G, \mathcal{P}, S)$ such that if $P, P' \in \mathcal{P}$, $g, g' \in G$ with $gP \neq g'P'$, then $\operatorname{diam}(\mathcal{N}_R(gPx) \cap \mathcal{N}_R(g'P'x)) \leq M$.*

The corollary follows immediately from Proposition 2.3.11 and the Milnor-Švarc Lemma.

Passing to a finite index subgroup induces a natural peripheral structure:

Proposition 2.3.13. *Let G be a group and let $\mathcal{P}$ be a finite collection of subgroups of G. Let $H \triangleleft G$ be a finite index normal subgroup. For each $P \in \mathcal{P}$, let $\mathcal{E}_0(P) = \{gPg^{-1} \cap H \,|\, g \in G\}$ and let $\mathcal{E}(P)$ be a set of representatives of H-conjugacy classes in $\mathcal{E}_0(P)$. Let $\mathcal{P}' = \bigsqcup_{\mathcal{P} \in P} \mathcal{E}(P)$.*

The pair $(G, \mathcal{P})$ is relatively hyperbolic if and only if $(H, \mathcal{P}')$ is relatively hyperbolic.

The elements of $\mathcal{P}'$ are each a finite index subgroup of an element of $\mathcal{P}$ and H is a finite index subgroup of G, so Proposition 2.3.13 follows from the fact that the inclusion of a Cayley graph for H into a Cayley graph for G is a quasi-isometry.

Another key step in understanding the geometry of spaces with a geometric action of a relatively hyperbolic group is developing an analogue to thin triangles.

Definition 2.3.14. *Let X be a geodesic metric space, and let $\mathcal{B}$ be a collection of subspaces.*

Let $\triangle abc$ be a geodesic triangle in X and let $\delta > 0$. Let $T(a,b,c)$ be the comparison triangle defined in Definition 2.3.3 and let $f : \triangle abc \to T(a,b,c)$ be the comparison map. If there exists $F \in \mathcal{B}$ such that for all $p \in T(a,b,c)$, either

1. $\operatorname{diam}(f^{-1}(p)) < \delta$ *or*
2. $f^{-1}(p) \subseteq \mathcal{N}_\delta(F)$,

*then $\triangle abc$ is **δ-thin relative to** F.*

The notion of fellow-traveling will be useful for describing behavior of geodesics that issue from the same point. Definitions of fellow-traveling may vary, so the one that will be used in this thesis is recorded here:

Definition 2.3.15. *Let $\alpha : [a_1, a_2] \to X$ and $\beta : [b_1, b_2] \to X$ be geodesics, and let $k \geq 0$. The geodesics α and β k**-fellow travel for distance** D if $d(\alpha(a_1 + t), \beta(b_1 + t)) \leq k$ for all $0 \leq t \leq D$. If α and β k-fellow travel for distance D, then α **and** β k**-fellow travel distance** D **from** x.*

Definition 2.3.16. *Let X be a geodesic metric space, $\delta \geq 0$ and let $\mathcal{B}$ be a collection of subspaces. The space X has the δ–**relatively thin triangle property relative to** $\mathcal{B}$ if each geodesic triangle Δ is δ-thin relative to some $B_\Delta \in \mathcal{B}$.*

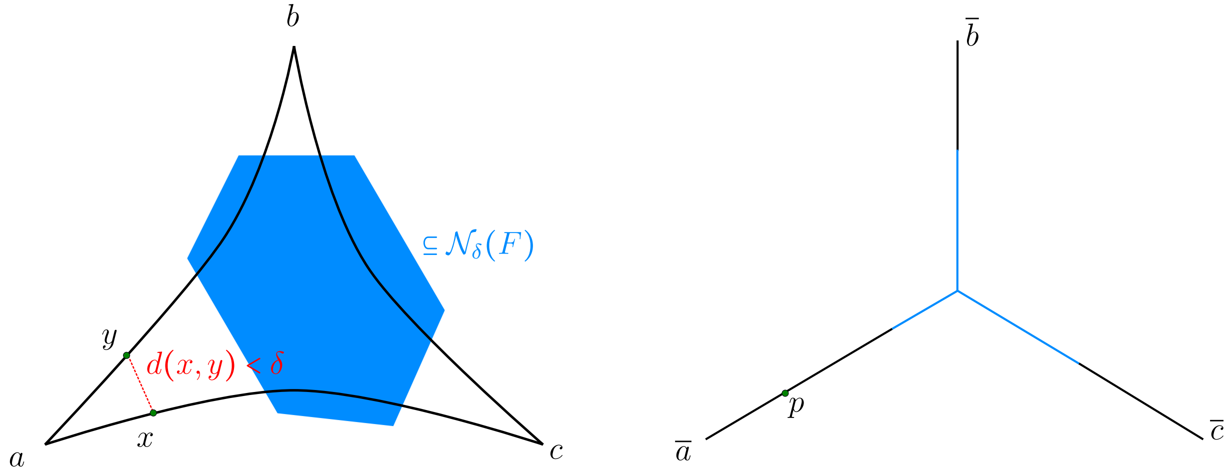

Figure 2.5: An example of a triangle which is δ-thin relative to some F with its comparison tripod. Points in the blue part of the tripod have preimages in the triangle which lie in the blue shaded region. All other points have preimages in the triangle with diameter δ like the point p whose preimates x, y have $d(x,y) < \delta$. The fat part (see Definition 2.3.18) of each side is the subsegment that intersects the blue shaded region.

The space X may contain triangles that are δ-thin. By definition, these triangles are δ-thin relative to every element of $\mathcal{B}$.

In the applications, X will usually be a CAT(0) space with a geometric action by a relatively hyperbolic group G where the elements of $\mathcal{B}$ are convex subspaces of X that lie in uniformly bounded neighborhoods of peripheral coset orbits. If $(G, \mathcal{P})$ is a relatively hyperbolic group pair, a CAT(0) space with a geometric action by G has the relatively thin triangle property relative to $\mathcal{F} = \{gPx \mid g \in G,\ P \in \mathcal{P}\}$:

Proposition 2.3.17 ([28] Theorem 4.1, Proposition 4.2, see also Section 8.1.3 of [10])**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let G act properly and cocompactly on a CAT(0) space X by isometries. Let $x \in X$ be a base point and set $\mathcal{F} = \{gPx \mid g \in G,\ P \in \mathcal{P}\}$.*

Then there exists $\delta > 0$ such that whenever $\triangle$ is a geodesic triangle in X, $\triangle$ is δ-thin relative to some $F \in \mathcal{F}$.

In particular, if $R \geq 0$ and $\mathcal{F}' = \{\mathcal{N}_R(F) : F \in \mathcal{F}\}$, then X still has the relatively thin triangle property relative to $\mathcal{F}'$.

Definition 2.3.18. *Let X be a CAT(0) geodesic metric space with triangles that are δ-thin relative to $\mathcal{B}$. Let $\triangle \subseteq X$ with vertices a, b, c with comparison map $f : \triangle abc \to T(a, b, c)$. Let L_a be the closure of the leg of the tripod $T(a, b, c)$ that contains $f(a)$. Let* $\textbf{\textit{Thin}}_a := \{f^{-1}(p) : \text{diam}(f^{-1}(p)) < \delta\} \cap f^{-1}(L_a)$. *The **corner segments** of $\triangle$ at a are the subsegments of the sides in* $\textbf{\textit{Thin}}_a$ *and the **corner length** is the length of a corner segment at a.*

*The **fat part** of the side $ab \subseteq \triangle$ in $\triangle$ is the closure of* $ab \smallsetminus (\textbf{\textit{Thin}}_a \cup \textbf{\textit{Thin}}_b)$.

The corner segments at a are subsegments of the sides issuing from a that δ-fellow travel. Each of these segments have the same length, which is defined to be the corner length. If $\triangle$ is δ-thin relative to $B_\triangle \in \mathcal{B}$, the fat part of each side of $\triangle$ is the subsegment that does not lie in any of the corner segments and hence lies in $\mathcal{N}_\delta(B_\triangle)$. Note that the fat part of a side may be empty. Since X is CAT(0), each corner segment or fat part of a side is connected.

Similarly, quasigeodesic triangles in the Cayley graph of a relatively hyperbolic group also satisfy a thinness condition which is used to obtain Proposition 2.3.17:

Theorem 2.3.19 ([28] Theorem 4.1, originally due to [10])**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair with Cayley graph Γ. For all $\lambda \geq 1$, $\epsilon > 0$ there exists a $\delta > 0$ such that if $\triangle$ is a (λ, ϵ)-quasigeodesic triangle in Γ with sides c_0, c_1, c_2, either:*

1. *there exists a point p that lies within $\frac{\delta}{2}$ of each side or*

2. *each side of c_i has a subpath c_i' such that the terminal endpoint of c_i' and the initial point of c_{i+1}' (indices mod 3) are within distance δ of each other.*

There is also a generalized version of quasiconvexity for relatively hyperbolic groups.

Definition 2.3.20 ([17] Definition 6.10)**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair. Let $H \leq G$. Let S be any finite set such that $S \cup \mathcal{P}$ generates G, and suppose there exists $\kappa(S, d_S)$ such that every vertex of $\hat{\Gamma}(G, S, \mathcal{P})$-geodesic γ in G lies within κ of H in d_S. Then H is* ***relatively quasiconvex in*** $(G, \mathcal{P})$.

Note that there are other equivalent definitions which are discussed in [17]. The definition is also independent of the choice of finite relative generating set (see [17] Theorem 7.10).

In this thesis, relative quasiconvexity will only be needed for the peripheral subgroups:

Proposition 2.3.21. *Let $(G, \mathcal{P})$ be a relatively hyperbolic group pair. Then every element of $\mathcal{P}$ is relatively quasiconvex in $\mathcal{P}$.*

Proof. In $\hat{\Gamma}(G, S, \mathcal{P})$ every $P \in \mathcal{P}$ has diameter 1. □

Lemma 2.3.22 is just several applications of the triangle inequality, but is instrumental for working with relatively thin triangles.

Lemma 2.3.22. *Let Δabc be a geodesic triangle that is δ-thin relative to F. Let ab, bc, ac denote the sides of Δabc. If the length of the fat part of ac in Δabc is bounded above by*

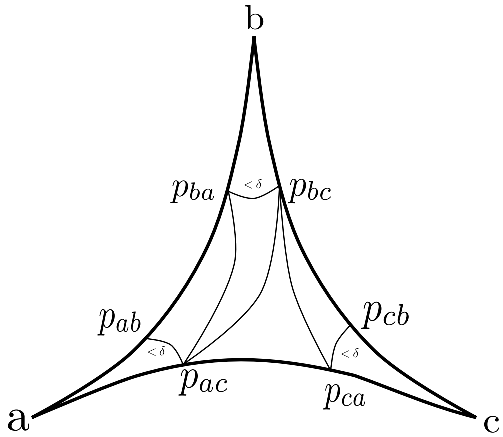

Figure 2.6: Applying the triangle inequality four times gives a bound on the difference between the length of $[p_{ab}, p_{ba}]$ and the length of $[p_{bc}, p_{cb}]$ in terms of $|[p_{ac}, p_{ca}]|, \delta$.

M, then the length of the fat part of bc and the length of the fat part of ab differ by at most $M + 3\delta$.

Proof. Choose points $p_{ab}, p_{ba}, p_{ac}, p_{ca}, p_{bc}, p_{cb}$ (see Figure 2.6) so that $[a, p_{ab}]$, $[a, p_{ac}]$ are corner segments at a, $[b, p_{ba}]$, $[b, p_{bc}]$ are the corner segments at b and $[c, p_{cb}], [c, p_{ca}]$ are corner segments at c. Then $[p_{ab}, p_{ba}] \subseteq ba$, $[p_{ac}, p_{ca}] \subseteq ac$, and $[p_{bc}, p_{bc}] \subseteq bc$ are the fat parts of each respective side and lie in $\mathcal{N}_\delta(F)$. The claim is that $|d(p_{ab}, p_{ba}) - d(p_{bc}, p_{cb})| \leq d(p_{ac}, p_{ca}) + 3\delta$. Refer to Figure 2.6. By the triangle inequality:

$$||[p_{ab}, p_{ba}]| - |[p_{ba}, p_{ac}]|| < \delta$$

$$||[p_{ba}, p_{ac}]| - |[p_{bc}, p_{ac}]|| < \delta$$

$$||[p_{bc}, p_{ac}]| - |[p_{bc}, p_{ca}]|| < |[p_{ac}, p_{ca}]|$$

$$||[p_{bc}, p_{ca}]| - |p_{bc}, p_{cb}|]|| < \delta.$$

Putting these four inequalities together:

$$|d(p_{ab}, p_{ba}) - d(p_{bc}, p_{cb})| < d(p_{ac}, p_{ca}) + 3\delta \qquad \square$$

When triangles that are thin relative to $\mathcal{B}$ where elements of $\mathcal{B}$ have bounded coarse intersections, the bounds on coarse intersections can be used to help bound the fat part of one side of a triangle that is δ-thin relative to $B \in \mathcal{B}$. By Lemma 2.3.22, a bound on the fat part of one side of a relatively thin triangle helps control the lengths of the fat parts of the other two sides. This technique will be used repeatedly, particularly in Chapter 4.

CHAPTER 3

GRAPHS OF GROUPS AND HIERARCHIES

3.1 Graphs of Groups

A graph of groups (together with an isomorphism from the fundamental group) is a way of decomposing a group along a finite number of splittings and HNN extensions. Further decomposing the vertex groups as graphs of groups, decomposing the resulting vertex groups as a graph of groups again and continuing this process a finite number of times yields a kind of "multilevel graph of groups" called a hierarchy which will be defined in Definition 3.2.1.

Definition 3.1.1. *A **graph of groups** (Γ, χ) consists of the following data:*

1. *a finite graph $\Gamma = \Gamma(V, E)$ where V is the vertex set of Γ and E is the oriented edge set of Γ with an involution $e \mapsto \overline{e}$ that switches the orientation of each edge,*
2. *an **assignment map** $\chi : V \sqcup E \to$ **Grp** that assigns a group to each vertex and edge,*
3. *for all $e \in E$, $\chi(e) = \chi(\overline{e})$,*
4. ***attachment homomorphisms** $\psi_e : \chi(e) \to \chi(t(e))$ where $t(e)$ is the terminal vertex of the edge e.*

*Γ is a **faithful** graph of groups if the attachment homomorphisms ψ_e are injective.*

An intuitive way of thinking about a graph of groups is to view a graph of groups as a way of splitting the fundamental group of a topological space X. The assignment map χ may sometimes be left implicit. A **graph of spaces** is constructed like a graph of groups, except that the assignment map χ assigns a topological space instead of a group to each edge and vertex. The attachment homomorphisms are replaced by continuous **attachment maps**, and a **faithful graph of spaces** has π_1-injective attachment maps. A **graph of spaces realization** of X is a way of identifying X with a space constructed by gluing together the mapping cylinders of the attachment maps:

Definition 3.1.2. *Let (Γ, χ) be a graph of spaces with vertex and edge sets V, E. Let X be a topological space, and let Y be the topological space constructed as follows:*

$$Y := \bigsqcup_{v \in V} \chi(v) \sqcup \left(\bigsqcup_{e \in E} \chi(e) \times [0,1] \right) / \{ \chi(e) \times \{t\} = \chi(\bar{e}) \times \{1-t\}, \\ \chi(e) \times \{1\} = \psi_e(\chi(e)) : \forall e \in E,\ 0 \le t \le 1 \}.$$

*A triple (Γ, χ, f) is a **graph of spaces realization for** X if Γ, χ, Y are as above and f is a homotopy equivalence between X and Y.*

Applying the Seifert-Van Kampen theorem to Y immediately yields a splitting of the fundamental group of X. Generally speaking, this thesis will use graphs of spaces to build intuition but will ultimately aim to derive results for graphs (and hierarchies) of groups.

Note that some authors such as Wise take faithfulness to be a part of the definition of a graph of groups. Not requiring faithfulness makes it easier to define graphs of groups in terms of graphs of spaces. For the applications in Chapter 6, graphs of groups will be constructed first without showing that they are faithful, but with some effort, these graphs of groups will turn out to be faithful.

The **fundamental group of a graph of groups** is analogous to the space Y constructed in Definition 3.1.2. For a tree of groups (i.e. a graph of groups where the underlying graph is a tree), the fundamental group is an amalgamated free product of the vertex groups where the image of an edge group in its terminal vertex is identified with its image in its initial vertex. To extend this construction to a general graph of groups (Γ, χ), choose a maximal tree, T in the underlying graph, construct the fundamental group of T and for each pair of edges $(e, \overline{e})$ of Γ not in T, perform an HNN extension with stable letter s_e so that conjugation of the image of $\chi(e)$ in the vertex group $\psi_{\overline{e}}(\chi(e))$ by s_e is identified with the image of $\chi(e)$ in the vertex group $\psi_e(\chi(e))$.

Definition 3.1.3. *Let (T, χ) be a tree of groups with vertex set V. The **fundamental group of** T is given by the following construction. Let:*

$$\pi_1(T) \coloneqq *_{v \in V} \chi(v) / \langle\langle \{\psi_e(\chi(e)) = \psi_{\overline{e}}(\chi(\overline{e})) : e \text{ is an edge of } T\} \rangle\rangle$$

*Let (Γ, χ) be a graph of groups and let T be a maximal tree in Γ with complementary edge set $E_T = \{e_1, \ldots, e_m, \overline{e_1}, \ldots, \overline{e_m}\}$ (where $\overline{e_i}$ has the same vertices as e_i but with the opposite orientation). The **fundamental group of** Γ **with respect to** T is the group:*

$$\pi_1(\Gamma, T) \coloneqq \langle \pi_1(T), s_1, \ldots, s_m \,|\, s_i^{-1} \psi_{\overline{e_i}}(\chi(\overline{e_i}) s_i = \psi_{e_i}(\chi(e_i)) \rangle$$

A choice of maximal tree to determine the fundamental group of a graph of groups is analogous to a choice of base point for determining the fundamental group of a space. Given a graph of groups (Γ, χ) and a maximal tree T, then $G \coloneqq \pi_1(\Gamma, T)$ acts on a **covering tree** $\hat{T}$ without inversions (i.e. for every edge e of $\widetilde{\Gamma}$ and $g \in G$, $g \cdot e \neq \overline{e}$) such that the quotient is $\widetilde{\Gamma}/G = \Gamma$. Choosing a different maximal tree T' will yield a G-equivariant isomorphism of covering trees and an isomorphism $\pi_1(\Gamma, T) \cong \pi_1(\Gamma, T')$ (See [30] for details). Some care will be needed to account for choices of maximal trees.

A **graph of groups structure** for a group is analogous to a graph of spaces realization for a topological space:

Definition 3.1.4. *Let G be a group, let (Γ, χ) be a graph of groups where T is a maximal tree and let $\phi : G \to \pi_1(\Gamma, T)$ be an isomorphism. The triple (Γ, ϕ, T) is a* ***graph of groups structure on*** *G.*

The structure (Γ, ϕ, T) is ***degenerate*** *if Γ is a single vertex labeled with G and ϕ is the identity.*

While a graph of groups structure determines a splitting of G, the choice of isomorphism and maximal tree affects the precise splitting. For example, given a graph of groups structure (Γ, ϕ, T) for G and any $g \in G \setminus 1$, there is another graph of groups structure (Γ', ϕ', T') where Γ' has the same underlying graph as Γ and T' is the same tree as T in the underlying graph, but each vertex group G_v has been replaced by $g^{-1}G_v g$ and each edge group G_e has been replaced by $g^{-1}G_e g$. In many cases, it suffices to give the splitting up to conjugacy which will be the case in the examples below. When the splitting is given up to conjugacy, the choice of maximal tree also becomes unnecessary.

Example 3.1.5. *Figure 3.1 shows a graph of spaces decomposition of a genus 2 surface and a graph of groups splitting of the fundamental group induced by the graph of spaces decomposition.*

More generally, the pants decomposition of a surface Σ_g will give a graph of groups decomposition of $\pi_1(\Sigma_g)$ where each vertex group is a free group of rank 2 and every edge group is an infinite cyclic group.

It is possible to construct a graph of spaces corresponding to a graph of groups.

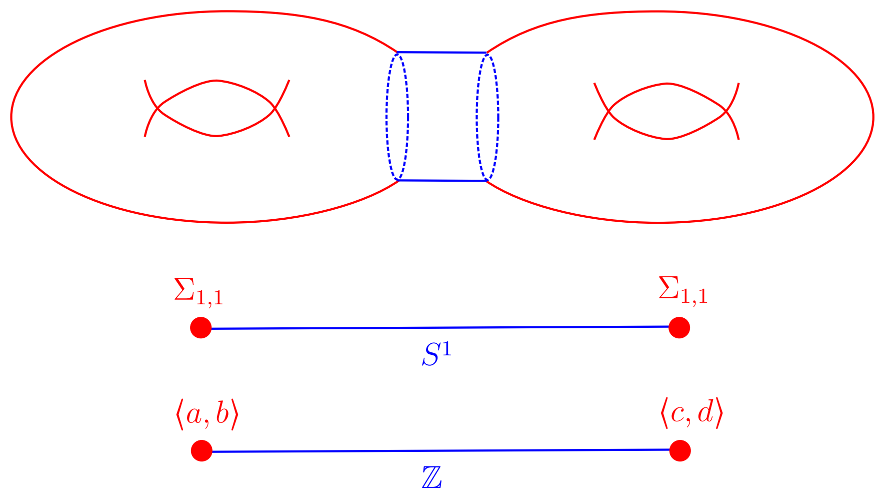

Figure 3.1: A graph of spaces realization of a genus 2 surface where $\Sigma_{1,1}$ is a punctured torus, together with the corresponding graph of groups obtained by applying the π_1 functor.

Example 3.1.6. *Let Γ be the graph of groups in Figure 3.1. Let $F_1 = \langle a, b \rangle$ and $F_2 = \langle c, d \rangle$ be the vertex groups and let $G_e = \langle z \rangle \cong \mathbb{Z}$ be the edge group. It is possible to construct a graph of spaces for $\pi_1 G$ (there is only one possible maximal tree here) by taking vertex spaces V_1 and V_2 where V_i is a rose with 2 petals, each petal is labeled a or b in V_1, and each petal is labeled c or d in V_2. The edge space is a copy of S^1 labeled c. The attaching maps will be $\psi_1 : S^1 \to V_1$ defined by $c \mapsto aba^{-1}b^{-1}$ and $\psi_2 : S^1 \mapsto (cdc^{-1}d^{-1})^{-1}$.*

A realization can be made by taking mapping cylinders for ψ_1 and ψ_2 and gluing them together in the natural way.

This construction can be made much more general whenever the graph of groups is faithful and the vertex and edge groups are finitely generated. For each vertex group G_v and edge group G_e, create vertex spaces $X_v \cong K(G_v, 1)$ and edge spaces $X_e \cong K(G_e, 1)$ with the attaching maps induced by the maps $X_e^{(1)} \to X_{t(e)}^{(1)}$.

Passing to finite index subgroups is a commonly used technique in geometric group theory. Graph of groups structures interact naturally with finite index normal subgroups. The following proposition is Proposition 3.18 of [3] but is originally due to Bass [5].

Proposition 3.1.7. *Suppose G has a graph of groups structure (Γ, ϕ, T), $H \triangleleft G$ and H is finite index in G. Then H has an induced graph of groups structure $(\widetilde{\Gamma}, \widetilde{\phi}, T')$ so that:*

1. *Every vertex group of $(\widetilde{\Gamma}, T')$ has the form $(K^g \cap H) \triangleleft K^g$ and is finite index in K^g for some vertex group K of (Γ, T) and some $g \in G$.*
2. *Every edge group of $(\widetilde{\Gamma}, T')$ has the form $(K^g \cap H) \triangleleft K^g$ and is finite index in K^g for some edge group K of (Γ, T) and some $g \in G$.*

Proposition 3.1.7 follows from considering the action of H on the covering tree of the graph of groups structure for G. Taking the graph of groups quotient provides the desired splitting.

3.2 Hierarchies

Hierarchies of groups are inductively defined multilevel graphs of groups:

Definition 3.2.1. *A **hierarchy of groups of length** 0 is a single vertex labeled by a group.*

*A **hierarchy of groups of length** n is a graph of groups (Γ_n, χ_n) together with hierarchies of length $n-1$ on each vertex of Γ_n.*

*If $\mathcal{H}$ is a length n hierarchy of groups, the n**th level** of $\mathcal{H}$ is the graph of groups Γ_n. For $1 \leq k \leq n$, the $(n-k)$th level of $\mathcal{H}$ is the disjoint union of the $(n-k)$th levels of the hierarchies on the vertices of Γ_n.*

*The **terminal groups** are the groups labeling the vertices at level 0.*

It will be useful to think of graphs of groups as length 1 hierarchies. Realizing a group as a hierarchy is similar to finding a graph of groups structure for that group:

Definition 3.2.2. *Let G be a group, $\mathcal{H}$ be a hierarchy of length n. Let (Γ_n, χ_n) be the level n graph of groups. A **hierarchy for** G is $\mathcal{H}$ together with a graph of groups structure (Γ_n, ϕ, T) for G. Let $\mathcal{P}$ be a collection of subgroups of G. The hierarchy structure **terminates in** $\mathcal{P}$ if every terminal group of $\mathcal{H}$ is conjugate to $\phi(\mathcal{P})$ for some $P \in \mathcal{P}$.*

It will often be convenient to forget the choice of maximal tree and only give a hierarchy structure for a group up to conjugacy. In general, hierarchies will be allowed to contain degenerate splittings, but in order to obtain non-trivial results, it will be necessary to ensure that at least one of the splittings in the hierarchy is non-degenerate.

Wise's hierarchies in [35] permit only one-edge splittings rather than allowing a graph of groups splitting for each vertex group in the hierarchy. The hierarchies in Definition 3.2.2 can be converted to hierarchies with one-edge splittings for each vertex group at the expense of increasing the length of the hierarchy. Wise's hierarchies also terminate in the trivial group while Definition 3.2.2 allows arbitary terminal groups. In practice, the goal in Chapter 6 will be to find a hierarchy for a relatively hyperbolic group $(G, \mathcal{P})$ that terminates in the peripheral subgroups $\mathcal{P}$. Chapter 7 will explore what happens to the hierarchy after quotienting out the peripheral subgroups.

A **hierarchy of spaces** and a **hierarchy realization for a space** X can be defined analogously by replacing groups in Definition 3.2.1 with topological spaces and replacing graph of groups structures by realizations in Definition 3.2.2.

Malnormality is an important group property which will play a role in Chapter 7 and is useful for amalgamating virtually special groups to make new virtually special groups (see [20]).

Definition 3.2.3. *Let G be a group and let $H \leq G$. The subgroup H is **malnormal in** G if for all $g \in G \setminus H$, $g^{-1}Hg \cap H = \{1\}$. Similarly, H is **almost malnormal in** G if for all $g \in G \setminus H$, $|g^{-1}Hg \cap H| < \infty$.*

Malnormality also extends to collections of subgroups. Let $\mathcal{P}$ be a collection of subgroups of G. The collection $\mathcal{P}$ is (almost) malnormal in G if for all $g \in G$ and $P, P' \in \mathcal{P}$ either $g^{-1}Pg \cap P'$ is trivial (finite) or $P = P'$ and $g \in P$.

For example, if $(G, \mathcal{P})$ is a relatively hyperbolic pair and G is finitely generated, then the collection $\mathcal{P}$ is almost malnormal in G by Proposition 2.3.11.

The height of a subgroup generalizes malnormality:

Definition 3.2.4. *Let $H \leq G$. The **height** of H in G is the the largest $n \geq 0$ such that there are n distinct cosets $\{g_1H, \ldots, g_nH\}$ where $\bigcap_i g_iHg_i^{-1}$ is infinite.*

Therefore, a finite group has height 0 while an almost malnormal group has height 1. If G is torsion free, then every height 1 subgroup is malnormal. Height will be important in Chapter 7.

Definition 3.1.1 (graphs of groups) and Definition 3.2.1 (hierarchies) are very flexible, but in practice, some further restrictions will be needed to ensure that graphs of groups and hierarchies produce useful splittings:

Definition 3.2.5. *Let (Γ, χ) be a faithful graph of groups and let (Γ, ϕ) be a graph of groups structure (up to conjugacy) for a group G.*

1. *Γ is **quasiconvex** if every edge attachment map is a quasi-isometric embedding into $\pi_1(\Gamma)$.*
2. *Γ is **(almost) malnormal** if for every $e \in E$, the image of the attachment homomorphism ψ_e in $\pi_1(\Gamma)$ is (almost) malnormal in $\pi_1(\Gamma)$.*

Let $\mathcal{H}$ be a hierarchy for G.

1. *$\mathcal{H}$ is **faithful** if every graph of groups at every level of $\mathcal{H}$ is faithful.*
2. *$\mathcal{H}$ is **quasiconvex** if every edge group of every graph of groups at every level of $\mathcal{H}$ quasi-isometrically embeds in G.*
3. *$\mathcal{H}$ is **(almost) malnormal** if every edge group of every graph of groups at every level of $\mathcal{H}$ is (almost) malnormal in G.*

It may be possible to give a reasonable weaker definition of quasiconvex (or malnormal) hierarchy by only requiring an edge group G_e of a graph of groups H in $\mathcal{H}$ to be quasi-isometrically embedded (malnormal) in each adjacent vertex group, but the stronger definition given here will be needed in Chapter 7.

Here are some examples to help illustrate the definition of a hierarchy:

Example 3.2.6. *Figure 3.2 shows a length 2 hierarchy of spaces realization for a genus 4 surface. By applying the π_1 functor, a faithful graph of groups structure is obtained.*

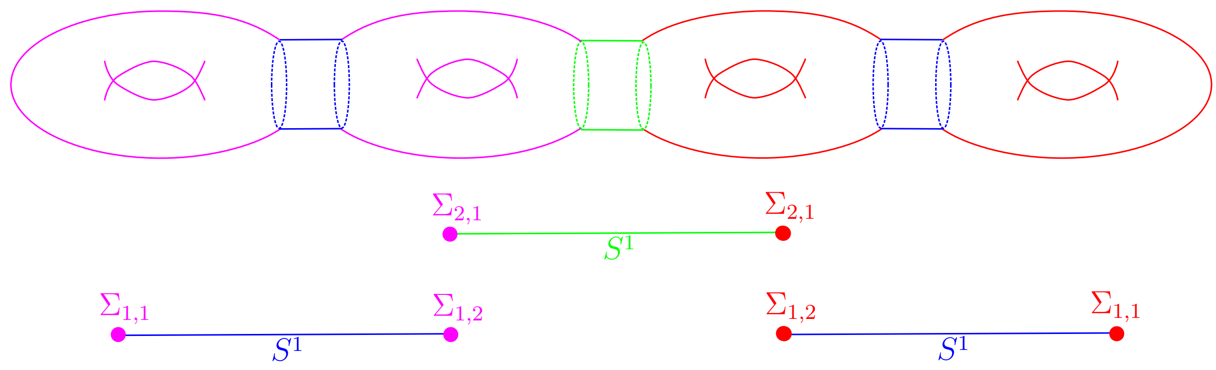

Figure 3.2: A length 2 hierarchy of spaces decomposition for a genus 4 surface Σ_4.

Applying the π_1 functor gives a hierarchy structure for $\pi_1(\Sigma_4)$ with terminal subgroups isomorphic to F_2 or F_3 (the free group on 2 or 3 letters). This hierarchy is quasiconvex. Indeed, if $\langle g \rangle$ is one of the infinite cyclic edge groups, $\langle g \rangle$ corresponds to a simple closed geodesic of Σ_4 which corresponds to a quasi-isometrically embedded hyperbolic axis fixed by g in $\mathbb{H}^2$ the universal cover of Σ_4 with quasi-isometry constants depending on the translation length of g. Further, the conjugates of g by elements in $G \smallsetminus \langle g \rangle$ cannot be powers of g because $\pi_1(\Sigma_4)$ is hyperbolic and does not contain an embedded copy of a Baumslag-Solitar group. Therefore, if $h \in G \smallsetminus \langle g \rangle$, then $\langle g \rangle \cap h\langle g \rangle h^{-1} = \{1_G\}$, so the hierarchy is also malnormal.

The hierarchy in Example 3.2.6 could be replaced by a graph of groups splitting (length 1 hierarchy) by performing all the splittings simultaneously rather than in two steps. However, there are some splittings that cannot be accomplished in a single step like this:

Example 3.2.7. *Figure 3.3 shows a length 2 hierarchy for the fundamental group of a genus 2 surface, Σ_2. Cuts are made along the both the blue and green simple closed curves which intersect, so the iterated splitting of the fundamental group cannot be accomplished by a graph of groups (length 1 hierarchy).*

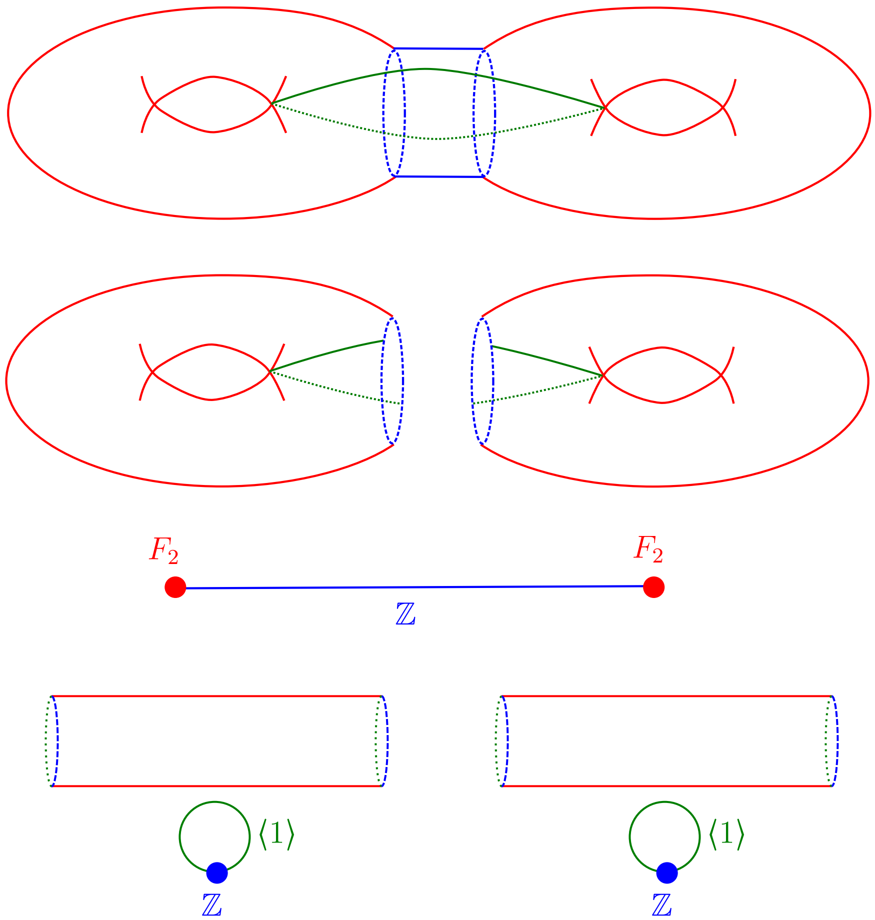

Figure 3.3: A hierarchy for $\pi_1(\Sigma_2)$, the fundamental group of a genus 2 surface Σ_2, where the iterated splitting of $\pi_1(S_2)$ cannot be realized by a graph of groups. The first splitting is over the infinite cyclic subgroup of $\pi_1(\Sigma_2)$ corresponding to one of the blue copies of S^1. The resulting vertex spaces are punctured tori whose fundamental groups are rank 2 free groups. Cutting along the green arc in each punctured torus makes an annulus. Then the fundamental group of a punctured torus splits as an HNN extension of the fundamental group of an annulus ($\mathbb{Z}$) over the trivial group (corresponding to the green arcs in each annulus which are glued together to make a punctured torus).

Other notable examples of hierarchies are the Haken Hierarchy for Haken 3-manifolds, see [23] Chapter 9.4, and the Magnus-Moldvanskii hierachy for one-relator groups, see [35] Section 18.

Proposition 3.1.7 extends to hierarchies by induction on the length of the hierarchy.

Corollary 3.2.8. *Supppose G has a hierarchy $\mathcal{H}$ and H is a finite index normal subgroup of G, then $\mathcal{H}$ has an induced hierarchy $\mathcal{H}'$ such that the length of $\mathcal{H}$ is the length of $\mathcal{H}'$ and:*

1. *every vertex group at level i of the hierarchy $\mathcal{H}'$ is of the form $K^g \cap H$ which is finite index and normal in K^g for some vertex group K of $\mathcal{H}$ at level i and some $g \in G$,*
2. *every edge group at level i of the hierarchy $\mathcal{H}'$ is of the form $K^g \cap H$ which is finite index and normal in K^g for some edge group K of $\mathcal{H}$ at level i and some $g \in G$.*

Lemma 3.2.9 follows from Corollary 3.2.8:

Lemma 3.2.9. *If $\mathcal{H}$ is a quasiconvex hierarchy for G and G_0 is a finite index normal subgroup of G, then the induced hierarchy on $\mathcal{H}_0$ on G_0 is quasiconvex.*

The definition of a quasiconvex hierarchy for a group G only requires that the edge groups are quasi-isometrically embedded in G; when a graph of groups (Γ, ϕ, T) structure for G is quasiconvex, the vertex groups are quasi-isometrically embedded as well.

Lemma 3.2.10. *Let (Γ, T) be a graph of groups structure for G. If the edge groups of G are quasi-isometrically embedded in G, then the vertex groups of Γ are quasi-isometrically embedded in G.*

Here is a rough sketch of the proof of Lemma 3.2.10. A Cayley graph $\Lambda(G, S)$ of G coarsely looks like a "tree of spaces" whose underlying (infinite) graph is the covering tree of (Γ, T) where the edge spaces are Cayley graphs of edge groups and the vertex spaces are Cayley graphs of vertex groups. If $\Lambda_v := \Lambda(G_v, S_v)$ is one of the vertex spaces, the coarse tree structure ensures that if a $\Lambda(G, S)$-geodesic shortcut γ between two points in Λ_v exits Λ_v through an edge space Λ_e, it must return through Λ_e and the geodesic in Λ_e between the entry and exit points p_e, p'_e in $\Lambda_v \cap \Lambda_e$ is a quasi-geodesic with constants set by the quasi-isometric embedding of Λ_e into $\Lambda(G, S)$. If γ enters and exists Λ_v at points $p_{e_1}, p'_{e_1}, \ldots, p_{e_m}, p'_{e_m}$, then a piecewise geodesic path ρ connecting the initial point of γ to p_{e_1}, the entry and exit points in order and the endpoint of γ to p'_{e_m} lies entirely in Γ_v and cannot be much longer than γ.

3.3 Fully $\mathcal{P}$-Elliptic Hierarchies

Given a relatively hyperbolic group pair $(G, \mathcal{P})$ and a hierarchy $\mathcal{H}$ for G, the goal in Chapter 7 will be to strategically find a quotient of G that has a hierarchy induced by $\mathcal{H}$ and inherits a relatively hyperbolic structure from $(G, \mathcal{P})$ that is also compatible with the induced hierarchy structure. To ensure that this happens, some additional restrictions must be imposed on the interactions between the edge and vertex groups of the hierarchy and the peripheral subgroups of G.

Definition 3.3.1. *Let $\mathcal{H}$ be a hierarchy for a group G and let $\mathcal{P}$ be a collection of subgroups of G. Let $\mathcal{V}$ be the vertex groups of $\mathcal{H}$. For each $H \in \mathcal{V}$, let $\pi_1(\Gamma_H, \phi_H, T_H)$ be the graph of groups structure for H induced by the hierarchy $\mathcal{H}$. The hierarchy $\mathcal{H}$ is $\mathcal{P}$-**elliptic** if the following holds: whenever there exists a $g \in G$ such that $P' := gPg^{-1} \subseteq H \in \mathcal{V}$, then there exists an $h \in H$ such that $hP'h^{-1}$ is contained in some* *vertex group of* Γ_H.

A $\mathcal{P}$*-elliptic hierarchy is* ***fully*** $\mathcal{P}$ ***elliptic*** *if the following holds: whenever* E *is an edge group in* $\mathcal{H}$*, then for all* $g \in G$*, either* $gPg^{-1} \cap E$ *is finite or* $gPg^{-1} \subseteq E$*.*

When $\mathcal{H}$ is a fully $\mathcal{P}$-elliptic hierarchy for G and G_0 is a finite index normal subgroup of G, the induced hierarchy from Corollary 3.2.8 for H is also fully $\mathcal{P}$-elliptic in the induced peripheral structure provided by Proposition 2.3.13:

Proposition 3.3.2. *Suppose that* G_0 *is finite index normal in* G *and let* $(G_0, \mathcal{P}_0)$ *be the peripheral structure induced on* G_0 *by Proposition 2.3.13. If* G *has a fully* $\mathcal{P}$ *elliptic hierarchy, then the induced hierarchy* $\mathcal{H}_0$ *of* G_0 *is fully* $\mathcal{P}_0$*-elliptic.*

Proposition 3.3.2 follows immediately from the explicit characterizations of the edge and vertex groups of the induced hierarchies in Corollary 3.2.8 and from the explicit description of the induced peripheral structure.

CHAPTER 4

A RELATIVELY HYPERBOLIC COMBINATION LEMMA

The construction of hierarchies in Chapter 6 is quite similar to the hierarchy constructed in [3]. However, new methods are needed to prove that the hierarchy is indeed quasiconvex. Let X be a NPC cube complex with fundamental group $G := \pi_1 X$ where $(G, \mathcal{P})$ is a relatively hyperbolic pair, and let $\widetilde{X}$ be the CAT(0) universal cover of X with base point $x \in \widetilde{X}$.

Let $\mathcal{A}$ and $\mathcal{B}$ be collections of convex subspaces of $\widetilde{X}$. Let E be a subspace of X so that an elevation of E to the universal cover is a union of elements of $\mathcal{A}$ and $\mathcal{B}$ and triangles in $\widetilde{X}$ are thin relative to $\mathcal{B}$. Suppose elements $B \in \mathcal{B}$ have embedded neighborhoods isometric to $B \times [0, \epsilon)$. If every $\mathcal{A} \in A$ that intersects B intersects this neighborhood as $A \cap B \times [0, \epsilon)$, then the main result of the chapter, Proposition 4.2.11 implies that an elevation of E to $\widetilde{X}$, the universal cover of X, is quasi-isometrically embedded in X. In the applications, the complex X can be modified using a trick called augmentation (see Definition 6.2.3) so that each $B \in \mathcal{B}$ has an image in X with an embedded neighborhood $B \times [0, \epsilon)$ in a homotopically equivalent cube complex $\hat{X}$. For example, $\mathcal{A}$ is a collection of hyperplanes of $\hat{X}$, then intersections between elements of $\mathcal{A}$ and $B \times [0, \epsilon)$ will have the desired form. In the application in Chapter 6, the spaces in $\mathcal{A}$ will be convex subsets of hyperplanes rather than hyperplanes themselves.

There are three major components in the proof of Proposition 4.2.11: the first step is to prove a path taming result, Proposition 4.1.5. Proposition 4.1.5 tames piecewise geodesic paths in E to piecewise geodesic paths in $\widetilde{X}$ so that the new path has length bounded below by a linear function of the original path, every other piece of the new piecewise geodesic path lies in some $B \in \mathcal{B}$

and pieces not contained in some $B \in \mathcal{B}$ depart from B at a rate bounded from below. The next step is to prove a combination lemma, Lemma 4.2.2 that uses the geometry of relatively thin triangles to show that certain paths produced by Proposition 4.1.5 are quasigeodesics. Proposition 4.2.11 finally shows how the results of Proposition 4.1.5 and Lemma 4.2.2 can be applied to paths of a more general form.

4.1 Improving Piecewise Geodesics

One special feature of CAT(0) spaces is that convex subspaces have nearest point projections maps:

Lemma 4.1.1 ([9] II.2.4)**.** *Let $\widetilde{X}$ be a CAT(0) space and let $Y \subseteq \widetilde{X}$ be convex and complete in the induced metric. There exists a unique closest point projection $\pi_Y : \widetilde{X} \to Y$.*

Let γ be a piecewise geodesic with geodesic pieces a_γ and b_γ where a_γ lies in some $A \in \mathcal{A}$ and $b_\gamma \subseteq B \in \mathcal{B}$. Suppose also that $\gamma = a_\gamma \cup b_\gamma$ is geodesic in the length metric on an edge space of a hierarchy. The goal of this subchapter is to reconfigure the pieces a_γ and b_γ so that a departs the peripheral space B quickly.

The key technical tool is Proposition 4.1.9 which says that either a_γ makes an angle of at least $\frac{\pi}{6}$ with B or otherwise, a_γ and b_γ make an angle of at least $\frac{2\pi}{3}$ with each other. When a_γ makes an angle less than $\frac{\pi}{6}$ with B, projecting a point p'' on a_γ close to B to the point p' and then creating a piecewise geodesic γ' by connecting p' with the endpoints of γ makes a new path γ' that cannot be shorter than a fixed linear function of the length of γ because a_γ and b_γ make an angle of $\frac{2\pi}{3}$. When γ has more than 2 pieces, the process is more complicated (see Construction 4.1.10), but the construction is similar.

Definition 4.1.2. *Let $\widetilde{X}$ be a CAT(0) space, let δ, $D \geq 0$ and let $\mathcal{B}$ be a collection of complete convex subspaces. The pair $(\widetilde{X}, \mathcal{B})$ is a (δ, D)-**Relatively Hyperbolic pair** if*

1. *every triangle in $\widetilde{X}$ is δ-thin relative to $\mathcal{B}$,*
2. *for all $B_1, B_2 \in \mathcal{B}$ with $B_1 \neq B_2$, $\operatorname{diam}(\mathcal{N}_{3\delta}(B_1) \cap \mathcal{N}_{3\delta}(B_2)) \leq D$.*

Definition 4.1.3. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, D)-relatively hyperbolic pair. Let $\mathcal{A}$ be a collection of complete convex subspaces of $\widetilde{X}$. Let E be a union of subspaces of $\widetilde{X}$ in $\mathcal{A}$ and $\mathcal{B}$ so that*

$$E = \left(\bigcup_{i \in I_A} A_i\right) \cup \left(\bigcup_{j \in I_B} B_j\right)$$

where each $A_i \in \mathcal{A}$ and each $B_j \in \mathcal{B}$. If there exists $\epsilon > 0$ such that:

1. *each B_j has an embedded neighborhood in $\widetilde{X}$ isometric to $B_j \times [0, \epsilon)$ where $B_j = B_j \times \{0\}$ and*
2. *whenever $A_j \cap B_j \neq \varnothing$, then $A_j \cap (B_j \times [0, \epsilon)) = (B_j \cap A_j) \times [0, \epsilon)$,*

*then E is a $(\mathcal{A}, \mathcal{B})$-**Local Peripheral Projection Closed subspace** of $\widetilde{X}$. For short, E will be called an $(\mathcal{A}, \mathcal{B})$-**LPPC**.*

Let $\widetilde{X}$ be a CAT(0) cube complex where triangles are δ-thin relative to some set $\mathcal{B}$. Let $\widetilde{X}_0$ be the mapping cylinder of the maps $F \to \widetilde{X}$ where $F \in \mathcal{B}$ and $F \times \{0\}$ is an embedded copy of F in $\widetilde{X}_0$ and $F \times \{1\}$ is the image of F in $\widetilde{X}$. The prototypical example of an $(\mathcal{A}, \mathcal{B})$-LPPC in $\widetilde{X}_0$ is a hyperplane H of $\widetilde{X}_0$ together with any elements of $\mathcal{B}$ that intersect H. Each of the intersections (locally) looks like the union of the green hyperplane H and the red cube face in Figure 4.1.

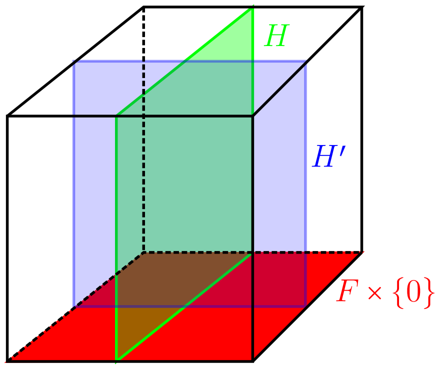

Figure 4.1: The green hyperplane and the red cube face look like an intersection of an element of $\mathcal{A}$ with an element of $\mathcal{B}$ in an $(\mathcal{A}, \mathcal{B})$-LPPC subspace.

The definition of an LPPC also accommodates the situation where a hyperplane has been cut-up by another hyperplane. For example, the closure of the union of the red face of the cube and the half of H in front of the blue hyperplane H' in Figure 4.1 is still an LPPC.

When E is an $(\mathcal{A}, \mathcal{B})$-LPPC subspace, intersections between elements of $\mathcal{A}$ and $\mathcal{B}$ contained in E interact nicely with the CAT(0) projection map.

Proposition 4.1.4. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, D)-relatively hyperbolic pair and let E be an $(\mathcal{A}, \mathcal{B})$-LPPC subspace.*

Let $A \in \mathcal{A}$ and $B \in \mathcal{B}$ be each contained in E and suppose $A \cap B$ is non-trivial. Let γ be a geodesic in the length metric on E, and suppose that γ can be written as $\gamma = ab$ where a is a geodesic in A and b is a geodesic in B, and a, b meet at a point $x \in A \cap B$. Let π be the CAT(0) projection onto B. Then there exists a neighborhood U of x in A such that $\pi(a \cap U)$ is the image of a geodesic in $A \cap B$.

Proof. The local product structure implied by an LPPC subspace ensures that for a sufficiently small neighborhood U_0 of x, $a \cap U_0$ can be written as $a(t) =$

$(a_1(t), a_2(t))$ where $a_1(t) : [0, \alpha) \to U_0 \cap A \cap B$ and $a_2(t) : [0, \alpha) \to [0, \epsilon)$. Then $\pi(a(t)) = a_1(t)$. By [9] Proposition I.5.3(3), for any $r < \alpha$, since $(a_1(t), a_2(t))$ is geodesic, on $[0, r]$, a_1 restricted to $[0, r]$ is a linearly reparameterized geodesic. By choosing a smaller $U \subseteq U_0$ so that $a \cap U \subseteq a([0, r])$, $\pi(a \cap U)$ is the image of a geodesic in $A \cap B$. □

The remainder of the subchapter will be devoted to proving the following proposition:

Proposition 4.1.5. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, D)-relatively hyperbolic pair and let E be an $(\mathcal{A}, \mathcal{B})$-LPPC subspace of $\widetilde{X}$.*

Let $\gamma \subseteq E$ be a geodesic in E. The E-geodesic γ can be written as a piecewise geodesic $a_1 b_1 \ldots a_n b_n$ in $\widetilde{X}$ where each a_i is a point or lies in $A_i \in \mathcal{A}$ and each $b_i \subseteq B_i \in \mathcal{B}$. Assume n is minimal. Then there exists a piecewise geodesic $\gamma' = a'_1 b'_1 \ldots a'_n b'_n$ such that:

1. *for all $1 \leq i \leq n$, $b'_i \subseteq B_i$,*
2. $\operatorname{diam}(a'_i \cap \mathcal{N}_{3\delta}(B_i)) \leq 2D + 29\delta$ *and for* $1 < i \leq n$, $\operatorname{diam}(a'_i \cap \mathcal{N}_{3\delta}(B_{i-1})) \leq 2D + 29\delta$.
3. $|b'_i| \geq \frac{1}{256}|b_i| - 8\delta$.
4. $|\gamma'| \geq \frac{1}{256}|\gamma| - 16n\delta$ *and γ' has the same endpoints as γ.*

Before proving Proposition 4.1.5, a number of intermediate results are needed about the geometry of angles in CAT(0) spaces. The first lemma shows that if two geodesic segments issuing from a point x make an angle of $\frac{\pi}{3}$, the length of the geodesic between the other two endpoints is bounded below by a linear function of the combined length of the segments.

Lemma 4.1.6. *Let $a : [0, \ell_1] \to \widetilde{X}$ and $b : [0, \ell_2] \to \widetilde{X}$ be distinct geodesics in a CAT(0) space $\widetilde{X}$ such that $x := a(0) = b(0)$. Let $0 \le t_1 \le \ell_1$ and $0 \le t_2 \le \ell_2$. Let $y = a(t_1)$ and $z = b(t_2)$. If $\angle_x(a, b) \ge \frac{\pi}{3}$, then $|[y, z]| \ge \frac{1}{16}(t_1 + t_2)$.*

Proof. Recall $\overline{\angle}_x(y, z)$ is the CAT(0) comparison angle at x in the triangle $\triangle xyz$. By [9] Proposition II.3.2, $\overline{\angle}_x(y, z) \ge \angle_x(a, b)$. If $\triangle x'y'z' \subseteq \mathbb{E}^2$ is a Euclidean triangle such that $d(x', y') = t_1$, $d(x', z') = t_2$ and $\overline{\angle}_{x'}(y', z') \ge \frac{\pi}{3}$, it follows that $d(y, z) \ge d(y', z')$.

Claim: $d(y', z') \ge \frac{1}{16}(t_1 + t_2)$. By the law of cosines:

$$d(y', z')^2 \ge t_1^2 + t_2^2 - t_1 t_2$$

Without loss of generality, assume $t_2 \ge t_1$.

Case 1: $t_1 \le \frac{1}{2} t_2$:

$$d(y', z')^2 \ge t_1^2 + t_2^2 - \frac{1}{2} t_2^2 \ge \frac{1}{256} t_1^2 + \frac{1}{8}\left(t_1^2 + t_2^2\right) + \frac{1}{256} t_2^2 \ge \frac{1}{256} t_1^2 + \frac{1}{8} t_1 t_2 + \frac{1}{256} t_2^2 = \left(\frac{1}{16} t_1 + \frac{1}{16} t_2\right)^2$$

Case 2: $\frac{1}{2} t_2 < t_1 \le t_2$:

$$d(y', z')^2 \ge t_1^2 + t_2^2 - t_1 t_2 \ge \frac{1}{256} t_1^2 + \frac{255}{256} t_1^2 + \frac{1}{256} t_2^2 + \frac{255}{256} t_2^2 - t_1 t_2$$

apply the inequality $\frac{1}{2} t_2 < t_1 \le t_2$ to the second and fourth terms:

$$\ge \frac{1}{256} t_1^2 + \frac{255}{512} t_1 t_2 + \frac{255}{256} t_1 t_2 + \frac{1}{256} t_2^2 - t_1 t_2 \ge \frac{1}{256} t_1^2 + \frac{1}{8} t_1 t_2 + \frac{1}{256} t_2^2.$$

It then follows that in all cases $d(y', z')^2 \ge (\frac{1}{16} t_1 + \frac{1}{16} t_2)^2$ so $d(y', z') \ge \frac{1}{16}(t_1 + t_2)$. $\square$

The **space of directions** at a point in a CAT(0) space is a space where the angle is a metric.

Definition 4.1.7. *Let $\widetilde{X}$ be a CAT(0) space and let $p \in \widetilde{X}$. The* ***space of directions*** *of $\widetilde{X}$ at p is:*

$$\Sigma_p^{\widetilde{X}} = \{\gamma | \gamma \text{ is a non-trivial geodesic with endpoint at } p\} / \sim$$

where $\gamma \sim \gamma'$ if and only if $\angle_p(\gamma, \gamma') = 0$.

Observe $\Sigma_p^{\widetilde{X}}$ has a natural metric, $d(\gamma, \gamma') = \angle_p(\gamma, \gamma')$, the Alexandrov angle at p. It is easy to see the metric is symmetric. That it satisfies the triangle inequality follows from [9] I.1.17.

In the Euclidean case, if a is a geodesic path with one endpoint in a convex subspace $B \subseteq \mathbb{R}^n$, then a and the projection of a onto B realize the infimum of angles between a and a geodesic in B. Lemma 4.1.8 shows that something similar happens with an $(\mathcal{A}, \mathcal{B})$-LPPC when $B \in \mathcal{B}$ and a is a geodesic in $A \in \mathcal{A}$ with one endpoint in B.

Lemma 4.1.8. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, D)-relatively hyperbolic pair and let E be an $(\mathcal{A}, \mathcal{B})$-LPPC subspace.*

Let $A \in \mathcal{A}$ and $B \in \mathcal{B}$ be each contained in E and suppose $A \cap B$ is non-trivial. Let γ be a geodesic in E, and suppose that γ can be written as $\gamma = ab$ where a is a geodesic in A and b is a geodesic in B, and a, b meet at a point $x \in A \cap B$. Let π be the CAT(0) projection to B and let U be a neighborhood of x in A such that $\pi(a \cap U)$ is the image of an $A \cap B$-geodesic $\sigma : [0, t_\sigma) \to A \cap B$ with $\sigma(0) = x$ as in Proposition 4.1.4.

Then, $\angle_x^E(\sigma, a) = \angle_x^{\widetilde{X}}(\sigma, a) = \angle_x^{\widetilde{X}}(B, a)$,

Proof. Since $\widetilde{X}$ is CAT(0) and A, B are convex, A, B are CAT(0) and hence uniquely geodesic. By Proposition 4.1.4, σ is an $A \cap B$-geodesic. Since $A \cap B$ is an intersection of convex subspaces, σ is a geodesic in $\widetilde{X}$. Therefore, $\angle_x^E(\sigma, a) = \angle_x^{\widetilde{X}}(\sigma, a)$. Since B is convex, σ is also a B-geodesic.

Let λ be a non-degenerate geodesic in B issuing from x and assume a is parameterized so that $a(0) = x$ and $a(s)$ is distance s from x. By definition of the projection map:

$$d(\lambda(s), a(s)) \geq d(\pi a(s), a(s)).$$

For $0 < t, t' \leq |\lambda|$ there exists $t'' > 0$ such that:

$$\overline{\angle}_x^{\widetilde{X}}(\lambda(t), a(t')) \geq \overline{\angle}_x^{\widetilde{X}}(\lambda(t''), a(t'')) \geq \overline{\angle}_x^{\widetilde{X}}(\pi(a(t'')), a(t'')) \geq \angle_x^{\widetilde{X}}(\sigma, a(t'')) = \angle_x^{\widetilde{X}}(\sigma, a)$$

which follows from the fact that comparison angles are non-decreasing functions of t, t' (see [9] Proposition 3.1). Therefore:

$$\angle_x^{\widetilde{X}}(\sigma, a) \leq \angle_x^{\widetilde{X}}(\lambda, a)$$

as desired. □

Since $\pi(a \cap U)$ coincides with the image of the geodesic σ (but is not geodesically parameterized), for notational convenience, let $\angle_x^Y(a, \pi(a)) := \angle_x(a, \sigma)$ where $Y = E$ or $\widetilde{X}$.

Proposition 4.1.9 uses Lemma 4.1.8 to show that the geodesic a either departs B quickly or that a and b make a large angle at x.

Proposition 4.1.9. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, D)-relatively hyperbolic pair and let E be a $(\mathcal{A}, \mathcal{B})$-LPPC subspace, and let $\gamma = ab$ and $x \in \widetilde{X}$ as in Proposition 4.1.8.*

Then at least one of the following occurs

1. $\angle_x^{\widetilde{X}}(a, B) \geq \frac{\pi}{6}$ *or*

2. $\angle_x^{\widetilde{X}}(a,b)$, *the angle between* a *and* b *at* x *is more than* $\frac{2\pi}{3}$ *in* $\widetilde{X}$.

Proof. Suppose that $\angle_x^{\widetilde{X}}(a,B) < \frac{\pi}{6}$.

Let π be the CAT(0) projection map onto B. By the triangle inequality for angles and Lemma 4.1.8:

$$\pi = \angle_x^E(a,b) \le \angle_x^E(a,\pi(a)) + \angle_x^E(\pi(a),b)$$

$$= \angle_x^{\widetilde{X}}(a,B) + \angle_x^B(\pi(a),b) < \frac{\pi}{6} + \angle_x^B(\pi(a),b)$$

Therefore $\angle_x^{\widetilde{X}}(\pi(a),b) = \angle_x^B(\pi(a),b) > \frac{5\pi}{6}$. Since

$$|\angle_x^{\widetilde{X}}(a,b) - \angle_x^{\widetilde{X}}(\pi(a),b)| \le \angle_x^{\widetilde{X}}(\pi(a),a) = \angle_x^{\widetilde{X}}(a,B) < \frac{\pi}{6},$$

where the last inequality follows by Lemma 4.1.8. Hence $\angle_x^{\widetilde{X}}(a,b) > \frac{2\pi}{3}$. □

For the E-geodesic γ in Proposition 4.1.5, Construction 4.1.10 gives the construction of the improved path γ'. The b_i' are constructed first and then for $i \ge 2$ the a_i' are constructed to connect an endpoint of b_{i-1}' to an endpoint of b_i'. The endpoints of a_1' are an endpoint of γ and an endpoint of b_1'. Figure 4.2 and Figure 4.4 illustrate two possible ways the endpoints q_i' and p_{i+1}' of b_i' can be chosen. Figure 4.3 and Figure 4.5 show the paths b_i' which will be part of γ' and the auxiliary paths b_i'', c_i and d_i which are needed in the proof of Proposition 4.1.5.

For the following construction, adapt the convention that if a_i is a point, then a_i makes an angle of π with B_i and B_{i-1}.

Recall the following hypotheses from Proposition 4.1.5: $(\widetilde{X},\mathcal{B})$ is a (δ,D)-relatively hyperbolic pair, E is a $(\mathcal{A},\mathcal{B})$-LPPC subspace of $\widetilde{X}$, and $\gamma \subseteq E$ is a geodesic in E. The E-geodesic γ can be written as a piecewise geodesic $a_1b_1\dots a_nb_n$ in $\widetilde{X}$ where each a_i is a point or lies in $A_i \in \mathcal{A}$ and each $b_i \subseteq B_i \in \mathcal{B}$.

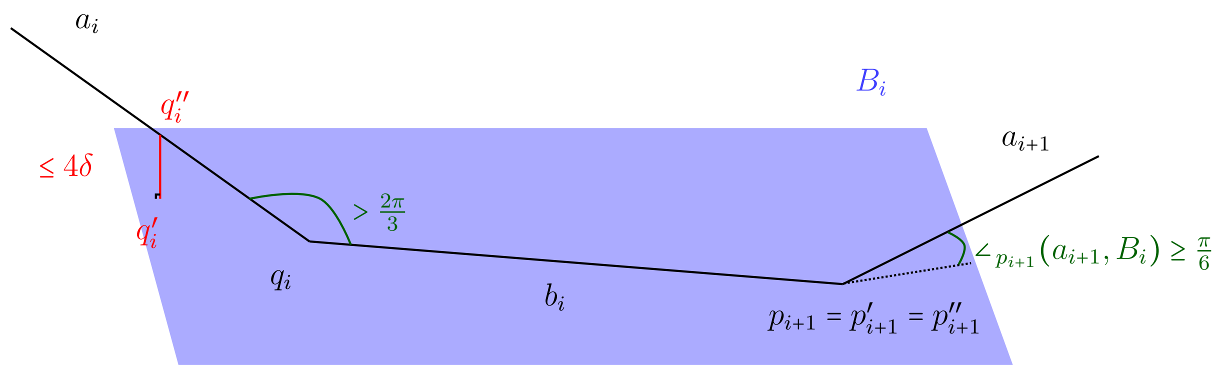

Figure 4.2: In the case where $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$ and $\angle_{p_{i+1}}(a_{i+1}, B_i) \geq \frac{\pi}{6}$, q_i', q_i'' are chosen as in the case shown in Figure 4.4, but here $p_{i+1} = p_{i+1}' = p_{i+1}''$.

Construction 4.1.10. *Label vertices so that* $a_i = [p_i, q_i]$ *and* $b_i = [q_i, p_{i+1}]$*. For each* $1 \leq i \leq n$*:*

1. *Choose an endpoint,* q_i'*, of* b_i'*:*
   (a) *If* $\angle_{q_i}(a_i, B_i) \geq \frac{\pi}{6}$*, set* $q_i'' := q_i' := q_i$*.*
   (b) *Otherwise, if* $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$ *and if* $a_i \notin \mathcal{N}_{4\delta}(B_i)$*, set* q_i'' *to be the point on* a_i *which is exactly* 4δ *from* B_i *and set* q_i' *to be the point in* B_i *closest to* q_i''*.*
   (c) *In all other cases, set* $q_i'' := p_i$ *and set* q_i' *to be the point in* B_i *closest to* q_i''*.*
2. *Choose the other endpoint* p_{i+1}' *of* b_i' *and set* p_1'*:*
   (a) $p_1'' = p_1 = p_1'$*.*
   (b) *If* $i = n$ *or* $\angle_{p_{i+1}}(a_{i+1}, B_i) \geq \frac{\pi}{6}$*, set* $p_{i+1}'' = p_{i+1}' = p_{i+1}$*.*
   (c) *If* $i < n$ *and* $a_{i+1} \notin \mathcal{N}_{4\delta}(B_i)$*, set* p_{i+1}'' *to be the point on* a_{i+1} *which is exactly* 4δ *from* B_i*. If* $i = n$*, set* $p_{i+1}'' = p_{i+1}$*. Let* p_{i+1}' *be the point in* B_i *closest to* p_{i+1}''*.*
   (d) *In all other cases: set* p_{i+1}'' *to be* q_{i+1} *and let* p_{i+1}' *be the point in* B_i *closest to* p_{i+1}''*.*

3. *For* $1 \le i \le n$*, set* $b_i' := [q_i', p_{i+1}']$ *and take* $a_i' := [p_i', q_i']$.

4. *Let* $c_i := [q_i'', q_i] \subseteq a_i$ *for all* i *and let* $d_i := [p_{i+1}'', p_{i+1}] \subseteq a_{i+1}$ *for all* $1 \le i < n$.

5. *Let* $\gamma' := a_1' b_1' \dots a_n' b_n'$.

Remark 4.1.11. *Some notes on the construction:*

1. *If* $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$*, then* $\angle_{q_i}(a_i, b_i) > \frac{2\pi}{3}$ *by Proposition 4.1.9. Similarly, if* $\angle_{q_i}(a_{i+1}, B_i) < \frac{\pi}{6}$*, then* $\angle_{p_{i+1}}(a_{i+1}, b_i) > \frac{2\pi}{3}$.

2. *The paths* $c_i \subseteq a_i$ *and* $d_i \subseteq a_{i+1}$ *are not always contained in* γ'*. Instead, they will be used in the proof of Proposition 4.1.5 to show that the length of* γ' *is bounded below by an affine linear function of the length of* γ.

3. *It is possible that* p_i'' *is closer to* q_i *than* q_i''*. In this case,* c_i *and* d_{i-1} *intersect in a non-trivial subsegment of* a_{i+1}*. However, their intersection must have short length. Indeed, for* $2 \le i \le n$*,* $|c_i \cap d_{i-1}| \le \operatorname{diam}(\mathcal{N}_{4\delta}(B_i) \cap \mathcal{N}_{4\delta}(B_{i-1}))$.

4. *To obtain a lower bound on the length of* γ' *in terms of the length of* γ*, the length of* b_i' *will be compared to the sum of the lengths of* b_i, c_i *and* d_i *and the length of* a_i *will be compared to the length of* $a_i'' := a_i \smallsetminus (c_i \cup d_{i-1})$ *except in the special case where* c_i *and* d_{i-1} *intersect non-trivially which will be addressed by Lemma 4.1.15.*

Lemma 4.1.12 compares the length of b_i' to the lengths of b_i, c_i and d_i in the case that exactly one of $\angle_{q_i}(a_i, B_i)$ or $\angle_{p_{i+1}}(B_i, a_{i+1})$ are less than $\frac{\pi}{6}$.

Lemma 4.1.12. *Suppose exactly one of the following holds:*

1. $q_i' = q_i$ *and* $\angle_{p_{i+1}}(B_i, a_{i+1}) < \frac{\pi}{6}$*,*

2. *or* $p_{i+1}' = p_{i+1}$ *and* $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$*,*

then $|b_i'| \ge \frac{1}{16}(|b_i| + |c_i| + |d_i|) - 4\delta$.

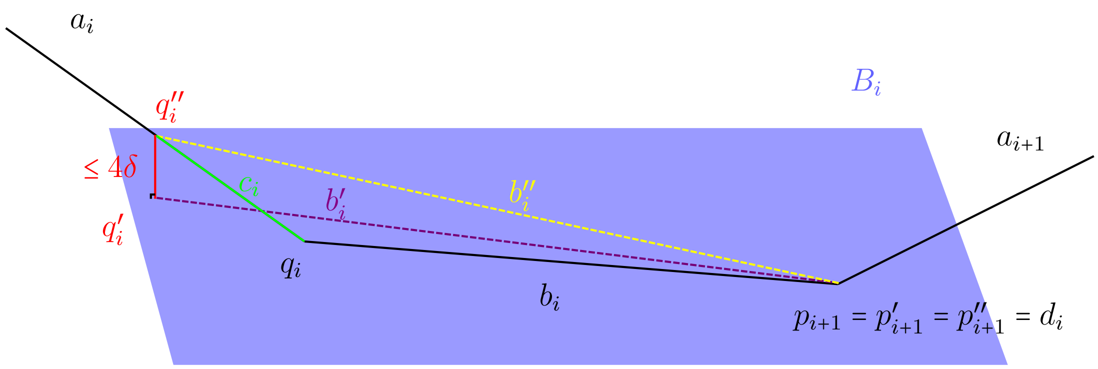

Figure 4.3: In the case where $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$ and $\angle_{p_{i+1}}(a_{i+1}, B_i) \geq \frac{\pi}{6}$, the segments b_i'' and b_i' are constructed as shown.

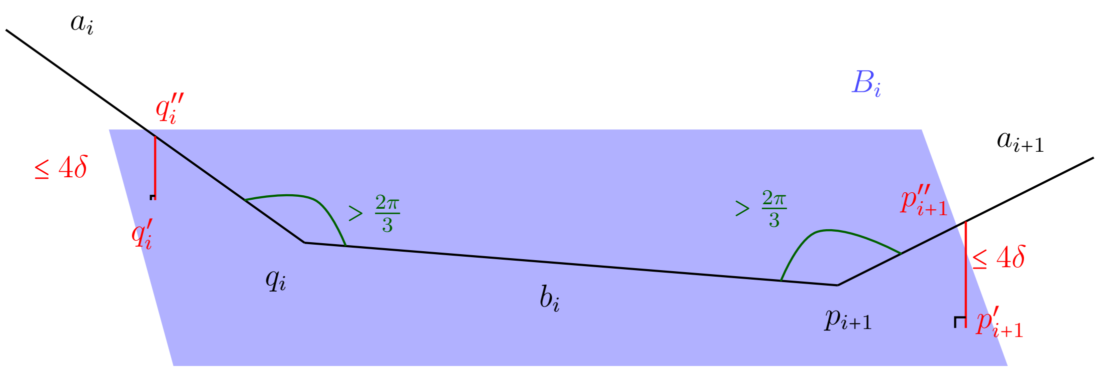

Figure 4.4: In this case, $\angle_{q_i}(a_i, B_i)$, $\angle_{p_{i+1}}(a_{i+1}, B_i) < \frac{\pi}{6}$. The point q_i'' is the point on a_i that is 4δ away from B_i or if no such point exists $q_i'' = p_i$. Then q_i' is the CAT(0) projection of q_i'' onto B_i. The points p_{i+1}', p_{i+1}'' are chosen similarly.

Proof. Assume that $p_{i+1} = p_{i+1}' = p_{i+1}''$ and $\angle_{q_i}(a_i, B_i) < \frac{\pi}{6}$. By Proposition 4.1.9, $\angle_{q_i}(a_i, b_i) > \frac{2\pi}{3}$. In this case, d_i is just a point. Refer to Figure 4.3 for an illustration. By Lemma 4.1.6, $|b_i''| \geq \frac{1}{16}(|b_i| + |d_i|) = \frac{1}{16}(|b_i| + |c_i| + |d_i|)$. By the triangle inequality, $||b_i'| - |b_i''|| \leq 4\delta$, so $|b_i'| \geq \frac{1}{16}(|b_i| + |c_i| + |d_i|) - 4\delta$. The other case is similar and is hence left to the reader. □

Lemma 4.1.13 is similar to Lemma 4.1.12, but deals with the case where both

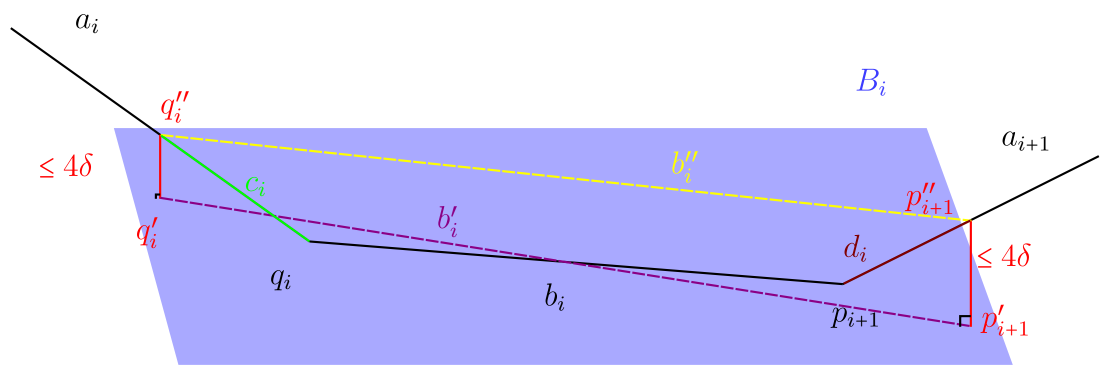

Figure 4.5: In the case where $\angle_{q_i}(a_i, B_i), \angle_{p_{i+1}}(a_{i+1}, B_i) < \frac{\pi}{6}$, set $b_i' = [q_i', p_{i+1}']$ and $b_i'' = [q_i'', p_{i+1}'']$.

$\angle_{q_i}(a_i, b_i) > \frac{2\pi}{3}$ and $\angle_{p_{i+1}}(a_{i+1}, b_i) > \frac{2\pi}{3}$.

Lemma 4.1.13. *Let γ, γ' be as in Proposition 4.1.5. Suppose both $\angle_{q_i}(a_i, b_i)$ and $\angle_{p_{i+1}}(a_{i+1}, b_i) \geq \frac{2\pi}{3}$. Let c_i and d_i be as in Construction 4.1.10. Then*

$$|b_i'| \geq \frac{1}{256}(|c_i| + |d_i| + |b_i|) - 8\delta$$

Proof. Refer to Figure 4.5 for an illustration of this case. Observe that by the triangle inequality (applied twice), $||b_i'| - |b_i''|| \leq 8\delta$, so it suffices to prove that $|b_i''| \geq \frac{1}{256}(|c_i| + |d_i| + |b_i|)$. Consider the quadrilateral bounded by vertices $q_i'', p_{i+1}'', p_{i+1}, q_i$ and draw an additional diagonal ρ (see Figure 4.6). In a CAT(0) space, the angle sum of a triangle is at most π, so $\angle_{q_i}(\rho, b_i) \leq \frac{\pi}{3}$. By the triangle inequality for angles $\angle_{q_i}(c_i, \rho) + \angle_{q_i}(b_i, \rho) > \frac{2\pi}{3}$, so $\angle_{q_i}(c_i, \rho) \geq \frac{\pi}{3}$. Therefore, by Lemma 4.1.6, $|\rho| \geq \frac{1}{16}(|b_i| + |d_i|)$ and $|b_i''| \geq \frac{1}{16}(|\rho| + |d_i|)$. Putting these two inequalities together:

$$|b_i''| \geq \frac{1}{256}(|b_i| + |d_i| + |c_i|)$$

as desired. □

When neither Lemma 4.1.12 nor Lemma 4.1.13 applies, then both $\angle_{q_i}(a_i, B_i)$

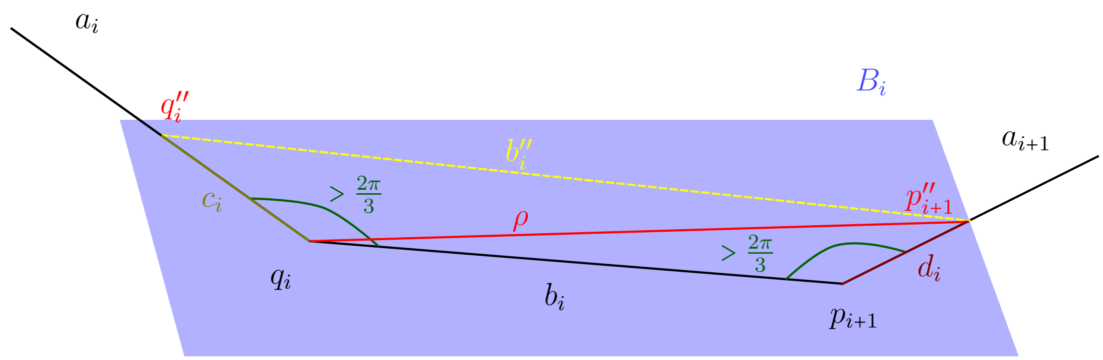

Figure 4.6: The quadrilateral for the proof of Lemma 4.1.13.

and $\angle_{p_{i+1}}(a_{i+1}, B_i)$ are both at least $\frac{\pi}{6}$ by Proposition 4.1.9. In these cases, $b_i = b_i'$, and both c_i and d_i are points.

The next task is to show that the a_i' leave B_i and B_{i-1} at a rate that is bounded below.

Proposition 4.1.14. *Let d_0 be the point p_1. Suppose $c_i \cap d_{i-1}$ is a point or empty. Let $\alpha_i := [p_i'', q_i'']$, $L > 0$. Recall:*

$$D \geq \max\{\mathrm{diam}(\mathcal{N}_{3\delta}(B) \cap \mathcal{N}_{3\delta}(B')) \mid B, B' \in \mathcal{B}\}.$$

Then:

1. $||a_i'| - |\alpha_i|| \leq 8\delta$,
2. $\mathrm{diam}(a_i' \cap \mathcal{N}_{3\delta}(B_i)) \leq 2D + 29\delta$.
3. *if* $i \geq 2$, $\mathrm{diam}(a_i' \cap \mathcal{N}_{3\delta}(B_{i-1})) \leq 2D + 29\delta$.

For the following proof, it may be useful to recall corner segments and fat parts of sides of relatively thin triangles from Definition 2.3.18.

Proof. That $||a_i| - |\alpha_i|| \leq 8\delta$ follows immediately by the triangle inequality.

Case: $p'_i \neq p_i$ **and** $q'_i \neq q_i$**.** Consider the quadrilateral with vertices p''_i, q''_i, q'_i, p'_i and let ℓ be the diagonal p''_i, q'_i. For convenience, let $\triangle_1 := \triangle p''_i p'_i, q'_i$ (the lower triangle in Figure 4.7) and let $\triangle_2 := \triangle p''_i q'_i q''_i$ (the upper triangle in Figure 4.7). Suppose that $\triangle_1, \triangle_2$ are thin relative to $F_1, F_2 \in \mathcal{B}$ respectively. Decompose a'_i and ℓ into subsegments using the fact that $\widetilde{X}$ has the δ-relatively thin triangle property (see Figure 4.7) as follows:

1. a'_i decomposes as a concatenation of three geodesics:
   (a) the corner segment of $\triangle_1$ at p'_i of length at most 4δ,
   (b) $(a'_i)_1 \subseteq \mathcal{N}_\delta(F_1)$, the fat part of a'_i in $\triangle_1$,
   (c) a corner segment of $\triangle_1$, $(a'_i)_2 \subseteq \mathcal{N}_\delta(\ell)$, at q'_i.

   Note that $(a'_i)_1$ may be empty.
2. ℓ decomposes as a concatenation of three geodesics as follows:
   (a) ℓ_1 is the union of a corner segment at p''_i in $\triangle_2$ and the intersection of a fat part of ℓ in $\triangle_2$ that is not in a corner segment of $\triangle_1$ at q'_i,
   (b) ℓ_2 is the possibly empty intersection of a corner segment of $\triangle_1$ at q'_i with the fat part of ℓ in $\triangle_2$, and
   (c) a corner segment in $\triangle_2$ of length at most 4δ.

By construction, $\ell_1 \subseteq \mathcal{N}_\delta(\alpha_i) \cup \mathcal{N}_\delta(F_2)$, the fat part of ℓ in $\triangle_1$ is contained in ℓ_1 and $\ell_2 \subseteq \mathcal{N}_\delta((a'_i)_2) \cap \mathcal{N}_\delta(F_2)$.

The segment α_i intersects $\mathcal{N}_{4\delta}(B_i)$ only at the common endpoint with ℓ. Either $F_2 = B_i$ (if $\triangle_2$ is δ-thin, assume $F_2 \neq B_i$) or $|\ell_1 \cap \mathcal{N}_\delta(B_i)| \leq D$. Indeed, the corner segment of $\triangle_2$ at p_i'' contained in ℓ_1 lies in $\mathcal{N}_\delta(\alpha_i)$ and cannot intersect $\mathcal{N}_\delta(B_i)$. The rest of ℓ_1 lies in the fat part of ℓ in $\triangle_2$, and $\operatorname{diam}(\mathcal{N}_\delta(B_i) \cap \mathcal{N}_\delta(F_2)) \leq D$. When $F_2 \neq B_i$, all but 4δ of $(a_i')_2$ δ-fellow travels ℓ_2 because they are both contained in corner segments of $\triangle_1$ at q_i'. Since ℓ_2 lies in the fat part of ℓ in $\triangle_2$, $\ell_2 \subseteq \mathcal{N}_\delta(F_2)$. Hence all but 4δ of $(a_i')_2$ lies in $\mathcal{N}_{2\delta}(F_2)$ where $F_2 \neq B_i$, so at most $D + 4\delta$ of $(a_i')_2$ lies in $\mathcal{N}_{3\delta}(B_i)$.

If $F_2 = B_i$, as before, the subsegment of ℓ_1 that is a corner segment in $\triangle_2$ at p_i'' cannot intersect $\mathcal{N}_\delta(B_i)$. The rest of ℓ_1 lies in the fat part of ℓ in $\triangle_2$, so $|\ell_1 \cap \mathcal{N}_\delta(B_i)| \leq 7\delta$ by Lemma 2.3.22 because α_i intersects $\mathcal{N}_{4\delta}(B_i)$ at a point. By a similar argument, $|\ell_2| \leq 7\delta$ because ℓ_2 lies in the fat part of ℓ in $\triangle_2$. Then, by δ-relative thinness, $|(a_i')_2| \leq D + 11\delta$.

In both cases, at most $D + 11\delta$ of $(a_i')_2$ lies in $\mathcal{N}_{3\delta}(B_i)$.

Subcase: $F_1 \neq B_i$. Here $|(a_i')_1 \cap \mathcal{N}_\delta(B_i)| \leq D$ because $F_1 \neq B_i$. Thus $|a_i' \cap \mathcal{N}_{3\delta}(B_i)| \leq 2D + 15\delta$.

Subcase: $F_1 = B_i$. At most $\max(D, 7\delta) \leq D + 7\delta$ of ℓ_1 lies in $\mathcal{N}_\delta(B_i)$ by the preceding. Since ℓ_1 contains the fat part of ℓ in $\triangle_1$, by Lemma 2.3.22, the length $(a_i')_1$, which is the fat part of a_i' in $\triangle_1$, is at most $D + 14\delta$. At most $D + 11\delta$ of $(a_i')_2$ lies in $\mathcal{N}_{3\delta}(B_i)$. In total, $|a_i' \cap \mathcal{N}_{3\delta}(B_i)| \leq 2D + 29\delta$.

Bounding $|a_i' \cap \mathcal{N}_{3\delta}(B_{i-1})|$ from above is similar and left to the reader.

Case: $q_i \neq q_i'$, $p_i = p_i'$. Let $\triangle := \triangle p_i' q_i'' q_i'$ where $\triangle$ is δ-thin relative to $F_\triangle \in \mathcal{B}$. Split a_i' into three segments (see Figure 4.8):

1. $(a_i')_1$ is a corner segment of $\triangle$ at p_i' contained in $\mathcal{N}_\delta(\alpha_i)$,
2. $(a_i')_2 \subseteq \mathcal{N}_\delta(F_\triangle)$ is the fat part of a_i' in $\triangle$,

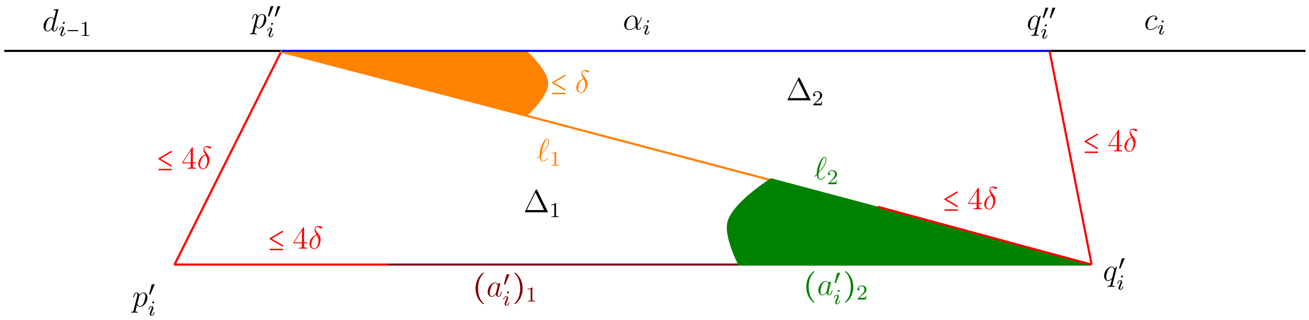

Figure 4.7: Accompanies the proof of Proposition 4.1.14 in the case $p_i \neq p'_i$ and $q_i \neq q'_i$.

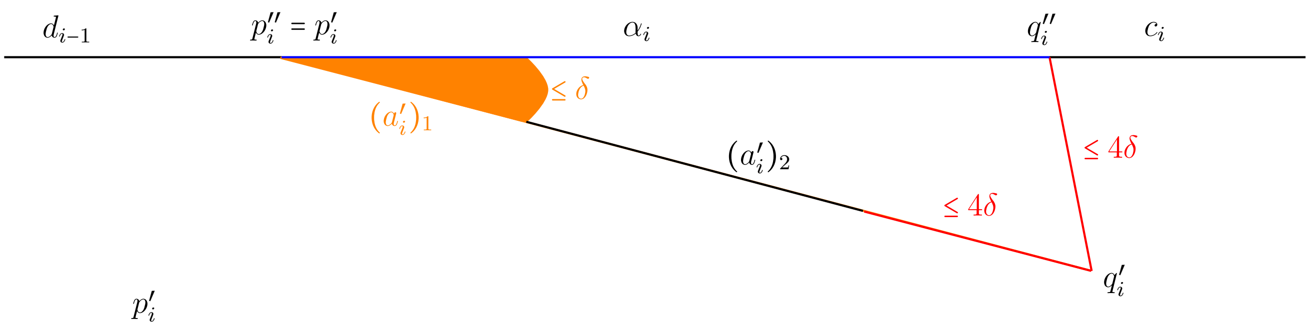

Figure 4.8: Accompanies the proof of Proposition 4.1.14 in the case $p_i = p'_i$ and $q_i \neq q'_i$. The segment a_i breaks down into a segment $(a'_i)_1 \subseteq \mathcal{N}_\delta(\alpha_i)$, a segment $(a'_i)_2 \subseteq F_\triangle \in \mathcal{B}$, and a segment of length at most 4δ.

3. a segment of length at most 4δ.

Immediately, $(a'_1)_1$ does not intersect $\mathcal{N}_{3\delta}(B_i)$ because α_i intersects $\mathcal{N}_{4\delta}(B_i)$ at a point.

If $F_\triangle = B_i$, then by Lemma 2.3.22, $|(a'_i)_2| \leq 7\delta$ because α_i intersects $\mathcal{N}_{3\delta}(B_i)$ at a point. Thus if $F_\triangle = B_i$, then $|a'_i \cap \mathcal{N}_{3\delta}(B_i)| \leq 11\delta$.

If $F_\triangle \neq F_i$, then $|(a'_i)_2 \cap \mathcal{N}_{3\delta}(B_i)| \leq D$, so $|a_i \cap \mathcal{N}_{3\delta}(B_i)| \leq D + 4\delta$.

For B_{i-1} the argument is similar but instead, α_i makes an angle of $\frac{\pi}{6}$ with B_{i-1}, so by trigonometry, $|\alpha_i \cap \mathcal{N}_{3\delta}(B_{i-1})| \leq 6\delta$. Hence at most 6δ of $(a_i')_1$ may lie in $\mathcal{N}_{3\delta}(B_{i-1})$. If $F_\triangle = B_{i-1}$, by Lemma 2.3.22, $|(a_i')_2| \leq 10\delta$, so $|a_i' \cap \mathcal{N}_{3\delta}(B_{i-1})| \leq 20\delta$. On the other hand, if $F_\triangle \neq B_{i-1}$, then $|(a_i')_2 \cap \mathcal{N}_{3\delta}(B_{i-1})| \leq D$, so $|a_i' \cap \mathcal{N}_{3\delta}(B_{i-1})| \leq D + 10\delta$.

Case: $p_i' = p_i$ **and** $q_i' = q_i$. In this case, it follows that $\alpha_i = a_i$ makes an angle of $\frac{\pi}{6}$ with B_i and B_{i-1}, and the result follows by trigonometry. □

There is a chance that the segments c_i and d_{i-1} overlap. Since the lengths of b_i' and b_{i-1} are bounded below in terms of the lengths of c_i and d_{i-1} respectively, some caution is needed. However, Lemma 4.1.15 shows that if c_i and d_{i-1} overlap, then a_i had to be relatively short to begin with.

Lemma 4.1.15. *Suppose $c_i \cap d_{i-1}$ is a non-trivial subsegment of a_i. Then $|c_i \cap d_{i-1}| \leq D$, $|a_i'| \leq D + 8\delta$ and $|d_{i-1}| + |c_i| + |a_i'| \geq |a_i|$.*

Proof. By construction, $d_{i-1} \subseteq \mathcal{N}_{4\delta}(B_{i-1})$ and $c_i \subseteq \mathcal{N}_{4\delta}(B_i)$. By the minimality of n and convexity, $B_i \neq B_{i-1}$. The intersection $\mathcal{N}_{4\delta}(B_i) \cap \mathcal{N}_{4\delta}(B_{i-1})$ has diameter bounded by D.

The second statement follows immediately from the fact that if c_i, d_{i-1} overlap, then their union covers a_i. □

Proposition 4.1.5 now follows:

Proof of Proposition 4.1.5: First observe that by construction, the endpoints of each b_i' are contained in B_i. By convexity, $b_i' \subseteq B_i$.

If $B_i = B_j$ for some $i \neq j$, then γ was not an E-geodesic because each B_i is convex and n is minimal.

By Proposition 4.1.14, a_i' spends at most $2D + 29\delta$ in $\mathcal{N}_{3\delta}(B_i)$ or has length at most $D + 8\delta$ by Proposition 4.1.15. Since $b_i' \subseteq B_i$, it is impossible for a_i' and b_i' to 2δ-fellow travel for more than $2D + 29\delta$. Similarly, b_i', a_{i+1}' do not 2δ-fellow travel for more than $2D + 29\delta$.

For Conclusion (3), either $b_i = b_i'$ in the case $q_i = q_i'$ and $p_i = p_{i+1}'$ or one of Lemmas 4.1.12 and 4.1.13 applies and yields the desired conclusion.

In the case $b_i = b_i'$, c_i and d_i are just points, so $|b_i| = |b_i'| + |c_i| + |d_i|$.

Otherwise, the conclusions of Lemmas 4.1.12 and 4.1.13 imply that

$$|b_i'| \geq \frac{1}{256}(|b_i| + |c_i| + |d_i|) - 8\delta$$

If $c_i \cap d_{i-1}$ is empty or a point, then by Proposition 4.1.14,

$$|a_i| = |\alpha_i| + |c_i| + |d_{i-1}| \leq |a_i'| + 8\delta + |c_i| + |d_{i-1}|$$

If $c_i \cap d_{i-1}$ is nontrivial subsegment of a_i, then by Proposition 4.1.15:

$$|a_i| \leq |a_i'| + |c_i| + |d_{i-1}|.$$

In any case:

$$|a_i'| \geq |a_i| - |c_i| - |d_{i-1}| - 8\delta$$

Therefore:

$$|\gamma'| = \sum_{i=1}^{n}(|a_i'| + |b_i'|) \geq \sum_{i=1}^{n}\left(\frac{1}{256}(|b_i| + |c_i| + |d_i|) - 8\delta + \frac{1}{256}|a_i'|\right)$$
$$\geq \sum_{i=1}^{n}\left(\frac{1}{256}(|b_i| + |c_i| + |d_i|) - \frac{1}{256}(|a_i| - |c_i| - |d_i|) - 8\delta\right)$$

$$\geq \left(\sum_1^n |a_i| + |b_i|\right) - 16n\delta = \frac{1}{256}|\gamma| - 16n\delta,$$

concluding the proof. □

4.2 The Combination Lemma

The goal of this subchapter is to prove a combination lemma for showing that piecewise geodesics satisfying the conclusions of Proposition 4.1.5 are uniformly quasigeodesic. Lemma 4.2.2 does most of the work, but special cases involving the initial and terminal geodesic pieces will have to be added later in Proposition 4.2.11 which will be used in the proof that the edge spaces of the hierarchy in Chapter 6 are π_1-injective, the proof that the hierarchy is quasiconvex and the proof that the hierarchy is fully $\mathcal{P}$-elliptic.

The following terminology will be used extensively in this chapter:

Definition 4.2.1. *Let $\gamma : [\alpha, \beta] \to \widetilde{X}$ be a length parameterized path with an endpoint $x \in \gamma$ and let $R \leq |\alpha - \beta|$. An R-**tail of** γ **at** x is the subpath $\gamma([\alpha, \alpha + R])$ if $x = \gamma(\alpha)$ and is the subpath $\gamma([\beta - R, \beta])$ if $x = \gamma(\beta)$.*

Lemma 4.2.2. *Let $(\widetilde{X}, \mathcal{B})$ be a (δ, M)-relatively hyperbolic pair with $M \geq 2\delta$.*

Let $\gamma := b_1 a_2 b_2 a_3 b_3 \dots a_n b_n$ be a piecewise geodesic path in $\widetilde{X}$ such that

1. *For each i, there exists $F_i \in \mathcal{B}$ such that each $b_i \subset F_i$ and $F_i \neq F_j$ for $i \neq j$,*
2. *for $2 \leq i \leq n$, diam$(a_i \cap \mathcal{N}_{3\delta}(F_i)) \leq M$ and diam$(a_i \cap \mathcal{N}_{3\delta}(F_{i-1})) \leq M$,*
3. *$|b_1| \geq 5M + 6\delta$ and for $1 < i < n$, $|b_i| \geq 13M + 15\delta$.*

Let γ_i be the geodesic running from the initial point of γ on b_1 to the terminal point of b_i.

Then for all $1 \le i \le n$, $|\gamma_i| \ge |\gamma_{i-1}| + |a_i| + |b_i| - 20M - 18\delta$, and for $1 \le i \le n$, a $5M + 6\delta$ tail of γ_i (at the endpoint on b_n) lies in $\mathcal{N}_{2\delta}(F_i)$.

Before proving Lemma 4.2.2, consider the case where $c = c_1 c_2 \dots c_n$ is a piecewise geodesic in a δ-hyperbolic metric space $\widetilde{Y}$. If each of the c_i are sufficiently long relative to the length that c_i and c_{i+1} 2δ-fellow travel from their concatenation point, then c and the geodesic connecting the endpoints of c remains within a bounded Hausdorff distance of each other. For example, see [3, Lemma 4.9] and [9, III.H.1.13].

The main applications of Lemma 4.2.2 will use the lower bound on the length of γ_n while the conclusion that γ_i has a $5M + 6\delta$-tail of γ_i in $\mathcal{N}_{2\delta}(F_i)$ will be used as part of the inductive proof of Lemma 4.2.2.

Proof. First establish some notation. See Figure 4.9 for a visual example of the notation:

1. Let ω_i be the geodesic from the initial point of γ, to the terminal point of a_i.
2. Let Δ_i^1 be the triangle formed by $\gamma_{i-1}, a_i, \omega_i$ and let Δ_i^2 be the triangle formed by γ_i, b_i, ω_i.
3. In all figures, shaded regions between two geodesics will denote δ-fellow traveling.
4. Fix F_i^j so that each Δ_j^i is δ-thin relative to F_i^j. Recall that if Δ_i^j is δ-thin, F_i^j can be any element of $\mathcal{B}$.

5. If ρ is a side of a triangle Δ_i^j, the superscript ρ^j denotes the fat part of ρ in Δ_i^j.
6. Let $p_1 = q_1$ be the initial point of γ, let $a_i = [p_i, q_i]$ and let $b_i = [q_i, p_{i+1}]$.
7. The **fat girth of** $\triangle_i^1$ is $\min\{|a_i^1|, |\gamma_{i-1}^1|\}$ and the **fat girth of** $\triangle_i^2$ is $\min\{|\omega_i^2|, |b_i^2|\}$.
8. The **corner girth of** $\triangle_i^1$ is the corner length of $\triangle_i^1$ at p_i and the **corner girth of** $\triangle_i^2$ is the corner length of $\triangle_i^2$.

See Figure 4.9 for an example of some of this notation. The definition of the fat girth for each triangle is motivated by Lemma 2.3.22 which says that a bound on the fat girth will give some control over the relative lengths of the fat parts of the sides of the triangle.

The following lemma outlines the strategy of the proof of Lemma 4.2.2:

Lemma 4.2.3. *Let e_1 be an upper bound on the corner girth of $\triangle_i^1$, let e_2 be an upper bound on the fat girth of $\triangle_i^1$. Then,*

$$|\omega_i| \geq |\gamma_{i-1}| + |a_i| - (2e_1 + 3\delta) - 2e_2.$$

Let e_4 be an upper bound on the corner girth of $\triangle_i^2$, let e_5 be an upper bound on the fat girth of $\triangle_i^2$. Then,

$$|\gamma_i| \geq |\omega_i| + |b_i| - (2e_4 + 3\delta) - 2e_5.$$

Combining these two inequalities:

$$|\gamma_i| \geq |\gamma_{i-1}| + |a_i| + |b_i| - (2e_1 + 2e_2 + 2e_4 + 2e_5 + 6\delta).$$

The notation e_3 will be used later for an upper bound on the corner length of $\triangle_i^1$ at q_i that will be used to find acceptable values for e_4.

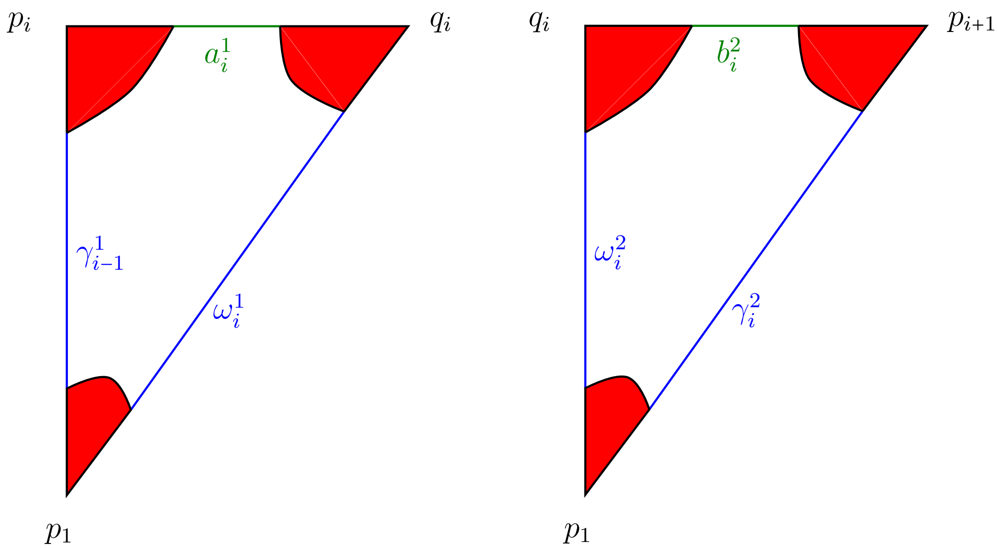

Figure 4.9: Triangles $\triangle_i^1$ and $\triangle_i^2$ with fat parts of sides marked $a_i^1, b_i^2, \gamma_i^1, \omega_i^1, \omega_i^2, \gamma_i^2$. The shaded regions denote δ-fellow travelling between subsegments of sides, in this case, between corner segments.

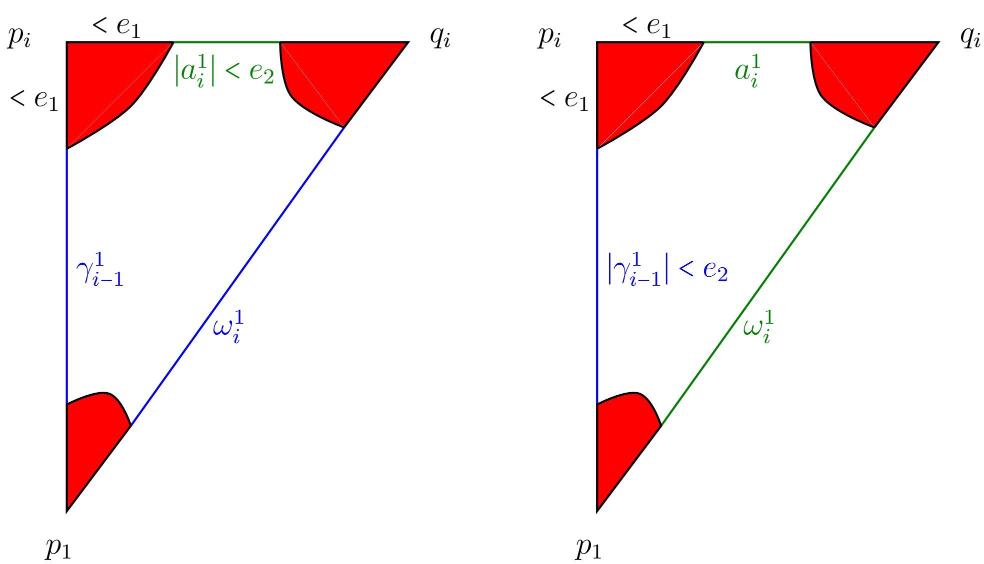

Figure 4.10: Let e_1 be an upper bound for the corner girth of $\triangle_i^1$ and let e_2 be an upper bound for the fat girth of $\triangle_i^1$. The left panel shows the case where e_2 bounds $|a_i^1|$ and the right panel shows the case where e_2 bounds $|\gamma_{i-1}^1|$. The notation e_3 will be reserved for an upper bound on the corner lengths of $\triangle_i^1$ at q_i.

Proof of Lemma 4.2.3. For the first inequality, refer to Figure 4.10. The corner segments of $\triangle_i^1$ at p_1 contained in γ_{i-1} and ω_i have the same length. Similarly, the corner segments of $\triangle_i^1$ at q_i contained in ω_i and a_i have the same length. If e_2 is an upper bound on the fat girth of $\triangle_i^1$, then either $|\omega_i^1|$ and $|\gamma_{i-1}^1|$ differ by at most $e_2 + 3\delta$ (see left panel of Figure 4.10) or $|\omega_i^1|$ and $|a_i^1|$ differ by at most $e_2 + 3\delta$ (see right panel of Figure 4.10) by Lemma 2.3.22. Therefore, $|\omega_i| \geq |\gamma_{i-1}| + |a_i| - 2e_1 - 2e_2 - 3\delta$ where the deficit of $2e_1 + 2e_2 + 3\delta$ comes from losing the length of both corner segments of $\triangle_i^1$ at p_i with combined length at most $2e_1$ (twice the corner girth of $\triangle_i^1$), the length of either γ_i^1 or a_i^1 depending on which side has length is bounded above by e_2, an upper bound on the fat girth, and the difference between the lengths of the fat parts of the other two sides which is bounded by $e_2 + 3\delta$.

The second inequality follows by relabeling Figure 4.10 to make Figure 4.11 and then using the exact same argument to show that $|\gamma_i| \geq |\omega_i| + |b_i| - 2e_4 - (2e_5 + 3\delta)$. Combining the first two inequalities:

$$|\gamma_i| \geq |\gamma_{i-1}| + |a_i| + |b_i| - 2e_1 - 2e_2 - 2e_4 - 2e_5 - 6\delta \tag{4.1}$$

immediately gives the final inequality. □

Lemma 4.2.3 shows that to prove Lemma 4.2.2, it suffices to find upper bounds for the corner girths and fat girths of $\triangle_i^1$ and $\triangle_i^2$ and to show that γ_i has a sufficiently long tail in $\mathcal{N}_{2\delta}(F_i)$.

The proof proceeds by induction on i. If $i = 1$, $b_1 = \gamma_1$ is geodesic and the argument is trivial, so suppose $i \geq 2$.

Lemma 4.2.4. *The corner girth of $\triangle_i^1$ is bounded above by M.*

Proof. An M-tail of γ_{i-1} at p_i lies in $\mathcal{N}_{2\delta}(F_{i-1})$. If the corner girth of $\triangle_i^1$ exceeds M, then a_i δ-fellow travels an M-tail of γ_{i-1} at p_i contained in $\mathcal{N}_{2\delta}(F_{i-1})$ for a length at least M which is impossible because a_i spends at most M in $\mathcal{N}_{3\delta}(F_{i-1})$. □

The following is used repeatedly:

Lemma 4.2.5. *Suppose that $\rho \subseteq \gamma_{i-1}$ has an endpoint a with $d(a, p_i) < 4M + 6\delta$. If $F \in \mathcal{B}$ with $F \neq F_{i-1}$ and either $\rho \subseteq \mathcal{N}_{3\delta}(F_i)$ or δ-fellow travels a segment in $\mathcal{N}_{3\delta}(F)$, then $|\rho| \leq M$.*

Proof. Since $F \neq F_{i-1}$, $\operatorname{diam}(\mathcal{N}_{3\delta}(F) \cap \mathcal{N}_{3\delta}(F_{i-1})) \leq M$. Suppose toward a contradiction that $|\rho| > M$. If $\rho \subseteq \mathcal{N}_{3\delta}(F_i)$, then ρ has an initial segment ρ' in the $5M + 6\delta$-tail of γ_{i-1} at p_i so that $\rho' \subseteq \mathcal{N}_{2\delta}(F_{i-1})$ and $|\rho'| > M$. If $\rho \subseteq \mathcal{N}_{2\delta}(F_{i-1})$, then $\rho' \subseteq \mathcal{N}_{3\delta}(F) \cap \mathcal{N}_{3\delta}(F_{i-1})$. If ρ δ-fellow travels a segment $\rho'' \subseteq \mathcal{N}_{3\delta}(F)$, then $|\rho''| = |\rho'| > M$ and $\rho'' \subseteq \mathcal{N}_{3\delta}(F) \cap \mathcal{N}_{3\delta}(F_{i-1})$, a contradiction. □

Lemma 4.2.6. *The fat girth of $\triangle_i^1$ is bounded above by M.*

Proof.

1. **Case:** $F_i^1 = F_{i-1}$. By hypothesis, a_i spends at most M in $\mathcal{N}_{3\delta}(F_{i-1})$, so $|a_i^1| \leq M$ because $a_i^1 \subseteq F_{i-1}$.
2. **Case:** $F_i^1 \neq F_{i-1}$. Since the corner girth of $\triangle_i^1$ is at most M, the corner length of $\triangle_i^1$ at p_i is at most M. Then by Lemma 4.2.5, $|\gamma_{i-1}^1| \leq M$ so the fat girth of $\triangle_i^1$ is bounded by M. □

The calculations for corner girth and fat girth of $\triangle_i^2$ depend on both $\triangle_i^1$ and $\triangle_i^2$.

Lemma 4.2.7. *The corner girth of $\triangle_i^2$ is at most $4M + 3\delta$.*

Proof. Let e_3 be the corner length of $\triangle^1_{i-1}$ at q_i.

Claim: if $e_3 > M$, then the corner girth of $\triangle^2_i$ is at most M. If $e_3 > M$ and the corner girth of $\triangle^2_i$ exceeds M, then the corner length of $\triangle^2_i$ at q_i exceeds M. If the corner girth of $\triangle^2_i$ is more than M as well, then a_i and b_i 2δ-fellow travel at q_i for more than M, so a_i spends more than M in $\mathcal{N}_{3\delta}(F_i)$ which is impossible, proving the claim.

Henceforth assume $e_3 \leq M$.

There are several cases:

1. **Case: ω^1_i is contained in a corner segment of $\triangle^2_i$ at q_i.**

   Subcase: $F^1_i = F_{i-1}$. Refer to Figure 4.12. First, the fat part of a_i in $\triangle^1_i$ is contained in $\mathcal{N}_\delta(F_{i-1})$ by definition. The combined length of a_i is at most $2M$ because the corner segment of $\triangle^1_i$ at p_i and the fat part of a_i in $\triangle^1_i$ lie in $\mathcal{N}_{3\delta}(F_{i-1})$ and $e_3 \leq M$.

   The corner segment of $\triangle^2_i$ at q_i contained in ω_i can be split into three subsegments: a corner segment of $\triangle^1_i$ of length at most M since $e_3 \leq M$, the intersection with ω^1_i which by the assumption in this case completely contains ω^1_i and its intersection with a corner segment of $\triangle^1_i$ at p_1.

   The intersection with ω^1_i by definition lies in $\mathcal{N}_\delta(F_{i-1})$, but also lies in $\mathcal{N}_{2\delta}(F_i)$ because it lies in a δ neighborhood of a corner segment of $\triangle^2_i$ at q_i contained in b_i, so this segment (and ω^1_i) has length at most M because $F_i \neq F_{i-1}$.

   Since ω^1_i lies in a corner segment of $\triangle^1_i$ at p_1 and a corner segment of $\triangle^2_i$ at q_i, let $c_{\gamma_{i-1}} \subseteq \gamma_{i-1}$ and $c_{b_i} \subseteq b_i$ be length $|c_{\omega_i}|$ δ-fellow traveling segments contained in corner segments.

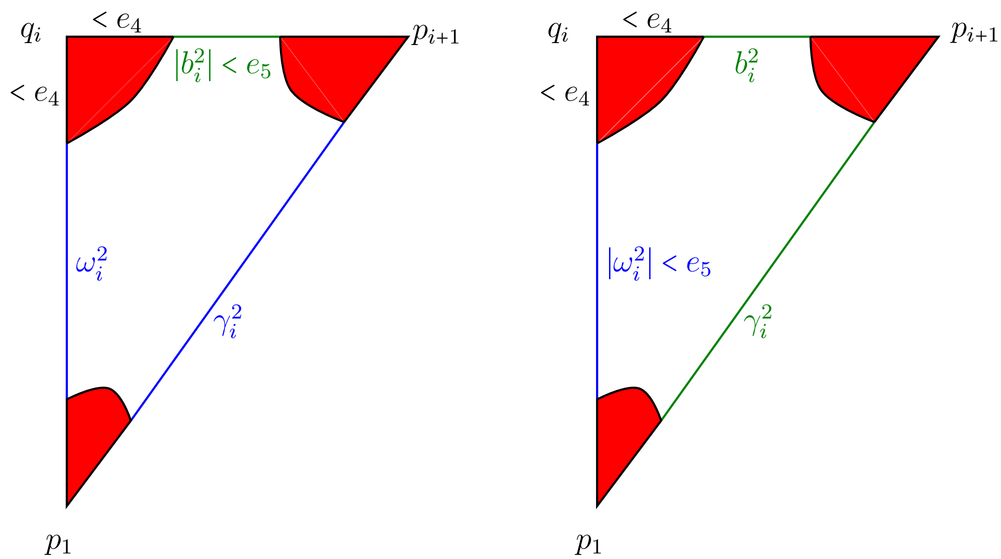

Figure 4.11: Let e_4 be an upper bound for the corner girth of $\triangle_i^2$ and let e_5 be an upper bound for the fat girth of $\triangle_i^2$. The left panel shows the case where e_5 bounds $|b_i^2|$ and the right panel shows the case where e_5 bounds $|\omega_i^2|$.

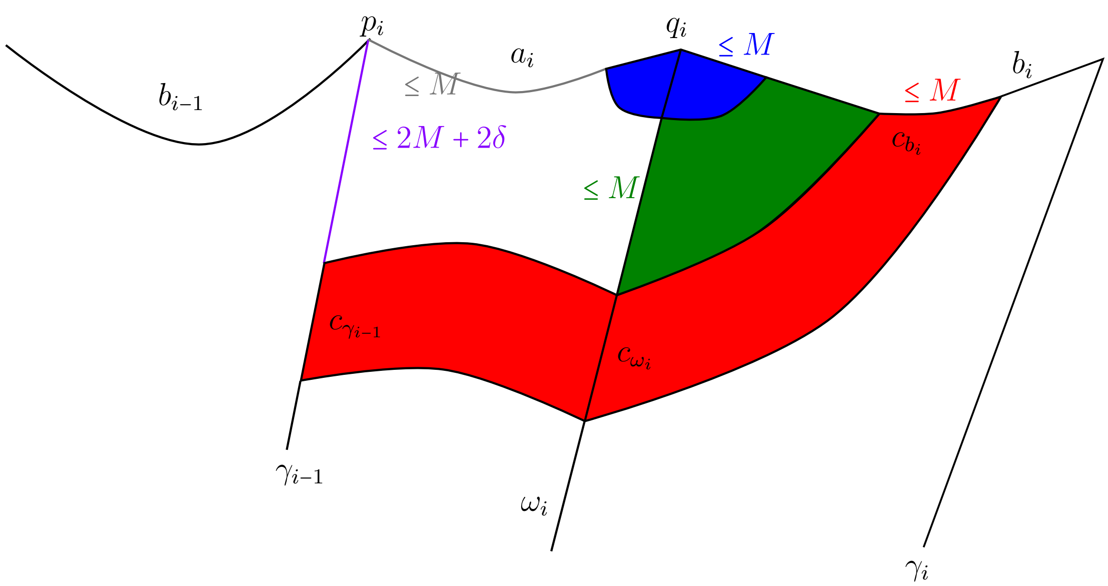

Figure 4.12: Here Δ_i^1 is δ-thin relative to F_{i-1}, ω_i^1 is contained in a corner segment of $\triangle_i^2$ at q_i and $e_3 \leq M$. Shaded regions indicate δ–fellow traveling between pairs of sides where the fellow traveling segments have equal length.

By the triangle inequality, the tail of γ_{i-1} from p_i to the closer endpoint of $c_{\gamma_{i-1}}$ has length at most $2M+2\delta$, so by Lemma 4.2.5, $|c_{\omega_i}| \leq M$. Thus in this subcase, the corner girth of $\triangle_i^2$ is at most $3M$.

Subcase: $F_i^1 \neq F_{i-1}$. Refer to Figure 4.13 for the following argument. Observe that $|\gamma_{i-1}^1| \leq M$; indeed, $\gamma_{i-1}^1 \subseteq \mathcal{N}_\delta(F_i^1)$ by definition, $\gamma_{i-1}^1 \subseteq \mathcal{N}_{2\delta}(F_{i-1})$ by the induction hypothesis and the fact that $\operatorname{diam}(\mathcal{N}_{3\delta}(F_i^1) \cap \mathcal{N}_{3\delta}(F_{i-1}) \leq M$. Label c_{ω_i}, $c_{\gamma_{i-1}}$ and c_{b_i} as in the previous case. As in the previous case $c_{\omega_i} \subseteq \mathcal{N}_{3\delta}(F_i)$. Since the corner girth of $\triangle_i^1$ is at most M and $|\gamma_{i-1}^1| \leq M$, the length of the tail of γ_{i-1} from p_i to the closer endpoint of $c_{\gamma_{i-1}}$ is at most $2M$. Thus by Lemma 4.2.5, $|c_{\omega_i}| \leq M$.

If $F_i^1 \neq F_i$, then $|\omega_i^1| \leq M$. Otherwise when $F_i^1 = F_i$, $|a_i^1| \leq M$. From before $|\gamma_{i-1}^1| \leq M$, so $|\omega_i^1| \leq 2M+3\delta$ by Lemma 2.3.22. Hence, $|\omega_i^1| \leq 2M+3\delta$. Consequently, since $e_3 \leq M$, $|c_{\omega_i}| \leq M$ and $|\omega_i^1| \leq 2M+3\delta$, the corner girth of $\triangle_i^2$ is at most $4M+3\delta$ in this subcase.

2. **Case: not all of ω_i^1 δ-fellow travels a subsegment of b_i.** Refer to Figure 4.14 for the following argument. Recall that $e_3 \leq M$ here. As in the previous case, the same argument shows that if $F_i^1 \neq F_{i-1}$, then $|\omega_i^1| \leq M$ and if $F_i^1 \neq F_i$, then $|\omega_i^1| \leq 2M+3\delta$. The corner segment of $\triangle_i^2$ at q_i contained in $\triangle_i^2$ is contained in the union of ω_i^1 and a corner segment of $\triangle_i^1$ at q_i whose length is bounded by $e_3 \leq M$, so in this case, the corner girth of $\triangle_i^2$ is at most $4M+3\delta$.

In summary, the corner girth of $\triangle_i^2$ is at most $4M+3\delta$ in all cases. □

Lemma 4.2.8. *If the fat girth of $\triangle_i^2$ is bounded above by $4M+3\delta$, then a $5M+6\delta$-tail of γ_i at p_{i+1} lies in $\mathcal{N}_{2\delta}(F_i)$.*

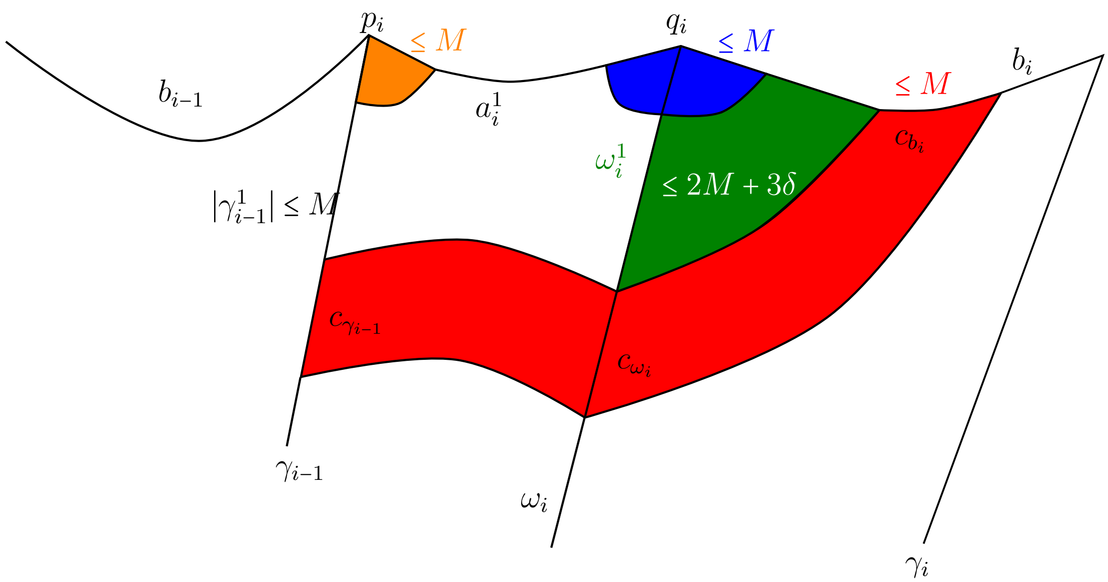

[p]

Figure 4.13: Here Δ_i^1 is δ-thin relative to $F \neq F_{i-1}$ and $e_3 \leq M$, the entire length of ω_i^1 δ–fellow travels a subsegment of b_i.

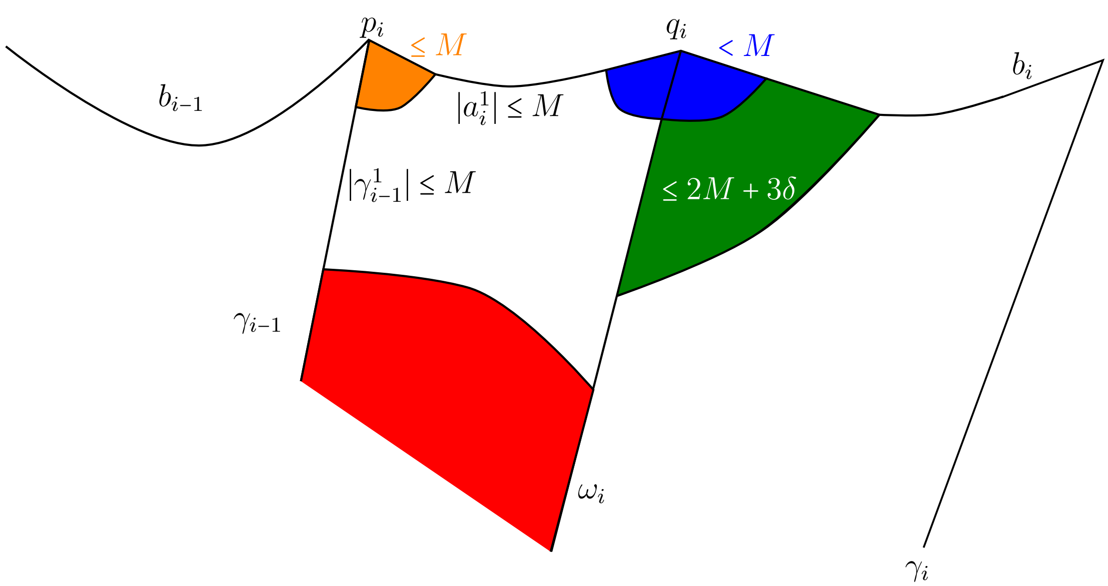

Figure 4.14: Calculating e_4 in the case that not all of ω_i^1 δ–fellow travels a subsegment of b_i.

Proof. **Case:** $F_i^2 \neq F_i$**.** The corner length of $\triangle_i^2$ at p_{i+1} can be obtained from $|b_i|$ by subtracting the corner length of $\triangle_i^2$ at q_i and $|b_i^2|$. The corner length of $\triangle_i^2$ at q_i is the corner girth of $\triangle_i^2$ which is at most $4M + 3\delta$. Here $b_i^2 \subseteq \mathcal{N}_{3\delta}(F_i) \cap \mathcal{N}_{3\delta}(F_i^2)$, so $|b_i^2| \leq M$. Since $|b_i| \geq 13M + 15\delta$ for $i \geq 2$, the corner length of $\triangle_i^2$ at q_i is at least $8M + 12\delta$. Therefore, a $8M + 12\delta$-tail of γ_i at q_i δ-fellow travels b_i. This tail therefore lies in $\mathcal{N}_{2\delta}(F_i)$, as desired.

Case: $F_i^2 = F_i$**.** Recall the corner girth of $\triangle_i^2$ is at most $4M+3\delta$ by Lemma 4.2.7. If $|b_i^2| \leq 4M + 3\delta$ and the corner girth of $\triangle_i^2$ at q_i is at most $4M + 3\delta$, then by an argument similar to the one in the previous case, the corner length of $\triangle_i^2$ at p_{i+1} is at least $5M + 9\delta$ and γ_i has an at least $5M + 6\delta$ tail at p_{i+1} contained in $\mathcal{N}_{2\delta}(F_i)$.

On the other hand, if $|b_i^2| > 4M + 3\delta$, then $|\omega_i^2|$ is bounded by $4M + 3\delta$ (the fat girth of $\triangle_i^2$). Let L_i be the corner length of $\triangle_i^2$ at p_{i+1}. By Lemma 2.3.22:

$$|\omega_i^2| \geq |b_i^2| - (4M + 3\delta) - 3\delta$$

Since the corner girth of $\triangle_i^2$ is at most $4M + 3\delta$:

$$|b_i^2| \geq |b_i| - L_i - (4M + 3\delta)$$

so that after combining the two inequalities and adding L_i to both sides:

$$|\omega_i^2| + L_i \geq |b_i| - L_i - (4M + 3\delta) - (4M + 3\delta) - 3\delta + L_i \geq 13M + 15\delta - 8M - 9\delta = 5M + 6\delta$$

At least a $L_i + |\gamma_i^2| \geq (5M + 6\delta)$-tail of γ_i at p_{i+1} lies in $\mathcal{N}_{2\delta}(F_i)$ because the corner segment of $\triangle_i^2$ at p_{i+1} contained in γ_i δ-fellow travels $b_i \subseteq \mathcal{N}_\delta(F_i)$ and $\gamma_i^2 \subseteq \mathcal{N}_\delta(F_i^2) = \mathcal{N}_\delta(F_i)$. □

The final step is to bound the fat girth of $\triangle_i^2$:

Lemma 4.2.9. *The fat girth of $\triangle_i^2$ is at most $4M + 3\delta$.*

Proof. If $F_i^2 \neq F_i$, the argument in the previous lemma showed that $|b_i^2| \leq M$ which means that the fat girth of $\triangle_i^2$ is at most M when $F_i^2 \neq F_i$, so assume that $F_i^2 = F_i$.

There are now 6 cases depending on the the order in which the endpoints of ω_i^1 and ω_i^2 appear on ω_i. Let $\omega_i^1 = [i_1, i_3]$ where i_1 is closer to p_1 than i_3 and let $\omega_i^2 = [i_2, i_4]$ where i_2 is closer to p_1 than i_4. Establish an ordering $i_j \leq i_k$ if i_j is closer to p_1 than i_k.

1. **Case:** $i_2 \leq i_1 \leq i_4 \leq i_3$.

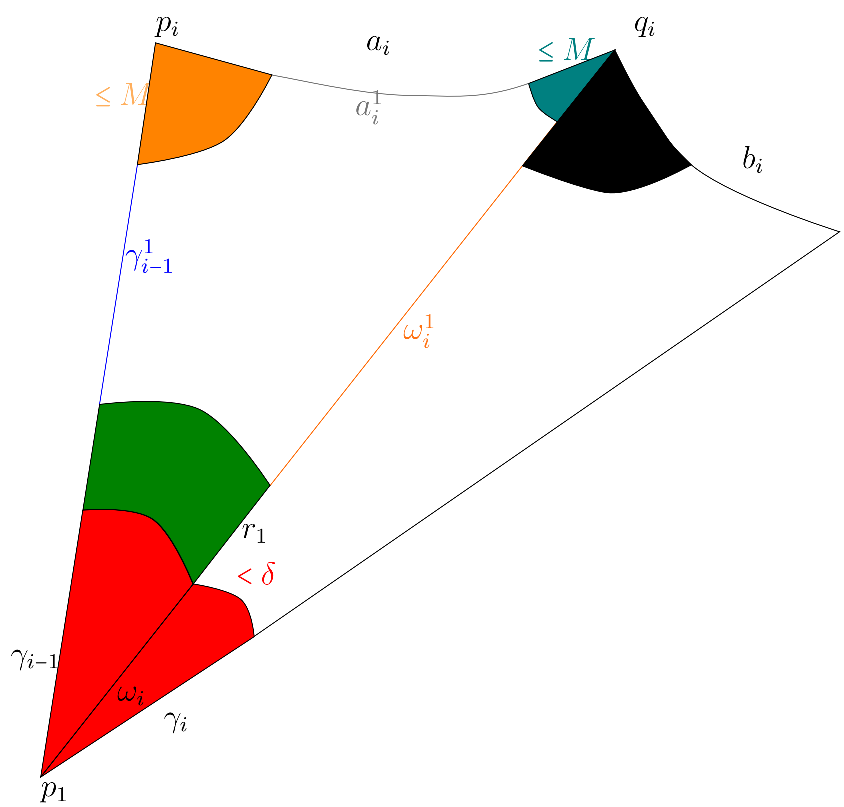

Figure 4.15: Bounding e_5 in the case $i_2 \leq i_1 \leq i_4 \leq i_3$.

Refer to Figure 4.15. If $F_i^1 \neq F_i$, then $|\omega_i^1| \leq M$ because $\omega_i^1 \subseteq \mathcal{N}_{2\delta}(F_i) \cap \mathcal{N}_{2\delta}(F_i^1)$. Since the fat girth of $\triangle_i^1$ is at most M, by Lemma 2.3.22, $|\gamma_{i-1}^1| \leq 2M + 3\delta$. Let $r_1 := [i_2, i_1] \subseteq \omega_i^2 \subseteq \omega_i$. Then r_1 lies in $\mathcal{N}_{2\delta}(F_i^2) = \mathcal{N}_{2\delta}(F_i)$.

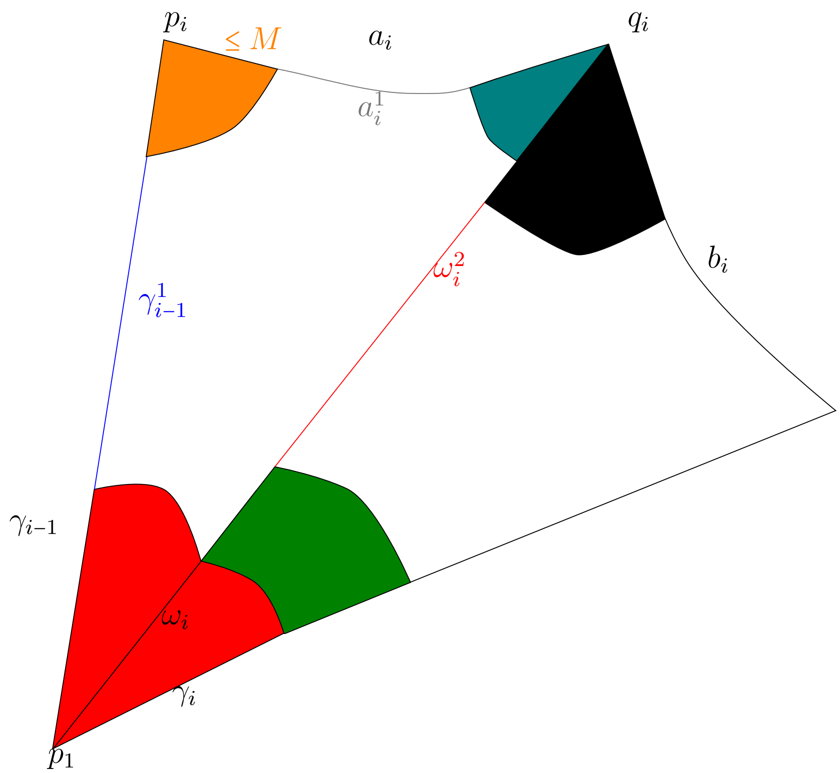

Figure 4.16: Bounding e_5 in the case $i_1 \leq i_2 \leq i_4 \leq i_3$.

Further, the corner girth of $\triangle_i^1$ is at most M and $|\gamma_{i-1}^1| \leq 2M + 3\delta$, so $|r_1|$ δ–fellow travels a segment of γ_{i-1} with an endpoint that is at most $|\gamma_{i-1}^1| + M \leq 3M + 3\delta$ from p_i, so $|r_1| \leq M$ by Lemma 4.2.5. Therefore, $|\omega_i^2| \leq |\omega_i^1| + |r_1| \leq 2M$.

On the other hand, if $F_i = F_i^1$: By hypothesis, a_i can spend at most M in $\mathcal{N}_{3\delta}(F_i)$, so $|a_i^1| \leq M$. Also, since the corner girth of $\triangle_i^1$ is at most M, $|\gamma_{i-1}^1| \leq M$ because at least $3M + \delta$ of the tail of γ_{i-1} at p_i must lie in $\mathcal{N}_{2\delta}(F_i)$. By construction $\gamma_{i-1}^1 \subseteq \mathcal{N}_\delta(F_i)$. By the triangle inequality (using the bounds on $|a_i^1|$ and $|\gamma_{i-1}^1|$), $|\omega_i^1| \leq 2M + 3\delta$. Here $|r_1| \leq M$, by using Lemma 4.2.5 and an argument similar to the one in the previous subcase, and so $|\omega_i^2| \leq 3M + 3\delta$. Therefore, in this case, the fat girth of $\triangle_i^2$ is at most $3M + 3\delta$.

2. **Case:** $i_1 \leq i_2 \leq i_4 \leq i_3$.

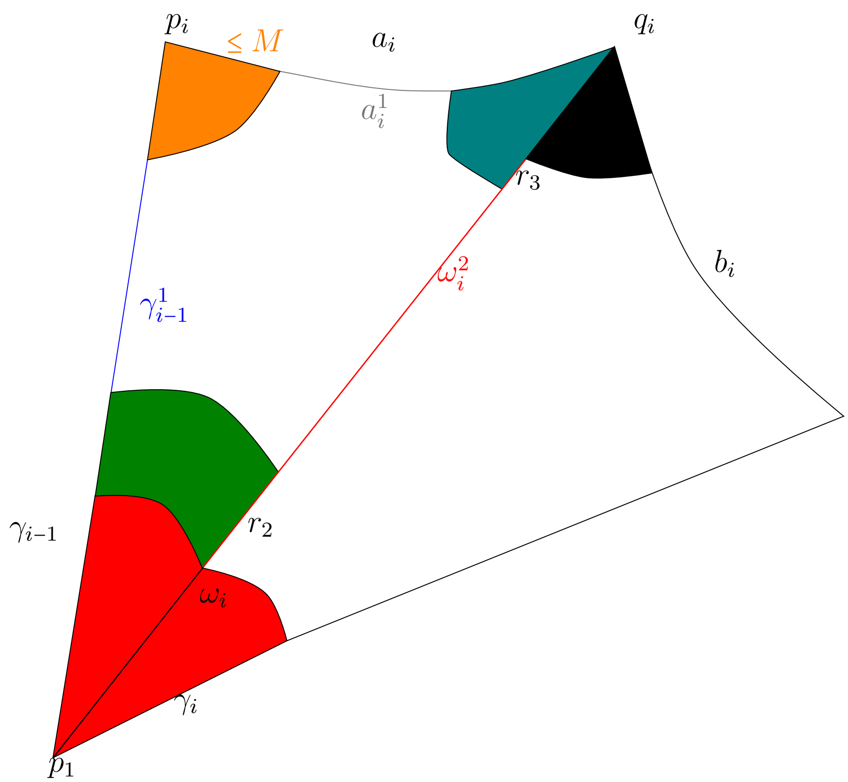

Figure 4.17: Bounding e_5 in the case $i_2 \leq i_1 \leq i_3 \leq i_4$.

If $F_i^1 \neq F_i$, then $\omega_i^2 \subseteq \mathcal{N}_\delta(F_i) \cap \mathcal{N}_\delta(F)$, so the fat girth of $\triangle_i^2$ is at most M. Thus assume $F_i^1 = F_i$. Refer to Figure 4.16. As in the previous case, $|a_i^1| \leq M$ and $|\gamma_i^1| \leq M$, so by the triangle inequality, $|\omega_i^2| \leq |\omega_i^1| \leq 2M + 3\delta$. Hence the fat girth of $\triangle_i^2$ is at most $2M + 3\delta$ in this case.

3. **Case:** $i_2 \leq i_1 \leq i_3 \leq i_4$.

   Refer to Figure 4.17. Let $r_2 := [i_2, i_1]$ and let $r_3 = [i_3, i_4]$. By an argument similar to the one in Case 1 when $F_i^1 \neq F_i$, then $|\omega_i^1| \leq M$ and the corner girth of $\triangle_i^1$ is at most M so that $|r_2| \leq M$ by Lemma 4.2.5.

   Since $r_3 \subseteq \omega_i^2 \subseteq \mathcal{N}_\delta(F_i^2) = \mathcal{N}_\delta(F_i)$ and $r_3 \subseteq \mathcal{N}_\delta(a_i) \cap \mathcal{N}_\delta(F_i)$, $|r_3| \leq M$ because a_i cannot spend more than M in $\mathcal{N}_{3\delta}(F_i)$. Thus if $F_i^1 \neq F_i$, then $|\omega_i^2| = |r_2| + |\omega_i^2| + |r_3| \leq 3M$, so in this subcase the fat girth of $\triangle_i^2$ is at most $3M$.

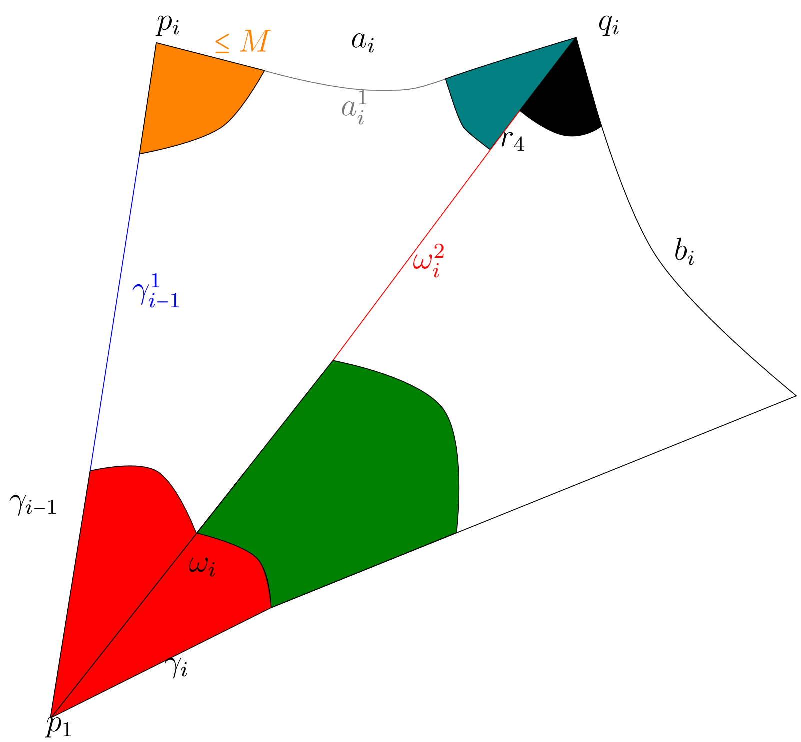

Figure 4.18: Bounding e_5 in the case $i_1 \leq i_2 \leq i_3 \leq i_4$.

On the other hand, suppose $F_i^1 = F_i$. As in the previous subcase $|r_3| \leq M$. Since γ_{i-1}^1 lies in $\mathcal{N}_{2\delta}(F_{i-1})$ and the corner girth of $\triangle_i^1$ is at most M, $|\gamma_{i-1}^1| \leq M$ by Lemma 4.2.5. Also $|a_i^1| \leq M$ because it lies in $\mathcal{N}_\delta(F_i) \cap a_i$. Thus by the bounds on $|a_i^1|$ and $|\gamma_{i-1}^1|$ and Lemma 2.3.22, $|\omega_i^1| \leq 2M + 3\delta$.

Since $|\gamma_{i-1}^1| \leq M$ and the corner girth of $\triangle_i^1$ is at most M, r_2 δ-fellow travels a subsegment of γ_{i-1} with an endpoint at most $2M$ from p_i, so $|r_2| \leq M$ by Lemma 4.2.5. Therefore $|\omega_i^2| = |r_2| + |\omega_i^1| + |r_3| \leq 4M + 3\delta$. This bound implies the fat girth of $\triangle_i^2$ is at most $4M + 3\delta$ in this case.

4. **Case:** $i_1 \leq i_2 \leq i_3 \leq i_4$**.**

   Refer to Figure 4.18. Let $r_4 := [i_3, i_4]$. Then $|r_4| \leq M$ because r_4 δ-fellow travels a segment in $\mathcal{N}_{2\delta}(F_i) \cap a_i$.

   If $F_i^1 \neq F_i$, then the part of $\omega_i^2 \setminus r_4 \subseteq \omega_i^1 \subseteq \mathcal{N}_\delta(F_i^1) \cap \mathcal{N}_\delta(F_i)$ and hence has

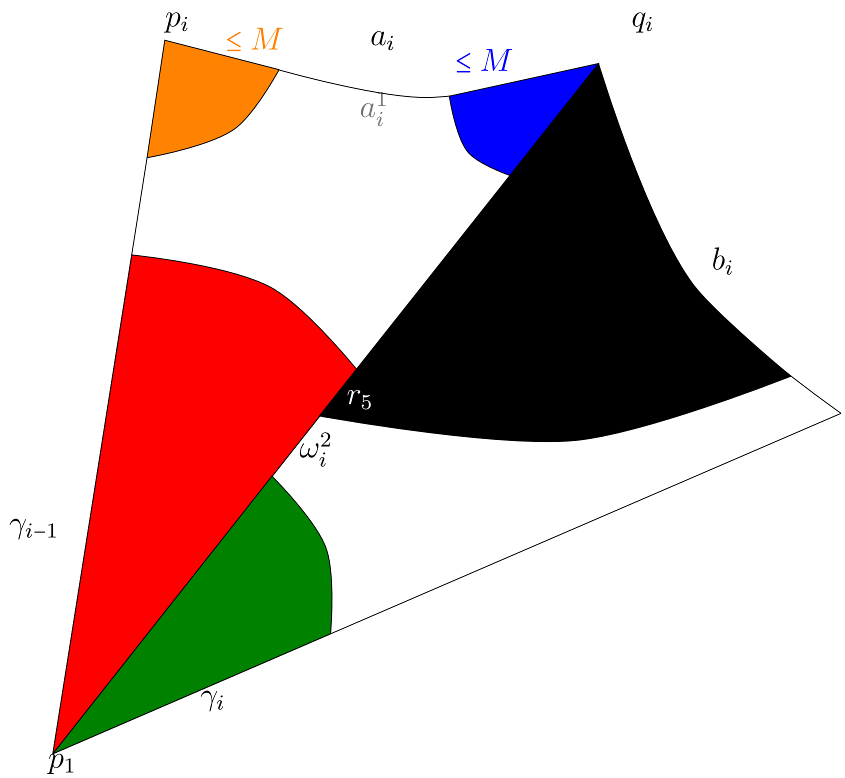

Figure 4.19: Bounding e_5 in the case $i_2 \le i_4 \le i_1 \le i_3$.

length at most M. Thus $|\omega_i^2| \le 2M$ and the fat girth of $\triangle_i^2$ is bounded above by $2M$ in this subcase.

On the other hand, if $F_i^1 = F_i$, arguments similar to ones in previous cases show that $|\gamma_i^1| \le 2M + 3\delta$, so $|\omega_i^2| \le 3M + 3\delta$ because $\omega_i^1 \subseteq \gamma_i^1$ so that the fat girth of $\triangle_i^2$ is at most $3M + 3\delta$ in this case.

5. **Case:** $i_2 \le i_4 \le i_1 \le i_3$**.**

   Refer to Figure 4.19. Let $r_5 := [i_4, i_1]$. First consider the case where $F_i^1 \neq F_i$. Then $|\omega_i^1| \le M$ because $\omega_i^1 \subseteq \mathcal{N}_{2\delta}(F_i) \cap \mathcal{N}_\delta(F)$, as in previous cases, $|a_i^1| \le M$, so by the triangle inequality $|\gamma_i^1| \le 2M + 3\delta$. On the other hand, if $\triangle_i^1$ is δ-thin relative to F_i, then $\gamma_{i-1}^1 \subseteq \mathcal{N}_\delta(F_i) \cap \mathcal{N}_{2\delta}(F_{i-1})$ and $|\gamma_{i-1}^1| \le M$ because a sufficiently long tail of γ_{i-1} lies in $\mathcal{N}_{2\delta}(F_{i-1})$ by the inductive hypothesis.

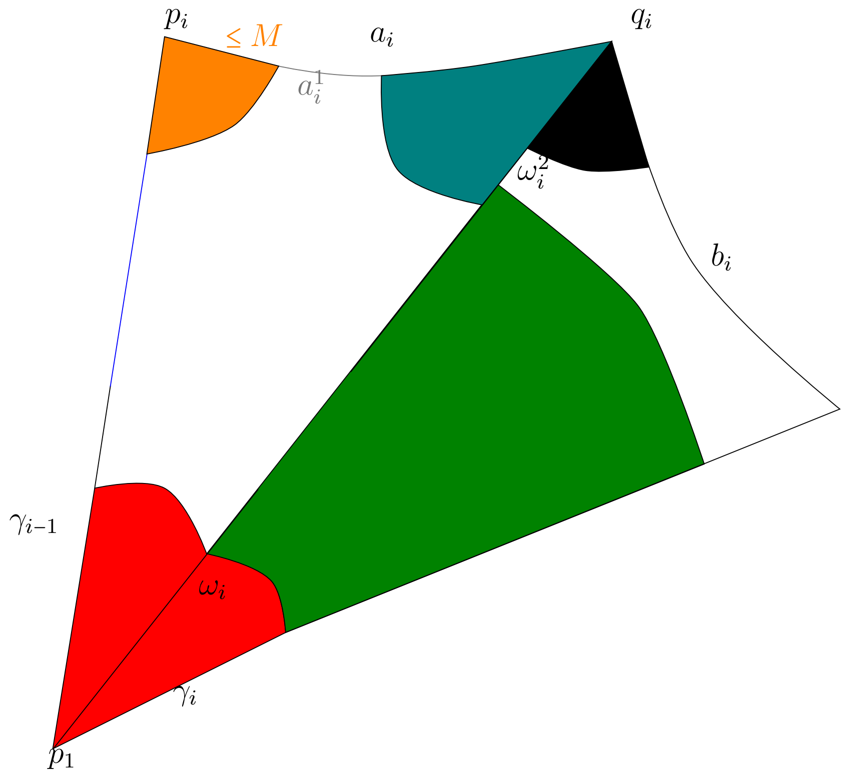

Figure 4.20: Bounding e_5 in the case $i_1 \leq i_3 \leq i_2 \leq i_4$.

In either of the above subcases, $|r_5| \leq M$ by Lemma 4.2.5 because r_5 δ-fellow travels a subsegment of γ_{i-1} with an endpoint at most $3M + 3\delta$ from p_i. By applying Lemma 4.2.5 again similarly, ω_i^2 must also have length at most M, so the fat girth of $\triangle_i^2$ is at most M in this case.

6. **Case:** $i_1 \leq i_3 \leq i_2 \leq i_4$.

   Refer to Figure 4.20. Immediately, $\omega_i^2 \subseteq \mathcal{N}_\delta(F_i)$ so that a_i has a subpath of length $|\omega_i^2|$ in $\mathcal{N}_{2\delta}(F_i)$ which can have length at most M by hypothesis. In this case, the fat girth of $\triangle_i^2$ is bounded above by M.

Hence the fat girth of $\triangle_i^2$ is at most $4M + 3\delta$ for all 6 subcases. □

In summary, the corner girth of $\triangle_i^1$ is at most M, the fat girth of $\triangle_i^1$ is at most

$e_2 = M$, the corner girth of $\triangle_i^2$ is at most $4M + 3\delta$, and the fat girth of $\triangle_i^2$ is at most $4M + 3\delta$. By Lemma 4.2.8, the $4M + 3\delta$-upper bound on the fat girth of $\triangle_i^2$ confirms that a $5M + 6\delta$ tail of γ_i lies in $\mathcal{N}_{2\delta}(F_i)$.

Recall Equation (4.1):

$$|\gamma_i| \geq |\gamma_{i-1}| + |a_i| + |b_i| - 2e_1 - 2e_2 - 2e_4 - 2e_5 - 6\delta$$

It then follows that:

$$|\gamma_i| \geq |\gamma_{i-1}| + |a_i| + |b_i| - 20M - 18\delta$$

completing the proof of Lemma 4.2.2. □

The next technical lemma is called the Arm Lemma because it is used twice in Proposition 4.2.11 to show that if γ satisfies the hypotheses of Lemma 4.2.2 and is a quasigeodesic, adding a geodesic segment or "**arm**" to each endpoint of γ will make a new quasigeodesic whose quasigeodesic constants differ from those of γ by at most some uniform amount.

Lemma 4.2.10 (Arm Lemma)**.** *Let $(\widetilde{Y}, \mathcal{B})$ be a (δ, R)-relatively hyperbolic pair, let $F_* \in \mathcal{B}$, and let $a' = [p', q']$, $\gamma = [q', q]$ and $\tau = [p', q]$ be geodesics in $\widetilde{Y}$ such that:*

1. $\operatorname{diam}(a' \cap \mathcal{N}_{4\delta}(F_*)) \leq R$, *and*
2. *more than a $2R$ tail of γ at q' lies in $\mathcal{N}_{3\delta}(F_*)$.*

Then $|\tau| \geq |a'| + |\tau| - 4R - 3\delta$.

If in addition:

1. *$H \in \mathcal{B}$ and a $5R + 6\delta$ tail of γ at q lies in $\mathcal{N}_{2\delta}(H)$,*

2. *a $5R$-tail of γ at q' lies in $\mathcal{N}_{2\delta}(F_*)$, and*

3. $|\gamma| \geq 10R + 3\delta$.

Then a tail of τ at q of length more than $4R$ either δ-fellow travels γ from q or lies in $\mathcal{N}_{3\delta}(H)$.

In the statement, the segment a' is supposed to be an arm and τ is the geodesic connecting the endpoints of the new segment formed by concatenating a' and γ. The second set of stronger hypotheses allows the attachment of a second arm a'' at q' so that if τ' is the geodesic joining the endpoints of the concatenation of τ with a'', the arm lemma can be used to obtain a lower bound for the length of τ'.

Proof. Let $\triangle$ denote the geodesic triangle $\gamma\tau a'$. The geodesics a' and γ can δ-fellow travel at q' for a distance of at most R because $\operatorname{diam}(a' \cap \mathcal{N}_{4\delta}(F_*)) \leq R$ and a more than $2R$ tail of γ at q' lies in $\mathcal{N}_{3\delta}(F_*)$, so the corner length of $\triangle$ at p' is at most R.

If $\triangle$ is δ-thin relative to $F \neq F_*$, the fat part of γ in $\triangle$ has length at most R because otherwise the fat part of γ has a subsegment of length more than R that coincides with the more than $2R$ tail of γ at p' that lies in $\mathcal{N}_{3\delta}(F_*)$. On the other hand, if $\triangle$ is δ-thin relative to F_*, then the fat part of a' in $\triangle$ has length at most R because $\operatorname{diam}(a' \cap \mathcal{N}_{4\delta}(F_*)) \leq R$. The argument that $|\tau| \geq |a'| + |\gamma| - 4R - 3\delta$ is now identical to the argument that $|\omega| \geq |\gamma_{i-1}| + |a| - 2e_1 - 2e_2 - 3\delta$ in the proof of Lemma 4.2.3 where e_1 is an upper bound on the corner girth of $\triangle_i^1$ and e_2 is an upper bound on the fat girth of $\triangle_i^1$.

Now suppose a $5R + 6\delta$ tail of γ at q lies in $\mathcal{N}_{2\delta}(H)$ for some $H \in \mathcal{B}$. If the corner length of $\triangle$ at q is at least $4R$, the second statement follows immediately, so suppose that the corner length of $\triangle$ at q less than $4R$. The corner length at q' is at most R by the preceding. Therefore, the length of the fat part of γ is at least $5R+3\delta$. Since the corner length of $\triangle$ at q' is less than R and a $5R$ tail of γ at q' lies in $\mathcal{N}_{2\delta}(F_*)$, the fat part of γ has a length more than R subsegment in $\mathcal{N}_{3\delta}(F_*)$, so $F = F_*$. Similarly, $F = H$, so $H = F_*$. In this case, the length of the fat part of a' is at most R, while the length of the fat part of γ is at least $5R + 3\delta$, so the length of the fat part of τ must be at least $4R$ by Lemma 2.3.22. The fat part of τ and the corner segment contained in τ at q together contain a length at least $4R$-tail of τ at q contained in $\mathcal{N}_{3\delta}(H)$. □

Proposition 4.2.11. *Let $(\widetilde{X}, \mathcal{B})$ be an (δ, M)–relatively hyperbolic pair. Let E be an $(\mathcal{A}, \mathcal{B})$-LPPC subspace of $\widetilde{X}$, and suppose that if $A_1, A_2 \in \mathcal{A}$, then $d(A_1, A_2) \geq 1024(8M + 6\delta + 2048(40M + 60\delta))$. If γ is a geodesic path in E, then γ embeds into $\widetilde{X}$ as a $(1024, 8M + 6\delta + 2048(40M + 60\delta))$–quasigeodesic.*

Proof. The path γ has one of the following piecewise geodesic forms:

1. $b_1 a_2 b_2 \ldots a_n b_n$ and $|b_1|, |b_n| \geq 1024M(40M + 60\delta)$
2. $a_1 b_1 a_2 b_2 \ldots a_{n-1} b_{n-1} a_n$
3. $a_1 b_1 a_2 b_2 \ldots a_n b_n$
4. $b_1 a_2 b_2 \ldots b_{n-1} a_n$
5. $b_1 a_2 b_2 \ldots a_n b_n$, where one or both of $|b_1|, |b_n|$ is less than $1024M(40M + 60\delta)$.

such that for all $1 \leq i \leq n$, $a_i \subseteq A_i \in \mathcal{A}$, for all $1 \leq i \leq n$, $b_i \subseteq B_i \in \mathcal{B}$, and for $1 \leq i \leq n$, $|b_i| \geq \max\{512(40M + 60\delta), 1024(8M + 6\delta + 2048(40M + 60\delta))\}$ unless specified otherwise above. The general strategy is to use Lemma 4.2.2 for paths of the first form and reduce the other possible forms to that case.

Note that since γ is an E-geodesic and each B_i is convex, if $B_i = B_j$, then $i = j$.

1. **Case:** $\gamma = b_1 a_2 b_2 \dots a_n b_n$ **and** $|b_1|, |b_n| \geq 1024(40M + 60\delta)$.

   By Proposition 4.1.5, there exists a path $\gamma' = b'_1 a'_2 b'_2 \dots a'_{n-1} b'_{n-1} a'_n b'_n$ such that γ' satisfies the hypotheses of Lemma 4.2.2, γ' has the same endpoints as γ, $|\gamma'| \geq \frac{1}{256}|\gamma| - 16n\delta$ and for $1 \leq i \leq n$, $|b'_i| \geq \frac{1}{256}|b_i| - 8\delta$. Let γ'' be the geodesic joining the endpoints of $b'_1 a'_2 b'_2 \dots a'_{n-1} a'_n b'_n$. By Lemma 4.2.2:

   $$|\gamma''| \geq |b'_1| + \sum_{i=2}^{n} (|a'_i| + |b'_i|) - n(20M + 18\delta)$$

   Note that for all $1 \leq i \leq n$, $|b'_i| \geq \frac{1}{256}|b_i| - 8\delta = 4(40M + 60\delta) - 8\delta \geq 2(20M + 18\delta + 16\delta)$. Then:

   $$|\gamma''| \geq \frac{1}{2}|\gamma'| + 16n\delta \geq \frac{1}{512}|\gamma|,$$

   so in this case, γ' is a $(2,0)$ quasigeodesic and γ is a $(512,0)$ quasigeodesic.

2. **Case:** $\gamma = a_1 b_1 a_2 b_2 \dots a_{n-1} b_{n-1} a_n$.

   Construct $\gamma' := a'_1 b'_1 a'_2 b'_2 \dots b'_{n-1} a'_n$ using Proposition 4.1.5. In this case, the subpath $\rho' := b'_1 \dots b'_{n-1}$ satisfies the hypotheses of the preceding case and is a $(512,0)$-quasigeodesic. Let $\rho'' := [q'_1, p'_n]$ be the geodesic connecting the endpoints of ρ'. By Lemma 4.2.2, the $5M + 6\delta$-tail of ρ'' at q'_1 lies in $\mathcal{N}_{2\delta}(F_1)$ and the $5M + 6\delta$-tail of ρ'' at p_n lies in $\mathcal{N}_{2\delta}(F_{n-1})$ where $F_1, F_{n-1} \in \mathcal{B}$ with $b_1 \subseteq F_1$ and $b_{n-1} \subseteq F_{n-1}$.

   By Proposition 4.1.5, $\operatorname{diam}(\mathcal{N}_{3\delta}(F_1) \cap a'_1) \leq M$ and $\operatorname{diam}(\mathcal{N}_{3\delta}(F_{n-1}) \cap a'_n) \leq M$. Observe that $p'_n \in \mathcal{N}_{\delta}(F_{n-1})$. Claim: $\operatorname{diam}(a'_1 \cap \mathcal{N}_{5\delta}(F_1)) \leq 2M$. Suppose not.

Then there would exist $x \in a_n'$ and $y \in \mathcal{N}_\delta(F_1)$ such that $d(x, p_n') > 2M$ and $d(x,y) \leq 4\delta$. However, the convexity of the CAT(0) metric and the convexity of $\mathcal{N}_\delta(F_{n-1})$ imply that more than M of $[p_n', x]$ lies in $\mathcal{N}_{3\delta}(F_{n-1})$ which is a contradiction.

Let $\rho''' = [p_1', p_n']$ so that ρ''', ρ'', a_1' form a geodesic triangle $\triangle_1$. Observe that $|b_1| \geq 1024(60M + 50\delta)$, so $|\rho''| \geq 2(60M + 50\delta)$. By applying Lemma 4.2.10 to $\triangle_1$ with $R = M$:

$$|\rho'''| \geq |\rho''| + |a_1'| - 4M - 3\delta$$

and that more than a $4M$ tail of ρ''' at p_n' lies in $\mathcal{N}_{3\delta}(F_{n-1})$. Now let $\tau = [p_1', q_n']$ which is the geodesic joining the endpoints of γ. Apply Lemma 4.2.10 to the triangle formed by τ, ρ'' and a_n' with $R = 2M$ to show that:

$$|\tau| \geq |\rho''| + |a_n'| - 4M - 3\delta \geq |\rho''| + |a_1'| + |a_n'| - 8M - 6\delta \geq \frac{1}{2}|\rho'| + |a_i'| + |a_n'| - 8M - 6\delta$$

$$\geq \frac{1}{2}|\gamma'| - 8M - 6\delta \geq \frac{1}{512}|\gamma| - 8M - 6\delta - 8n\delta \geq \frac{1}{1024}|\gamma| - 8M - 6\delta$$

using the fact that ρ' is a $(\frac{1}{2}, 0)$-quasigeodesic, the fact that $|\gamma'| \geq \frac{1}{256}|\gamma| - 16n\delta$ by Proposition 4.1.5 and that each $|b_i'| \geq 4(40M + 60\delta)$. Therefore γ is a $(1024, 8M + 6\delta)$-quasigeodesic.

3. **Case:** $\gamma = a_1 b_1 a_2 b_2 \ldots a_n b_n$.

   Let τ be the geodesic connecting the endpoints of γ.

   If $|b_n| \geq 1024(40M + 60\delta)$, apply the first case of this argument to $b_1 a_2 \ldots a_n b_n$. Then an argument similar to the preceding case using Lemma 4.2.10 shows that:

$$|\tau| \geq \frac{1}{1024}|\gamma| - 8M + 6\delta$$

If $|b_n| < 1024(40M + 60\delta)$, then apply the second case to $a_1 b_1 a_2 b_2 \ldots a_n$ of this argument which must be a $(1024, 8M + 6\delta)$–quasigeodesic. It then follows by the triangle inequality that γ is a $(1024, 8M + 6\delta + 1024(40M + 60\delta))$-quasigeodesic.

4. **Case:** $\gamma = b_1 a_2 b_2 \ldots a_n$**.**

   The argument is the same as in the immediately preceding case (this can be seen immediately rewriting γ as $a_n b_{n-1} a_{n-1} \ldots a_2 b_1$).

5. **Case:** $\gamma = b_1 a_2 b_2 \ldots a_n b_n$ **where one or both of** $|b_1|, |b_n|$ **is less than** $1024(40M + 60\delta)$**.**

   For the subcase $|b_1|$ and $|b_n| < 1024(40M + 60\delta)$, apply the second case to $a_2 b_2 \ldots b_{n-1} a_n$ and conclude that γ is a $(1024, 8M + 6\delta + 2048(40M + 60\delta))$-quasigeodesic.

   In the other two subcases, similar arguments show that γ is a $(1024, 8M + 6\delta + 1024(40M + 60\delta))$-quasigeodesic. □

CHAPTER 5

THE GEOMETRY OF SPECIAL CUBE COMPLEXES

5.1 Special cube complexes and separability

Non-positively curved cube complexes were introduced in Chapter 2. **Special** and **virtually special** cube complexes are NPC cube complexes with additional restrictions on the hyperplanes:

Definition 5.1.1. *Let X be a NPC cube complex. Suppose that:*

1. *Every hyperplane of X is embedded.*
2. *Every hyperplane of X is two-sided.*
3. *Every hyperplane of X does not **self-osculate**: let H be a two-sided embedded hyperplane of X corresponding to the parallelism class of edges $[e]$. Since H is two-sided, each codimension-1-midcube in H has a top and bottom side which induces an orientation on the edges of $[e]$. The hyperplane H **self-osculates** if there exist distinct oriented $e_1, e_2 \in [e]$with the same orientation such that $e_1 \cap e_2 = \{v\}$, v is the initial vertex of both e_1 and e_2 and (e_1, e_2) is not a pair of sides of a 2-cube of X.*
4. *No two hyperplanes **interosculate**: let H_1, H_2 be two hyperplanes of X with corresponding parallelism classes of edges $[e_1], [e_2]$. If there exist $e, f \in [e_1]$ and $e', f' \in [e_2]$ such that e and e' form the corner of a two-cube in X and f and f' intersect but do not form a corner of a two-cube in X, then H_1 and H_2 **interosculate**.*

*Then X is said to be **special**.*

*When X has a finite cover that is special, then X is **virtually special**.*

*A group G is **special** if it is the fundamental group of a special cube complex. Naturally, G is **virtually special** if G has a finite index special subgroup.*

See Figure 5.1 for visual examples of self-osculation and interosculation.

Finite covers of special cube complexes are special:

Proposition 5.1.2. *Let X be a special cube complex and let $\hat{X} \to X$ be a finite covering. Then $\hat{X}$ is special.*

The proof is straightforward: if $\hat{X}$ is not special, then the pathology that prevents $\hat{X}$ from being special has an image under the covering map that prevents X from being special.

Right angled Artin groups provide important examples of special cube complexes.

Definition 5.1.3. *Let Γ be a graph with vertices V. The **right angled Artin Group** on Γ is the group:*

$$A(\Gamma) := \langle v \in V \mid [v_1, v_2] \text{ if } v_1 \in V \text{ and } v_2 \in V \text{ are connected by an edge} \rangle.$$

*The **Salvetti complex of** Γ is the complex $S(\Gamma)$ which has 1 vertex and an edge for each $v \in V$. For every $n \geq 2$ inductively add an n-cube for each n-clique in Γ so that each parallelism class of edges has a single element corresponding to the vertices of that n-clique and every $(n-1)$-clique contained in an n-clique corresponds to a face of the n-cube corresponding to that n-clique.*

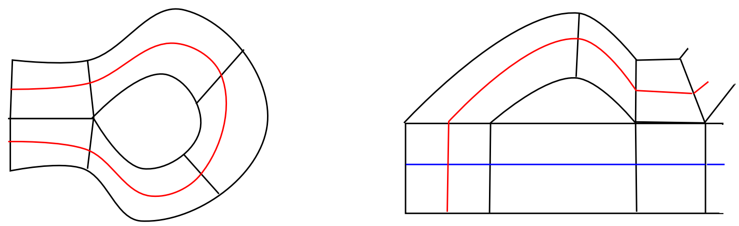

Figure 5.1: Left: an example of a self osculating hyperplane (red). Right: an example of two hyperplanes (red and blue) that interosculate.

By examining the 2-skeleton, $\pi_1(S(\Gamma)) \cong A(\Gamma)$. In fact, if $G = A(\Gamma)$, then $S(\Gamma)$ is a $K(G,1)$. By construction, every hyperplane of $S(\Gamma)$ corresponds to a vertex of Γ because each parallelism class of edges contains a single element.

Example 5.1.4. *The Salvetti complex of a RAAG is special. Indeed, every hyperplane H has a single dual edge e_H. Since each parallelism class of edges contains a single element, no hyperplane self-osculates, and no pair of hyperplanes inter-osculate.*

In fact, Salvetti complexes are prototypical examples of special cube complexes. If X is a non-positively curved cube complex, the **crossing graph of** X is a graph Γ_X that has a vertex for each hyperplane of X and for each pair of vertices v_1, v_2, there is an edge between v_1 and v_2 if and only if the corresponding hyperplanes cross.

Every CAT(0) cube complex is also special, see [34] for details.

Theorem 5.1.5 (Haglund-Wise [15], see [34] Theorem 4.4)**.** *Let X be a NPC cube complex. Then X is special if and only if there is a local isometry from X to $S(\Gamma_X)$ where Γ_X is the crossing graph of X.*

Special cube complexes were defined specifically to determine which cube complexes admit local isometries into Salvetti complexes of RAAGs. The first application of Theorem 5.1.5 is to show that the fundamental group of a special cube complex is residually finite:

Definition 5.1.6. *A group G is **residually finite** if for all $g \in G \smallsetminus \{1\}$, there exists a homomorphism $\psi_g : G \to F$ such that F is a finite group and $\psi_g(g) \neq 1$.*

It is well known that right angled Artin groups are linear groups which are residually finite. Subgroups of residually finite groups are residually finite, so Proposition 2.1.6 yields the following corollary:

Corollary 5.1.7. *Let X be a non-positively curved virtually special cube complex. Then $\pi_1 X$ is residually finite.*

When G is the fundamental group of a special cube complex, G also has certain subgroups which satisfy a useful residual property called separability:

Definition 5.1.8. *Let G be a group and let $H \leq G$. The subgroup H is **separable in** G if for all $g \in G \smallsetminus H$, there exists a homomorphism $\psi_g : G \to F$ such that F is finite, $H \leq \ker \psi_g$ and $\psi_g \neq 1$.*

Equivalently, H is separable if it is the intersection of the finite index subgroups of G that contain H. Separability is also equivalent to being closed in the profinite topology on G. A group G is residually finite if and only if the trivial subgroup is separable in G.

Example 5.1.9. *Let F be a free group and let H be a finitely generated subgroup of F. The subgroup H is separable in G by Marshall Hall's theorem (see [31] for an intuitive proof). Free groups are therefore called **locally extended residually finite** because their finitely generated subgroups are separable.*

Passing to finite index subgroups is compatible with separability:

Lemma 5.1.10. *Let G be a group, let G_0 be a finite index subgroup of G and let $H \leq G$. Then H is separable in G if and only if $H \cap G_0$ is separable in G_0.*

Proof. If H is separable, then H is closed in the profinite topology on G. Therefore, $H \cap G_0$ is closed in the induced profinite topology on G_0.

On the other hand, suppose that $H \cap G_0$ is closed in the induced profinite topology on G_0. Since G_0 is finite index, G_0 and its complement are open in the profinite toplogy on G. The complement of $H \cap G_0$ in G_0 is open in the profinite topology on G_0 and hence in the profinite topology on G as well because G_0 is open in that topology. Thus $G \setminus H = (G \setminus G_0) \cup (G_0 \setminus H)$ is open in the profinite topology on G, so H is closed in the profinite topology on G. □

Separability has an important geometric consequence:

Theorem 5.1.11 (Scott's Criterion, [29])**.** *Let X be a connected complex, $G = \pi_1 X$ and $H \leq G$. Let $p : X^H \to X$ be the cover corresponding to H. The subgroup H is separable in G if and only if for every compact subcomplex $Y \subseteq X^H$, there exists an intermediate finite cover $X^H \to \hat{X} \to X$ such that $Y \hookrightarrow \hat{X}$.*

In other words, an immersion to X whose image has separable fundamental group in $\pi_1 X$ can be promoted to an embedding into a finite cover of X. Stallings' proof of Marshall Hall's theorem in [31] shows explicitly how immersions of finite graphs can be promoted to embeddings in finite covers.

In special cube complexes, the hyperplanes are separable:

Proposition 5.1.12. *Let X be a virtually special compact and non-positively curved cube complex. Let W be a hyperplane of X. Then $\pi_1(W)$ is separable in $\pi_1(X)$.*

Haglund and Wise's Canonical Completion and Retraction from [15] which extend Stallings' methods for finite groups are needed to prove separability of the hyperplanes:

Lemma 5.1.13 ([34, Construction 4.12])**.** *Let X be a non-positively curved special cube complex. Let $f : Y \to X$ be a local isometry where Y is compact. Then there exists a finite sheeted cover $\hat{X}$ such that f lifts to an embedding $\hat{f} : Y \to \hat{X}$. Further, there exists a retraction $r : \hat{X} \to Y$.*

Proof of Proposition 5.1.12. Recall that being separable is equivalent to being closed in the profinite topology. Let $f : W \to X$ be a local isometric immersion. By Corollary 2.2.5, f induces an injective map $f_* : \pi_1(W) \to \pi_1(X)$. Passing to a finite special cover and taking a canonical completion $\hat{f} : Y \to \hat{X}$ yields an induced map $\hat{f}_* \pi_1(W) \to H$ where H is a finite index subgroup of $\pi_1(W)$ in $\pi_1(X)$. Since $\hat{X}$ is a special cube complex, $\pi_1(\hat{X})$ is residually finite and hence has Hausdorff profinite topology. By point set topology, since $r_* : \pi_1(\hat{X}) \to \pi_1(W)$ is a retraction on the Hausdorff space $\pi_1(\hat{X})$, H is closed in the profinite topology on $\pi_1(\hat{X})$ and is hence separable $\pi_1(\hat{X})$. Thus $\pi_1(W)$ is separable in $\pi_1(X)$ by Lemma 5.1.10. □

5.2 Elevations and R-embeddings

This thesis relies heavily on the fact that the hyperplanes of a special cube complex have separable fundamental group. When X is a special NPC cube complex, Scott's criterion can be used to pass to a finite cover $\hat{X}$ so that elevations of hyperplanes in $\hat{X}$ to the CAT(0) universal cover $\widetilde{X}$ are sufficiently separated. This will be instrumental in ensuring that geodesics in the edge spaces of the hierarchy in Chapter 6 satisfy the hypotheses of Proposition 4.2.11. This subchapter builds up the technical tools and terminology used to obtain finite covers whose hyperplanes elevate to sufficiently separated images in the universal cover.

The first step is to formalize the notion of an elevation:

Definition 5.2.1. *Let W be a connected topological space and let $\phi : W \to Z$ be a continuous map. Let $p : \hat{Z} \to Z$ be a covering map. An **elevation of** W **to** $\hat{Z}$ is a minimal covering $\hat{p} : \hat{W} \to W$ such that the map $\hat{\phi} := \phi \circ \hat{p}$ lifts to a map $\hat{W} \to \hat{Z}$.*

Often, the map $\hat{W} \to \hat{Z}$ will be implied and an elevation of ϕ will instead refer to the image of some $\hat{\phi}$ such that $\hat{\phi}$ is some elevation.

*Elevations may not be unique: two elevations of the same map are **distinct** if they have different images.*

When $\phi : W \to Z$ is an inclusion map, then the distinct elevations of ϕ are precisely the components of $p^{-1}(W)$. Lemma 5.2.2 illustrates this principle in the important case where p is the universal covering map:

Lemma 5.2.2. *Let $i : Y \hookrightarrow X$ be an inclusion of connected topological spaces and let $W = i(Y)$. Let $\widetilde{X}$ be the universal cover of X and let $p : \widetilde{X} \to X$ be the canonical covering map. If $f : \widetilde{Y} \to \widetilde{X}$ is an elevation of the map i to $\widetilde{X}$, then the image of f is a connected component of $p^{-1}(W)$.*

Proof. Let $A = \operatorname{Im} f$ and let $p_Y : \widetilde{Y} \to Y$ be the canonical covering map. By definition of an elevation, $A \subseteq p^{-1}(W)$. Since A is connected, A is contained in some connected component of $p^{-1}(W)$. Suppose y is in the same connected component of $p^{-1}(W)$ as A. Then there exists a path γ in $p^{-1}(W)$ from y to a point $x \in A$. Let $x = f(z)$. The path $p(\gamma)$ lifts to a path γ_Y in $\widetilde{Y}$ issuing from z, let w be the other endpoint of the path. Then $i \circ p_Y w = p(y)$. Since f is a lift of $p \circ p_Y$, $f(w) = y$, so $y \in A$ and A is a connected component of $p^{-1}(W)$. □

An R**-embedded** subspace is a connected subspace with elevations that have embedded R-neighborhoods:

Definition 5.2.3. *Let X be a metric space, $R \geq 0$ and let $Y \subseteq X$ be connected. Let $p : X^Y \to X$ be the covering space associated to $\pi_1(Y)$ so that the inclusion $Y \hookrightarrow X$ lifts canonically to X^Y. The subspace Y is* R***-embedded in*** X *if p is injective on $\mathcal{N}_R(Y) \subseteq X^Y$.*

The following lemma is straightforward but will be important:

Lemma 5.2.4. *Let $p : \hat{X} \to X$ be a finite cover. If A is R-embedded in X, then each component of $p^{-1}(X)$ is R-embedded in $\hat{X}$.*

Separability of the hyperplane subgroups is used in this thesis to prove the following: every compact virtually special cube complex has a finite cover where every hyperplane is R-embedded. Moreover, these R-embedded hyperplanes have elevations to the universal cover that are $2R$-separated:

Proposition 5.2.5. *Let X be a compact virtually special non-positively curved cube complex. Given $R \geq 0$, then there exists a finite regular (compact) special cover C such that for every hyperplane $W \subseteq X$, W is R-embedded in C.*

If $\widetilde{W_1}, \widetilde{W_2}$ are distinct elevations of a hyperplane W of X to the universal cover $\widetilde{X}$, then $d_{\widetilde{X}}(\widetilde{W_1}, \widetilde{W_2}) \geq 2R$.

Proof. Since X is compact and therefore finite, every hyperplane is finite. In particular, if W is a hyperplane of X, then $\mathcal{N}_R(W)$ is compact and connected.

Since X is special, $\pi_1(W)$ is separable by Proposition 5.1.12. By theorem 5.1.11, there exists a finite covering $\hat{p} : \hat{X} \to X$ such that there is an embedding $i_W : \mathcal{N}_R(W) \hookrightarrow \hat{X}$.

Let $\widetilde{p} : \widetilde{X} \to X$, $p^W : \widetilde{X}^W \to X$ and $p : \widetilde{X} \to X^W$ be canonical covering maps so that $\widetilde{p} = p^W \circ p$. Let $\widetilde{W} \to \widetilde{W_1}, \widetilde{W} \to \widetilde{W_2}$ be distinct elevations of W to $\widetilde{X}$ with respect to the covering map $p : \widetilde{X} \to X^W$ and let $\widetilde{w_1} \in \widetilde{W_1}$ and $\widetilde{w_2} \in \widetilde{W_2}$. Note that $\widetilde{W_1}$, $\widetilde{W_2}$ are disjoint because they are connected components of $\widetilde{p}^{-1}(W)$ and have distinct images in $\widetilde{X}$.

Let γ be a path in $\widetilde{X}$ of length less than $2R$ connecting $\widetilde{w_1}$ and $\widetilde{w_2}$. Suppose toward a contradiction that there exists a point $\widetilde{x} \in \gamma$ such that $d(\widetilde{x}, \widetilde{w_1}) < R$ and $d(\widetilde{x}, \widetilde{w_2}) < R$.

Since $\widetilde{W_1}$ and $\widetilde{W_2}$ are elevations of the same hyperplane, there exists $g \in \pi_1(X) \smallsetminus \pi_1(W)$ such that $g \cdot \widetilde{w_1} \in \widetilde{W_2}$, and $g \notin \pi_1(W)$ because otherwise $g \cdot \widetilde{w_1} \in W_1 \cap W_2$ in which case, $\widetilde{w_1} \in \widetilde{W_2}$ but $\widetilde{w_1} \notin \widetilde{W_2}$. Now $d(g \cdot x, \widetilde{W_2}) \leq R$. Since $g \notin \pi_1(W)$, $p(x) \neq p(g \cdot x)$. By definition of an elevation, $p(W_2)$ is contained in the image of the inclusion of W into X^W. Also $p(x)$, $mp(g \cdot x)$ lie in an R-neighborhood of the image of W in X^W. However,

$$p^W \circ p(x) = \widetilde{p}(x) = \widetilde{p}(g \cdot x) = p^W \circ p(g \cdot x)$$

contradicting the fact that W is R-embedded in X^W.

The cover X^W corresponds to some finite index subgroup $H \leq \pi_1(X)$ and by passing to a finite index subgroup of H which is normal in $\pi_1(X)$, one may assume that X^W is regular. Further, the number of orbits of hyperplanes under deck transformations of X^W is at most the number of hyperplanes of X, and every hyperplane in the orbit of an elevation of W to X^W is R-embedded.

If there exists a hyperplane W' which is not R-embedded in X^W, repeat the same process to produce a finite regular cover of X^W (and hence X) where elevations of W and W' and any hyperplanes in their deck transformation orbits are R-embedded. Every hyperplane of W' lies in a deck transformation orbit of an elevation of a hyperplane of W, so passing to a finite cover finitely many times will produce a finite covering $C \to X$ where every hyperplane of C is R-embedded because each hyperplane is in the orbit of an elevation of a hyperplane of X. □

The separation between hyperplane elevations will be used later in Section 6 to produced piecewise geodesics whose pieces connecting two elevations of a hyperplane are sufficiently long – that is, at least length $2R$.

5.3 Convex Cores

Specialness also plays a role in building a geometric representation of the peripheral structure. Haglund and Wise proved that quasiconvex subgroups of hyperbolic virtually special compact groups have "convex cores:"

Theorem 5.3.1 ([14],[27], see also [15] Proposition 7.2)**.** *Let $\widetilde{X}$ be a CAT(0) cube complex. Let G be a hyperbolic group acting geometrically on $\widetilde{X}$ with compact virtually special quotient. If H is a quasiconvex subgroup of G and K is any compact subcomplex of $\widetilde{X}$, then there exists a proper H-cococompact convex subcomplex of $\widetilde{X}$ containing K.*

Theorem 5.3.1 with canonical completion and retraction can be used to show that hyperbolic special groups are **QCERF** or **quasiconvex extended residually**

finite meaning that if G is hyperbolic and special, then every quasiconvex subgroup of G is separable.

A similar result exists in the relatively hyperbolic case. One might imagine that replacing the quasiconvex subgroup H by a relatively quasiconvex subgroup might yield a generalization; however, some care is required. Consider the following example:

Example 5.3.2. *Take the standard action of $\mathbb{Z}^2 = \langle(1,0),(0,1)\rangle$ on $\mathbb{R}^2$ by translation. The diagonal $D := \{(r,r) : r \in \mathbb{R}\}$ is a subspace stabilized by $L := \langle(1,1)\rangle \leq \mathbb{Z}^2$. The subgroup L is $(2,0)$-quasi-isometrically embedded in the given presentation of $\mathbb{Z}^2$, but the convex hull of D is all of $\mathbb{R}^2$.*

A relatively hyperbolic group G with virtually abelian peripheral subgroups of rank at least 2 acting geometrically on a CAT(0) cube complex $\widetilde{X}$ could have a cyclic parabolic subgroup that stabilizes a proper subset of a flat stabilized by a peripheral subgroup of G in $\widetilde{X}$.

Full relatively quasiconvex subgroups eliminate these pathologies:

Definition 5.3.3. *Let $(G,\mathcal{P})$ be a relatively hyperbolic group pair and let H be a relatively quasiconvex subgroup of G. The subgroup H is a **full relatively quasiconvex subgroup** of G if for each $g \in G$ and $P \in \mathcal{P}$, either $gPg^{-1} \cap H$ is finite or $gPg^{-1} \cap H$ is finite index in gPg^{-1}.*

By restricting to full quasicovex subgroups, Sageev and Wise successfully generalized Theorem 5.3.1.

Theorem 5.3.4 ([28, Theorem 1.1])**.** *Let X be a compact non-positively curved cube complex with $G = \pi_1(X)$ hyperbolic relative to subgroups $P_1, \dots, P_n$. Let $\widetilde{X}$ be the*

CAT(0) universal cover of X. If H is a full relatively quasiconvex subgroup of G, then for any compact $U \subseteq \widetilde{X}$, then there exists an H-cocompact convex subcomplex $\widetilde{Y} \subseteq \widetilde{X}$ with $U \subseteq \widetilde{Y}$.

By Proposition 2.3.21, if $(G, \mathcal{P})$ is a relatively hyperbolic group pair, the elements of $\mathcal{P}$ and their conjugates are relatively quasiconvex. By Proposition 2.3.11, the elements of $\mathcal{P}$ and their conjugates are fully relatively quasiconvex. Therefore:

Lemma 5.3.5. *Let X be a non-positively curved cube complex with CAT(0) universal cover $\widetilde{X}$ and $G := \pi_1(X)$. Let $(G, \mathcal{P})$ be a relatively hyperbolic pair. Let $x \in \widetilde{X}$ be a base point in the universal cover. For each $P \in \mathcal{P}$, there exists a $Z'(P, x)$ such that $Z'(P, x)$ is a P-cocompact convex subcomplex of $\widetilde{X}$ containing x.*

It follows immediately that there exists a $Q \geq 0$ such that the convex hull of Px is contained in $\mathcal{N}_Q(Px)$.

CHAPTER 6

A MALNORMAL QUASICONVEX FULLY $\mathcal{P}$-ELLIPTIC HIERARCHY

6.1 Relative Fellow Traveling

For this subsection, fix a relatively hyperbolic pair $(G, \mathcal{P})$. In [18] Proposition 4.1.6, Hruska and Kleiner proved a version of quasigeodesic stability called the **relative fellow traveling property** for CAT(0) spaces with isolated flats. When G acts geometrically on a CAT(0) space with isolated flats, the peripheral subgroups $\mathcal{P}$ are virtually free abelian of rank at least 2. The remainder of this subsection consists of several lemmas to show in Theorem 6.1.11 that when a relatively hyperbolic group (with arbitrary peripheral subgroups) acts geometrically on a CAT(0) space $\widetilde{X}$ and orbits of peripheral cosets are uniformly quasiconvex, then $\widetilde{X}$ also has the **relative fellow traveling property**:

Definition 6.1.1. *Let $\widetilde{X}$ be a CAT(0) space and let G act geometrically on $\widetilde{X}$ with basepoint $x \in \widetilde{X}$. If for all $\lambda \geq 1$ and $\epsilon \geq 0$, there exists $\ell(\lambda, \epsilon) \geq 0$ such that for all pairs of (λ, ϵ)-quasigeodesics $\gamma : [a, b] \to \widetilde{X}$ and $\gamma' : [a', b'] \to \widetilde{X}$ with the same endpoints, there exist partitions:*

$$a \leq s_0 \leq s_1 \leq \ldots \leq s_m \leq b \text{ and } a' \leq t_0 \leq t_1 \leq \ldots t_m \leq b'$$

such that for all i, $d(\gamma(s_i), \gamma'(t_i)) \leq \ell$ and:

1. *either $d_{Haus}(\gamma((s_i, s_{i+1})), \gamma'(t_i, t_{i+1})) \leq \ell$ or*
2. *$\gamma((s_i, s_{i+1})), \gamma'(t_i, t_{i+1}) \subseteq \mathcal{N}_\ell(g_i P_i x)$,*

then $\widetilde{X}$ has the ***relative fellow traveling property relative to*** $\{gPx \mid g \in G,\, P \in \mathcal{P}\}$.

For the remainder of this subsection, assume G is acting geometrically on a CAT(0) space $\widetilde{X}$ with base point $x \in \widetilde{X}$. Drutu and Sapir's Lemma 8.10 of [10] gives an analogy of quasigeodesic stability for the Cayley graph of a relatively hyperbolic group.

Definition 6.1.2. *Let S be a finite generating set for G and let $\lambda \geq 1$ and $\epsilon \geq 0$. Let γ be a (λ, ϵ)-quasigeodesic in $\Gamma(G,S)$. For any $\mu \geq 0$ and any quasigeodesic, let* ***the μ-saturation of*** *γ in $\Gamma(G,S)$ be*

$$CaySat_\mu(\gamma) := \bigcup \{gP \,|\, g \in G,\ P \in \mathcal{P},\ \gamma \cap \mathcal{N}_\mu(gP) \neq \varnothing\}.$$

Lemma 6.1.3 ([10, Lemma 8.10])**.** *Let G, S be as in Definition 6.1.2. Fix $\lambda \geq 1$ and $\epsilon \geq 0$. There exists $\mu_0(\lambda, \epsilon)$ such that for all $\mu \geq \mu_0$, there exists $D = D(\mu, \lambda, \epsilon)$ such that for all (λ, ϵ) quasigeodesics γ, γ' in $\Gamma(G,S)$ with the same endpoints:*

$$\gamma' \subseteq \mathcal{N}_D(\gamma) \cup \mathcal{N}_D(CaySat_\mu(\gamma)).$$

A similar statement holds in a space being acted upon geometrically by G:

Definition 6.1.4. *Let G act geometrically on a CAT(0) space $\widetilde{X}$ with base point $x \in \widetilde{X}$. Let $\lambda \geq 1$ and $\epsilon \geq 0$ and let γ be a (λ, ϵ)-quasigeodesic in $\widetilde{X}$. Define*

$$Sat_\mu(\gamma) := \bigcup \{gPx \,|\, g \in G,\ P \in \mathcal{P},\ \gamma \cap \mathcal{N}_\mu(gP) \neq \varnothing\}$$

Corollary 6.1.5. *Let $G, \widetilde{X}$ and $x \in \widetilde{X}$ be as in Definition 6.1.4. Fix $\lambda \geq 1$ and $\epsilon > 0$. There exists $\mu_0(\lambda, \epsilon) \geq 0$ such that for all $\mu \geq \mu_0$, there exists a $D \geq 0$ depending on λ, μ, ϵ and the action of G on $\widetilde{X}$ such that if γ, γ' are (λ, ϵ)-quasigeodesics in $\widetilde{X}$ with the same endpoints, then*

$$\gamma' \subseteq \mathcal{N}_D(\gamma) \cup \mathcal{N}_D(Sat_\mu(\gamma)).$$

The corollary follows immediately from the fact that there is a quasi-isometry $\Gamma(G,S) \to \widetilde{X}$.

These two statements and Theorem 2.3.19 give some notion of quasigeodesic stability for a CAT(0) space with a geometric action by a relatively hyperbolic group.

The first lemma (together with Corollary 6.1.5) says that when γ and γ' are (λ, ϵ)–quasigeodesics, if $gP \neq g'P'$ are two peripheral cosets and γ' enters gPx, then γ' cannot enter $g'P'x$ without passing within a bounded distance of γ where the bound depends only on $\widetilde{X}$, the action of G and the quasigeodesic constants λ, ϵ.

Lemma 6.1.6. *Fix $\lambda \geq 1$ and $\epsilon \geq 0$ and let μ, D be the constants specified by λ, ϵ and Corollary 6.1.5. Suppose there exists $L' \geq 0$ so that for every peripheral coset gP, gPx is L'-quasiconvex in $\widetilde{X}$. Then for all (λ, ϵ)-quasigeodesics γ, γ' with the same endpoints, there exists $D'_0 = D'_0(G, \mathcal{P}, \lambda, \epsilon, \widetilde{X})$ such that for all $D' \geq D'_0$, if $\gamma'(t) \in \mathcal{N}_{D+\mu+\lambda+\epsilon}(g_1P_1x) \cap \mathcal{N}_{D+\mu+\lambda+\epsilon}(g_2P_2x)$ where g_1P_1 and g_2P_2 are peripheral cosets in the μ-saturation of γ, then $\gamma'(t) \in \mathcal{N}_{D'}(\gamma)$.*

Proof. Since g_1P_1x and g_2P_2x are in the μ saturation of γ, there exist points $a \in \gamma \cap \mathcal{N}_\mu(g_1P_1x)$ and $b \in \gamma \cap \mathcal{N}_\mu(g_2P_2x)$. Let $\triangle$ be a quasigeodesic triangle with sides $[\gamma'(t), a] \subseteq \mathcal{N}_{L'+D+\mu+\lambda+\epsilon}(g_1P_1x)$, $[\gamma'(t), b] \subseteq \mathcal{N}_{L'+D+\mu+\lambda+\epsilon}(g_2P_2x)$ and $(ab) \subseteq \gamma$ where (ab) is a subpath of γ between a and b.

The quasigeodesic triangle $\triangle$ δ-thin relative to some gPx by Theorem 2.3.19 (and the fact that $\widetilde{X}$ is CAT(0) in that:

1. there exists a point $p \in \widetilde{X}$ that is $\frac{\delta}{2}$ from all three sides of $\triangle$ or
2. there exist **corner segments** that are δ-fellow traveling subsegments of the geodesic sides of $\triangle$ at $\gamma'(t)$ and **fat segments** that lie in each of the geodesic

sides so that each fat segment has one endpoint on a corner segment and one endpoint that is distance δ from (ab).

The length of the corner segments is at most $A := \max \operatorname{diam}(\mathcal{N}_{\delta+L'+D+\mu+\lambda+\epsilon}(g_1P_1x) \cap \mathcal{N}_{\delta+L'+D+\mu+\lambda+\epsilon}(g_2P_2x))$ and at least one of the fat segments has length at most A in gPx because at least one of g_1P_1 and g_2P_2 does not equal gP. By Theorem 2.3.19, a path of length $2A+\delta$ from $\gamma'(t)$ to (ab), the third side of the quasigeodesic triangle $\triangle$, and $(ab) \subseteq \gamma$, so $d(gPx, \gamma) \leq 2A + \delta$. Setting $D_0' = 2A+\delta$ gives the desired bound. □

If $\gamma'(t)$ lies near a peripheral coset orbit gPx, then there is an interval I containing t such that the endpoints of $\gamma'(t)$ are close to γ and the remaining points lie in a uniformly bounded neighborhood of a single peripheral coset orbit:

Lemma 6.1.7. *Fix $\lambda \geq 1$ and $\epsilon \geq 0$. Choose sufficient μ, D, D' with $D' \geq D$ that make the conclusions of Corollary 6.1.5 and Lemma 6.1.6 hold. Let $L = D' + \mu + \lambda + \epsilon$ and let γ, γ' be continuous (λ, ϵ)–quasigeodesics with the same endpoints.*

If $\gamma'(t_0) \in \mathcal{N}_{D+\mu}(gPx)$, then the interval $I := [t_-, t_+]$ where $t_- := \inf\{t \leq t_0 : \gamma'(t) \subseteq \mathcal{N}_{D+\mu}(gPx)\}$ and $t_+ := \sup\{t \geq t_0 : \gamma'(t) \subseteq \mathcal{N}_{D+\mu}(gPx)\}$ has the following properties:

1. *$t_0 \in I$,*
2. *$\gamma'(I) \subseteq \mathcal{N}_{D+\mu}(gPx)$,*
3. *$\gamma'(t_-), \gamma'(t_+) \in \mathcal{N}_L(\gamma)$.*

Proof. By construction $t_0 \in I$. By continuity, $\gamma'(I) \subseteq \mathcal{N}_{D+\mu}(gPx)$.

For some small fixed $0 < \rho < 1$, $\gamma'(t_+ + \rho) \notin \mathcal{N}_{D+\mu}(gPx)$. If $\gamma'(t_+ + \rho) \in \mathcal{N}_{D+\mu}(\gamma)$, then $\gamma'(t_+) \in \mathcal{N}_{D+\mu+\lambda+\epsilon}(\gamma)$. Otherwise by Corollary 6.1.5, $\gamma'(t_+ + \rho) \in$

$\mathcal{N}_{D+\mu}(g'P'x)$ for some peripheral coset $g'P' \neq gP$. Therefore, $\gamma'(t_+ + \rho) \in \mathcal{N}_{D+\mu+\lambda+\epsilon}(gPx) \cap \mathcal{N}_{D+\mu+\lambda+\epsilon}(g'P'x)$. By Lemma 6.1.6, then $\gamma'(t_+ + \rho) \in \mathcal{N}_{D'}(\gamma)$, so $\gamma'(t_+) \in \mathcal{N}_{D'+\mu+\lambda+\epsilon}(\gamma) = \mathcal{N}_L(\gamma)$. Similarly, $\gamma'(t_-) \in \mathcal{N}_L(\gamma)$. □

The following lemma shows how to partition the domain of γ' for relative fellow traveling:

Lemma 6.1.8. *Let L, D, D', μ be as in Lemma 6.1.7. Let $\gamma' : [a', b'] \to \widetilde{X}$ be a (λ, ϵ)-quasigeodesic, let γ be a (λ, ϵ)-quasigeodesic with the same endpoints and let $R > 0$, then there exists $L_0 = L_0(R, \lambda, \epsilon) \geq 0$ and a partition $a' \leq t_0 \leq t_1 \leq \ldots \leq t_m \leq b'$ such that*

1. *$\gamma(t_j) \in \mathcal{N}_L(\gamma)$ for all j,*
2. *$|t_{2i} - t_{2i+1}| \geq R$,*
3. *$\gamma'([t_{2i}, t_{2i+1}]) \subseteq \mathcal{N}_{D+\mu}(g_iP_ix)$ for some peripheral coset g_iP_i, and*
4. *if $t_{2i+1} \leq t_- \leq t_+ \leq t_{2i+2}$, $|t_- - t_+| \geq R$, then $\gamma'([t_-, t_+]) \notin \mathcal{N}_{D+\mu}(gPx)$ for all peripheral cosets gP,*
5. *$\gamma'([a, t_0]), \gamma'([t_{2i}, t_{2i+1}]), \gamma'([t_m, b]) \subseteq \mathcal{N}_{L_0}(\gamma)$.*

Further, there exist s_j so that

1. *$d(\gamma(s_j), \gamma'(t_j)) \leq L$,*
2. *$|s_j - s_{j+1}| \geq \frac{(|t_j - t_{j+1}| - 2L - \epsilon)}{\lambda}$.*

Proof. Let $t_{-1} := a'$. Partition $[a, b]$ inductively as follows:

- let $t_{2j} = \inf\{t \in [t_{2j-1}, b] : \gamma'([t, t+R]) \subseteq \mathcal{N}_{D+\mu}(gPx)$ for some $gPx\}$, halt if no such t_j exists,

- let $t_{2j+1} = \sup\{t \leq b' : \gamma'([t, t+R]) \subseteq \mathcal{N}_{D+\mu}(gPx)\}$.

Let m be the largest subscript for which t_j was determined and let $t_{m+1} := b'$. For each t_{2j+1}, let I_{2j} be the interval specified by Lemma 6.1.7 with $\gamma'(t_{2j+1}) \in \mathcal{N}_{D+\mu}(gPx)$. Immediately, $\gamma'(t_{2j+1}) \in \mathcal{N}_L(gPx)$. Either t_{2j} is the left hand endpoint of I_{2j} in which case $\gamma'(t_{2j}) \in \mathcal{N}_L(gPx)$ by Lemma 6.1.7, or $t_{2j} = t_{2j-1}$, so $t_{2j} \in \mathcal{N}_L(gPx)$ by Lemma 6.1.6. Either way, by construction $\gamma'([t_{2j}, t_{2j+1}]) \subseteq \mathcal{N}_{D+\mu}(g_jP_jx)$ for some peripheral coset g_jP_j.

Observe that $[t_{2i+1}, t_{2i+2}]$ cannot contain any length at least R subintervals whose image under γ' lies in $\mathcal{N}_{D+\mu}(gPx)$ for some peripheral coset gP because otherwise that subinterval would have t_{2i+1} as its left endpoint.

Let $L_0 := L + R\lambda + \epsilon$. For $t \in [t_{2i+1}, t_{2i+2}]$, if $\gamma(t) \notin \mathcal{N}_{D+\mu}(gP)$, then there exists t' such that $|t - t'| \leq R$ and $\gamma'(t') \in \mathcal{N}_L(\gamma)$ by Lemma 6.1.7, so $\gamma'(t) \in \mathcal{N}_{L+R\lambda+\epsilon}(\gamma) = \mathcal{N}_{L_0}(\gamma)$ because γ' is a $(\lambda, \epsilon)-$quasigeodesic.

The final two assertions follow immediately from the fact that $\gamma'(t_-)$, $\gamma'(t_+) \in \mathcal{N}_L(\gamma)$ and the fact that γ is a $(\lambda, \epsilon)-$quasigeodesic. $\square$

Note that the constant L_0 can be made arbitrarily larger if desired.

The next two lemmas are devoted to showing that when γ' is (λ, ϵ)-geodesic and γ is geodesic, and γ' does not have a long subpath in any peripheral coset orbit, then they lie within a bounded Hausdorff distance of each other.

Lemma 6.1.9. *Fix $\lambda \geq 1, \epsilon \geq 0$. Suppose there exists an $L' \geq 0$ so that for any peripheral coset gP, the orbit gPx is L'-quasiconvex in $\widetilde{X}$. Let $\delta \geq 0$ be sufficiently large so that Theorem 2.3.19 holds, and let $D+\mu$ be as before, but possibly enlarged so that $D+\mu \geq \delta$. Suppose there exists $R = R(D, \mu, \lambda, \epsilon)$ such that whenever $a((t_-, t_+)) \subseteq \mathcal{N}_{D+\mu}(gPx)$,*

$|t_- - t_+| \le R$. Let a be a continuous (λ, ϵ)–quasigeodesic, let b be a geodesic with $|b| \le L$ and let c be geodesic such that $\triangle := \triangle abc$ is a quasigeodesic triangle.

Then there exists S not depending on the choice of (λ, ϵ)-quasigeodesic triangle $\triangle$ such that if $y \in c$, $d(y, a) \le S$. In other words, $c \subseteq \mathcal{N}_S(a)$.

Proof. If $|c| \le \delta$, then $S = \delta \le D + \mu + R\lambda + \epsilon + 4\delta$ suffices.

Now suppose $S = D + \mu + R\lambda + \epsilon + 4\delta$ suffices for $|c| \le (n\delta)$ for $n \in \mathbb{N}$.

Suppose $|c| \le (n+1)\delta$. Assume c is parameterized $c : [0, \alpha] \to \widetilde{X}$ with $c(0)$ on a, $\alpha > 0$ and $c(\alpha)$ on b. Let $P = \{t \in [0, \alpha] : d(c(t), a) = L\}$.

There are four cases:

1. If $P = \{\alpha\}$, since $d(c(0), a) = 0$, then for all $y \in c$, $d(y, a) \le L$ by continuity.
2. There exists $q \in P \cap [\alpha - 2\delta, \alpha - \delta]$. Then $S = L + R\lambda + \epsilon + 4\delta$ suffices because every point in $c([0, q])$ lies in an S neighborhood of a by the assumption and $c([q, \alpha])$ is in an $L + 2\delta$ neighborhood of a because $|b| \le L$.
3. $P \cap [\alpha - 2\delta, \alpha) = \emptyset$. This case requires some additional argument. See below.
4. There exists $q \in P \cap (\alpha - \delta, \alpha)$ and $P \cap [\alpha - 2\delta, \alpha - \delta] = \emptyset$. For $t \in [q, \alpha]$, $d(c(t), a) \le \delta + L$ because $|b| \le L$. If $t \in [0, q)$, consider the triangle formed by $c|_{[0,q]}$, the geodesic of length L connecting c to a and the subsegment of a connecting their other two endpoints. This is a triangle satisfying one of the three preceding cases, so $d(c(t), a) \le L + R\lambda + \epsilon + 4\delta$.

Thus proving $S = L + R\lambda + \epsilon + 4\delta$ suffices in the third case is all that remains. Choose $0 < t \le \alpha - 2\delta$ so that $t = \sup\{0 \le t' \le \alpha - 2\delta \,|\, c(t') \in P\}$. Then for all

$y \in c([0,t])$, $d(y,a) \leq S$ because $|t| \leq \alpha - 2\delta \leq (n-1)\delta$. Let $r \in (t,\alpha)$. The goal is to show $d(c(r),a) \leq L + R\lambda + \epsilon + 4\delta$. By Theorem 2.3.19 either:

1. there exists a point $z \in \widetilde{X}$ such that z is $\frac{\delta}{2}$ from each side of $\triangle$ or
2. each side a,b,c of $\triangle$ has a subpath $a',b',c' \subseteq N_\delta(gPx)$ for some peripheral coset gP such that the terminal endpoint of c' closer to $c(0)$ is close to one endpoint of a' and the terminal endpoint of c' close to $c(\alpha)$ is within $2\delta + L$ of the other endpoint of c' (because $|b'| \leq L$).

In the first case, let $s \in (0,\alpha)$ so that $d(c(s),b) \leq \delta$, and $d(c(s),a) \leq \delta$. If $t \leq r \leq s \leq \alpha$, then $d(c(s),\alpha) \leq \delta \leq L$, so $d(c(r),a) \leq L$ by continuity because no points of $c((t,\alpha))$ are distance L from a. If $s \leq r$, since $\widetilde{X}$ is CAT(0), $d(c(t),b) \leq \delta$, so $d(c(t),a) \leq L + \delta$.

In the second case, let $c' = [c(t_-),c(t_+)]$ so that $0 \leq t_- \leq t_+ \leq \alpha$ and $d(c(t_+),b) \leq \delta$. When $r \geq t_+$, let $s = t_+$ and the argument is similar to the previous case where $s \leq r$. When $t \leq r \leq t_-$, then $d(c(t_-),a) \leq \delta$ and the argument is similar to the previous case where $s \geq r$. On the other hand if $t_- \leq r \leq t_+$, by hypothesis the distance between the endpoints of $a' \subseteq \mathcal{N}_\delta(gPx)$ is at most $\lambda R + \epsilon$ and by Lemma 2.3.22, $|c'| \leq \lambda R + \epsilon + 3\delta + T$ because $|b'| \leq T$. Therefore $|t_- - t_+| \leq \lambda R + \epsilon + 3\delta + L$, and $d(c(t_-),a) \leq \delta$, so $d(c(t),a) \leq \lambda R + \epsilon + 4\delta + L$. Therefore, by induction on n, the statement holds. □

Lemma 6.1.10. *Let $\mu, D, D' \geq 0$ be large enough so that the conclusions of Lemma 6.1.6, Lemma 6.1.7 and Lemma 6.1.9 hold for (λ,ϵ)-quasigeodesics. Let gPx be $L' \geq 0$-quasiconvex in $\widetilde{X}$ for all peripheral cosets gP. Let $L \geq 0$ satisfy the conclusions of Lemma 6.1.7 for (λ,ϵ) quasigeodesics, let $\gamma : [0,a] \to \widetilde{X}$ be a geodesic and let $\gamma' : [0,b] \to \widetilde{X}$ be a (λ,ϵ)–quasigeodesic such that $d(\gamma'(0),\gamma(0))$, $d(\gamma'(b),\gamma(a)) \leq L$.*

Suppose there exists $R \geq 0$ such that whenever I is some interval and either $\gamma'(I)$ or $\gamma(I) \subseteq \mathcal{N}_{D+\mu}(gPx)$ for some peripheral coset gP, $|I| \leq R$. Then there exists S not depending on the choice of γ or γ' such that $d_{Haus}(\gamma, \gamma') \leq S$.

Proof. Since $\widetilde{X}$ is CAT(0), γ and the geodesic connecting the endpoints of γ' are Hausdorff distance L apart. Therefore, it suffices to assume γ and γ' have the same endpoints.

Let $p = \gamma'(t)$. Let D' (depending on D, μ) be as in the conclusion of Lemma 6.1.6. By Corollary 6.1.5, either $d(p, \gamma) \leq D$ or $p \in \mathcal{N}_{D+\mu}(gPx)$ for some peripheral coset gP. In the latter case, by Lemma 6.1.6, there exists t_0 such that $|t - t_0| \leq R + 1$ and $d(\gamma'(t_0), \gamma) \leq D'$. Therefore, $d(p, \gamma) \leq D + \lambda(R + 1) + \epsilon + D'$ because γ' is (λ, ϵ)–quasigeodesic.

On the other hand suppose there exists a point $q \in \gamma$ such that $d(q, \gamma') \geq L$. Let $q' \in \gamma'$ with $d(q, q') \geq L_0$. Then the geodesic $\psi := [q, q']$ splits the quasigeodesic bigon formed by γ and γ' into two quasigeodesic triangles satisfying the hypotheses of Lemma 6.1.9. Therefore, there exists S' such that for every $y \in \gamma$, then $d(y, \gamma') \leq S'$. □

Theorem 6.1.11 shows that $\widetilde{X}$ has the relative fellow traveler property. Theorem 6.1.11 is presumed to be known based on the works of [10], [17] and others, but the exact formulation used here proved difficult to find in the literature. The preceding lemmas are combined to partition the domain of γ' into intervals where γ' is close to γ and where γ' lies in a neighborhood of some peripheral coset orbit. The endpoints of each interval have images close to γ which suggests a way to partition γ using the projections of the endpoints of each interval onto γ. However, the projections of these points may not appear in the correct order, so some adjustments will need to be made.

Theorem 6.1.11. *Let gPx be L'-quasiconvex in $\widetilde{X}$ for every peripheral coset gP.*

For all $\lambda \geq 1$ and $\epsilon > 0$ there exists $\ell = \ell(\lambda, \epsilon) \geq 0$ such that for all pairs of (λ, ϵ)-quasigeodesics $\gamma : [a, b] \to \widetilde{X}$ and $\gamma' : [a', b'] \to \widetilde{X}$ with the same endpoints, there exist partitions:

$$a \leq s_0 \leq s_1 \leq \ldots \leq s_m \leq b \text{ and } a' \leq t_0 \leq t_1 \leq \ldots t_m \leq b'$$

such that for all i, $d(\gamma(s_i), \gamma'(t_i)) \leq \ell$ and:

1. *either $d_{Haus}(\gamma((s_i, s_{i+1})), \gamma'((t_i, t_{i+1}))) \leq \ell$ or*
2. *$\gamma((s_i, s_{i+1})), \gamma'(t_i, t_{i+1}) \subseteq \mathcal{N}_\ell(g_i P_i x)$.*

Proof. It suffices to assume that γ, γ' are continuous quasigeodesics because γ and γ' are each within a fixed Hausdorff distance (depending only on λ, ϵ) of continuous quasigeodesics γ_0, γ_0' with the same endpoints whose quasigeodesic constants depend only on (λ, ϵ) (see [9] Lemma III.H.1.11).

Let L, L_0, D, μ be as in Lemma 6.1.7 and Lemma 6.1.8 and let D' satisfy the conclusions of Lemma 6.1.6. Let $M \geq 0$ be a constant so that

$$M \geq \max_{gP, g'P' \text{ peripheral cosets}} \operatorname{diam}(\mathcal{N}_{D+D'+100L_0+L'+1}(gPx) \cap \mathcal{N}_{D+D'+100L_0+L'+1}(g'P'x),)$$

and so that $M > 2\lambda L_0 + \epsilon$.

First assume γ is geodesic. Let:

$$a' \leq t_0' \leq t_1' \leq t_2' \leq \ldots \leq t_m' \leq b'$$

be the partition specified by Lemma 6.1.8 with $R = 3\lambda M + \epsilon + 2\lambda L$. F

By Lemma 6.1.8, there exist $a \leq s'_1, s'_2, \dots, s'_m \leq b$ such that $d(\gamma'(t'_i), \gamma(s'_i)) \leq L$ and $d(s'_{2i}, s'_{2i+1}) \geq 3M$.

By Lemma 6.1.8, whenever $|t_- - t_+| \geq 3\lambda M + \epsilon + 2L$, then $\gamma'((t_-, t_+)) \notin \mathcal{N}_{D+\mu}(gPx)$ for all peripheral cosets gP.

By Lemma 6.1.10, there exists $S \geq 0$ such that $d_{\text{Haus}}(\gamma|_{(t'_{2i-1}, t'_{2i})}, \gamma'|_{(s'_{2i-1}, s'_{2i})}) \leq S$. The L'-quasiconvexity of g_iP_ix implies that $\gamma((s'_{2i}, s'_{2i+1}))$ lies in $\mathcal{N}_{L'+2L}(g_iP_ix)$.

This nearly completes the proof; however, there is the possibility that $j > i$ but that $s'_j < s'_i$.

Then there exist $i < k_1 < k_2$ where $k_2 = k_1 + 1$ so that $s'_i \in [s'_{k_1}, s'_{k_2}]$. There are three possibilities:

1. **Case:** $\gamma([s'_{k_1}, s'_{k_2}])$ **lies in** $\mathcal{N}_L(\gamma'([t_{k_1}, t_{k_2}]))$. Then $|t_{k_1} - t_i| \leq \lambda(2L) + \epsilon$ because there exists some $t_{k_1} \leq t \leq t_{k_2}$ with $d(\gamma'(t), \gamma(s_i)) \leq L$ so that $d(\gamma'(t_i), \gamma'(t)) \leq 2L$ and $t_i \leq t_{k_1} \leq t$. Therefore, by construction, i is odd, $j = i + 1$ and $|s_i - s_j| \leq \lambda^2(2L) + \lambda\epsilon + \epsilon$.
2. **Case:** $\gamma([s'_{k_1}, s'_{k_2}])$ **lies in** $\mathcal{N}_{D+\mu}(gPx)$ **where** gP **is a peripheral coset. Either** $\gamma([s'_{i-1}, s'_i])$ **or** $\gamma([s'_i, s'_{i+1}])$ **lies in** $\mathcal{N}_{2L+L'+D+\mu}(hPx)$ **where** hP **is a peripheral coset and** $hP \neq gP$**.** If $gP \neq hP$, then either $|s'_i - s'_{k_1}| \leq M$ or $|s'_i - s'_{k_2}| \leq M$ because M bounds the diameter of $\mathcal{N}_{2L+L'+D+\mu}(gPx) \cap \mathcal{N}_{2L+L'+D+\mu}(hPx)$. In either case, there exists $t \in \{t'_{k_1}, t'_{k_2}\}$ with $d(\gamma'(t), \gamma(s'_i)) \leq L + M$ so that $|t'_{k_1} - t'_i| \leq \lambda(2L + M) + \epsilon$. As in the previous case, then i is odd, $j = i + 1$ and $|s'_i - s'_j| \leq \lambda^2(2L + M) + \lambda\epsilon + \epsilon$.
3. **Case: same as case** (2) **except that** $gP = hP$**.** If $gPx = hPx$, then $\gamma([s'_i, s'_{k_2}]) \subseteq \mathcal{N}_{L+L'}(gPx)$, so whenever $i \leq 2k_3 \leq 2k_3 + 1 \leq k_2$, then

$g_{k_3}P_{k_3} = gPx$ because $\gamma|_{[s'_{k_3}, s'_{k_3+1}]}$ must have a length $3M$ subsegment in $\mathcal{N}_L(g_{k_3}P_{k_3})$. Therefore, $\gamma'([t'_i, t'_{k_2}]) \subseteq \mathcal{N}_{2L+L'+S}(gPx)$.

The partitions $a' \leq t_0 \leq t'_1 \leq \ldots \leq t'_m \leq b'$ and $a \leq s'_0, s'_1, \ldots, s'_m \leq b$ can now be reworked. First set $s_0 := s'_0$ and $t_0 := t'_0$. Given that s_{2i} was set equal to some s'_{2j}, set s_{2i+1}, s_{2i+2}, t_{2i+1} and t_{2i+2} as follows:

1. If $g_{j+1}P_{j+1} = g_jP_j$, reset $s_{2i} := s'_{2j+2}$ and $t_{2i} := s'_{2j+2}$ and repeat this process to determine s_{2i+1} and s_{2i+2}.
2. If $2j = m - 1$, then set $s_{i+1} := s'_m$, set $t_{i+1} := t'_m$ and stop.
3. If $s'_{2j+1} \geq s'_{2j}$, set $s_{2i+1} := s'_{2j+1}$, set $s_{2i+2} := s'_{2j+2}$, set $t_{2i+1} := t'_{2j+1}$ and set $t_{2i+2} := t'_{2j+2}$.
4. If $s'_{2j+1} < s'_{2j}$, set $s_{2i+1} := s'_{2j}$, set $s_{2i+2} := s'_{j+2}$, set $t_{2i+1} := t'_{2j+1}$ and set $t_{2i+2} = t'_{2j+2}$.

First observe that $s_{2i+2} \geq s_{2i}$. Indeed, if $s'_{2j+2} \leq s'_{2j}$, then by the argument in case (3) above, there exist some s'_{2k}, s'_{2k+1} where $2k > 2j$ such that $g_kP_k = g_jP_j$. However, then g_jP_j, $g_{j+1}P_{j+1}, \ldots, g_kP_k$ are all equal, so step (1) precludes $s_{2i} > s_{2i+2}$.

By construction and cases (1) and (2) above, $[s_{2i+1}, s_{2i+2}]$ is either a point s'_{2j} where $|s'_{2j} - s'_{2j+1}| \leq \lambda^2(2L + M) + \epsilon$ or an interval $[s'_{2j}, s'_{2j+1}]$. In either case, $d_{\text{Haus}}(\gamma([s_{2i}, s_{2i+1}]), \gamma'([t_{2i}, t_{2i+1}])) \leq L + \lambda^2(2L + M) + \epsilon$. By the argument in Case (3) above and Step (4) in the selection of s_{2i+1} and t_{2i+1}, both $\gamma([s_{2i}, s_{2i+1}]) \subseteq \mathcal{N}_{2L+L'+S+L+\lambda^2(2L+M)+\epsilon}(gPx)$ and $\gamma'([t_{2i}, t_{2i+1}]) \subseteq \mathcal{N}_{2L+L'+S+L+\lambda^2(2L+M)+\epsilon}(gPx)$ for some (common) peripheral coset gP. Thus setting $\ell := 2L+L'+S+L+\lambda^2(2L+M)+\epsilon$ completes the case that γ is geodesic.

The case where both γ and γ' are quasigeodesics can be handled by taking ρ to be the geodesic joining the common endpoints of γ and γ' and using the case where one path is geodesic on the pairs (γ, ρ) and (γ', ρ). Let $M' = \max \operatorname{diam} \mathcal{N}_{2L+L'+S+L+\lambda^2(2L+M)+\epsilon}(gPx) \cap \mathcal{N}_{2L+L'+S+L+\lambda^2(2L+M)+\epsilon}(g'P'x)$ where $gP \neq g'P'$ are peripheral cosets. After repartitioning γ, γ' and ρ similarly, the constant $\ell := 2(2L + L' + S + L + \lambda^2(2L + M) + \epsilon + M')$ suffices. □

6.2 Superconvexity

If $\widetilde{X}$ is a δ-hyperbolic metric space then there exists a $R = R(\lambda, \epsilon)$ such that if $\alpha, \beta \subseteq \widetilde{X}$ are biinfinite (λ, ϵ)quasigeodesics, and for some $D > 0$, $d_{\text{Haus}}(\alpha, \beta) \leq D$, then $d_{\text{Haus}}(\alpha, \beta) < R$. A similar statement holds if α is replaced by a quasiconvex subset of $\widetilde{X}$. **Superconvexity** is a generalization of this property for a general geodesic metric space. A superconvex biinfinite geodesic is a biinfinite geodesic α with a fixed $R = R(\alpha)$ so that every biinfinite geodesic β that lies in some bounded Hausdorff distance of α has $d_{\text{Haus}}(\alpha, \beta) \leq R$. The notion of superconvexity will be extended in Definition 6.2.1 to certain immersions whose elevations to the universal cover exhibit a similar property.

Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let X be a NPC compact cube complex with $\pi_1 X := G$. Recall $\{Z'(P, X) | P \in \mathcal{P}\}$, a collection of proper P-cocompact convex subcomplexes of $\widetilde{X}$, the CAT(0) unviersal cover of X, constructed in Lemma 5.3.5.

The edge spaces of the hierarchy constructed in the subsequent subsections will contain the images of the immersed complexes $P \backslash Z'(P, x) \to X$. To obtain a fully $\mathcal{P}$-elliptic hierarchy (recall Definition 3.3.1) in the following construction, it will be necessary to enlarge the $Z'(P,x)$ to create a new convex subcomplex of $\widetilde{X}$ so that every element $g \in \pi_1 X$ conjugate in G into P has a positive power g^k that is an element of $P \cong \pi_1(P\backslash Z'(P,x))$. Geometrically, this means replacing $Z'(P,x)$ by some $Z(P,x) := \mathcal{N}_L(Z(P,x))$ for some $L \geq 0$ so that every biinfinite geodesic in $\widetilde{X}$ that lies in a finite Hausdorff neighborhood of $Z(P,x)$ lies in $Z(P,x)$. The image of an immersed complex and its accompanying immersion whose elevations to the universal cover have this property are called **superconvex**.

Definition 6.2.1. *Let X be a non-positively curved cube complex and let $\phi : Z \to X$ be a local isometry. The map ϕ is **superconvex** if for any elevation $\widetilde{\phi} : \widetilde{Z} \to \widetilde{X}$ of Z to the universal cover $\widetilde{X}$ of X and any bi-infinite geodesic γ in $\widetilde{X}$ such that $d_{Haus}(\gamma, \widetilde{Z})$ is bounded, then $\gamma \subseteq \widetilde{Z}$.*

If the immersion $\phi : Z \to X$ is superconvex, then Z is said to be superconvex in X (with respect to ϕ).

Proposition 6.2.2. *Let X, G, $(G, \mathcal{P})$, x and $Z'(P,x)$ be as in Lemma 5.3.5. Then there exists an $L \geq 0$ such that $Z(P,x) := \mathcal{N}_L(Z'(P,x))$ has a superconvex immersed quotient $\overline{Z}(P,x) := P\backslash Z(p,x)$.*

Proof. It suffices to prove that there exists $L \geq 0$ such that every bi-infinite geodesic within a fixed Hausdorff distance of $Z'(P,x)$ lies in $\mathcal{N}_L(Z'(P,x))$ because any elevation of $\overline{Z}(P,x)$ is a translate of $Z(P,x)$. Let γ be a bi-infinite geodesic in $\widetilde{X}$ such that $d_{\mathrm{Haus}}(\gamma, Z'(P,x)) \leq R$ for some $R \geq 0$. Recall that $\widetilde{X}$ has the δ-relatively thin triangle property for some $\delta > 0$ and by Theorem 6.1.11 $(\widetilde{X}, G, \mathcal{P})$ has the relative fellow traveler property for some $\sigma > 0$ relative to the peripheral coset orbits of x. Let a, b be points on γ and let $\alpha = [a,b] \subseteq \gamma$. Let $a', b' \in Z'(p,x)$ be closest points to a, b respectively. Consider the path

$\beta = [a,a'] \cup [a',b'] \cup [b',b]$. This path is a $(1,2R)$ quasigeodesic. Then there exist partitions:

$$0 = s_0 \le t_0 \le s_1 \le t_1 \le \ldots \le s_n \le t_n = \ell$$

$$0 = s'_0 \le t'_0 \le s'_1 \le t'_1 \le \ldots \le s'_n \le t'_n = \ell'$$

and $\sigma > 0$ such that $d_{\text{Haus}}(\alpha([s_i,t_i]), \beta([s'_i,t'_i])) \le \sigma$ and $\alpha([t_i,s_{i+1}]), \beta([t'_i,s'_{i+1}]) \subseteq \mathcal{N}_\sigma(g_i P_i x)$ for some $g_i \in G$ and $P_i \in \mathcal{P}$.

In particular, $d(\alpha(t_i),\beta(t'_i)), d(\alpha(s_{i+1}),\beta(s'_{i+1})) \le \sigma$. $\mathcal{N}_R(Z'(P,x))$ is convex, as is α, so by CAT(0) geometry, any point on γ lies within 2σ of the geodesic segment $\beta([t_i,s'_{i+1}]) \subseteq \mathcal{N}_R(Z'(P,x))$.

Therefore, setting $L = R + 2\sigma$ suffices. □

The immersed complexes $\overline{Z}(P,x)$ constructed in Proposition 6.2.2 are called **peripheral complexes**.

Taking a quotient $\overline{Z}(P,x) := P\backslash Z(P,x)$ gives an immersion $\phi_P : \overline{Z}(P,x) \to X$. There is a convenient way to upgrade the immersion to an embedding:

Definition 6.2.3. *Let X be a non-positively curved cube complex with CAT(0) universal cover $\widetilde{X}$ and $G := \pi_1(X)$. Let $(G,\mathcal{P})$ be a relatively hyperbolic pair. Let $\mathcal{Z} := \bigsqcup_{P\in\mathcal{P}} \overline{Z}(P,x)$. The* ***augmented cube complex for the pair*** *$(X,\mathcal{Z})$ is the complex:*

$$C(X,\mathcal{Z}) := (X \cup \bigsqcup_{P\in\mathcal{P}} \overline{Z}(P,x) \times [0,1]) / (\overline{Z}(P,x) \times \{1\}) \sim \phi_{P,x}(Z(P,x)),$$

the mapping cylinders of the $\phi_{P,x}$.

Note that every non-peripheral hyperplane is an extension of a hyperplane of X.

The hyperplanes $Z_i \times \frac{1}{2}$ are called ***peripheral hyperplanes*** *while the remaining hyperplanes of $C(X,\mathcal{P},x)$ are* ***non-peripheral****.*

The advantage of the augmented complex is that $\overline{Z}(P, x) \times \{0\}$ is now an embedded subcomplex of $C(X, \mathcal{P}, x)$. From a group theoretic standpoint, it is easy to see that $\pi_1 G \cong \pi_1(C(X, \mathcal{Z}))$, so a hierarchy for $\pi_1(C(X, \mathcal{Z}))$ determines a hierarchy of $\pi_1 G$. Moreover, if X is special, then $C(X, \mathcal{Z})$ is as well. Technically, the definition of $C(X, \mathcal{Z})$ depends on the base point, but for the group theoretic applications, this will just mean that the result is determined up to conjugacy.

Proposition 6.2.4. *Let $C(X, \mathcal{Z})$ be the augmented cube complex for the pair $(X, \mathcal{Z})$ described in Definition 6.2.3. Let $\widetilde{C}$ be the universal cover of $C(X, \mathcal{Z})$. Let $\mathcal{B}$ be the collection of elevations of components of $\mathcal{Z}$ to $\widetilde{C}$.*

Then there exists $M > 0$ so that $(\widetilde{C}, \mathcal{B})$ is a (δ, M) relatively hyperbolic pair.

Proof. By construction each $B \in \mathcal{B}$ is distance 1 from the convex hull of a peripheral coset orbit which lies in a uniformly bounded neighborhood of that peripheral coset orbit. Recall that for all $t \geq 0$, there exists $f(t)$ such that for all pairs of peripheral cosets gP, $g'P'$ $\operatorname{diam} \mathcal{N}_t(gPx) \cap \mathcal{N}_t(g'P'x) \leq f(t)$. Therefore, there exists M so that for all pairs $B_1, B_2 \in \mathcal{B}$, $\operatorname{diam} \mathcal{N}_{3\delta}(B_1) \cap \mathcal{N}_{3\delta}(B_2) \leq M$.

Similiarly, since triangles are δ_0-thin relative to the collection of peripheral coset orbits by Proposition 2.3.17, so triangles in $\widetilde{X}$ are $\delta := \delta_0 + 1$-thin relative to $\mathcal{B}$. □

6.3 The Double Dot Hierarchy

The construction of a hierarchy will use a finite cover called the **double dot cover**.

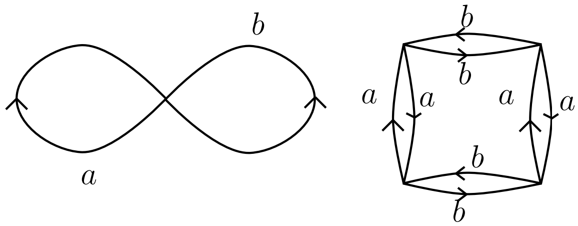

Figure 6.1: The figure 8 loop on the left whose two hyperplanes are the two edge midpoints and the double dot cover of the figure 8 loop on the right.

Definition 6.3.1 ([35] Construction 9.1)**.** *Let X be a cube complex, let $W \subseteq X$ be a hyperplane of X. Let γ be a based loop and let $[\gamma] \in \pi_1 X$. Then $[\gamma]$ has a well defined (mod 2) intersection number with W.*

Let $\mathcal{W}$ be the set of embedded, 2-sided, non-separating hyperplanes of X. Then there exist maps $i_W : \pi_1 X \to \mathbb{Z}/2\mathbb{Z}$ a map:

$$\Phi : \pi_1 X \to \bigoplus_{W \in \mathcal{W}} \mathbb{Z}/2\mathbb{Z} \qquad \Phi = \bigoplus_{W \in \mathcal{W}} i_W$$

The ***double dot cover*** *of X is the cover corresponding to the subgroup $\ker \Phi \leq \pi_1 X$.*

The double dot cover of a cube complex is usually a high degree cover. Therefore, constructing examples can be quite difficult. Fortunately, the double dot cover of a rose with 2 petals is easy to construct:

Example 6.3.2. *See Figure 6.1 for an example of the double dot cover of the figure 8 loop.*

An important feature of the double dot cover is that the cover is taken over non-separating hyperplanes. This serves two purposes: first, making sure that double dot cover is not trivial and second, making sure that the double dot hierarchy constructed later has non-trivial splittings.

Fortunately, there is a way to obtain a complex where every hyperplane is non-separating:

Theorem 6.3.3 ([8] Proposition 2.12)**.** *Let X be a compact special NPC cube complex, then X is homotopy equivalent to a compact special NPC cube complex whose hyperplanes are all non-separating.*

Let X be a special cube complex with finitely many hyperplanes $\mathcal{W} := \{W_1, \ldots, W_n\}$ where every hyperplane is non-separating and let $p : \ddot{X} \to X$ be the double dot cover of X. The hyperplanes of $\ddot{X}$ are elevations of hyperplanes of X, and they divide $\ddot{X}$ in a natural way. Let $x \in \ddot{X} \smallsetminus \bigcup p^{-1}(\mathcal{W})$. Any two paths γ_1, γ_2 between x and a lift of $p(x)$ to $\ddot{X}$ represent the same element of $\ker \phi$ precisely when the number of times γ_1 and γ_2 cross elevations of W agree $(\bmod\ 2)$ for every $W \in \mathcal{W}$. Therefore, $\ddot{X} \smallsetminus \cup p^{-1}(\mathcal{W})$ has components labeled by an element of $\bigoplus_{W \in \mathcal{W}} \mathbb{Z}/2$.

The labeling helps construct the double dot hierarchy whose construction follows in Construction 6.3.4. Given a compact special non-positively curved cube complex X whose hyperplanes are all non-separating, a local isometric immersion $\Phi : \mathcal{Z} \to X$ and an ordering on the hyperplanes of C_Φ, the augmented cube complex induced by Φ (recall Definition 6.2.3), the double dot hierarchy will produce a hierarchy of spaces for $\ddot{C}_\phi$. When these inputs satisfy certain criteria discussed in Section 6.4, the double dot hierarchy gives rise to a quasi-convex, and fully $\mathcal{P}$-elliptic hierarchy of groups for $\pi_1(\ddot{C}_\Phi)$ which is isomorphic to a finite index subgroup of $\pi_1 X$. Passing to a particular finite cover will produce an induced hierarchy that is also malnormal. An important feature of the double dot hierarchy is that it is not unique: the construction that follows explicitly depends on an ordering of the hyperplanes of X (which induces an order on the non-peripheral hyperplanes of C_Φ).

Construction 6.3.4 (cf. [3] Section 5)**.** *Let X be a compact special NPC cube complex whose hyperplanes are all non-separating. Let $\mathcal{Z} = \bigsqcup_{i=1}^{n} \overline{Z}_i$ and let $\Phi : \mathcal{Z} \to X$ be a local isometric immersion of NPC cube complexes. Let $C := C_\Phi$ be the augmented cube complex. Recall that Φ can now be canonically regarded as an embedding $\Phi : \mathcal{Z} \hookrightarrow C$. Let $\mathcal{W}$ be the set of non-peripheral hyperplanes of C and let $W_1, \ldots, W_n$ be an ordering of the elements of $\mathcal{W}$. Let $p : \ddot{C} \to C$ be the double dot cover and let $\ddot{Z} := p^{-1}(\mathcal{Z})$.*

As above, choose a basepoint $x \in \ddot{C}$ with $p(x) \notin \bigcup \mathcal{W}$ so that each component of $C \smallsetminus p^{-1}(\bigcup \mathcal{W})$ is labeled by a vector $\hat{t} \in \bigoplus_{i=1}^{n} \mathbb{Z}/2\mathbb{Z}$. For each $1 \leq i \leq n$, let $\mathcal{W}_i$ be the first i hyperplanes and let $M_i = \bigoplus_1^i \mathbb{Z}/2\mathbb{Z}$, and as before, the complementary components of $\bigcup \mathcal{W}_i$ are labeled by elements of M_i. For each $\hat{t} \in M_i$, let $K_{\hat{t}}$ be the closure of the union of components labeled by $\hat{t}$.

*For each $\hat{t} \in M_i$, a $\hat{t}$**-vertex space at level** $n-i+1$ consists of components of $K_{\hat{t}} \cup \ddot{\mathcal{Z}}$ that intersect $K_{\hat{t}}$. In the construction of the double dot hierarchy, the set of components of $\hat{t}$-vertex spaces at level $n-i+1$ specifies all of the vertex spaces at each level, but the actual graph of spaces structure at each level must be described.*

*If A is the closure of a component of $c^{-1}(W_i) \smallsetminus \bigcup_{j<i} p^{-1}(W_j)$, then A is called a **partly-cut-up elevation of** W_i. The idea is that the hierarchy should be constructed by cutting along an elevation of a hyperplane W_i to $\ddot{C}$ and any elements of $\ddot{\mathcal{Z}}$ that intersect W_i, but the elevation of the hyperplane W_i may have already been cut by one of the other hyperplane elevations of W_j with $j < i$.*

By construction, any two $\hat{t}$-vertex spaces at level $n-i+1$ are either disjoint or intersect in a union of components of $\ddot{\mathcal{Z}}$ and disjoint partly-cut-up elevations of W_i.

Now it is time to construct the graph of spaces structures at each level. Let V be a *vertex space at level $n-i+1$ so that V is the $\hat{t}$-vertex space for some $\hat{t} \in M_i$. Consider the canonical projection $\pi : M_{i+1} \to M_i$, let $\hat{t}^+$ and $\hat{t}^-$ be the preimages of $\hat{t}$ under π. Let $\hat{V}^+ := \{V_1^+, \ldots, V_p^+\}$ and $\hat{V}^- := \{V_1^-, \ldots, V_m^-\}$ be the components labeled by $\hat{t}^+$ and $\hat{t}^-$ respectively. Then $V = \bigcup \hat{V}^+ \cup \bigcup \hat{V}^-$. By construction, elements of $\hat{V}^+$ are pairwise disjoint and similarly, elements of $\hat{V}^-$ are pairwise disjoint. The elements of $\hat{V}^+$ and $\hat{V}^-$ are the vertex spaces in the graph of spaces for V and this graph of spaces will have a bipartite underlying graph Γ. For convenience, the edges representing multiple components can be repeated so that each edge space is connected and the edges of Γ are in one-to-one correspondence with components of $\bigcup \left(\hat{V}^+\right) \cap \left(\bigcup \hat{V}^-\right)$. The attaching maps are the inclusion maps of edge spaces into vertex spaces while the realization is provided by a homotopy equivalence collapsing the mapping cylinders of the edge spaces onto the images of the edge spaces.*

Let $\hat{t} \in M_n$. Then the components of the $\hat{t}$-vertex spaces are the vertex spaces of level 1 of the hierarchy, so the terminal spaces of the hierarchy are precisely these spaces.

The hierarchy $\mathcal{H}$ built in Construction 6.3.4 is called the **double dot hierarchy for the pair** $(X, \mathcal{Z})$. While this is not technically correct because the hierarchy depends on the ordering of the hyperplanes, the applications that follow will only depend on the existence of a hierarchy given some local isometric immersion $\mathcal{Z} \to X$.

In general, the double dot hierarchy does not have nice properties. It may fail to be faithful and even if it is, may fail to be quasiconvex or malnormal. Also, the terminal spaces may not be useful. However, when hyperplanes are embedded, nonseparating and two-sided, the terminal spaces are easy to understand:

Lemma 6.3.5. *Let $\Phi : \mathcal{Z} \hookrightarrow X$ be a local isometric immersion of NPC special compact cube complexes. Let $C := C_\Phi$ be the augmented cube complex and let $p : \ddot{C} \to X$ be the* *double dot cover. Let* $\ddot{\mathcal{Z}} = p^{-1}(\mathcal{Z})$*. Suppose that every hyperplane of* X *is nonseparating. If* Y *is a terminal space of the double dot hierarchy for* $(X, \mathcal{Z})$*, then* Y *has a graph of spaces structure* (Γ, χ) *such that:*

1. Γ *is bipartite with vertex set* $V = V^+ \sqcup V^-$*,*
2. *if* $v \in V^+$*,* $\chi(v)$ *is contractible,*
3. *if* $v \in V^-$*,* $\chi(v)$ *is a component of* $\ddot{\mathcal{Z}}$ *and*
4. *every edge space is contractible.*

Proof. Recall that $\mathcal{W}$ is the set of non-peripheral hyperplanes of C in Construction 6.3.4. At the bottom level of the hierarchy, the vertex spaces that remain are pieces of $\ddot{C}$ that are cut along all elevations of non-peripheral hyperplanes of C which are joined by parts of the elevations of mapping cylinders of $\ddot{\mathcal{Z}} \to X$ to components of $\ddot{\mathcal{Z}}$. Cutting along the partly-cut-up elevations of peripheral hyperplanes yields the desired graph of spaces structure.

The parts that do not intersect $\ddot{\mathcal{Z}}$ are contractible because all the hyperplanes of $\ddot{C}$ have been cut up, while the other parts deformation retract onto components of $\ddot{Z}$. The edge spaces are cubical subspaces of the peripheral hyperplanes that have also been completely cut up and are therefore also contractible. □

Corollary 6.3.6. *Under the same assumptions as Lemma 6.3.5, the fundamental group of a terminal space of the double dot hierarchy is a free product of the form* $(*_{i=1}^{p} G_i) * F$ *where* F *is a finitely generated free group and each* $G_i := \pi_1(Z_i)$ *where* Z_i *is a component of* $\mathcal{Z}$*.*

6.4 A fully $\mathcal{P}$-elliptic malnormal quasiconvex hierarchy

In the special case that $\mathcal{Z}$ is a union of the complexes constructed in Proposition 6.2.2 and X is compact special, strategically passing to finite covers and building the double dot hierarchy will produce a faithful, quasiconvex and fully $\mathcal{P}$–elliptic hierarchy. Let $C := C_\Phi$ be the augmented complex for the pair $(X, \mathcal{Z})$ and let $\widetilde{C}$ be its universal cover. Each edge space of the double dot hierarchy consists of unions of components of $\ddot{\mathcal{Z}}$ and partly-cut-up hyperplane elevations of a single hyperplane of C. When elevations of the edge spaces are uniformly quasi-isometrically embedded, the hierarchy will be faithful and quasiconvex. Malnormality and full $\mathcal{P}$-ellipticity will follow. The strategy is to use Proposition 4.2.11 where $\mathcal{A}$ consists of elevations of a partly-cut-up hyperplane elevation to $\widetilde{C}$ and $\mathcal{B}$ consists of elevations of components of $\ddot{Z}$ to $\widetilde{C}$. Proposition 4.2.11 requires that $B \in \mathcal{B}$ when b is a geodesic connecting two elevations $A_1, A_2 \in \mathcal{A}$ is sufficiently long. Proposition 5.2.5 can be applied to X to pass to a finite cover of X where the hyperplanes are sufficiently separated.

For the following, let X_0 be a NPC compact special cube complex. There exists a homotopy equivalent compact special cube complex X whose hyperplanes are all non-separating. Let $\widetilde{X}$ be the universal cover of X. Let $G := \pi_1 X \cong \pi_1 X_0$ and suppose that $(G, \mathcal{P})$ is a relatively hyperbolic group pair. Let $x \in \widetilde{X}$ be a base point in $\widetilde{X}$ not in any hyperplane of $\widetilde{X}$.

By Proposition 6.2.2, for each $P \in \mathcal{P}$ there exists a complex Z_P and superconvex local isometric immersions $\phi_P : Z_P \to X$ such that $\pi_1 Z_P \cong P$ and the image of $\pi_1 Z_P$ in G is conjugate to P in G. Let $\Phi : \bigsqcup_{P \in \mathcal{P}} Z_P \to \widetilde{X}$ so that $\Phi|_{Z_i} = \phi_i$. The map Φ is still a superconvex local isometric immersion.

Let $C_1 := C_\Phi$ be the augmented cube complex of the pair $(X, \mathcal{Z})$.

Lemma 6.4.1. *Let C' be a finite regular cover of C_1. Then:*

1. *There exists a finite cover X' of X with $G' := \pi_1 X'$ and a superconvex local isometric immersion $\Phi' : \mathcal{Z}' \to X'$ such that $(G', \mathcal{P}')$ is the induced relatively hyperbolic group pair (see Proposition 2.3.13) and C' is the augmented cube complex of the pair $(X', \mathcal{Z}')$. The components of $\mathcal{Z}'$ have fundamental group isomorphic to elements of $\mathcal{P}'$ and for each component Z of $\mathcal{Z}'$, the image of $\pi_1 Z$ is conjugate to an element of $\mathcal{P}'$ in G'.*
2. *Every nonperipheral hyperplane of C_2 is nonseparating.*

Let $Z_1', \dots, Z_q'$ be the components of $\mathcal{Z}'$.

Let $\mathcal{B}$ be the collection of images of elevations of $Z_1', \dots, Z_q'$ to $\widetilde{X}$. These elevations are P_i-cocompact, so each $B \in \mathcal{B}$ lies in a bounded neighborhood of $gP_i x$ for some $1 \leq i \leq q$ and $g \in G$. Further, $\widetilde{X}$ has triangles that are $(\delta - 2)$-thin relative to $\mathcal{B}$ for some $\delta > 0$. Let D be an upper bound on $\operatorname{diam}(\mathcal{N}_{3\delta}(B_1) \cap \mathcal{N}_{3\delta}(B_2))$ where B_1 and B_2 are distinct elements of $\mathcal{B}$. Let $M = 2(2D + 12\delta)$. Let T be some constant so that such that $(1024, 8M + 9\delta + 2048(40m + 60\delta))$–quasigeodesics $T - 2$-relatively fellow travel in $\widetilde{X}$ as in Theorem 6.1.11, and let $f(T) \geq \operatorname{diam}(\mathcal{N}_{T+2M+2\delta+2}(B_1) \cap \mathcal{N}_{T+2M+2\delta+2}(B_2))$ where $B_1, B_2 \in \mathcal{B}$ are distinct.

Using Proposition 5.2.5, let C_2 be a finite regular cover of C_1 such that every non-peripheral hyperplane of C_2 is $\max(f(T), 512(40M + 60\delta), 1024(8M + 6\delta(8M + 6\delta + 2048(40M + 60\delta)))) + 2$-embedded and nonseparating. Let C_2 be the augmented cube complex of a pair $(X_2, \mathcal{Z}'')$ where X_2 is a finite cover of X. Note that $\widetilde{X}$ naturally embeds in $\widetilde{C}$, the universal cover of C_2 so that $\widetilde{C}$ has δ-relatively thin triangles relative to $\mathcal{B}$ and $(1024, 8M + 9\delta + 2048(40m + 60\delta))$–quasigeodesics

T-relatively fellow travel in $\widetilde{C}$. Let $G_2 = \pi_1(C_2)$ and let $(G_2, \mathcal{P}'')$ be the induced peripheral structure. Let $(\ddot{G}_2, \ddot{\mathcal{P}}'')$ be the induced peripheral structure $\ddot{G}_2 := \pi_1 \ddot{C}_2$.

The next few statements will show that the double dot hierarchy on $c : \ddot{C}_2 \to C_2$, the double dot cover of C_2, is faithful, quasiconvex and fully $\ddot{\mathcal{P}}''$-elliptic hierarchy for $\pi_1 \ddot{C}_2$. Passing to a finite regular cover will later yield a hierarchy which is also malnormal.

Let E be an edge space of the double dot hierarchy on $\ddot{C}_2$. Let W be a partly-cut-up elevation of a non-peripheral hyperplane to $\ddot{C}_2$ so that E is a union of W and elements of $c^{-1}(\mathcal{Z}'')$. let $\mathcal{A}$ be the set of elevations of W to $\widetilde{C}$. To use Proposition 4.2.11, it must be shown that an elevation of E to $\widetilde{C}$ is a $(\mathcal{A}, \mathcal{B})$-LPPC subspace.

Lemma 6.4.2. *Let E be an edge space of the double dot hierarchy for $\ddot{C}_2$. Then any elevation $\widetilde{E} \to \widetilde{C}$ is a $(\mathcal{A}, \mathcal{B})$-LPPC subspace.*

Proof. The augmented structure of $\ddot{C}_2$ means that each $B \in \mathcal{B}$ has an (embedded) neighborhood N isometric to $B \times [0, \frac{1}{2})$ where $B = B \times \{0\}$. Each non-peripheral hyperplane H intersects $B \times [0, \frac{1}{2})$ as $(B \cap H) \times [0, \frac{1}{2})$. If $\widetilde{W}$ is an elevation of a partly-cut-up hyperplane, then $\widetilde{W} \cap N$ is the intersection of N with a hyperplane $\widetilde{W}_0$ of $\widetilde{C}$ and half spaces of other hyperplanes so E is a $(\mathcal{A}, \mathcal{B})$-LPPC subspace. ∎

Faithfulness and quasiconvexity of the double dot hierarchy on $\ddot{C}_2$ now follow from Proposition 4.2.11.

Proposition 6.4.3. *Let E be an edge space of the double dot hierarchy on $\ddot{C}_2$. Let $\widetilde{E}$ be* *the universal cover of E. Then any elevation $\widetilde{E} \to \widetilde{C}$ of E to $\widetilde{C}$ is a* $(1024, 8M + 6\delta + 2048(40M + 60\delta))$*–quasi-isometric embedding.*

Proof. The edge space E is a union of a partly-cut-up hyperplane elevation $\hat{W}$ with elements of $\mathcal{B}$. Since $\hat{W}$ is contained in an elevation of a hyperplane that is $L := \max(f(T), 512(40M + 60\delta), 1024(8M + 6\delta + 2048(40M + 60\delta)))$-embedded, by Proposition 4.2.11 every geodesic in $\widetilde{E}$ is a quasigeodesic of the specified type. □

Proposition 6.4.4. *Let E be an edge space of the double dot hierarchy on $\ddot{C}_2$. Then the map $E \to \ddot{C}_2$ is π_1 injective.*

Proof. Suppose not toward a contradiction. Then there exists a loop γ in E such that γ is essential in E but has trivial image in $\pi_1(\ddot{C}_2)$.

Since γ is π_1 trivial in $\pi_1(C)$, there exists a loop $\widetilde{\gamma}$ in $\widetilde{C}$ which is the image of a geodesic $\hat{\gamma}$ in $\widetilde{E}$ such that $\widetilde{\gamma}$ projects to γ under the covering map $\widetilde{C} \to \ddot{C}_2$. Recall $\widetilde{E}$ is a $(\mathcal{A}, \mathcal{B})$-LPPC subspace so that $\widetilde{\gamma}$ embeds in $\widetilde{C}$ as a piecewise geodesic whose pieces lie in $\mathcal{A}$ where $\mathcal{A}$ consists of elevations of a partly-cut-up hyperplane to $\widetilde{C}$ and $\mathcal{B}$ consists of elevations of $\ddot{Z}''$ to $\widetilde{C}$.

By Proposition 6.4.3, $\widetilde{\gamma}$ is a $(1024, 8M + 6\delta + 2048(40M + 60\delta))$-quasigeodesic loop in $\widetilde{C}$. However, because elements of $\mathcal{A}$ and $\mathcal{B}$ are convex, $\widetilde{\gamma}$ contains a subpath of the form $a_j b_j a_{j+1}$ where a_j, a_{j+1} are in distinct elements of $\mathcal{A}$ and $b_j \subseteq B \in \mathcal{B}$, so the length of $\hat{\gamma}$ is at least $|b_j| \geq 2048(8M + 6\delta + 2048(40M + 60\delta) + 2)$ by the L-embeddedness of every hyperplane of $\ddot{C}_2$.

However, by Proposition 4.2.11 $\widetilde{\gamma}$ is a $(1024, 8M+6\delta+2048(40M+60\delta))$ quasigeodesic, but there cannot be a $(1024, 8M + 6\delta + 2048(40M + 60\delta))$-quasigeodesic loop γ whose length is at least $2048(8M + 6\delta + 2048(40M + 60\delta) + 2)$. □

Corollary 6.4.5. *The double dot hierarchy induced on $\pi_1\ddot{C}_2$ is faithful and quasiconvex.*

Proof. The hierarchy is faithful by Proposition 6.4.4. An elevation $\widetilde{E} \to \widetilde{C}$ is $(1024, 8M + 6\delta + 2048(40M + 60\delta))$–quasi-isometric embedding by Proposition 6.4.3. Therefore, $\pi_1 E$ is $(1024, 8M+6\delta+2048(40M+60\delta))$–quasi-isometrically embedded in $\pi_1\ddot{C}_2$. □

The next step is to prove that the double dot hierararchy on $\ddot{C}_2$ is fully $\ddot{\mathcal{P}}''$-elliptic. Definition 6.4.6 introduces geometric terminology for the situation where a subgroup of a relatively hyperbolic group pair $(G, \mathcal{P})$ contains an element g conjugate into a peripheral subgroup P such that no positive power of g lies in $E \cap P$.

Definition 6.4.6. *Let $(\widetilde{X}, \mathcal{B})$ be an (δ, M)-relatively hyperbolic pair and let $\widetilde{X} \to X$ be a covering. Let B be a locally convex subspace of X. Let $E \subseteq X$. The subspace E has an* ***accidental*** *B* ***loop*** *if there exists a homotopically essential loop, γ, which is both freely homotopic to a geodesic loop in B and has no positive power homotopic in E to a geodesic loop in B.*

The next few statements will show that the edge spaces of the double dot hierarchy for $\ddot{C}_2$ have no accidental $\ddot{\mathcal{Z}}''$-loops. This will imply the hierachy is fully $\ddot{\mathcal{P}}''$-elliptic. The first step is to show that elevations of partly-cut-up hyperplanes do not have accidental $\ddot{\mathcal{Z}}''$-loops.

Lemma 6.4.7 ([3, Lemma 5.15])**.** *Let $(X, \mathcal{Z})$ be a superconvex pair where each component of $\mathcal{Z}$ is embedded and let C be the corresponding augmented cube complex. For $n \geq 1$, let $\{W_1, \ldots, W_n\}$ be a collection of embedded, 2-sided, nonseparating hyperplanes of C. Let Q be a component of $W_n \smallsetminus \cup_{i<n} W_i$. Then Q has no accidental $\mathcal{Z}$-loops.*

Proof. Let $\alpha \subseteq Q$ be a geodesic loop homotopic in C to a geodesic loop α' in $Z_j \subseteq \mathcal{Z}$. Recall that X embeds into C. If $\alpha \subseteq C \setminus X$, then α is homotopic into $\mathcal{Z}$ because $C \setminus X$ is isometric to $\mathcal{Z} \times [0,1)$.

By construction, the intersection of Q with $C \setminus X$ is equal to $(\mathcal{Z} \cap Q) \times [0,1)$ in $\mathcal{Z} \times [0,1)$. Therefore, the homotopy of α into $\mathcal{Z}$ can be arranged so that it happens entirely in Q.

Hence assume α intersects X.

Claim: $\alpha \subseteq X$: Up to a reparameterization, assume α is parameterized as a geodesic loop $\alpha : [0,L] \to C$ and suppose $\alpha(0) = \alpha(L) \in X$ and that for small $0 < t < \epsilon$, $\alpha(t) \notin X$. Let R be the smallest positive number so that $\alpha(R) \in X$, and let $\beta = \alpha|_{[0,R]}$. Then β lies in $Z_j \times [0,1]$ so write $\beta(t) = (z(t), s(t))$ for $z(t) \in Z_j$, $s(t) \in [0,1]$ with $s(0) = s(t) = 1$. Since $s(t)$ is locally geodesic by [9] Proposition I.5.3, $s'(t)$ is constant. Therefore, $s(t) = 1$ for all t, so in fact $\alpha(t) \subseteq X$.

Therefore, α elevates to the universal cover $\widetilde{X}$ of X as a biinfinite geodesic $\widetilde{\alpha}$ which lies in a bounded neighborhood of the image of an elevation of the immersion $\phi_j : Z_j \to X$ to $\widetilde{X}$. By the superconvexity of ϕ_j, $\widetilde{\alpha}$ lies in the elevation of ϕ_j, so $\alpha \subseteq \phi_j(Z_j) \cap Q = Q \cap (Z_j \times \{1\})$. Therefore, by using the product structure, α is homotopic in Q into Z_j and hence $\mathcal{Z}$, so α is not an accidental $\mathcal{Z}$-loop. $\square$

Proposition 6.4.8. *Let E be an edge space of the double dot hierarchy for $\ddot{C}_2$. Then E has no accidental $\ddot{\mathcal{Z}}''$–loops.*

Proof. Recall that E is a union of a partly-cut-up hyperplane elevation Q and components of $\ddot{\mathcal{Z}}''$ that intersect Q. By Lemma 6.4.7, Q has no accidental $\ddot{\mathcal{Z}}''$-loops.

Suppose there exists a $\ddot{C}_2$-essential loop γ in E such that γ is freely homotopic in $\ddot{C}_2$ into $\ddot{\mathcal{Z}}''$. Then a representative of the homotopy class of γ lifts to a bi-infinite $\widetilde{E}$-geodesic $\hat{\gamma}$ where $\widetilde{E}$ is an elevation of E to $\widetilde{C}$, and a representative of the homotopy class of γ lifts to a bi-infinite $\widetilde{C}$-geodesic $\rho \subseteq \widetilde{Z}$, an elevation of a component of $\ddot{\mathcal{Z}}''$ and there exists $R \geq 0$ so that $\hat{\gamma} \subseteq \mathcal{N}_R(\rho)$.

Let $\hat{\gamma}_0$ be a subsegment of $\hat{\gamma}$ with $|\hat{\gamma}_0| = |\gamma|$ (e.g. take $\hat{\gamma}_0$ to be the subsegment between two consecutive lifts of a point of γ to $\hat{\gamma}$). If $\hat{\gamma}_0 \subseteq \widetilde{Z}'$ where $\widetilde{Z}'$ is an elevation of a component of $\ddot{\mathcal{Z}}''$, then $\hat{\gamma} \subseteq \widetilde{Z}'$ and $\hat{\gamma}$ is geodesic in $\widetilde{C}$. Then $\widetilde{Z} = \widetilde{Z}'$ because $\mathrm{diam}(\mathcal{N}_R(\widetilde{Z}) \cap \mathcal{N}_R(\widetilde{Z}_i)) = \infty$ in which case γ was not an accidental $\ddot{\mathcal{Z}}$ loop.

On the other hand, if $\hat{\gamma}_0 \subseteq \widetilde{Q}$ where $\widetilde{Q}$ is some elevation of Q to $\widetilde{C}$, then Q has an accidental $\mathcal{Z}$-loop, contradicting the fact that there are no such accidental $\mathcal{Z}$ loops.

Therefore, there exist subsegments $\hat{\gamma}$ of the form $\gamma_m = a_{m,1}b_{m,1}a_{m,2}b_{m,2}\dots a_{m,m}b_{m,m}$ such that $\cup_1^\infty \gamma_m = \gamma$, $|\gamma_m| \to \infty$ as $m \to \infty$, a_i lies in an elevation $\widetilde{Q}_i$ of Q to $\widetilde{C}$, $b_{m,i} \subseteq \widetilde{Z}_i$ where $\widetilde{Z}_i$ is an elevation of a component of $\ddot{\mathcal{Z}}''$, and if $i \neq j$, $b_{m,i} \subseteq \widetilde{Z}_i$ and $b_{m,j} \subseteq \widetilde{Z}_j \neq \widetilde{Z}_i$ (otherwise, by convexity of $\mathcal{Z}$, γ_m could be written as a concatenation of fewer geodesic segments).

Let τ_m be the $\widetilde{C}$-geodesic connecting the endpoints of γ_m.

Claim: there exists $S > 0$ such that all but the S neighborhoods of the endpoints of τ_m lies in $\mathcal{N}_{2\delta}(\widetilde{Z})$.

Recall τ_m has endpoints on $\hat{\gamma} \subseteq \mathcal{N}_R(\rho)$, so R-neighborhood of ρ, so there exist length R geodesics connecting the endpoints of τ_m to ρ. Write $\tau_m = [p, q]$. Let ℓ

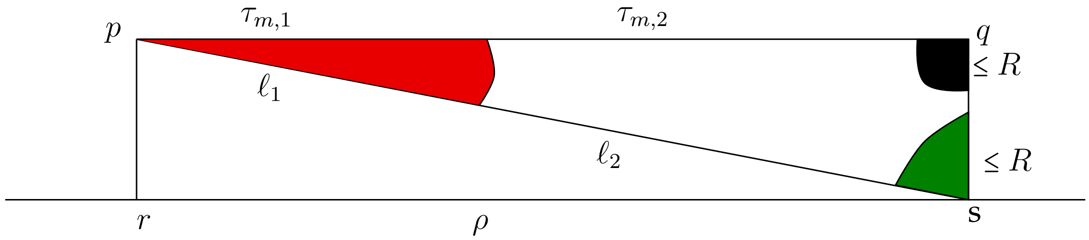

Figure 6.2: A quadrilateral consisting of τ_m, a subsegment of ρ and two geodesics of length at most R with diagonal $\ell = [p, s]$.

be the diagonal starting at p of the resulting quadrilateral, giving two triangles (refer to Figure 6.2), and let $[p, r]$ be a minimal length geodesic of length at most R connecting p to ρ.

Recall that $\mathcal{B}$ is the set of components of elevations of elements of $\ddot{\mathcal{Z}}''$ and that triangles in $\widetilde{C}$ are δ-thin relative to $\mathcal{B}$. The bottom triangle in Figure 6.2 is either δ-thin relative to $\widetilde{Z}$ or some other $F' \in \mathcal{B}$. In the first case, all but the corner segment of the bottom triangle at p contained in ℓ either is in the corner segment of the bottom triangle at s and hence δ-fellow travels a subsegment of ρ or lies in the fat part of ℓ in the bottom triangle and hence lies in $\mathcal{N}_\delta(\widetilde{Z})$. Therefore, all but an at most R-tail of ℓ at p is not in $\mathcal{N}_\delta(\widetilde{Z})$. In the second case, the length of the fat part of ρ in $\mathcal{N}_\delta(F')$ is at most M, so by Lemma 2.3.22 and the fact that $|[p, r]| \leq R$, at most a $M + 2R + 3\delta$ tail of ℓ at p is not in the corner segment of the bottom triangle at s.

The top triangle in Figure 6.2 is δ-thin relative to some $F'' \in \mathcal{B}$. Decompose ℓ as a corner segment of the top triangle at p, ℓ_1, the fat part of ℓ in the top triangle,ℓ_2 and the corner segment at s of length at most R. Then ℓ_1 δ-fellow travels $\tau_{m,1}$, the corner segment at p in τ. The fat part of ℓ, ℓ_2, lies in $\mathcal{N}_\delta(F'')$. Let $\tau_{m,2}$ be the fat part of τ in the top triangle so that $\tau_{m,2} \subseteq \mathcal{N}_\delta(F'')$. If $F'' = \widetilde{Z}$, then all but a $M + 2R + 3\delta$ tail of $\tau_{m,1}$ at p must 2δ-fellow travel a subsegment of ρ and all but a length R tail of the remainder of τ at q lies in $\mathcal{N}_\delta(F'')$, so $S = M + 2R + 3\delta$ suffices, and the claim is satisfied in this case.

If $F'' \neq F', \widetilde{Z}$, then at most $2M + 2R + 3\delta$ of ℓ_2 lies in $\mathcal{N}_\delta(\widetilde{Z})$ since at most $M + 2R + 3\delta$ of ℓ does not δ-fellow travel a subsegment of ρ. By Lemma 2.3.22, $|\tau_{m,2}| \leq 2M + 3R + 6\delta$, so at most a $2M + 4R + 6\delta$ tail of τ at q does not lie in $\mathcal{N}_{2\delta}(\widetilde{Z})$ and at most $M + R + 3\delta$ of $\tau_{m,1}$ at p does not lie in $\mathcal{N}_{2\delta}(\widetilde{Z})$. Hence it suffices to take $S = 2M + 4R + 6\delta$.

Recall the constant T defined above so that $(1024, 8M + 9\delta + 2048(40m + 60\delta))$–quasigeodesics $T - 2$-relatively fellow travel in $\widetilde{C}$.

By Proposition 4.2.11, each γ_m is a $(1024, 8M + 6\delta + 2048(40M + 60\delta))$–quasigeodesic, so if τ_m is the geodesic connecting the endpoints of γ_m, $|\tau_m| \to \infty$ as $m \to \infty$. Choose $m >> 0$ such that there exist $b_{m,i}$ and $b_{m,j}$ so that $b_{m,i}$, $b_{m,j}$ are at least $S + T + 2M + 2\delta + 2$ from the endpoints of γ_m and τ_m.

At most M of $b_{m,i}$ can lie in $\mathcal{N}_\delta(\widetilde{Z})$ unless $\widetilde{Z}_i = \widetilde{Z}$. Therefore, if $\widetilde{Z}_i \neq \widetilde{Z}$, by the relative fellow traveler property (Theorem 6.1.11) $b_{m,i} \subseteq \mathcal{N}_{T+2M}(\tau_m)$ because $b_{m,i}$ is sufficiently far from the endpoints of τ_m and $m >> 0$, in fact $b_{m,i}$ is near the part of τ_m which lies in $\mathcal{N}_{2\delta}(\widetilde{Z})$, so $b_{m,i} \subseteq \mathcal{N}_{T+2M+2\delta+2}(\widetilde{Z})$. However, since elevations of every hyperplane are at least $f(T) + 2M + 2\delta$ apart, $|b_{m,i}| \geq f(T) + 2M + 2\delta$ and $\widetilde{Z} = \widetilde{Z}_i$ because $f(T)$ bounds the diameter of $\mathcal{N}_{T+2M+2\delta+2}(B_1) \cap \mathcal{N}_{T+2M+2\delta+2}(B_2)$ when $B_1, B_2 \in \mathcal{B}$ and $B_1 \neq B_2$. Similarly, $\widetilde{Z}_j = \widetilde{Z}$, contradicting the fact that $\widetilde{Z}_i \neq \widetilde{Z}_j$. Therefore, γ could not have been an accidental $\ddot{\mathcal{Z}}$ loop. □

Corollary 6.4.9. *The double dot hierarchy on $\ddot{C}_2$ is fully $\ddot{\mathcal{P}}''$-elliptic.*

Faithfulness, quasiconvexity and full $\mathcal{P}$-ellipticity are preserved by taking the induced hierarchy of a finite regular cover of $\ddot{C}_2$. The final step is to show that there exists a finite cover of $\ddot{C}_2$ whose induced hierarchy is also a malnormal hierarchy.

The following lemma is straightforward:

Lemma 6.4.10. *Suppose $H \leq G$ and G_0 is a finite index subgroup of G and let $H_0 = H \cap G_0$. If H is malnormal in G, then H_0 is malnormal in H.*

Proposition 6.4.11. *Let G be the fundamental group of a relatively hyperbolic special compact NPC cube complex, and let $H \leq G$ be full relatively quasiconvex. Then H is separable in G.*

The idea is to follow the proof of Theorem 7.3 of [15] except to use Theorem 5.3.4 to produce a cocompact convex core for H. Once we have a convex core for H, there is a compact cube complex A, special cube complex X with $\pi_1(X) = G$ and local isometry $f : A \to X$. Then use canonical completion and retraction as in 5.1.12 to show that H is a virtual retract of a residually finite group.

Proposition 6.4.12. *(Hruska-Wise [19, Theorem 9.3]) If G is relatively hyperbolic and $H \leq G$ is relatively quasiconvex and separable, then there exists a finite index subgroup $K \leq G$ containing H such that for every $g \in K$ either gHg^{-1} is finite or gHg^{-1} is conjugate into a peripheral subgroup of G.*

Therefore, if H is also full relatively quasiconvex, then H is almost malnormal in K.

The following is based on [3] (Corollary 3.29) and follows immediately from the two preceding statements and the fact that when G is virtually special, G is linear and hence virtually torsion free.

Corollary 6.4.13. *If G is hyperbolic relative to $\mathcal{P}$ and special, and $H \leq G$ is full relatively quasiconvex, then H is virtually malnormal.*

Theorem 6.4.14. *Let G be special, virtually torsion-free and let $(G, \mathcal{P})$ be a relatively hyperbolic pair. Let $\mathcal{H}$ be a fully $\mathcal{P}$-elliptic quasiconvex hierarchy for G, then there exists a finite index subgroup $G_0 \leq G$ with induced fully $\mathcal{P}$–elliptic quasiconvex hierarchy $\mathcal{H}_0$ of G_0 which is malnormal and fully $\mathcal{P}$–elliptic.*

The proof here is nearly the same as in Theorem 3.30 of [3].

Proof. Because $\mathcal{H}$ is fully $\mathcal{P}$-elliptic, the edge subgroups are full. By [3] 3.24, if $H \leq G$ is almost malnormal in G and G_0 is a finite index normal subgroup of G, then $H \cap G_0$ is malnormal in G_0. Since there are finitely many edge groups, by Corollary 6.4.13 there exists some G_0 such that for every edge group E of $\mathcal{H}$, $E \cap G_0$ is malnormal in G_0. Since G_0 is normal, conjugation by $g \in G$ is an automorphism of G_0, so in particular, these edge groups $E \cap G_0$ are malnormal in G. □

Theorem 6.4.15. *Let X be a NPC virtually special compact cube complex. If $G = \pi_1 X$ and $(G, \mathcal{P})$ is a relatively hyperbolic pair, then G has a finite index normal subgroup G_0 with induced peripheral structure $\mathcal{P}_0$ such that G_0 has a faithful, quasiconvex, fully $\mathcal{P}_0$-elliptic and almost malnormal hierarchy terminating in $\mathcal{P}_0$.*

Proof. First, pass to a finite index regular cover of X, X_1 that is special. By applying a homotopy equivalence, X_1 is homotopy equivalent to a cube complex where every hyperplane gives a nontrivial splitting of $\pi_1 X_1$ (see [3] Lemma 5.17).

By Corollary 6.4.5, there exists a special cube complex X_1' homotopy equivalent to X_1 with a finite regular cover X_2 such that $G_2 := \pi_1 X_2$ with induced peripheral structure $(G_2, \mathcal{P}_2)$ has a faithful, quasiconvex, fully $\mathcal{P}_2$–elliptic hierarchy terminating in $\mathcal{P}_2 * F_k$ where F_k is a free group.

By Theorem 6.4.14, there exists a finite regular cover X_0 with $G_0 := \pi_1 X_0$ and induced peripheral structure $(G_0, \mathcal{P}_0)$ such that the induced hierarchy on G_0 is malnormal as well and terminates in free products of free groups and elements of $\mathcal{P}_0$. The hierarchy can then be continued to a malnormal, quasiconvex, fully-$\mathcal{P}_0$-elliptic one that terminates in $\mathcal{P}_0$. □

CHAPTER 7

A RELATIVELY HYPERBOLIC VERSION OF THE MALNORMAL SPECIAL QUOTIENT THEOREM

Wise's malnormal special quotient theorem roughly says the following:

Theorem 7.0.1 ([35])**.** *Let G be hyperbolic and virtually special and suppose $(G, \mathcal{P})$ is a relatively hyperbolic pair. Then G has a virtually special quotient with kernel normally generated by subgroups of elements of $\mathcal{P}$ that are finite index normal in G.*

The purpose of this chapter is to apply Theorem 6.4.15 to obtain a relatively hyperbolic version of Wise's malnormal special quotient theorem using techniques from Sections 6-9 of [3].

Wise's theorem characterizing the virtually special hyperbolic groups [35] Theorem 13.3 has the following useful consequence:

Corollary 7.0.2. *Let G be a hyperbolic group with a quasiconvex hierarchy terminating in finite groups. Then G is virtually special.*

The technique for proving a relatively hyperbolic analog of Theorem 7.0.1 will be to start with the hierarchy provided by Theorem 6.4.15 and strategically take quotients using group theoretic Dehn fillings (see Definition 7.1.1). These quotients can be constructed to be hyperbolic, and with some care, the hierarchy structure can be passed down to the quotient so that Corollary 7.0.2 can be used. In [3], the authors avoided using Corollary 7.0.2 because their account aimed to give a new proof of auxiliary results used to prove Corollary 7.0.2. Consequently, they needed to ensure that the hierarchy structure on the quotient is also a malnormal hierarchy. Here, by using Corollary 7.0.2, it will only be necessary to produce a quasiconvex hierarchy for such a quotient.

7.1 Group Theoretic Dehn Filling

For this chapter, let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let $\mathcal{P} = \{P_1, \ldots, P_m\}$ unless stated otherwise. When M is a finite volume hyperbolic 3–manifold with torus cusps, a **Dehn filling** of M is a gluing of solid tori $T_i \cong D \times S^1$ by a diffeomorphism to the boundary components. The result of the gluing depends only on the isotopy class of the curve $\gamma_i \subseteq \partial M$ that each copy of $\partial D \times \{p\} \subseteq T_i$ is glued to (see e.g. [23] Chapter 10.1). In this situation $\pi_1 M$ is hyperbolic relative to a collection of copies of $\mathbb{Z}^2$, one for each boundary component of M.

The next definition is a group theoretic analog of Dehn filling

Definition 7.1.1. *Let $\{N_i \triangleleft P_i : 1 \leq i \leq m\}$. Then there exists a **group theoretic Dehn filling** of G with **filling map** π defined by the quotient:*

$$\pi : G \to G(N_1, \ldots, N_m) := G/\langle\langle \bigcup N_i \rangle\rangle.$$

*The subgroups N_i are called **filling kernels**.*

*A filling is called **peripherally finite** if each filling kernel N_i is finite index in P_i.*

For a classical filling, if every T_i is filled by gluing along the curves γ_i that are sufficiently long, Thurston's Dehn filling theorem says that the resulting manifold is hyperbolic. The group theoretic analog of a sufficiently long classical

Dehn filling is a group theoretic Dehn filling where the filling kernels avoid a finite set of elements:

Definition 7.1.2. *A statement $\mathfrak{P}$ holds for all sufficiently long fillings if there exists a finite $B \subseteq G \setminus 1$ such that whenever $B \cap N_i = \varnothing$ for all $1 \leq i \leq m$, the filling $G(N_1, \ldots N_m)$ has $\mathfrak{P}$.*

Osin showed that sufficiently long Dehn fillings of relatively hyperbolic groups are relatively hyperbolic, have kernels which intersect each peripheral subgroup P_i precisely in N_i and can be manipulated so that any finite set of elements are not killed by the filling map.

Theorem 7.1.3 ([25])**.** *Let $F \subseteq G$ be any finite subset of G. Then for all sufficiently long Dehn fillings:*

1. $\ker(\phi|_{P_i}) = N_i$ *for* $i = 1, 2, \ldots, m$,
2. *the pair* $(G(N_1, \ldots, N_m), \{\phi(P_1), \ldots, \phi(P_m)\}$ *is a relatively hyperbolic pair, and*
3. $\phi|_F$ *is injective.*

The edge subgroups of the hierarchy from Theorem 6.4.15 will need to be full relatively quasiconvex subgroups of G. The quasiconvexity of the hierarchy will ensure that these subgroups are relatively quasiconvex.

Theorem 7.1.4 ([17] Theorem 1.5)**.** *Let $H \leq G$ be a quasi-isometrically embedded subgroup. Then H is relatively quasiconvex in G.*

Theorem 7.1.5 ([17])**.** *Let $H \leq G$ be relatively quasiconvex. Then there exists a relatively hyperbolic structure $(H, \mathcal{D})$ where $\mathcal{D}$ is finite and every element of $\mathcal{D}$ is conjugate into an element of $\mathcal{P}$.*

For the following, it is convenient to introduce equivalent formulations of relative hyperbolicity and relative quasiconvexity. Let $\Gamma(G,S)$ be a Cayley graph for G and let $\mathcal{P} = \{P_1,\dots,P_m\}$ be a finite collection of subgroups of G. For each i, let T_i be a left transversal for P_i.

Definition 7.1.6 (see [2, Section 1]). *For each pair (i,t) with $1 \leq i \leq m$ and $t \in T_i$, let $\Gamma(i,t)$ be the full subgraph of $\Gamma(G,S)$ containing tP_i. Define the **combinatorial horoball** $\mathcal{H}(i,t)$ as follows: the zero skeleton of $\mathcal{H}(i,t)$ is the set $\Gamma(i,t)\times(\{0\}\cup\mathbb{N})$. Add a single edge between the unordered pair $\{(g_1,y),(g_2,y)\}$ whenever $0 < d_S(g_1,g_2) \leq 2^y$, and for each $g \in tP_i$ and $y \geq 0$, add an edge connecting the unordered pair $\{(g,y),(g,y+1)\}$.*

*For $n \in \mathbb{Z}^{\geq 0}$, a point (g,n) is n-**deep** in $\mathcal{H}(i,t)$.*

*The **cusped graph** $(X,\mathcal{P},S)$ is the space formed by taking $\Gamma(G,S)\cup\bigsqcup_{i,t}\Gamma(i,t)$ and gluing each $(tP_i,0)\subseteq\mathcal{H}(i,t)$ to $tPi \subseteq \Gamma(G,S)$.*

*The pair $(G,\mathcal{P})$ is **AGM-Relatively Hyperbolic** if and only if $(X,\mathcal{P},S)$ is a hyperbolic graph.*

By Lemma 3.1 of [2], if $H \leq G$ and $(H,\mathcal{D})$ is a relatively hyperbolic pair so that every $D \in \mathcal{D}$ is conjugate into some $P \in \mathcal{P}$, then there exists a generating set T such that the inclusion $\phi : H \hookrightarrow G$ extends to an H-equivariant Lipschitz map $\hat{\phi} : X(H,\mathcal{D},T)^{(0)} \to X(G,\mathcal{P},S)^{(0)}$.

Definition 7.1.7. *Fix a relatively hyperbolic pair $(G,\mathcal{P})$. The subgroup $H \leq G$ is **AGM-quasiconvex** if there exists a relatively hyperbolic structure $(H,\mathcal{D})$ such that if $\phi : H \hookrightarrow G$ is the inclusion map, the map $\hat{\phi}$ has quasiconvex image.*

Theorem 7.1.8 ([2, Proposition 1.9] , [22, Theorem A.10] , [17],). *The pair $(G,\mathcal{P})$*

is AGM-relatively hyperbolic if and only if (G, P) is relatively hyperbolic in the sense of Definition 2.3.8.

A subgroup $H \leq G$ is relatively quasiconvex if and only if H is AGM-relatively quasiconvex.

Proposition 7.1.9. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair. Let $(H, \mathcal{D})$ be a relatively hyperbolic pair such that there exists a relatively hyperbolic pair $(H, \mathcal{D})$ such that the embedding $H \hookrightarrow G$ induces a map $\hat{\phi} : X(H, \mathcal{D}, T)^{(0)} \to X(G, \mathcal{P}, S)^{(0)}$ with quasiconvex image. Then there are finitely many H-conjugacy classes of infinite intersections of H with a conjugate of some $P \in \mathcal{P}$.*

Proof. Let $H \cap gPg^{-1}$ be infinite for some $P \in \mathcal{P}$ and $g \in G$. Then $H \cap gPg^{-1}$ stabilizes an infinite diameter subset of gP in $\Gamma(G, S)$. Let $d_{\mathcal{P}}$ denote distance in $X(G, \mathcal{P}, S)$. Then there exists $g_n \in H \cap gPg^{-1}$ and $x \in H \cap gP$ such that $d_S(x, g_n x), d_T(x, g_n x) \to \infty$ and there is a geodesic path γ_n connecting x to $y_n := g_n x$ so that x, y_n are the only points in γ_n that are 0-deep and γ_n has a point that is at least n-deep in the combinatorial horoball glued to gP. By quasiconvexity, for $n >> 0$, then there exists a peripheral coset $h_n D_n$ for some $h_n \in H$, $D_n \in \mathcal{D}$ such that $x, y_n \in h_n D_n$. There exist only finitely many $\mathcal{D}$ cosets containing x, so $H \cap gPg^{-1}$ has infinite intersection with hDh^{-1} for some $h \in H$ and $D \in \mathcal{D}$. Therefore, hDh^{-1} is conjugate into P by g and hDh^{-1} and gPg^{-1} fix the same parabolic point in $\partial X(G, \mathcal{P}, S)$. By the relative quasiconvexity of H, there are only finitely many H-orbits of such points, so there exist only finitely many H-conjugacy classes of conjugates of elements of $\mathcal{P}$ that have infinite intersection with H. □

Corollary 7.1.10. *The collection $\mathcal{D}$ can be modified so that:*

1. *Every element of $\mathcal{D}$ is infinite.*

2. *Every infinite intersection of H with a conjugate of some $P \in \mathcal{P}$ is conjugate in H to some element of $\mathcal{D}$.*

Proof. For the first statement, simply remove all finite elements of $\mathcal{D}$. Since $\mathcal{D}$ was finite, removing finite sized elements of $\mathcal{D}$ will affect distances in the coned-off Cayley graph by at most a fixed constant.

The second statement follows immediately from Proposition 7.1.9 □

When a filling of G interacts nicely with a subgroup H, it is possible to induce a filling on the subgroup H.

Definition 7.1.11. *Let $H \leq G$. A filling $G \to G(N_1, \ldots, N_m)$ is an H-**filling** if whenever $gP_ig^{-1} \cap H$ is infinite for some $P_i \in \mathcal{P}$, then $gN_ig^{-1} \subseteq H$.*

Definition 7.1.12. *If $H \leq G$ is a relatively quasiconvex subgroup and $(H, \mathcal{D})$ is the relatively hyperbolic structure from Theorem 7.1.5 and Corollary 7.1.10. Let $\pi : G \to G(N_1, \ldots, N_m)$ be an H-filling. Let $D_j \in \mathcal{D}$. Then there exists some $P_i \in \mathcal{P}$ and $g \in G$ with $g^{-1}D_jg \subseteq P_i$. Let $K_j := gN_ig^{-1}$. Since π is an H-filling, $K_j \lhd D_j$, so the groups K_j determine a filling:*

$$\pi_H : H \to H(K_1, \ldots, K_N)$$

*called the **induced filling of** H with respect to $G(N_1, \ldots, N_m)$.*

Since N_i is normal in P_i, then groups K_j (and hence the filling) do not depend on the choice of $g \in G$.

The following theorem appears as stated in [3] as Theorem 7.11 and collects results about induced Dehn fillings from [2]:

Theorem 7.1.13. *Let $H \leq G$ be a full relatively quasiconvex subgroup and let $F \subseteq G$ be a finite subset. For all sufficiently long H-fillings, $\phi : G \to G(N_1, \ldots, N_m)$ of G:*

1. *$\phi(H)$ is a full relatively quasiconvex subgroup of $G(N_1, \ldots, N_m)$,*
2. *$\phi(H)$ is isomorphic to the induced filling in that if $\phi_H : H \to H(K_1, \ldots, K_m)$ is the induced filling map, then $\ker \phi_H = \ker \phi \cap H$, and*
3. *$\phi(F) \cap \phi(H) = \phi(F \cap H)$.*

7.2 The filled hierarchy

Let $\mathcal{H}$ be a quasiconvex fully $\mathcal{P}$–elliptic hierarchy. By Lemma 3.2.10, Theorem 7.1.4 and the full $\mathcal{P}$-ellipticity of the hierarchy, the edge and vertex groups of the hierarchy are fully relatively quasiconvex. Let $\pi : G \to \overline{G}$ be a filling and let $(\overline{G}, \overline{\mathcal{P}})$ be the relatively hyperbolic structure induced on the filling by Theorem 7.1.3. The goal of this subchapter is to build an induced hierarchy $\overline{\mathcal{H}}$ (which may not be faithful) for $\overline{G}$ based on $\mathcal{H}$ where the vertex and edge groups of $\overline{\mathcal{H}}$ are induced fillings of vertex and edge groups of $\mathcal{H}$. The hierarchy $\overline{\mathcal{H}}$ will be called a **filled hierarchy** for $(\overline{G}, \overline{\mathcal{P}})$.

The filled hierarchy is built by starting at the top level and building the hierarchy inductively downward.

At the top level, let $\overline{\mathcal{H}}$ have the degenerate graph of groups decomopsition for $\overline{G}$ consisting of a single vertex labeled $\overline{G}$. Let n be the length of $\mathcal{H}$. Suppose the filled hierarchy has been filled down to the $(n-i)$th level and let $\overline{A}$ be a vertex group at level $n - i$ so that A is the induced filling of a vertex group A at level $n - i$ of $\mathcal{H}$. Let (Γ, α^{-1}, T) be the graph of groups structure for A provided by $\mathcal{H}$ and let χ be the assignment map for the graph of groups (Γ, T).

If x is a vertex or edge of Γ, let $\mathcal{A}_x := \chi(x)$, the corresponding vertex or edge group. Let $\overline{\chi}(x) := \overline{\mathcal{A}}_x$ where $\overline{A}_x$ is the induced filling $\pi_x : \mathcal{A}_x \to \overline{\mathcal{A}}_x$. The problem is that the pair $(\Gamma, \overline{\chi})$ still needs attachment homomorphisms to be a graph of groups.

Let $\phi_e : \mathcal{A}_e \to \mathcal{A}_v$ be an attachment homomorphism of an edge group $\mathcal{A}_e$ to a vertex group $\mathcal{A}_v$. Two details need to be checked: first there need to be attachment maps $\overline{\phi}_e : \overline{\mathcal{A}}_e \to \overline{\mathcal{A}}_v$ such that $\overline{\phi}_e \circ \pi_e = \pi_v \phi_e$. Then there will need to be an isomorphism $\overline{\alpha} : \pi_1(\Gamma, \overline{\chi}, T) \to \overline{A}$ so that $(\Gamma, \overline{\alpha}, T)$ is a graph of groups structure for $\overline{A}$ where $\overline{\alpha} \circ \pi_\Gamma = \pi_A \circ \alpha$.

Completing the square:

$$\begin{array}{ccc} \mathcal{A}_e & \xrightarrow[\pi_e]{} & \overline{\mathcal{A}}_e \\ \downarrow{\scriptstyle \phi_e} & & \downarrow \\ \mathcal{A}_v & \xrightarrow[\pi_v]{} & \overline{\mathcal{A}}_v \end{array}$$

with a map $\overline{\phi}_e : \overline{\mathcal{A}}_e \to \overline{\mathcal{A}}_v$ is straightforward because π_e is surjective and $\ker \pi_e \subseteq \ker \pi_v \circ \phi_e$. Even when $\mathcal{H}$ is a faithful hierarchy, the map ϕ_e may fail to be injective.

Constructing the desired isomorphism $\overline{\alpha} : \pi_1(\Gamma, \overline{\chi}, T) \to \overline{G}$ amounts to completing the square:

$$\begin{array}{ccc} \pi_1(\Gamma, \chi, T) & \xrightarrow[\pi_\Gamma]{} & \pi_1(\Gamma, \overline{\chi}, T) \\ \downarrow{\scriptstyle \alpha} & & \downarrow \\ A & \xrightarrow[\pi_A]{} & \overline{A} \end{array}$$

Lemma 7.2.1. *There exists an isomorphism $\overline{\alpha} : \pi_1(\Gamma, \overline{\chi}, T) \to \overline{G}$ that completes the diagram.*

Proof. It suffices to show that $\ker \pi_\Gamma = \ker(\pi_A \circ \alpha)$.

The first step is to show that $\ker \pi_\Gamma \subseteq \ker(\pi_A \circ \alpha)$. Let $k \in \ker \pi_\Gamma$. Then k can be written as:

$$k = \prod_i k_i^{g_i}$$

where each $g_i \in \pi_1(\Gamma, \chi, T)$ and $k_i \in \ker \pi_{v_i}$ where v_i is a vertex of Γ.

It then suffices to show that $\ker \pi_v \subseteq \ker \pi_A \circ \alpha$ for each vertex v, so assume $k \in \ker \pi_v$.

The vertex group $\mathcal{A}_v$ is fully relatively quasiconvex in $(G, \mathcal{P})$ as noted above, and by Corollary 7.1.10 there is an induced peripheral structure $\mathcal{D}_v := \{D_1, \ldots, D_l\}$ on $\mathcal{A}_v$ such that each (infinite) $D_i \subseteq (P_{j_i})_i^g$, and $D_i = (P_{j_i})^{g_i}$ by fullness. The element k can be written as:

$$k = \prod_\beta n_\beta^{a_\beta}$$

where n_β lies in a filling kernel $N_\beta \triangleleft D_\beta \subseteq (P_{j_\beta})^{g_\beta}$. By fullness, $(P_{j_\beta})^{g_\beta}$ is conjugate in A to an element of the peripheral structure on A induced by $(G, \mathcal{P})$, so n_β is conjugate to an element of some filling kernel of the induced filling π_A. Therefore $k \in \ker \pi_A \circ \alpha$.

On the other hand, if $k \in \ker(\pi_A \circ \alpha)$, then $k = \prod_{i=1}^l \alpha^{-1}(k_i^{g_i})$ where each $k_i \in K_{j_i} \triangleleft D_{j_i}$ and K_{j_i} is a filling kernel for the induced filling π_A. By full $\mathcal{P}$-ellipticity, $\alpha^{-1}(D_{j_i})$ is conjugate into some vertex group $\mathcal{A}_v$ of (Γ, χ), so $\alpha^{-1}(k_i)$ is conjugate into $\ker \pi_v$ for some $v \in V$. Therefore, $\alpha^{-1}(k_i) \in \ker \pi_\Gamma$, and $\ker(\pi_A \circ \alpha) \subseteq \ker \pi_\Gamma$. □

For the following, let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let $\mathcal{H}$ be a quasiconvex fully $\mathcal{P}$–elliptic hierarchy for G. The next lemma ties together some definitions:

Lemma 7.2.2. *If $A \leq G$ is an edge or vertex group of $\mathcal{H}$, then A is fully relatively quasiconvex subgroup of $(G, \mathcal{P})$ and every filling is an A-filling.*

Proof. That A is fully relatively quasiconvex follows immediately from the definition of full $\mathcal{P}$-ellipticity and Theorem 7.1.4.

Whenever $gP_ig^{-1} \cap A$ is infinite, then $gP_ig^{-1} \subseteq A$, so if $N_i \triangleleft P_i$, then $gN_ig^{-1} \triangleleft A$. □

Lemma 7.2.3. *Let A be an edge or vertex group of $\mathcal{H}$. Then for all sufficiently long fillings:*

$$\pi : (G, \mathcal{P}) \to (\overline{G}, \overline{\mathcal{P}})$$

the following hold:

1. *The subgroup $\overline{A} := \phi(A)$ is fully relatively quasiconvex in $(\overline{G}, \overline{P})$,*
2. *If $\overline{G}$ is hyperbolic, then $\overline{A}$ is quasiconvex in $\overline{G}$,*
3. *The subgroup $\overline{A}$ is isomorphic to the induced filling of A.*

Proof. There are only finitely many edge and vertex groups, so the first and third statements follow from Theorem 7.1.13.

If $\overline{A}$ is fully relatively quasiconvex in $(\overline{G}, \overline{P})$, then $\overline{A}$ is undistorted in $\overline{G}$ by [17, Theorem 10.5] and by [9, Corollary III.Γ.3.6], $\overline{A}$ is quasiconvex in $\overline{G}$. □

The third point also makes the filled hierarchy $\overline{H}$ faithful:

Corollary 7.2.4. *For all sufficiently long fillings $\pi : (G, \mathcal{P}) \to (\overline{G}, \overline{\mathcal{P}})$, the filled hierarchy $\overline{\mathcal{H}}$ on $\overline{G}$ is faithful.*

Proof. Let $\phi_e : A_e \to A_v$ be an attachment homomorphism mapping an edge group A_e to a vertex group A_v. Since $\pi(A_e)$ and $\pi(A_v)$ are isomorphic to the induced filling, so $\ker \pi|_{A_v} = \ker \pi \cap \ker \pi \cap A_v$ and $\ker \pi|_{A_e} = \ker \pi \cap \ker A_e$. Let $g_e \in A_v$ and let $\overline{\phi}_e : \overline{A}_e \to \overline{A}_v$ be the induced edge homomorphism. Then $\overline{\phi}_e \pi(g_e) = \pi\phi_e(g_e)$. If $\pi\phi_e(g_e) = 1$, then $\phi_e(g_e) \in \ker \pi$, so $g_e \in \ker \pi$. Therefore, $\pi(g_e) = 1$. Therefore, $\overline{\phi}_e$ is injective. □

The preceding results combine to produce a quasiconvex hierarchy:

Theorem 7.2.5 (see [3, Theorem 2.12])**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let $\mathcal{H}$ be a quasiconvex fully $\mathcal{P}$-elliptic hierarchy terminating in $\mathcal{P}$. For all sufficiently long peripherally finite fillings $\pi : (G, \mathcal{P}) \to (\overline{G}, \overline{\mathcal{P}})$ so that every $\overline{P} \in \overline{\mathcal{P}}$ is hyperbolic, the group $\overline{G}$ is hyperbolic and has a quasiconvex hierarchy terminating in $\overline{\mathcal{P}}$.*

Proof. By Corollary 7.2.4, the quotient $\overline{G}$ has a faithful hierarchy $\overline{H}$ where the underlying graphs and every vertex or edge group of $\overline{H}$ is the image of a vertex or edge group (respectively) of $\overline{H}$ under π.

By Lemma 7.2.3 (2), every edge and vertex group of $\overline{H}$ is quasiconvex in $\overline{G}$ and is hence also quasi-isometrically embedded in $\overline{G}$, so the hierarchy $\overline{H}$ is quasiconvex.

By construction, the terminal groups are fillings of the terminal groups of $\mathcal{H}$, so the terminal groups of $\overline{\mathcal{H}}$ are in $\overline{\mathcal{P}}$. □

Theorem 7.2.5 works for a group with a quasiconvex hierarchy, but Theorem 6.4.15 only gives a hierarchy for a finite index subgroup. When the filling kernels are chosen carefully, a filling of a finite index subgroup $G' \triangleleft G$ can be promoted to a filling of G.

Definition 7.2.6. *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair and let $G' \triangleleft G$ be a finite index normal subgroup with induced peripheral structure $(G', \mathcal{P}')$. Let $\{N'_j \triangleleft P'_j \mid P'_j \in \mathcal{P}'_j\}$ be a collection of filling kernels. The collection $\{N'_j\}$ is **equivariantly chosen** if*

1. *whenever gP'_jg^{-1} and hP'_kh^{-1} both lie in P_i, then $gN'_jg^{-1} = gN'_kh^{-1}$ and*
2. *every such gN'_jg^{-1} is normal in P_i.*

*An **equivariant filling** of $(G', \mathcal{P}')$ is a filling with equivariantly chosen filling kernels.*

An equivariant filling of $(G', \mathcal{P}')$ will induce a nice equivariant filling of $(G, \mathcal{P})$:

Proposition 7.2.7. *An equivariant filling $(G', \mathcal{P}') \to (\overline{G}', \overline{\mathcal{P}}')$ determines a filling $(G, \mathcal{P}) \to (\overline{G}, \overline{\mathcal{P}})$ so that $\overline{G}'$ is finite index normal in $\overline{G}$ and $(\overline{G'}, \overline{\mathcal{P}}')$ is the peripheral structure induced by $(\overline{G}, \overline{\mathcal{P}})$.*

Theorem 7.2.8 (Theorem B Restated)**.** *Let $(G, \mathcal{P})$ be a relatively hyperbolic pair with $\mathcal{P} = \{P_1, \ldots, P_m\}$. There exist subgroups $\{\dot{P}_i \triangleleft P_i\}$ where $\dot{P}_i$ is finite index in P_i such that if $\overline{G} = G(N_1, \ldots, N_m)$ is any peripherally finite filling with $N_i \triangleleft \dot{P}_i$, then $\overline{G}$ is hyperbolic and virtually special.*

Proof. By Theorem 6.4.15, there exists $(G', \mathcal{P}')$ such that $G' \triangleleft G$ is finite index and $(G', \mathcal{P}')$ has a quasiconvex, malnormal fully $\mathcal{P}'$–elliptic hierarchy terminating in $\mathcal{P}'$. Since G is virtually special and hence residually finite, there exist arbitrarily long peripherally finite fillings of $(G', \mathcal{P}')$ that are sufficiently long for Theorem 7.2.5 to hold.

Let $G(K_1, \ldots, K_M)$ be such a peripherally finite filling which is also sufficiently long so that Theorem 7.1.3 holds. Now pass to subgroups of the filling kernels to obtain an equivariant filling; let:

$$K'_i = \bigcap\{K_j^g \,|\, g \in G,\ \#(K_j^g \cap P_i) = \infty\}.$$

The new filling kernels $K'_i \le K_i$, so the new filling $G'(K'_1, \ldots, K'_M)$ is still sufficiently long and remains peripherally finite. By Proposition 7.2.7, the filling $G'(K'_1, \ldots, K'_M)$ determines a filling of G.

Consider a filling $G(N_1, \ldots, N_m)$ so that for each i:

1. $N_i \lhd P_i$
2. $N_i \le \dot{P}_i$ and
3. P_i/N_i is virtually special and hyperbolic.

with an induced equivariant filling:

$$G' \to G'(N'_1, \ldots, N'_M)$$

so that $N'_j \le K'_j$ and $N'_j \lhd P'_j$ for each j. Such a filling is sufficiently long so that Theorem 7.2.5 holds.

Therefore $G'(N'_1, \ldots, N'_M)$ has a quasiconvex hierarchy terminating in $\overline{\mathcal{P}}' = \{P'_j/N'_j\}$. By Theorem 7.1.3, the pair $(G', \overline{\mathcal{P}}')$ is relatively hyperbolic, so G' is hyperbolic because the elements of $\mathcal{P}'$ are finite.

Then $G'(N'_1, \ldots, N'_M)$ is a hyperbolic group with an malnormal quasiconvex hierarchy that terminates in finite groups (which are hence hyperbolic and virtually special). So by Corollary 7.0.2 (see [35] Theorem 13.3), $G'(N'_1, \ldots, N'_M)$ is virtually special. By Proposition 7.2.7, $G'(N'_1, \ldots, N'_M)$ is finite index normal in $G(N_1, \ldots, N_m)$, so the filling $G(N_1, \ldots N_m)$ is also virtually special. □

BIBLIOGRAPHY

[1] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[2] Ian Agol, Daniel Groves, and Jason Fox Manning. Residual finiteness, QCERF and fillings of hyperbolic groups. *Geom. Topol.*, 13(2):1043–1073, 2009.

[3] Ian Agol, Daniel Groves, and Jason Fox Manning. An alternate proof of Wise's malnormal special quotient theorem. *Forum Math. Pi*, 4:e1, 54, 2016.

[4] Mark Baker and Daryl Cooper. A combination theorem for convex hyperbolic manifolds, with applications to surfaces in 3-manifolds. *J. Topol.*, 1(3):603–642, 2008.

[5] Hyman Bass. Covering theory for graphs of groups. *Journal of Pure and Applied Algebra*, 89(1):3 – 47, 1993.

[6] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[7] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[8] Corey Bregman. Automorphisms and homology of non-positively curved cube complexes. *arXiv:1609.03602*, 2016.

[9] 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.

[10] Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Mark Sapir.

[11] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.

[12] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[13] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.

[14] Frédéric Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.

[15] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.

[16] G. Christopher Hruska. Geometric invariants of spaces with isolated flats. *Topology*, 44(2):441–458, 2005.

[17] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[18] G. Christopher Hruska and Bruce Kleiner. Hadamard spaces with isolated flats. *Geom. Topol.*, 9:1501–1538, 2005. With an appendix by the authors and Mohamad Hindawi.

[19] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.

[20] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. *Invent. Math.*, 199(2):293–331, 2015.

[21] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[22] Jason Fox Manning and Eduardo Martínez-Pedroza. Separation of relatively quasiconvex subgroups. *Pacific J. Math.*, 244(2):309–334, 2010.

[23] B. Martelli. *An Introduction to Geometric Topology*. CreateSpace Independent Publishing Platform, 2016.

[24] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.

[25] Denis V. Osin. Peripheral fillings of relatively hyperbolic groups. *Inventiones mathematicae*, 167(2):295–326, Feb 2007.

[26] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[27] Michah Sageev and Daniel T. Wise. Periodic flats in CAT(0) cube complexes. *Algebr. Geom. Topol.*, 11(3):1793–1820, 2011.

[28] Michah Sageev and Daniel T. Wise. Cores for quasiconvex actions. *Proc. Amer. Math. Soc.*, 143(7):2731–2741, 2015.

[29] Peter Scott. Subgroups of surface groups are almost geometric. *J. London Math. Soc. (2)*, 17(3):555–565, 1978.

[30] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin-New York, 1980. Translated from the French by John Stillwell.

[31] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[32] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.

[33] Daniel T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

[34] 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; by the American Mathematical Society, Providence, RI, 2012.

[35] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy. 2012.