ISRAEL JOURNAL OF MATHEMATICS **168** (2008), 317–429
DOI: 10.1007/s11856-008-1070-6

DEHN FILLING IN RELATIVELY HYPERBOLIC GROUPS

BY

DANIEL GROVES*,†

California Institute of Technology
e-mail: groves@caltech.edu

AND

JASON FOX MANNING**,†

University at Buffalo, SUNY
e-mail: j399m@buffalo.edu

ABSTRACT

We introduce a number of new tools for the study of relatively hyperbolic groups. First, given a relatively hyperbolic group G, we construct a nice combinatorial Gromov hyperbolic model space acted on properly by G, which reflects the relative hyperbolicity of G in many natural ways. Second, we construct two useful bicombings on this space. The first of these, **preferred paths**, is combinatorial in nature and allows us to define the second, a relatively hyperbolic version of a construction of Mineyev.

As an application, we prove a group-theoretic analog of the Gromov-Thurston 2π Theorem in the context of relatively hyperbolic groups.

CONTENTS

* The first author was supported in part by NSF Grant DMS-0504251.
** The second author was supported in part by an NSF Mathematical Sciences Postdoctoral Research Fellowship.
† Both authors thank the NSF for their support. Most of this work was done while both authors were Taussky-Todd Fellows at Caltech.
Received September 14, 2006 and in revised form June 18, 2007

1. Introduction

A finitely generated group is **word hyperbolic** [16] if it acts properly and cocompactly on a metric space (e.g., its Cayley graph) satisfying a certain coarse geometric property called Gromov hyperbolicity (Definition 2.6 below). Two spaces acted on properly and cocompactly by the same group will have the same coarse geometry, so word hyperbolicity depends only on the group in question. For example, the fundamental group of a compact hyperbolic n-manifold acts properly and cocompactly on the hyperbolic space $\mathbb{H}^n$, and so is word hyperbolic. If G is the fundamental group of a non-compact but finite volume hyperbolic n-manifold, then G also has an apparently natural action on $\mathbb{H}^n$. We say that groups with such actions on Gromov hyperbolic spaces are **relatively hyperbolic**. Considering the case of fundamental groups of finite area hyperbolic surfaces, we see that it is necessary first to specify a collection of "peripheral" subgroups to determine an action up to some kind of coarse equivalence. Thus it makes no real sense to ask whether a group is relatively hyperbolic, but only to ask whether a group is hyperbolic *relative* to a collection of subgroups. There are several competing ways to say what it means for a group to be (strongly) hyperbolic relative to a collection of subgroups [16, 13, 4]. These definitions are now all known to be equivalent.[1] (See Section 2.9 for precise definitions and more examples of relatively hyperbolic groups.)

This paper has three main purposes. First, we introduce a new space (the "cusped space") for studying relatively hyperbolic groups (Section 3). Second, we construct a pair of useful bicombings on this space (Sections 5 and 6). Third, we extend Thurston's Hyperbolic Dehn Surgery Theorem to the context of

[1] In [35], Osin gives a more general definition where it is not assumed that parabolic subgroups are finitely generated, or that there are finitely many conjugacy classes of parabolic subgroups. Since the current paper appeared as a preprint, Chris Hruska [23] has extended the definition proposed here (Theorem 3.25.(5)) to the infinitely generated setting, and proved its equivalence to Osin's definition.

(torsion-free) relatively hyperbolic groups (Part 2). We discuss these now in turn.

1.1. A NEW GEOMETRY FOR RELATIVELY HYPERBOLIC GROUPS. Roughly speaking, a group G is hyperbolic relative to a collection $\mathcal{P}$ of subgroups if the geometry of G is δ-hyperbolic, except for that part corresponding to the subgroups in $\mathcal{P}$. In Subsection 3.1 we define **combinatorial horoballs** which are a way of embedding any graph into a δ-hyperbolic graph. We build the cusped space by gluing onto the Cayley graph of G a collection of combinatorial horoballs, one for each coset of each subgroup in $\mathcal{P}$. The resulting space is Gromov hyperbolic if and only if G is hyperbolic relative to $\mathcal{P}$. In case the space is Gromov hyperbolic, the action of G on the space satisfies the conditions given in Gromov's original definition of relative hyperbolicity (Definition 2.38). The cusped space thus combines the most useful combinatorial and coarse geometric aspects of relatively hyperbolic groups. A closely related construction appears in work of Cannon and Cooper [7] (also see [38] for another related construction). As in [7], if G is the fundamental group of a finite volume hyperbolic n-manifold, and $\mathcal{P}$ consists of the fundamental groups of the cusps, the space we build is quasi-isometric to $\mathbb{H}^n$, though we do not provide a proof of this here.

Part of the novelty of the cusped space compared to the one in [7] is that it is a graph (metrized so the length of each edge is 1), so that the metric and combinatorial aspects harmonize with each other more easily. This allows us to more easily adapt a number of constructions and results in word hyperbolic groups to the relative setting. In particular, we consider combinatorial isoperimetric inequalities (Subsection 2.6), homological isoperimetric inequalities (Subsection 2.7) and a construction of Mineyev from [29]. Considering the different types of isoperimetric inequalities and spaces, we get a number of new characterizations of relatively hyperbolic groups.

Let G be a finitely generated group which is finitely presented relative to a collection $\mathcal{P} = \{P_1, \ldots, P_n\}$ of finitely generated subgroups (see Subsection 2.4). Let $\hat{\Gamma}$ be the coned-off Cayley graph for G with respect to $\mathcal{P}$ (see Definition 2.41), $\hat{C}$ the coned-off Cayley complex (see Definition 2.47), and let X be the cusped space associated to G and $\mathcal{P}$ (defined in Section 3).

Then we have

THEORME 3.25: *The following are equivalent:*

(1) *G is hyperbolic relative to $\mathcal{P}$ in the sense of Gromov;*
(2) *G is hyperbolic relative to $\mathcal{P}$ (i.e., $\hat{\Gamma}$ is Gromov hyperbolic and fine);*
(3) *$\hat{C}$ satisfies a linear combinatorial isoperimetric inequality;*
(4) *$\hat{C}$ satisfies a linear homological isoperimetric inequality;*
(5) *$X^{(1)}$ is Gromov hyperbolic;*
(6) *X satisfies a linear combinatorial isoperimetric inequality;*
(7) *X satisfies a linear homological isoperimetric inequality.*

See Subsection 2.6 and Definition 2.28 for definitions of linear isoperimetric inequalities (both combinatorial and homological).

The equivalence of (1) and (2) is by now well-known (see, for example, [10, Appendix]). As far as we are aware, the equivalence of (2) and (3) has not appeared elsewhere, though it is implicit in [35]. What is really novel about Theorem 3.25 is the space X and the use of homological isoperimetric inequalities.

1.2. Bicombings on relatively hyperbolic groups. The second main purpose of this paper is the construction in Sections 4–6 of a pair of useful bicombings on the cusped space. In Section 4, we prove a general result about convex sets and between-ness in a Gromov hyperbolic space Υ. Given a family $\mathcal{G}$ of 'sufficiently separated' convex sets we construct, for any pair of points $a, b \in \Upsilon$, a canonical collection of sets in $\mathcal{G}$ which are 'between' a and b. These collections satisfy a number of axioms, and allow a great deal of combinatorial control over triangles in Υ built from quasi-geodesics using our construction.

This analysis is carried out in Section 5, where we define **preferred paths** for our space X. These give a G-equivariant bicombing of X by uniform quasi-geodesics, whose intersection with horoballs is very controlled (this is where the results from Section 4 are used). In particular (possibly partially ideal) triangles whose sides are preferred paths have very well controlled combinatorial structure (see Subsection 5.2 for details on this).

We expect that the construction in Section 4 and that of preferred paths in Section 5 will have many applications. The first is the bicombing q which is defined in Section 6. This gives a relatively hyperbolic version of a construction of Mineyev from [29]. Applications of Mineyev's construction are myriad (see, for example, [29, 30, 31, 32, 33, 42]). It can reasonably be expected that many of these results can be extended to the relatively hyperbolic setting using the bicombing from Section 6 of this paper, or variations on it. In particular, in [19],

we define a homological bicombing on the coned-off Cayley graph of a relatively hyperbolic group (using the bicombing from this paper in an essential way) in order to investigate relative bounded cohomology and relatively hyperbolic groups, in analogy with [29] and [30]. Also, in [17], the first author proves that if the parabolic subgroups of G act 'nicely' on a strongly bolic metric space (as defined by Lafforgue [28]) then so does G. Using the work of Lafforgue [28] and Druţu and Sapir [11], this has implications for the Baum-Connes conjecture for certain relatively hyperbolic groups.

It is also worth noting that in Part 2 of this paper, the major tool is preferred paths. The only time we need the homological bicombing (which is the analogue of Mineyev's construction) is in the proof of Theorem 11.11. Otherwise, we use only the results from Section 5, which have no relation to Mineyev's construction.

1.3. Relatively hyperbolic Dehn filling. Part 2 of this paper is devoted to another application of the constructions of Part 1. We investigate a group theoretic analogue of Dehn filling, which is the third and final purpose of this paper.

We first briefly remind the reader what is meant by "Dehn filling" in the context of 3-manifolds. Suppose that M is a compact 3-manifold, with some component T of ∂M homeomorphic to a torus. Let α be an essential simple closed curve in T. Let W be a solid torus, and μ a meridian for W (a curve which bounds an embedded disk in W), and let $\phi\colon (\partial W, \mu) \to (T, \alpha)$ be a homeomorphism of pairs. The 3-manifold

$$M(\alpha) = M \cup_\phi W$$

obtained by gluing using ϕ is called the **Dehn filling of** M **along** α, and depends up to homeomorphism only on the homotopy class of α in T. Thurston's Hyperbolic Dehn Surgery Theorem [40] says that if the interior of M admits a hyperbolic metric, then so does the interior of $M(\alpha)$, for all but finitely many choices of α. The number of curves to be excluded and the relationship between the geometry of M and that of $M(\alpha)$ can be made quite precise (see for example [22]).

On the level of fundamental groups, $\pi_1(M(\alpha)) = \pi_1(M)/\langle\langle a \rangle\rangle$, where a is an element of $\pi_1(M)$ whose conjugacy class is represented by α. One of the group theoretic statements implied the Hyperbolic Dehn Surgery Theorem is

THEOREM 1.1 ([40]): *Let $G = \pi_1(M)$ where M is a hyperbolic 3-manifold with a single torus cusp C, and let $P = \pi_1(C) \cong \mathbb{Z} \oplus \mathbb{Z}$ be the cusp subgroup. Then for all but finitely many $a \in P$, $G/\langle\langle a \rangle\rangle$ is infinite, non-elementary, and word hyperbolic.*

This group theoretic statement finds its strongest quantitative formulation in the "6 Theorem" independently due to Lackenby and Agol [26, 1].

Theorem 1.1 puts the conclusion of the Hyperbolic Dehn Surgery Theorem in an algebraic context. Koji Fujiwara asked whether there was an algebraic analogue of this theorem which also puts the hypotheses into an algebraic context. We learned of this question from Danny Calegari.

As an answer to this question, we provide the following result (where $|K_i|_{P_i}$ denotes the minimal length of a nontrivial element of K_i using the word metric on P_i with respect to the generating set $S \cap P_i$; see Definition 7.1):

THEOREM 7.2: *Let G be a torsion-free group, which is hyperbolic relative to a collection $\mathcal{P} = \{P_1, \ldots, P_n\}$ of finitely generated subgroups. Suppose that S is a generating set for G so that for each $1 \leq i \leq n$ we have $P_i = \langle P_i \cap S \rangle$.*

There exists a constant B depending only on $(G, \mathcal{P})$ so that for any collection $\{K_i\}_{i=1}^n$ of subgroups satisfying

- *$K_i \trianglelefteq P_i$; and*
- *$|K_i|_{P_i} \geq B$,*

then the following hold, where K is the normal closure in G of $K_1 \cup \cdots \cup K_n$.

(1) *The map $P_i/K_i \xrightarrow{\iota_i} G/K$ given by $pK_i \mapsto pK$ is injective for each i.*
(2) *G/K is hyperbolic relative to the collection*

$$\mathcal{Q} = \{\iota_i(P_i/K_i) \colon 1 \leq i \leq n\}.$$

It is well-known (see, for example, [13, Theorem 3.8]) that a group which is hyperbolic relative to a collection of word hyperbolic subgroups is word hyperbolic. Thus, an immediate corollary of Theorem 7.2 is the following:

COROLLARY 1.2: *Under the hypotheses of Theorem 7.2, if each of the P_i/K_i are themselves word hyperbolic, then G/K is word hyperbolic.*

Together with Theorem 11.12 (non-elementariness), Corollary 1.2 unifies a number of known results:

(1) We have already remarked that Theorem 7.2 generalizes Theorem 1.1. However, since we make no assumptions about the subgroups P_i, Theorem 7.2 also generalizes the results of Lackenby from [27] about filling 3-manifolds some of whose boundary components are higher genus surfaces (although we do not obtain such nice quantitative bounds as Lackenby).
(2) Corollary 1.2 also generalizes some known results about hyperbolic groups. For example, modulo the extra torsion-free hypothesis, Theorem 7.2 is a generalization of statements (1)–(3) of [16, Theorem 5.5.D, p. 149]. Of interest in this context is that we make no use whatsoever of small cancellation techniques.
(3) Much of the group-theoretic content of many "CAT(−1)" or "CAT(0) with isolated flats" filling constructions on hyperbolic manifolds with torus cusps is also contained in Theorem 7.2. Examples of this are in [39, 34]. (See also Remark 1.6.)

We now make a few more remarks about Theorem 7.2.

Remark 1.3: The "short" fillings really must be excluded in Theorem 7.2, as can be seen, for instance, from the many examples of exceptional fillings of hyperbolic 3-manifolds. By considering fillings of the Hopf link, we can see that it is also important that each of the lengths $|K_i|_{P_i}$ is large.

An even simpler example is given by G equal to the free group $\langle x, y\rangle$, and $\mathcal{P} = \{P_1 = \langle x\rangle, P_2 = \langle y\rangle, P_3 = \langle xy\rangle\}$. Choosing $K_1 = \langle x^p\rangle$, $K_2 = \langle y^q\rangle$, and $K_3 = \langle (xy)^r\rangle$, the quotient G/K will be infinite and word hyperbolic if and only if $1/p + 1/q + 1/r < 1$. This occurs if all three of p, q, and r are at least 4, but of course if $p = q = 2$, then r can be arbitrarily large while G/K remains finite.

Remark 1.4: Since we have been, from the very beginning, working in the coarse world of δ-hyperbolic spaces, we have no hope of obtaining the fine control over constants, as obtained in [26, 1, 22]. Therefore, we have made very little attempt throughout this paper to make our constants optimal.

However, it is worth remarking that there are some delicate interdependencies between some of the constants we use.

Remark 1.5: Denis Osin has independently proved Theorem 7.2; see [36]. In fact, Osin works in a somewhat more general setting, in two respects.

First, Osin has a more general notion of relative hyperbolicity, which allows infinitely many (possibly infinitely generated) parabolics. It can be shown (and this is really implicit in Osin's proof) that the appropriate statement for infinitely many infinitely generated parabolics follows from the statement for finitely many finitely generated parabolics.[2]

Second, and more seriously, Osin makes no assumption of torsion-freeness in [36]. We believe that our methods apply (with appropriate modification) to the non-torsion-free case, at least to prove the analogue of Theorem 7.2, but at such a cost in clarity and brevity that we have elected to deal only with the torsion-free case. We have tried to make explicit our use of torsion-freeness and how one might go about avoiding it (see Remarks 3.29, 5.2, 5.8, 5.45, 6.12, 5.18,8.8 and 9.10; on first reading, we recommend ignoring all of these comments; to facilitate this, they are all labelled as 'Remark (about torsion)'). Certain of our results, which are not used in the proof of Theorem 7.2, must be considerably rephrased in the presence of torsion (see especially Theorem 9.3).

Note also that Osin's main theorem [36, Theorem 1.1] states that given a finite set $\mathcal{F} \subset G$, there is a B so that under the conditions of Theorem 7.2 the map $G \to G/K$ is injective on $\mathcal{F}$. This is Corollary 9.7 below.

Remark 1.6: It is worth noting that, even when starting with a rank one locally symmetric space, our group-theoretic version of filling produces hyperbolic groups in many situations where the existence a locally CAT(-1) filling is not at all clear; see the non-existence results of [25]. The advantage of a CAT(-1) filling is that information about the fundamental group can be obtained from local information about the locally CAT(-1) model. For more on this, see [14], which is in preparation.

2. Preliminaries

2.1. Coarse geometry. All metric spaces will be assumed to be complete geodesic metric spaces, and the distance between two points x and y will usually be denoted $d(x, y)$.

Definition 2.1: If X is a metric space, $A \subset X$ and $R \geq 0$, then let $N_R(A)$ be the R-neighborhood of A in X.

[2] For more on this, see [21].

Definition 2.2: If X and Y are metric spaces, $K \geq 1$ and $C \geq 0$, a (K, C)-**quasi-isometric embedding** of X into Y is a function $q\colon X \to Y$ so that for all $x_1, x_2 \in X$

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

If in addition the map q is C-**coarsely onto**, i.e., $N_C(q(X)) = Y - q$, it is called a (K, C)-**quasi-isometry**. The two metric spaces X and Y are then said to be **quasi-isometric** to one another. This is a symmetric condition.

Definition 2.3: A (K, C)-**quasi-geodesic** in X is a (K, C)-quasi-isometric embedding $\gamma\colon \mathbb{R} \to X$. We will occasionally abuse notation by referring to the image of γ as a quasi-geodesic.

A (K, C)-**quasi-geodesic ray** is a (K, C)-quasi-isometric embedding

$$p\colon \mathbb{R}_{\geq 0} \to X.$$

2.2. Gromov hyperbolic spaces. Given a geodesic triangle $\Delta(x, y, z)$ in any geodesic metric space, there is a unique **comparison tripod**, Y_{xyz}, a metric tree so that the distances between the three extremal points of the tree, $\overline{x}$, $\overline{y}$ and $\overline{z}$, are the same as the distances between x, y and z (see Figure 1). There

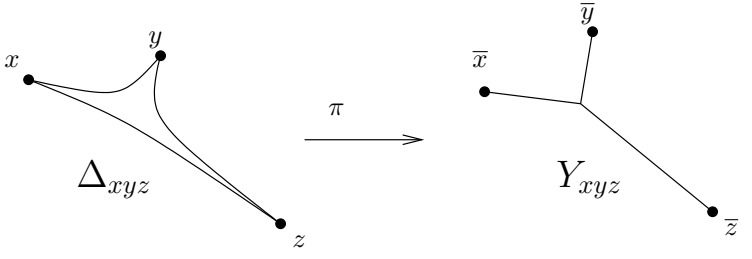

Figure 1. A triangle and its comparison tripod

is a unique map

$$\pi\colon \Delta(x, y, z) \to Y_{xyz}$$

which takes x to $\overline{x}$, y to $\overline{y}$ and z to $\overline{z}$, and which restricts to an isometric embedding on each side of $\Delta(x, y, z)$.

Definition 2.4: Let $\delta \geq 0$. The triangle $\Delta(x, y, z)$ is δ-**thin** if the diameter of $\pi^{-1}(p)$ is at most δ for every point $p \in Y_{xyz}$.

Remark 2.5: Any triangle which is δ-thin is also **δ-slim**, i.e., any side of the triangle is contained in the δ-neighborhood of the union of the other sides. More generally, a geodesic n-gon is $(n-2)\delta$-slim. Also, for every δ there is a δ' so that δ-slim triangles are δ'-thin (see [5, Proposition III.H.1.17, p. 408]).

Definition 2.6: A geodesic metric space X is **δ-hyperbolic** if every geodesic triangle in X is δ-thin. If δ is unimportant we may simply say that X is **Gromov hyperbolic**.

See [5, Chapter III.H] for the background and many basic results about Gromov hyperbolic spaces.

Definition 2.7: Let $x,\ y,\ z \in X$. The **Gromov product** of x and y with respect to z is $(x,y)_z = \frac{1}{2}(d(x,z) + d(y,z) - d(x,y))$. Equivalently, $(x,y)_z$ is the distance from $\overline{z}$ to the central vertex of the comparison tripod Y_{xyz} for any geodesic triangle $\Delta(x,y,z)$.

Definition 2.8: Fix some $z \in X$, where X is some Gromov hyperbolic metric space. We say that a sequence $\{x_i\}$ **tends to infinity** if $\liminf_{i,j\to\infty}(x_i,x_j)_z = \infty$. On the set of such sequences we may define an equivalence relation: $\{x_i\} \sim \{y_i\}$ if $\liminf_{i,j\to\infty}(x_i,y_j)_z = \infty$. The **Gromov boundary** of X, also written ∂X, is the set of equivalence classes of sequences tending to infinity. The Gromov boundary does not depend on the choice of z (see [5, Proposition III.H.3.7]).

Remark 2.9: We may topologize $X \cup \partial X$ so that if $\{x_i\}$ tends to infinity, then $\lim_{i\to\infty} x_i = [\{x_i\}]$. Furthermore, if $\gamma\colon [0,\infty) \to X$ is a quasi-geodesic ray, then for any sequence $\{t_i\}$ with $\lim_{i\to\infty} t_i = \infty$, the sequence $\{\gamma(t_i)\}$ tends to infinity. The point $\{\gamma(t_i)\} \in \partial X$ does not depend on the choice of $\{t_i\}$. An isometric action on X extends to a topological action on the boundary.

Remark 2.10: We will implicitly assume, whenever we say that a space is δ-hyperbolic, that triangles are δ-thin and δ-slim.

This can be achieved by replacing δ by some larger constant.

We will also assume that δ is an integer.

One can also consider geodesic "triangles" in a Gromov hyperbolic space X whose vertices are **ideal**, i.e., points in ∂X. The following is a simple exercise in δ-slim triangles.

LEMMA 2.11: *A geodesic triangle in a δ-hyperbolic space with some or all vertices ideal is 3δ-slim.*

2.3. CAYLEY GRAPH. Although there is little difference between various definitions of the Cayley graph, it is convenient to fix one here. By "G is generated by S", we mean that there is a surjection

$$\pi\colon \mathcal{F}(S) \to G,$$

where $\mathcal{F}(S)$ is the free group on the set S. If $S \subset G$, then we implicitly assume π is the homomorphism induced by inclusion. The **Cayley graph of G with respect to S**, written $\Gamma(G, S)$, is the graph with vertex set G, and edge set $G \times S$. The edge (g, s) connects the vertices g and $g\pi(s)$.

2.4. RELATIVE PRESENTATIONS. We recall the following definitions of Osin. (We change the notation slightly.)

Definition 2.12 ([35, Definition 2.1]): Let G be a group, $\{H_\lambda\}_{\lambda\in\Lambda}$ a collection of subgroups of G, and A a subset of G. We say that A is *a relative generating set for G with respect to $\{H_\lambda\}_{\lambda\in\Lambda}$* if G is generated by

$$A \cup \left(\bigcup_{\lambda\in\Lambda} H_\lambda \right).$$

We will be concerned with situations where the index set Λ is finite.

Definition 2.13 (Osin): Suppose that G is generated by $\mathcal{A}$ with respect to $\{H_\lambda\}_{\lambda\in\Lambda}$. Then G is a quotient of

$$F = F(\mathcal{A}) * (*_{\lambda\in\Lambda} H_\lambda),$$

where $F(\mathcal{A})$ is the free group on the alphabet $\mathcal{A}$. Suppose that N is the kernel of the canonical quotient map from F to G. If N is the normal closure of the set $\mathcal{R}$ then we say that

$$\langle \mathcal{A}, \{H_\lambda\}_{\lambda\in\Lambda} \mid \mathcal{R} \rangle,$$

is a **relative presentation** for G with respect to $\{H_\lambda\}_{\lambda\in\Lambda}$.

We say that G is **finitely presented relative to** $\{H_\lambda\}_{\lambda\in\Lambda}$ if we can choose $\mathcal{R}$ to be finite.

The following lemma is essentially contained in Theorem 2.44 of [35].

LEMMA 2.14: *If G is a finitely generated group, finitely presented relative to a collection of nontrivial subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$, then Λ is finite and each H_λ is finitely generated.*

Proof. By Theorem 2.44 of [35], there is a finite subset $\Lambda_0 \subseteq \Lambda$ so that

$$G \cong (*_{\lambda\in\Lambda\smallsetminus\Lambda_0} H_\lambda) * G_0,$$

where G_0 is generated by $\bigcup_{\lambda\in\Lambda_0} H_\lambda$ and the finite relative generating set for G.

By Grushko's Theorem, and the fact that G is finitely generated, $\Lambda \smallsetminus \Lambda_0$ is also finite. Moreover, each H_λ for $\lambda \in \Lambda \smallsetminus \Lambda_0$ is finitely generated, as is G_0.

Osin's theorem further asserts that G_0 has the structure of a tree of groups, where each edge group is finitely generated. The collection of vertex groups is $\{H_\lambda\}_{\lambda\in\Lambda_0} \cup \{Q\}$ where Q is some finitely generated group. Since G_0 is finitely generated, an elementary application of Bass-Serre theory shows that each H_λ is finitely generated. ∎

Definition 2.15: Suppose that G is finitely generated, and also finitely presented relative to $\{H_1, \ldots, H_m\}$. By Lemma 2.14, each of the H_i is finitely generated. In this situation, we usually fix a finite generating set S for G so that for each $1 \leq i \leq m$ we have $\langle S \cap H_i\rangle = H_i$. We call such a set S a **compatible** generating set for G.

Definition 2.16 (Relative Cayley complex): Suppose that G is finitely presented relative to the finitely generated subgroups $\{H_1, \ldots, H_m\}$, and that S is a finite compatible generating set for G. Let

$$\langle \mathcal{A}, \{H_1, \ldots, H_m\} \mid \mathcal{R}\rangle,$$

be a finite relative presentation for G, where $\mathcal{A} \subset S$ is a set of relative generators for G.

Let $\Gamma = \Gamma(G, S)$ be the Cayley graph of G with respect to S. The elements of $\mathcal{R}$ correspond to loops in Γ. We glue a 2-cell to each such loop, in a manner equivariant under the G-action on Γ. The resulting 2-complex is called the **relative Cayley complex** of G with respect to $\langle S, \{H_1, \ldots, H_m\} \mid \mathcal{R}\rangle$, and is denoted $C(G, S, \mathcal{R})$.

In general, the relative Cayley complex of G will not be simply-connected. This will only be the case if $\langle S \mid \mathcal{R}\rangle$ is already a presentation for G.

In Subsection 2.9 below, we define a simply-connected 2-complex $\hat{C}$ associated to a finite relative presentation of relatively hyperbolic group, starting with the relative Cayley complex.

Also, in Section 3, we construct another simply-connected 2-complex X, again starting with a finite relative presentation of a relatively hyperbolic group.

The spaces X and $\hat{C}$ both contain copies of the relative Cayley complex.

2.5. Combinatorial maps, chains, etc.

Definition 2.17: Let Y be a cell complex, and let $\omega = \sum_i \alpha_i \omega_i$ be a real cellular n-chain. The **1-norm** of ω is

$$|\omega|_1 = \sum_i |\alpha_i|.$$

Definition 2.18: Suppose that Y is a 2-complex and that $p\colon I \to Y^{(1)}$ is a combinatorial path, in the sense that there is a cell structure on I so that p sends each edge either to an edge or a vertex of Y.

The **length**, or **1-norm** of p is then the number of 1-cells in I which are mapped onto edges of Y. We denote the 1-norm of p by $|p|_1$.

The map p induces an obvious cellular 1-chain $\underline{p}$ on Y (using the orientation on I).

Definition 2.19: Suppose that Y is a 2-complex and Σ an oriented, cellulated surface. A **combinatorial map** is a map $f\colon \Sigma \to Y$ which sends each vertex of Σ to a vertex of Y, each edge of Σ to an edge or a vertex of Y and each 2-cell to a 2-cell, an edge, or a vertex of Y. Furthermore, if σ is a cell in Σ which is sent to a cell of the same dimension then the interior of σ is mapped homeomorphically by f onto its image.

The **area** of a combinatorial map $f\colon \Sigma \to Y$ is the number of 2-cells in Σ which are mapped onto 2-cells. When we refer to the '1-norm' of a map between 2-complexes, we mean the area.

As in the 1-dimensional setting, a combinatorial map $f\colon \Sigma \to Y$ induces an integral, cellular 2-chain $\underline{f}$ on Y.

Remark 2.20: If f is a combinatorial map of an interval or a surface into a 2-complex Y then $|\underline{f}|_1 \leq |f|_1$.

Remark 2.21: Throughout this paper we are somewhat cavalier about the difference between paths as maps, paths as subsets, and 1-chains. This should not cause any confusion.

2.6. Combinatorial isoperimetric inequalities. Any Gromov hyperbolic space satisfies a linear coarse isoperimetric inequality (see Proposition III.H.2.7 in [5]). However, when we work with homological isoperimetric inequalities below, we need to use simply connected spaces.

Thus we pause in this paragraph to consider combinatorial isoperimetric inequalities. There is little novel here, but it is worth noting that our spaces are not always locally finite, are not uniformly locally finite even when they are locally finite, and there will rarely be a proper and cocompact action on the Gromov hyperbolic spaces in this paper. Thus we need to be slightly careful about the hypotheses in the results below.

Proposition 2.22: *Let X be a simply-connected 2-complex and suppose that $X^{(1)}$ is δ-hyperbolic. Suppose further that for some $K \geq 0$, any combinatorial loop of length at most 16δ can be filled with a combinatorial disk of area at most K. Then any combinatorial loop c in X can be filled with a combinatorial disk whose area is at most $K|c|_1$.*

Proof. This is essentially Dehn's algorithm. See the proof of [5, Proposition III.H.2.7]. ∎

Proposition 2.23: *Suppose that X is a simply connected 2-complex and that there is a constant $M \geq 0$ so that the length of the attaching map of any 2-cell is at most M.*

If X satisfies a linear combinatorial isoperimetric inequality, then $X^{(1)}$ is δ-hyperbolic for some δ.

Furthermore, δ can be calculated in terms of M and the isoperimetric constant of X.

Proof. Follows from the proof of [5, Theorem III.H.2.9], or from a combinatorial version of the proof of Theorem 2.30 below. ∎

2.7. Homological things. In [15], Gersten proves that a group is hyperbolic if and only if it has a linear homological filling function. In this section we recall those notions and extend them slightly in order to account for the fact that the actions on Gromov hyperbolic spaces in this paper are rarely proper

and cocompact (though they are usually one or the other). In this section we prefer to work with rational coefficients. However, real or complex coefficients would also work.

We recall some standard definitions:

Definition 2.24: Let Y be a cell complex, A a normed abelian group (with norm $\|\cdot\|$), and let $\mathcal{C}_n$ be the set of n-cells of Y. A **locally finite n-chain in Y with coefficients in A**, ω is a formal sum

$$\omega = \sum_{c \in Y} \omega_c c,$$

where each $\omega_c \in A$. The chain ω is said to be **summable** if

$$|\omega|_1 := \sum_{c \in Y} \|a_c\| < \infty.$$

The quantity $|\omega|_1$ is the **norm of** ω. The **support** of ω is the union in Y of the n-cells c for which ω_c is nonzero.

The chains of compact support form the standard cellular chain group $C_n(Y, \mathbb{Z})$.

We first recall a result from [30] about expressing 1-chains as sums of paths. Note that this result is a generalization of [3, Theorem 3.3].

Definition 2.25: Suppose that Γ is a graph and T is a collection of vertices. A **T-path** is a directed path in Γ whose initial and terminal vertices are in T (or are equal).

A path is **simple** if it has no repeated vertices.

Let f be a summable 1-chain in a graph Γ. Let $\Gamma(f)$ denote the directed graph which is Γ with an orientation chosen so that $f(e) \geq 0$ for each edge e. Let $\Gamma_+(f)$ be the minimal subgraph of $\Gamma(f)$ containing all the edges e with $f(e) \neq 0$.

THEOREM 2.26 (Mineyev, Theorem 6, [30]): *Let Γ be a graph, T a set of vertices in Γ and f a summable 1-chain on Γ with coefficients in $\mathbb{Q}$ and $\mathrm{supp}(\partial f) \subset T$.*

(a) *There is a countable family $P = \{p_1, p_2, \ldots\}$ of simple T-paths in $\Gamma_+(f)$ and a sequence $\{\alpha_i\}$ in $\mathbb{Q} \cap [0, \infty)$ so that* (i) $f = \sum_i \alpha_i p_i$; *and* (ii) $|f|_1 = \sum_i \alpha_i |p_i|_1$.

(b) *If f has finite support then P can be chosen to be finite.*

Definition 2.27: Suppose that $f = \sum_i \alpha_i p_i$ expresses the 1-chain f as a sum of 1-chains $\alpha_i p_i$, where $\alpha_i \in \mathbb{Q}_{\geq 0}$. This sum is called **coherent** if $|f|_1 = \sum_i \alpha_i |p_i|_1$.

Definition 2.28: Suppose that Z is a simply-connected 2-complex. We say that Z satisfies a **linear homological isoperimetric inequality** if there is a constant $K \geq 0$ so that for any combinatorial loop c in Z there is some $\sigma \in C_2(Z; \mathbb{Q})$ with $\partial\sigma = \underline{c}$, satisfying

$$|\sigma|_1 \leq K|c|_1.$$

See [30, Theorem 7] for (many) other notions of what it might mean for a space to have a 'linear isoperimetric inequality'. In this paper, we exclusively use the notion from Definition 2.28 above.

The next result follows immediately from the definitions.

LEMMA 2.29: *If a simply-connected 2-complex Z satisfies a linear combinatorial isoperimetric inequality, then it satisfies a linear homological isoperimetric inequality.*

In [15], Gersten proves that the Cayley complex of a finitely presented group G satisfies a linear homological isoperimetric inequality if and only if G is hyperbolic. This is essentially a converse to Lemma 2.29 above. In this section we slightly generalize Gersten's result. However, most of our work is in noting that the proof of Gersten's result from [30] works in our setting.

THEOREM 2.30: *Suppose that Z is a simply-connected 2-complex and that there is a constant M so that the attaching map for each 2-cell in Z has length at most M.*

Suppose further that Z satisfies a linear homological isoperimetric inequality. Then $Z^{(1)}$, the 1-skeleton of Z, is δ-hyperbolic for some δ.

Proof. The proof consists of noting that a number of other proofs in the literature do not rely essentially on finite valence.

The first step is to prove that if Z is not Gromov hyperbolic then for any $\epsilon > 0$, there are ϵ-thick geodesic bigons in Z. This is essentially [37, Theorem 1.4], modified in the obvious way. Namely, for each M, define $f(r) = \inf\{d(\gamma(R+r)), \gamma'(R+r))\}$ where the infimum is taken over all positive integers R and all γ, γ', geodesics such that $\gamma(0) = \gamma'(0)$ and $d(\gamma(R), \gamma'(R)) \geq 2M^2$. This is not quite the function that Papasoglu uses, but it suffices for the proof. The remainder of the proof of this first step is identical to that in [37].

Now follow the proof of [30, Proposition 8] to prove that if there are ϵ-thick bigons in Z for $\epsilon > 0$, then Z does not have a linear homological isoperimetric inequality. Mineyev's argument relies only on a bound on the length of attaching maps for 2-cells (and not, for instance, on vertex transitivity or local finiteness). It is certainly worth remarking that Mineyev relies on a result of Gersten [15], which similarly requires only a bound on the attaching maps of 2-cells. ∎

2.8. Homological bicombings. There are various notions of "bicombing" for graphs. In particular, one can define combings made up of paths or of 1-chains.

Definition 2.31: Let Γ be a graph, and let $\mathrm{Geod}(\Gamma)$ be the set of (oriented) geodesics in Γ. That is, each element of $\mathrm{Geod}(\Gamma)$ is a path $\sigma\colon I \to \Gamma$, where $I \subset \mathbb{R}$ and σ is parametrized by arc length. A **geodesic bicombing on** Γ is a function

$$\gamma\colon \Gamma^{(0)} \times \Gamma^{(0)} \to \mathrm{Geod}(\Gamma)$$

so that $\gamma(x, y)$ is a geodesic which begins at x and ends at y.

Definition 2.32: Let Γ and $\mathrm{Geod}(\Gamma)$ be as in Definition 2.31. Let $\overline{\Gamma}$ be some compactification (or bordification) of $\Gamma^{(0)}$, and let $\Lambda \subseteq \Gamma$. A **geodesic bicombing on** Λ is a function

$$\gamma\colon \Lambda \times \Lambda \smallsetminus \Delta \to \mathrm{Geod}(\Gamma),$$

where $\Delta = \{(x, x)\colon x \in \Lambda\}$ and

(1) If $x \in \Gamma^{(0)}$, then $\gamma(x, y)$ starts at x; otherwise, $\lim_{t\to-\infty} \gamma(x, y)(t) = x \in \Lambda \smallsetminus \Gamma$.
(2) If $y \in \Gamma^{(0)}$, then $\gamma(x, y)$ ends at y; otherwise, $\lim_{t\to\infty} \gamma(x, y)(t) = y \in \Lambda \smallsetminus \Gamma$.

Definition 2.33 ([29]): Suppose that Γ is a graph, and $\mathbb{A}$ is a ring. Let $C_1(\Gamma; \mathbb{A})$ be the group of finite formal sums of 1-cells in Γ with coefficients in $\mathbb{A}$. A **homological bicombing on** Γ is a function

$$q\colon \Gamma^{(0)} \times \Gamma^{(0)} \to C_1(\Gamma; \mathbb{A})$$

so that $\partial q(a, b) = b - a$.

Remark 2.34: It is clear that a geodesic bicombing as in Definition 2.31 gives rise to a homological bicombing as in Definition 2.33. It is slightly less obvious (but also true) that we can use a bicombing as in Definition 2.32 to produce something homological.

Definition 2.35: Let Γ be a graph and $\mathbb{A}$ a ring, and let $\overline{C}_1(\Gamma;\mathbb{A})$ be the group of (locally finite) formal sums of 1-cells in Γ. Let $\overline{\Gamma}$ be some compactification (or bordification) of $\Gamma^{(0)}$, and let $\Lambda \subseteq \Gamma$. Let q be a function

$$q\colon\ \Lambda \times \Lambda \to \overline{C}_1(\Gamma;\mathbb{A}),$$

which is zero precisely on $\Delta = \{(x,x)\colon x \in \Lambda\}$. Given $R \in \mathbb{N}$, $x, y \in \Lambda$, and $z \in \Gamma^{(0)}$, let $q_{z,R}(x,y)$ be the 1-chain which is equal to $q(x,y)$ on the ball of radius R about z, and zero outside it. The function q is a **homological bicombing on** Λ if it satisfies the following condition: For every $x, y \in \Lambda$, and $z \in \Gamma^{(0)}$, there is an R_0 so that for every integer $R > R_0$, there exist 0-chains $\xi_{R,+}$ and $\xi_{R,-}$, each with coefficients summing to 1 so that:

(1) $\partial q_{z,R}(x,y) = \xi_{R,+} - \xi_{R,-}$ for all $R > R_0$,
(2) any sequence $\{y_i\}_{i=R_0}^{\infty}$ with $y_i \in \xi_{i,+}$ satisfies $\lim_{i\to\infty} y_i = y$, and any sequence $\{y_i\}_{i=R_0}^{\infty}$ with $x_i \in \xi_{i,-}$ satisfies $\lim_{i\to\infty} x_i = x$.

Definition 2.36: Let Γ be a graph with a compactification $\overline{\Gamma}$ of $\Gamma^{(0)}$, and let $\Lambda \subseteq \overline{\Gamma}$. Let $\epsilon > 0$. A homological bicombing $q : \Lambda \times \Lambda \to \overline{C}_1(\Gamma;\mathbb{R})$ is **ϵ-quasi-geodesic** if both

(1) $q(a,b)$ has support in the ϵ-neighborhood of some geodesic between a and b, and
(2) If a, $b \in \Gamma^{(0)}$, then $|q(a,b)|_1 \leq \epsilon d(a,b)$.

Remark 2.37: In general, one may also want to place constraints on the 1-norms of finite "subsegments" of $q(a,b)$, where a and b are ideal points. However, we do not need this refinement in this paper.

2.9. Relatively hyperbolic groups. Relatively hyperbolic groups were first defined by Gromov in [16]. Alternative definitions were given by Farb [13] and Bowditch [4]. These definitions are all equivalent. See [10, Appendix].

Further characterizations of relatively hyperbolic groups are given by Osin [35], in terms of relative Dehn functions, and Yaman [41], in terms of convergence group actions.

Recently there has been a large amount of interest in relatively hyperbolic groups. (See [6], [8], [12], [18], [35], among many others).

Here is the original definition of Gromov's [16, Section 8.6]:

Definition 2.38: Suppose that G acts isometrically and properly on a proper, geodesic, Gromov hyperbolic metric space X, so that the quotient is quasi-isometric to a wedge of n rays. Let $\gamma_1, \ldots, \gamma_n$ be unit-speed geodesic rays in X/G tending to distinct points in the Gromov boundary of X/G, and choose lifts $\widetilde{\gamma}_1, \ldots, \widetilde{\gamma}_n$ to X. For each i, let e_i be the point in ∂X to which $\widetilde{\gamma}_i$ limits, and let P_i be the stabilizer in G of e_i. For each i define a **horofunction** $h_i \colon X \to \mathbb{R}$ by

$$h_i(x) = \limsup_{t \to \infty} d(x, \widetilde{\gamma}_i(t)) - t.$$

The R**-horoballs** of X are the sub-level sets $B_i(R) = h_i^{-1}(-\infty, R)$ and their G-translates. Assume that there exists a constant R so that for any $g \in G$ and any i, j, either $gB_i(R) \cap B_j(R)$ is empty or $i = j$ and $g \in P_i$. Finally, suppose that G acts cocompactly on the complement of the union of the horoballs. Then we say that G is **hyperbolic relative to** $\mathcal{P} = \{P_1, \ldots, P_n\}$ **in the sense of Gromov**.

Definition 2.39: Suppose that G is a relatively hyperbolic group acting on the δ-hyperbolic space X as in Definition 2.38. An element $g \in G$ is called **hyperbolic** if it does not have a bounded orbit in X, and it fixes exactly two points in ∂X.

We say that G is **non-elementary relatively hyperbolic** if there are hyperbolic elements g, h in G so that $\mathrm{Fix}_{\partial X}(g) \cap \mathrm{Fix}_{\partial X}(h) = \emptyset$.

Remark 2.40: By the usual Ping-Pong argument, if g, h are as in Definition 2.39, then there is some $j \geq 1$ so that g^j and h^j generate a free group.

The following is another definition of relatively hyperbolic groups, which is a hybrid of Farb's [13] and of Bowditch's [4].

Definition 2.41 (Coned-off Cayley graph): Suppose that G is a finitely generated group, with finite generating set S. Let $\Gamma(G, S)$ be the Cayley graph of G with respect to S.

Suppose that $\mathcal{P} = \{P_1, \ldots, P_k\}$ is a finite collection of finitely generated subgroups of G. We form a new graph containing $\Gamma(G, S)$, called the **coned-off Cayley graph** and denoted $\widehat{\Gamma}(G, \mathcal{P}, S)$ as follows:

For each $i \in \{1, \ldots, k\}$ and each coset gP_i we add a new vertex, $v_{g,i}$ to $\Gamma(G, S)$. We also add a vertex from each element of gP_i to $v_{g,i}$.

Definition 2.42 (Fine graphs; see [4, page 11]): A (not necessarily locally finite) graph $\mathcal{K}$ is **fine** if for every edge e in $\mathcal{K}$ and each integer $L > 0$, the number of simple simplicial loops of length at most L which contain e is finite.

Definition 2.43: Suppose that G is a finitely generated group and that $\mathcal{P} = \{P_1, \ldots, P_k\}$ is a collection of finitely generated subgroups. We say that G is **hyperbolic relative to** $\mathcal{P}$ if the coned-off Cayley graph $\widehat{\Gamma}(G, \mathcal{P}, S)$ is fine and δ-hyperbolic for some $\delta > 0$.

Remark 2.44: By now the class of groups which we call 'relatively hyperbolic' is standard. However, we should point out that, in the terminology of [13], G is hyperbolic relative to $\mathcal{P}$ if and only if the coned-off Cayley graph is δ-hyperbolic. Farb's hypothesis of Bounded Coset Penetration is equivalent to fineness of $\widehat{\Gamma}$ (see, for instance, [10, Appendix]). It is shown in [6] and [10, Appendix] that Definitions 2.38 and 2.43 are equivalent.

Whenever G is hyperbolic relative to $\mathcal{P}$, we will always assume that our (finite) generating set for G is compatible, in the sense of Definition 2.15.

We briefly list some examples of relatively hyperbolic groups (and the subgroups they are hyperbolic relative to):

(1) Hyperbolic groups are hyperbolic relative to the empty collection of subgroups;
(2) Fundamental groups of geometrically finite hyperbolic manifolds are hyperbolic relative to the cusp subgroups;
(3) Free products are hyperbolic relative to the free factors;
(4) A group which acts properly and cocompactly on a CAT(0) space with isolated flats is hyperbolic relative to the stabilizers of maximal flats (see [24]; and [12]);
(5) Limit groups are hyperbolic relative to maximal noncyclic abelian subgroups (see [2] and [9]).

In Section 3, we introduce a 'cusped' space, $X(G, \mathcal{P}, S)$, associated to a group G and finite collection $\mathcal{P}$ of finitely generated subgroups. We prove that G is hyperbolic relative to $\mathcal{P}$ if and only if $X(G, \mathcal{P}, S)$ is δ-hyperbolic for some δ.

Remark 2.45: One of the important features of the space X defined in Section 3 below is that if G is hyperbolic relative to $\mathcal{P}$, then the action of G on $X(G, \mathcal{P}, S)$ satisfies the requirements of Definition 2.38 (see Theorem 3.25).

There are a number of previous constructions of cusped spaces (for example, Cannon and Cooper [7], Bowditch [4] and Rebbechi [38]). The novelty in our space X is that it is a graph metrized so that edge lengths are 1. Thus we can apply many combinatorial constructions directly, such as isoperimetric inequalities (combinatorial and homological). Importantly, in Section 6 below, we can also apply a construction of Mineyev from [29].

Another important feature of the space X is that there is a bound on the lengths of the attaching maps of 2-cells.

By Osin [35, Theorem 1.5], relatively hyperbolic groups are always finitely presented relative to their parabolics. (This also follows from the construction of the **relative Rips complex** in [8].)

THEOREM 2.46 ([35, Theorem 1.5]): *Let G be a finitely generated group, $\{H_1, \dots, H_m\}$ a collection of subgroups of G. The following are equivalent:*

(1) *G is finitely presented with respect to $\{H_1, \dots, H_m\}$ and the corresponding relative Dehn function is linear;*
(2) *G is hyperbolic relative to $\{H_1, \dots, H_m\}$.*

Recall by Lemma 2.14 that if G is finitely generated, and finitely presented relative to $\{H_1, \dots, H_m\}$, then each of the H_i is finitely generated.

In this paper, we have no need for the 'relative' Dehn functions of [35]. Rather, we construct simply-connected 2-complexes with linear combinatorial isoperimetric inequalities. (They also have linear homological isoperimetric inequalities; see Theorem 3.25 below.)

Definition 2.47 (Coned-off Cayley complex): Suppose that G is a finitely generated group, with a collection $\mathcal{P}$ of finitely generated subgroups, and that $\langle \mathcal{A}, \mathcal{P} \mid \mathcal{R} \rangle$ is a finite relative presentation for G.

Let S be a (finite) compatible generating set for G containing $\mathcal{A}$. Form a 2-complex $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$, called the **coned-off Cayley complex** as follows:

Let $C = C(G, \mathcal{P}, S)$ be the coned-off Cayley graph (which contains a copy of the Cayley graph $\Gamma(G, S)$. Attach 2-cells to C in a G-equivariant way, corresponding to the relations $\mathcal{R}$. Also, add a 2-cell to each loop of length three in C which contains an infinite valence vertex.

LEMMA 2.48: *The coned-off Cayley complex is simply-connected.*

PROPOSITION 2.49: *Suppose that G is hyperbolic relative to $\mathcal{P}$, and let S be a finite generating set for G containing generating sets for each subgroup in $\mathcal{P}$. Also, let $\langle S, \mathcal{P} \mid \mathcal{R} \rangle$ be a finite relative presentation for G.*

Form the coned-off Cayley complex $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$. Then $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$ has a linear combinatorial isoperimetric inequality.

Proof. Since G is hyperbolic relative to $\mathcal{P}$, the coned-off Cayley graph is δ-hyperbolic and fine. Therefore, there are only finitely many orbits of simple loops of length at most 16δ. Since $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$ is simply-connected, this implies that there exists K so that every combinatorial loop of length at most 16δ can be filled with a combinatorial disk in $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$ of area at most K.

The result now follows immediately from Proposition 2.22. ∎

We now prove the converse to Proposition 2.49.

PROPOSITION 2.50: *Suppose that G is finite presented relative to*

$$\mathcal{P} = \{P_1, \ldots, P_n\}$$

and that $\langle S, \mathcal{P} \mid \mathcal{R} \rangle$ is a finite relative presentation for G.

If the coned-off Cayley complex $\hat{C} = \hat{C}(G, \mathcal{P}, S, \mathcal{R})$ satisfies a linear combinatorial isoperimetric inequality then the coned-off Cayley graph $\hat{\Gamma} = \hat{\Gamma}(G, \mathcal{P}, S)$ is Gromov hyperbolic and fine.

Proof. That $\hat{\Gamma}$ is δ-hyperbolic for some δ follows from Proposition 2.23, since there is certainly a bound on the length of the attaching maps of 2-cells in $\hat{C}$.

It remains to prove that $\hat{\Gamma}$ is fine. Take an edge e and a simple simplicial loop c containing e of length L. Now, there are only finitely many 2-cells adjacent to each edge in $\hat{C}$. A simple loop may be filled by a topological disk, and the isoperimetric function supplies a bound on the area of such a disk in terms of the length. In particular, suppose a simple loop of length at most L can be filled by a topological disk of area at most L'. There are only finitely many topological disks of area at most L' containing e on the boundary, so there are only finitely many simple loops of length at most L containing e. Therefore $\hat{\Gamma}$ is fine. ∎

It is a consequence of Theorem 3.25 that the hypothesis in Proposition 2.50 can be weakened to that of having a **homological** linear isoperimetric inequality, but we do not know a direct proof. This raises a natural question (if the answer is positive, then it can be used to give a direct proof):

QUESTION 2.51: *Let X be a simply connected 2-complex with a homological (linear?) isoperimetric inequality, a bound on the length of attaching maps of 2-cells and finitely many 2-cells adjacent to any edge. Must X be fine?*

Part 1. The cusped space and preferred paths

3. The cusped space

The purpose of this section is to construct a space X from a finitely generated group G, and a finite collection $\mathcal{P}$ of finitely generated subgroups. The utility of X is that it is Gromov hyperbolic if and only if G is hyperbolic relative to $\mathcal{P}$ (see Theorem 3.25).

3.1. Combinatorial horoballs.

Definition 3.1: Let Γ be any 1-complex. The **combinatorial horoball based on** Γ, denoted $\mathcal{H}(\Gamma)$, is the 2-complex formed as follows:

- $\mathcal{H}^{(0)} = \Gamma^{(0)} \times (\{0\} \cup \mathbb{N})$
- $\mathcal{H}^{(1)}$ contains the following three types of edges. The first two types are called **horizontal**, and the last type is called **vertical**.
  (1) If e is an edge of Γ joining v to w, then there is a corresponding edge $\overline{e}$ connecting $(v, 0)$ to $(w, 0)$.
  (2) If $k > 0$ and $0 < d_\Gamma(v, w) \leq 2^k$, then there is a single edge connecting (v, k) to (w, k).
  (3) If $k \geq 0$ and $v \in \Gamma^{(0)}$, there is an edge joining (v, k) to $(v, k + 1)$.
- $\mathcal{H}^{(2)}$ contains three kinds of 2-cells:
  (1) If $\gamma \subset \mathcal{H}^{(1)}$ is a circuit composed of three horizontal edges, then there is a 2-cell (a **horizontal triangle**) attached along γ.
  (2) If $\gamma \subset \mathcal{H}^{(1)}$ is a circuit composed of two horizontal edges and two vertical edges, then there is a 2-cell (a **vertical square**) attached along γ.

(3) If $\gamma \subset \mathcal{H}^{(1)}$ is a circuit composed of three horizontal edges and two vertical ones, then there is a 2-cell (a **vertical pentagon**) attached along γ, unless γ is the boundary of the union of a vertical square and a horizontal triangle.

Remark 3.2: As the full subgraph of $\mathcal{H}(\Gamma)$ containing the vertices $\Gamma^{(0)} \times \{0\}$ is isomorphic to Γ, we may think of Γ as a subset of $\mathcal{H}(\Gamma)$.

Remark 3.3: Whenever $\mathcal{H}(\Gamma)$ is to be thought of as a metric space, we will always implicitly ignore the 2-cells, and regard $\mathcal{H}(\Gamma)^{(1)}$ as a metric graph with all edges of length one.

Definition 3.4: Let Γ be a graph and $\mathcal{H}(\Gamma)$ the associated combinatorial horoball. Define a **depth** function

$$D : \mathcal{H}(\Gamma) \to [0, \infty)$$

which satisfies:

(1) $D(x) = 0$ if $x \in \Gamma$,
(2) $D(x) = k$ if x is a vertex (v, k), and
(3) D restricts to an affine function on each 1-cell and on each 2-cell.

Definition 3.5: Let Γ be a graph and $\mathcal{H} = \mathcal{H}(\Gamma)$ the associated combinatorial horoball. For $N \geq 1$, let $\mathcal{H}_N \subset \mathcal{H}$ be the full sub-graph with vertex set $\Gamma^{(0)} \times \{0, \ldots, N\}$.

The following observation will be important in Section 10.

Observation 3.6: Let Γ be a graph, $\mathcal{H}_N$ as in Definition 3.5 above, and $\partial\mathcal{H}_N$ be the full sub-graph with vertex set $\Gamma^{(0)} \times \{N\}$.

Let $\mathcal{H}' = \mathcal{H}(\partial\mathcal{H}_N)$. Identify the copies of $\partial\mathcal{H}_N$ in $\mathcal{H}_N$ and $\mathcal{H}'$. The resulting complex is isomorphic to $\mathcal{H}$.

PROPOSITION 3.7: *Let Γ be a connected 1-complex so that no edge joins a vertex to itself. Then $\mathcal{H}(\Gamma)$ is simply-connected and satisfies a linear combinatorial isoperimetric inequality with constant at most* 3.

Proof. Let c be a combinatorial loop in $\mathcal{H}(\Gamma)$. To prove the proposition it suffices to show that c can be filled by a disk of area at most $3|c|_1$.

Let j be minimal so that there is some vertex (v, j) in c. Since we may clearly suppose c has no backtracking, there is at least one horizontal edge at depth j

in c. By gluing pentagons to edges in c at depth j, and squares when pentagons are not possible, we can reduce the length of c, and increase the minimal depth. Repeating this procedure, we eventually end up with a path of length 3, or 4, which is entirely horizontal. A path of length 3 can be filled be a horizontal triangle, whilst a path of length 4 can be filled by two pentagons beneath it. Being slightly careful about counting shows that the isoperimetric constant is at most 3, as required. ∎

THEOREM 3.8: *Let Γ be any 1-complex. Then $\mathcal{H}(\Gamma)^{(1)}$ is δ-hyperbolic, where δ is independent of Γ.*

Proof. This follows from Propositions 3.7 and 2.23. ∎

Remark 3.9: By studying the geometry of geodesics in combinatorial horoballs as in the results below, it is possible to directly prove that any combinatorial horoball is 20-hyperbolic (and 20 is not optimal).

Geodesics in combinatorial horoballs are particularly easy to understand.

LEMMA 3.10: *Let $\mathcal{H}(\Gamma)$ be a combinatorial horoball. Suppose that $x, y \in \mathcal{H}(\Gamma)$ are distinct vertices. Then there is a geodesic $\gamma(x, y) = \gamma(y, x)$ between x and y which consists of at most two vertical segments and a single horizontal segment of length at most 3.*

Moreover, any other geodesic between x and y is Hausdorff distance at most 4 from this geodesic.

Proof. Let γ' be any geodesic joining x to y.

We observe that if $h = [\gamma'(t_1), \gamma'(t_2)]$ is a maximal horizontal segment of length greater than 1, then $D(\gamma'(t_1 - 1))$ and $D(\gamma'(t_2 + 1))$ are both smaller than $D(h)$ (see Figure 2). It is easy to see that no geodesic in A can contain a horizontal segment of length longer than 5 (Figure 3). Indeed, the geodesic γ' can contain at most 5 horizontal edges in total.

We next observe that there can be at most two horizontal segments in γ' other than the one at maximal depth, at most one on each side of the deepest one. In fact, there can be at most *one* horizontal segment other than the one at maximal depth. Figure 4 shows representative paths with three horizontal segments together with ways to shorten them (there are other possibilities, which are easy to deal with).

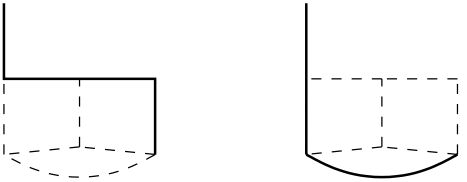

Figure 2. In a geodesic, a horizontal segment not at maximal depth must have length 1.

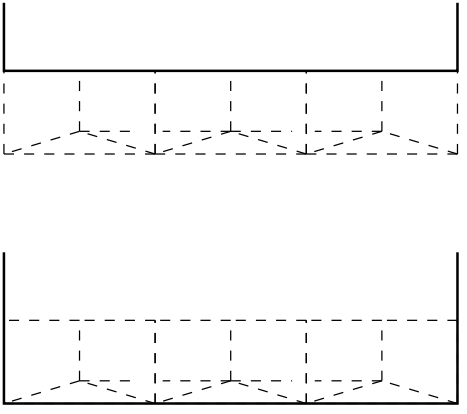

Figure 3. A geodesic can contain no horizontal segment of length greater than 5.

Let M' be the maximum depth achieved by γ'. There is a geodesic γ'' obtained from γ' by pushing all horizontal segments of γ' down to $D^{-1}(M')$. The Hausdorff distance between γ' and γ'' is at most $1\frac{1}{2}$. The geodesic γ'' consists of at most two vertical segments and one horizontal segment of length at most 5. If the horizontal segment $h \subset \gamma''$ has length 4 (respectively 5) then there is another path with the same endpoints as h, and the same length, consisting of two vertical edges and a horizontal path of length 2 (respectively 3). Replacing h with this new path if necessary, we obtain a geodesic of the form required by the lemma.

Now let γ be any geodesic satisfying the conclusion of the lemma. We argue that the Hausdorff distance between γ and γ'' is at most $2\frac{1}{2}$ which, combined

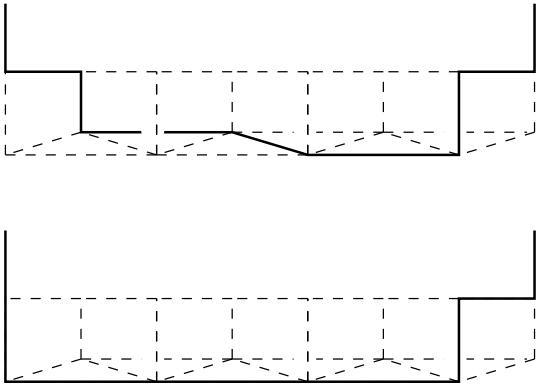

Figure 4. A path with three horizontal segments and a way to shorten the path.

with the earlier estimate on the Hausdorff distance between γ' and γ'', completes the proof.

Let M be the maximum depth of γ. The reader may readily verify that $0 \leq M - M' \leq 1$. If $M = M'$, then γ'' lies within Hausdorff distance at most $1\frac{1}{2}$ of γ. If $M = M' + 1$, then the Hausdorff distance between $\gamma(x, y)$ and γ'' is at most $2\frac{1}{2}$. ∎

LEMMA 3.11: *If A is a combinatorial horoball, then the Gromov boundary consists of a single point, denoted e_A. Moreover, for any $x \in A$, there is a geodesic ray from x to e_A consisting entirely of vertical edges. Any geodesic ray from x to e_A is Hausdorff distance at most $1\frac{1}{2}$ from the vertical ray.*

Proof. A geodesic ray has at most one horizontal edge. Given this observation, the proof is similar to that of Lemma 3.10. ∎

3.2. THE AUGMENTATION. Let G be a finitely generated group, with a finite collection $\mathcal{P}$ of subgroups. Let S be a compatible generating set for G. We define an augmentation of the relative Cayley graph complex of G by combinatorial horoballs. This augmentation will be hyperbolic exactly when G is hyperbolic relative to $\mathcal{P}$ (see Theorem 3.25).

Definition 3.12: Let G be a finitely generated group, let $\mathcal{P} = \{P_1, \ldots, P_n\}$ be a (finite) family of finitely generated subgroups of G, and let S be a generating set

for G so that $P_i \cap S$ generates P_i for each $i \in \{1, \dots, n\}$. For each $i \in \{1, \dots, n\}$, let T_i be a left transversal for P_i (i.e., a collection of representatives for left cosets of P_i in G which contains exactly one element of each left coset).

For each i, and each $t \in T_i$, let $\Gamma_{i,t}$ be the full subgraph of the Cayley graph $\Gamma(G, S)$ which contains tP_i. Each $\Gamma_{i,t}$ is isomorphic to the Cayley graph of P_i with respect to the generators $P_i \cap S$. Then we define

$$X = \Gamma(G, S) \cup (\cup\{\mathcal{H}(\Gamma_{i,t})^{(1)} \colon 1 \leq i \leq n, t \in T_i\}),$$

where the graphs $\Gamma_{i,t} \subset \Gamma(G, S)$ and $\Gamma_{i,t} \subset \mathcal{H}(\Gamma_{i,t})$ are identified as suggested in Remark 3.2.

Remark 3.13: The vertex set of X can naturally be identified with the set of 4-tuples (i, t, p, k), where $i \in \{1, \dots, n\}$, $t \in T_i$, $p \in P_i$, and $k \in \mathbb{N}$.

We will use this identification without comment in the sequel.

Remark 3.14: The group G acts isometrically and properly on the graph $X(G, \mathcal{P}, S)$.

A path in X starting at 1 determines an element of G in the following manner: Each horizontal edge is naturally labelled by a group element. Take the product of the labels of the horizontal edges in the path (and ignore the vertical edges) with the order coming from the path.

Paths in quotients X/H, where $H \leq G$, and also paths which start at a point $(i, 1, 1, k)$ lying directly 'beneath' 1 also naturally determine elements of G.

Supposing further that G is finitely presented relative to $\mathcal{P}$, there is a natural locally finite simply connected 2-complex with skeleton $X(G, \mathcal{P}, S)$:

Definition 3.15: Let G be a finitely generated group, let $\mathcal{P} = \{P_1, \dots, P_n\}$ be a (finite) family of finitely generated subgroups of G, and let S be a generating set for G so that $P_i \cap S$ generates P_i for each $i \in \{1, \dots, n\}$. Let $S' = S \setminus \bigcup \mathcal{P}$, and suppose that

$$G = \langle \mathcal{A}, \mathcal{P} \mid \mathcal{R} \rangle$$

is a finite *relative* presentation of G, in the sense of Definition 2.13 above (where $\mathcal{A} \subset S$). Then we may form a locally finite 2-complex $X(G, \mathcal{P}, S, \mathcal{R})$, whose one-skeleton is the space $X(G, \mathcal{P}, S)$ from Definition 3.12, and which contains the following 2-cells.

- 2-cells from $\mathcal{R}$: Each relator in $\mathcal{R}$ determines a loop γ in $\Gamma(G,S)$ beginning at 1. For each $g \in G$, there is a 2-cell attached along $g\gamma$.
- 2-cells from combinatorial horoballs: For each i and each $t \in T_i$ we have an embedding of $\mathcal{H}(\Gamma_{i,t})^{(1)}$ into $X(G,\mathcal{P},S)$. If $\gamma \in \mathcal{H}(\Gamma_{i,t})$ bounds a 2-cell, then there is a 2-cell attached along its image in $X(G,\mathcal{P},S,\mathcal{R})$.

Remark 3.16: The 2-complex X can also be obtained from the relative Cayley complex $C(G,S,\mathcal{R})$ as in Definition 2.16 by attaching combinatorial horoballs to the cosets gP_i (including the 2-cells in the horoballs). In particular, and this will be important several times throughout this paper, the relative Cayley complex is canonically embedded in X.

As in Definition 3.4, we define a depth function for the space $X(G,\mathcal{P},S,\mathcal{R})$:

Definition 3.17: A **depth** function is

$$D : X(G,\mathcal{P},S,\mathcal{R}) \to [0,\infty)$$

and it satisfies:

(1) $D(x) = 0$ if $x \in G$;
(2) $D(x) = n$ if x is a vertex (i,t,p,n);
(3) D is equivariant; and
(4) D restricts to an affine function on each 1-cell and on each 2-cell.

Remark 3.18: Because D is G-equivariant, it induces a depth function on the quotient space X/H, for any subgroup H of G. We refer to the depth function on the quotient by D also.

Remark 3.19: We observe that $D^{-1}(0) = \widehat{C}$, the relative Cayley complex of G.

Remark 3.20: In Section 6, we need to choose $\mathcal{R}$ carefully. We will need all sufficiently short loops in $\widehat{C}$ to be able to be filled in $\widehat{C}$. This can be ensured by including all of the short relations in G in $\mathcal{R}$.

We remark that X naturally breaks up into $\widehat{C}$ and an equivariant family of combinatorial horoballs.

Definition 3.21: Let $L > 0$. An **L-horoball** is a component of $D^{-1}[L,\infty)$. A **0-horoball** is the maximal subcomplex of $X(G,\mathcal{P},S,\mathcal{R})$ with vertices

$$tP_i \cup \{(i,t,p,k) \colon p \in P_i, k \in \mathbb{N}\}$$

for some $i \in \{1, \dots, n\}$ and some $t \in T_i$.

Remark 3.22: The space X is given the path metric, where each edge has length 1. We have not specified what kind of metric to put on the 2-cells. Thus, whenever we are discussing the metric of X, we will simply pretend that X is a 1-complex. In particular, a **geodesic** in X will always refer to a geodesic path in the 1-skeleton.

The purpose of the 2-cells is that we want a simply-connected space for which we will prove a linear isoperimetric inequality. This will imply that X is a δ-hyperbolic space for some δ, in case G is hyperbolic relative to $\mathcal{P}$. See Theorem 3.25 below.

3.3. Notions of relative hyperbolicity. The main result of this subsection is Theorem 3.25, which gives a collection of statements which are equivalent to relative hyperbolicity. In particular, G is hyperbolic relative to $\mathcal{P}$ if and only if $X(G, \mathcal{P}, S)$ is Gromov hyperbolic for any appropriate choice of S as in Definition 3.12 above.

The next result will allow us to translate the hyperbolicity of the coned-off Cayley complex into hyperbolicity of the space X. In order to make the proof of Theorem 3.25 easier, we choose to phrase it in terms of linear homological isoperimetric inequalities, though there is an analogous version with combinatorial isoperimetric inequalities (and the proof of this analogue is somewhat easier).

Theorem 3.23: *Suppose that G is finitely presented relative to $\mathcal{P}$, that S is a finite generating set for G so that $P = \langle S \cap P \rangle$ for each $P \in \mathcal{P}$, and suppose that $\langle S, \mathcal{P} \mid \mathcal{R} \rangle$ is a finite relative presentation for G.*

If the coned-off Cayley complex $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$ satisfies a linear homological isoperimetric inequality, then $X(G, \mathcal{P}, S, \mathcal{R})$ satisfies a linear homological isoperimetric inequality.

Proof. Let $\Gamma = \Gamma(G, S)$ be the Cayley graph of G with respect to S. Let $\hat{\Gamma} = \hat{\Gamma}(G, \mathcal{P}, S)$ be the coned-off Cayley graph and $\hat{C} = \hat{C}(G, \mathcal{P}, S, \mathcal{R})$ the coned-off Cayley complex.

By assumption, $\hat{C}$ satisfies a linear homological isoperimetric inequality. Let K be the isoperimetric constant for $\hat{C}$.

Take a 1-cycle c in X^1. By Theorem 2.26, we may assume without loss of generality that c is a simple loop. If the support of c lies entirely within a single

horoball, then c may be filled with a combinatorial disk of area at most $3|c|_1$, by Proposition 3.7. Therefore, we may suppose that the support of c does not lie entirely in a horoball.

Decompose c into pieces which lie in Γ, and pieces which lie entirely in a single horoball.

For each maximal sub-path of c which lies in a horoball, there is a path of length 2 in $\hat{\Gamma}$ with the same endpoints. This gives a loop $\hat{c}$ in $\hat{\Gamma}$. Clearly $|\hat{c}| \leq |c|$.

There is a rational 2-chain ω in $\hat{C}$ with boundary $\hat{c}$ so that

$$|\omega|_1 \leq K|\hat{c}|_1.$$

Consider an infinite-valence vertex $v \in \hat{\gamma}$, and the 2-cells in $\mathrm{supp}(\omega)$ which intersect v. Each such 2-cell is a triangle, and has a single edge in Γ. Thus to each infinite valence vertex v in $\hat{c}$ is associated (via the 2-chain ω) a 1-chain p_v whose support is contained in the link of v and whose 1-norm is exactly the 1-norm of ω restricted to the star of v.

The 1-chains p_v (of which there are finitely many for the given path $\hat{c}$) have total length bounded by $|\omega|_1$, and give a method of 'surgering' c. We have the decomposition of c into paths q_i in the Cayley graph and paths r_v, where r_v is a path through a horoball corresponding to the infinite valence vertex v. This induces a decomposition of c as

$$c = \sum_v (r_v + p_v) + \left(\sum_i q_i - \sum_v p_v \right),$$

where the paths are oriented so that $r_v + p_v$ is a 1-cycle, and so is $(\sum_i q_i - \sum_v p_v)$.

Now,

$$|c|_1 \leq \sum_v |r_v + p_v|_1 + \sum_i |q_i|_1 + \sum_v |p_v|_1 \leq 2K|\hat{c}|_1 + |c|_1 \leq (2K+1)|c|_1.$$

(The second inequality holds because $\sum_v |p_v|_1 \leq |\omega|_1$, and $\sum_v |r_v|_1 + \sum_i |q_i|_1 = |c|_1$.)

The 1-cycles $r_v + p_v$ lie entirely in a horoball, and can be filled efficiently, as described above (this follows from Theorem 2.26 and Proposition 3.7).

Therefore, it suffices to fill the 1-cycle $c_1 = (\sum_i q_i - \sum_v p_v)$ efficiently. The 1-cycle c_1 has support entirely in the Cayley graph Γ. Therefore, c_1 can be interpreted as a 1-cycle in $\hat{\Gamma}$. Hence there is a 2-chain ω_1 in $\hat{\Gamma}$ so that $\partial\omega_1 = c_1$ and $|\omega_1|_1 \leq K|c_1|_1$.

Let ω_1' be the 2-chain which is equal to ω_1 on the Cayley complex, and equal to zero on 2-cells adjacent to infinite valence vertices. Then $\partial\omega_1'$ decomposes into a sum of c, and a collection of 1-cycles d_i each of which lies entirely on the link of an infinite valence vertex. Also, $|\omega_1'|_1 \leq |\omega_1|_1$, so $|\partial\omega_1'|_1 \leq M|\omega_1'|_1 \leq M|\omega_1|_1$, where M is the maximum length of an attaching map of a 2-cell in $\hat{\Gamma}$.

Since the support of ω_1' is entirely contained in the Cayley complex, it can be considered as a 2-chain in X, whose boundary is the sum of c_1 and a collection of 1-cycles, d_i each of which lives entirely in a single parabolic coset. Each d_i can be filled by a 2-chain ν_i whose support is entirely contained in the appropriate combinatorial horoball, so that $|\nu_i|_1 \leq 3|d_i|_1$. We also have that $\sum_i |d_i|_1 \leq |\partial\omega_1'|_1$, by choosing the d_i to have distinct supports (choose one d_i for each parabolic coset).

Now, by choosing appropriate orientations, we have $\partial(\omega_1' + \sum_i \nu_i) = c_1$. We also have

$$|\omega_1' + \sum_i \nu_i|_1 \leq |\omega_1'|_1 + \sum_i |\nu_i|_1 \leq |\omega_1|_1 + 3\sum_i |d_i|_1 \leq K|c_1|_1 + 3|\partial\omega_1'|_1$$
$$\leq K|c_1|_1 + M|\omega_1|_1 \leq (KM + K)|c_1|_1.$$

This finishes the proof that X satisfies a linear homological isoperimetric inequality. ∎

We state the combinatorial version of Theorem 3.23 below for completeness. The proof is entirely analogous to that of Theorem 3.23 with loops playing the part of 1-cycles and disks the part of filling 2-chains.

THEOREM 3.24: *Let $G, \mathcal{P}, S$ and $\mathcal{R}$ be as in the statement of Theorem 3.23 above.*

If $\hat{C}(G, \mathcal{P}, S, \mathcal{R})$ satisfies a linear combinatorial isoperimetric inequality with constant K, then $X(G, \mathcal{P}, S, \mathcal{R})$ satisfies a linear combinatorial isoperimetric inequality, and we can take the constant to be $K_1 = 3K(2K + 1)$.

The following is the main result of this section, and gathers together a few notions of relative hyperbolicity, including some that are new in this paper.

THEOREM 3.25: *Suppose that G is a finitely generated group, $\mathcal{P} = \{P_1, \ldots, P_n\}$ is a finite collection of finitely generated subgroups of G, $G = \langle \mathcal{A}, \mathcal{P} \mid \mathcal{R} \rangle$ is a finite relative presentation for G, and S is a compatible generating set containing $\mathcal{A}$.*

Let $\hat{\Gamma}$ be the coned-off Cayley graph for Γ with respect to S and $\mathcal{P}$, and let $\hat{C}$ be the coned-off Cayley complex. Let $X(G, \mathcal{P}, S, \mathcal{R})$ be as defined in Definition 3.15 above. The following are equivalent:

(1) *G is hyperbolic relative to $\mathcal{P}$ in the sense of Gromov;*
(2) *G is hyperbolic relative to $\mathcal{P}$ (i.e., $\hat{\Gamma}$ is Gromov hyperbolic and fine);*
(3) *$\hat{C}$ satisfies a linear combinatorial isoperimetric inequality;*
(4) *$\hat{C}$ satisfies a linear homological isoperimetric inequality;*
(5) *$X^{(1)}$ is Gromov hyperbolic;*
(6) *X satisfies a linear combinatorial isoperimetric inequality;*
(7) *X satisfies a linear homological isoperimetric inequality.*

Proof.

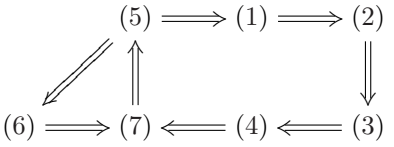

The main result of [6] is that (1) implies (2) (cf. Remark 2.44). By Proposition 2.49, (2) implies (3).

By Lemma 2.29, (3) implies (4) and (6) implies (7).

By Theorem 3.23, (4) implies (7).

Now, (7) implies (5), by Theorem 2.30.

Proposition 2.22 gives that (5) implies (6), provided we can find a bound on the area of fillings of short loops in X (where "short" means at most 16δ). Any combinatorial loop in X of length at most 16δ which lies in a horoball can be filled by a combinatorial disk of length 48δ, by Proposition 3.7. Up to the G-action, there are only finitely many loops of length less than 16δ which do not lie in a horoball. Since X is simply connected, these can all be filled, and so there is some universal constant $C(X)$ so that any loop in X of length less than 16δ can be filled by a disk of area at most $C(X)$.

It remains to observe that the space X satisfies the conditions of Definition 2.38, and so (5) implies (1).

Gathering together these implications, the theorem is proved. ∎

3.4. Metric properties of X. We now suppose that G is hyperbolic relative to $\mathcal{P}$, and that $X = X(G, \mathcal{P}, S)$ as in Definition 3.12 (for the moment we are

only concerned with the metric properties of X, and we restrict our attention to the 1-complex version of X).

By Theorem 3.25, the graph X is δ-hyperbolic for some δ. We assume that this δ satisfies the conditions of Remark 2.10

LEMMA 3.26: *If $L > \delta$, the L-horoballs are convex in X.*

Proof. Let H_1 be a 1-horoball, and let H_L be the L-horoball contained in H_1. We observe that H_L is convex in H_1 (where H_1 is endowed with its path metric). Thus if H_L fails to be convex in X, then two points in H_L are connected by a geodesic which passes through $D^{-1}(0)$. Let p and q be two such points in H_L, chosen to have minimal distance from one another in the path metric on H_1. (Note that this distance must be at least $2L$ so that the geodesic between them actually leaves H_1. They must also satisfy $D(p) = D(q) = L$.) Choose another point $r \in H_L$ so that $\max\{d_{H_1}(p,r), d_{H_1}(q,r)\} < d_{H_1}(p,q)$ and $D(r) = L$. There are then X-geodesics $[p,r]$ and $[q,r]$ which lie entirely in H_L, whereas p and q are joined by a geodesic which includes points in $D^{-1}(0)$, i.e., points which are of distance at least L from H_L. Since the triangle formed by these geodesics is δ-slim, this implies that $L \leq \delta$, a contradiction. ∎

We extend the choice of geodesics in Lemma 3.10 to a choice of geodesics between any two points in X.

LEMMA 3.27: *If G is torsion-free, then there is an antisymmetric, G-equivariant geodesic bicombing γ on X, so that if x and y lie in the same L-horoball for $L > 2\delta$, then $\gamma(x,y)$ is as described in Lemma 3.10.*

Proof. Choose a complete set $\{o_i\}$ of representatives for the orbits of vertices of X under the G-action. For each o_i and each $x \in X$, choose a geodesic $\gamma(o_i, x)$ as in Lemma 3.10. Extend equivariantly and antisymmetrically. Because G is torsion free, there is no element of G which exchanges two points of X. Thus equivariance and antisymmetry may coexist. ∎

Remark 3.28: We say the bicombing in Lemma 3.27 is 'antisymmetric' because we consider paths to be maps. If paths are considered as subsets, it would be symmetric. We will sometimes blur this distinction, but it will not introduce any confusion. For more about the parametrizations of the bicombing, see Paragraph 5.1.1.

When we define the homological bicombing in Section 6, it will be *antisymmetric*, of course.

Remark (about torsion) 3.29: If G is not torsion-free, then there may be no bicombing as in Lemma 3.27, because of the possible presence of 2-torsion. The way to proceed in the presence of torsion is to let $\gamma(x, y)$ be the collection of all of the geodesics between x and y. This makes a number of the arguments later in this paper more awkward. At times it is also useful to consider the "average" of all of the geodesics between x and y.

The following is a slight generalization of the usual notion of cone types.

Definition 3.30 (Cone types): Suppose that G is a group, that Ξ is a graph equipped with a free G-action and that $x \in \Xi$ is a vertex.

For a combinatorial path $\gamma\colon I \to \Xi$, let $[\gamma]$ denote the G-orbit of γ. If v is a vertex in Ξ, the **cone type of** v **viewed from** x is the collection of classes $[\gamma]$ of paths for which

(1) there exists $g \in G$ so that $g \cdot \gamma(0) = v$; and
(2) $d(x, g \cdot \gamma(1)) = d(x, v) + |\gamma|$.

Note that if v_1 and v_2 have the same cone type then, in particular, v_1 and v_2 lie in the same G-orbit.

In case Ξ is the Cayley graph of a finitely generated group G, and $x = 1$, the above definition is equivalent to the usual notion of cone types (see [5, III.Γ.2.16], for example).

The following result follows directly from the proof of [5, Theorem III.Γ.2.18].

LEMMA 3.31: *Let G be a finitely generated group acting freely on the locally compact δ-hyperbolic graph Ξ, and suppose that $x \in \Xi$ is a vertex. Each orbit of vertices in Ξ contains only finitely many cone types viewed from x.*

LEMMA 3.32: *Suppose that G is hyperbolic relative to $\mathcal{P} = \{P_1, \ldots, P_n\}$, where no P_i is equal to G. Suppose further that S is a compatible generating set for G with respect to $\mathcal{P}$, and that $X = X(G, \mathcal{P}, S)$ is δ-hyperbolic. Finally, suppose that* $\operatorname{diam}(\Gamma(P_i, S \cap P_i)) \geq 2^{15\delta+1}$ *for each i.*

For each $L > 0$ there exists a 10δ-local geodesic $\gamma \subset X$ of length at least L so that γ does not intersect any $(15\delta + 1)$-horoball in X.

Proof. Proceed as follows. Let H_1 and H_2 be distinct 10δ-horoballs in X. Let ν be a shortest path between H_1 and H_2. Note that the length of ν is at least 20δ.

Let $x = \nu \cap H_1$. For some i, some $t \in T_i$ and some $p \in P_i$, we have $x = (i, t, p, 10\delta)$. By hypothesis, there exists $q \in P_i$ so that the distance in $\Gamma(P_i, S \cap P_i)$ between p and q is exactly $2^{15\delta+1}$. This implies that the geodesic $\gamma(x, y)$ between x and $y = (i, t, q, 10\delta)$ intersects $D^{-1}(15\delta)$ but not $D^{-1}(15\delta + 1)$. This geodesic consists of two vertical segments, and a single horizontal segment of length 2.

Let $\nu' = qp^{-1} \cdot \nu$. Our 10δ-local geodesic segment is then constructed as follows: Start with ν, followed by $\gamma(x, y)$. Then take ν' to the 10δ-horoball $H_3 = qp^{-1}H_2$. From the endpoint x' of ν' in H_3, construct a a path $\gamma(x', y')$ exactly as above (see Figure 5). Continue in this manner. This construction can

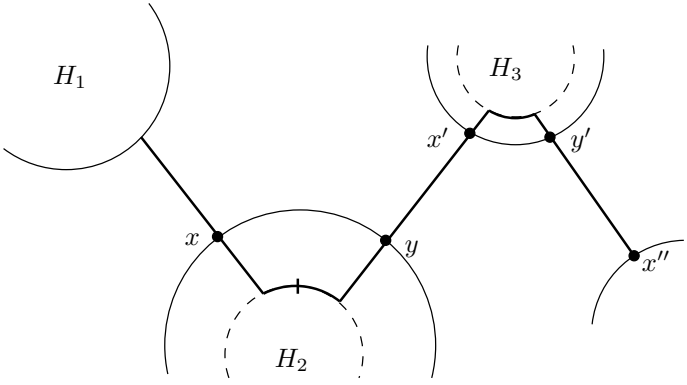

Figure 5. How to make an arbitrarily long 10δ-local geodesic which stays in the "thick" part of X.

be iterated as many times as necessary to ensure the path has at least length L, and it is not difficult to see that the ensuing path is a 10δ-local geodesic (note that the path ν, and its translates, must begin and end with a vertical path of length 10δ, since it is a shortest path between horoballs). ∎

THEOREM 3.33: *Suppose that G is hyperbolic relative to $\mathcal{P} = \{P_1, \ldots, P_n\}$, where no P_i is equal to G. Suppose further that S is a compatible generating*

set for G with respect to $\mathcal{P}$, and that $X = X(G, \mathcal{P}, S)$ is δ-hyperbolic. Finally, suppose that $\operatorname{diam}(\Gamma(P_i, S \cap P_i)) \geq 2^{15\delta+1}$ *for each i.*

Then there exists a hyperbolic element $g \in G$ which has an axis $\gamma \in X$ so that $\gamma \subset D^{-1}[0, 19\delta]$.

Proof. Let $K_{19\delta}$ be the total number of cone types of all points in $D^{-1}[0, 19\delta]$ as viewed from $1 \in X$.

Let σ be a 10δ-local geodesic of length longer than $\frac{7}{3}(K_{19\delta}) + 2\delta$ as in Lemma 3.32. In particular, σ does not penetrate any $(15\delta + 1)$-horoball. By translating σ by an element of G, we suppose that σ begins at $1 \in X$.

By [5, Theorem III.H.1.13(1)], any 10δ-local geodesic is contained in the 2δ-neighborhood of any geodesic joining the endpoints. A simple argument then shows that in fact the geodesic is contained in the 4δ-neighborhood of the k-local geodesic. Also, any 10δ-local geodesic is a $(\frac{7}{3}, 2\delta)$-quasi-geodesic (by [5, Theorem III.H.1.13(3)]).

Let ρ be a geodesic segment joining the endpoints of σ. We have ensured that the length of ρ is greater than $K_{19\delta}$, that ρ does not intersect any $(19\delta + 1)$-horoball, and that ρ starts at 1. [3]

Therefore, there exist vertices $v_1, v_2 \in \rho$ which have the same cone type as viewed from 1. Suppose that v_1 occurs before v_2 on ρ. Let ρ_1 be the subpath of ρ from the 1 to v_1, let ρ_2 be that part between v_1 and v_2, and let ρ_3 be the remainder of ρ.

Since v_1 and v_2 have the same cone type, they are in the same G-orbit. Let $g \in G$ be so that $g.v_1 = v_2$. Since $\rho_1\rho_2\rho_3$ is a geodesic, and v_1 and v_2 have the same cone type as viewed from 1, the path $\rho_1\rho_2(g \cdot \rho_2)(g \cdot \rho_3)$ is also a geodesic.

In turn, this implies that $\rho_1\rho_2(g \cdot \rho_2)(g^2 \cdot \rho_2)(g^2 \cdot \rho_3)$ is a geodesic.

Iterating this argument, we see that the path

$$\rho_2(g \cdot \rho_2)(g^2 \cdot \rho_2) \cdots,$$

is a geodesic ray, starting at v_1. Denote this ray by r. Note that $r \subset g^{-1} \cdot r$. Let γ be the union of the paths $g^{-i} \cdot r$ as $i \to \infty$, parametrized in the obvious way. This is a bi-infinite geodesic line, contained in $D^{-1}[0, 19\delta]$, upon which g acts by translation. ∎

[3] The following argument is very similar to that of [5, Proposition III.Γ.2.22].

THEOREM 3.34: *Suppose that G is hyperbolic relative to $\mathcal{P} = \{P_1, \ldots, P_n\}$, that $G \neq P_i$ for each i and that P_1 is infinite. Then G is non-elementary relatively hyperbolic.*

Proof. Suppose first that some parabolics are finite. Let X' be the space obtained from Definition 3.12 using all of the parabolics, and X the space obtained using only the infinite parabolics. Up to quasi-isometry, X' is X with a locally finite collection of rays attached to the cosets of the finite parabolics. In particular, X is Gromov hyperbolic if and only if X' is. Moreover, an element $g \in G$ acts hyperbolically on X' if and only if it acts hyperbolically on X.

Therefore, we are free to assume that all of the parabolic subgroups of G are infinite.

Let γ be the geodesic from Theorem 3.33, which is an axis of a hyperbolic element $g \in G$. Let γ^+ and γ^- be the points in ∂X at either end of γ.

Let H be a 25δ-horoball in X, and let e_H be the point in ∂X coming from H. Consider an ideal geodesic triangle, T, with vertices γ^-, γ^+ and e_H, and edge γ between γ^- and γ^+. Suppose furthermore, that the geodesics with endpoint e_H are vertical after depth 2δ. By Lemma 2.11, this triangle is 3δ-slim.

The triangle T intersects $H \cap D^{-1}(25\delta)$ in a pair of points $\{x^+, x^-\}$ which are at most 3δ apart. Since the parabolic P which stabilizes H is infinite, there exists $p \in P$ so that $d_X(\{p \cdot x^-, p \cdot x^+\}, \{x^-, x^+\}) \geq 10\delta$.

Now, pT is another ideal triangle, with one vertex e_H and the opposite side an axis for pgp^{-1}. We claim that $\mathrm{Fix}_{\partial X}(g) \cap \mathrm{Fix}_{\partial X}(pgp^{-1}) = \emptyset$.

An easy argument shows that the Hausdorff distance between two geodesics with the same endpoints at infinity is at most 2δ.

Suppose, for instance, that $p \cdot \gamma^+ = \gamma^-$. Then the fact that the geodesics from $p \cdot \gamma^+$ to e_H and from γ^- to e_H are vertical below the 2δ level of H, easily implies that $d_X(p \cdot x^+, x^-) \leq 4\delta$, in contradiction to the choice of p.

Thus, the elements g and pgp^{-1} are hyperbolic elements with disjoint fixed sets in ∂X, and G is non-elementary relatively hyperbolic, as required. ∎

3.5. CONSTANTS. By virtue of Theorem 3.25, we may assume that $X(G, \mathcal{P}, S)$ is δ-hyperbolic for some δ; we are free to assume that δ is an integer and that $\delta \geq 100$. It is useful to have a few different scales to work at, and so we choose the following constants. The choices made will be justified by the results of the subsequent sections.

Definition 3.35: We set $K = 10\delta$, $L_1 = 100K$, and $L_2 = 3L_1$.

We have made no attempt to make this choice of constants optimal.

4. Convex sets and betweenness

In this section we prove a theorem about collections of convex sets in an arbitrary proper Gromov hyperbolic space. The construction in this section is crucial to the construction of **preferred paths** in the next section. In turn, preferred paths are the key to our construction of the bicombing in Section 6.

Definition 4.1: Let Υ be a geodesic metric space. A collection $\mathcal{G}$ of convex sets is **N-separated** if for all $A, B \in \mathcal{G}$ the distance between A and B is at least N.

The elements of $\mathcal{G}$ will be called **globules**.

Let $\mathrm{Isom}(\Upsilon; \mathcal{G})$ be the collection of all isometries g of Υ so that for all $A \in \mathcal{G}$ we have $gA \in \mathcal{G}$.

Throughout the remainder of this section, we will suppose that $\mathcal{G}$ is a 50δ-separated collection of convex subsets of a δ-hyperbolic space Υ.

Remark 4.2: In subsequent sections of this paper, we will apply the results of this section in case $\Upsilon = X(G, \mathcal{P}, S, \mathcal{R})$ and $\mathcal{G}$ is the collection of all L_1-horoballs in X.

Another interesting example is $\mathbb{H}^n$ with some collection of (sufficiently separated) horoballs. It is worth remarking that the existence of a function $\mathcal{D}^\infty$ satisfying the properties (A1)–(A7) below is not obvious even in the case of horoballs in $\mathbb{H}^n$.

The main purpose of this section is to define, for any pair of points $a, b \in \Upsilon$, a family of globules $\mathcal{D}^\infty_{a,b}$, which will be the collection of globules which are 'between' a and b (in case a is contained in a globule A we will have $A \in \mathcal{D}^\infty_{a,b}$ for all b). We want our collections $\mathcal{D}^\infty_{a,b}$ to satisfy various conditions, listed as Axioms (A1)-(A7) below. Most of these are quite straightforward to ensure but Axioms (A5)–(A7), the most important for the applications, are much more difficult to guarantee.

Remark 4.3: The construction performed in this section can be done if the globules are only quasi-convex. How far the globules must be separated depends on the constants of quasi-convexity.

The following two lemmas will be useful:

LEMMA 4.4: *If C is a convex subset of Υ then for any $R \geq 0$, the set $N_R(C)$ is δ-quasi-convex.*

LEMMA 4.5: *Let A, B be disjoint closed δ-quasi-convex sets in the proper δ-hyperbolic space Υ. Let a, $a' \in A$, and b, $b' \in B$. If q is a geodesic between a and b, and q' a geodesic between a' and b', then q' lies in the 3δ-neighborhood of $A \cup B \cup q$.*

Moreover, for any $4\delta < R < \frac{1}{2}d(A, B)$, the Hausdorff distance between the (nonempty) sets $q \smallsetminus N_R(A \cup B)$ and $q' \smallsetminus N_R(A \cup B)$ is at most 5δ.

Proof. Let $[a, a']$ be any geodesic between a and a', and let $[b, b']$ be any geodesic between b and b'; since A and B are δ-quasi-convex, $[a, a'] \subset N_\delta(A)$ and $[b, b'] \subset N_\delta(B)$. Consider the quadrilateral $[a, a'] \cup q \cup [b, b'] \cup [b, a']$, pictured in Figure 6. Since Υ is δ-hyperbolic, this quadrilateral is 2δ-slim. In particular, if x is

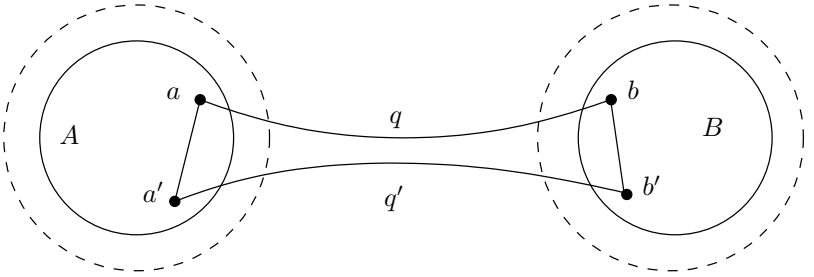

Figure 6. Tube lemma.

any point on q', then

$$x \in N_{2\delta}\left([a, a'] \cup q \cup [b, b']\right) \subset N_{3\delta}\left(A \cup q \cup B\right).$$

For the second assertion of the lemma, let R be as in the statement; since $R < \frac{1}{2}d(A, B)$, the set $q \smallsetminus N_R(A \cup B)$ is non-empty. Let $x \in q \smallsetminus N_R(A \cup B)$. We must show that

$$d(x, q' \smallsetminus N_R(A \cup B)) \leq 5\delta.$$

As above, there is a point $x' \in [a, a'] \cup q' \cup [b, b']$ so that $d(x, x') \leq 2\delta$. Since $R > 3\delta$, and x is at least R from $A \cup B$, the point x' must be on q'. If x' is

outside the R-neighborhood of $A \cup B$, then we are done, so suppose without loss of generality that x' lies in $N_R(A)$.

Let y be the point on q' which is distance exactly R from A. Let t be a point in A with $d(y,t) = d(y,A) = R$. Consider the geodesic triangle pictured in Figure 7 with vertices s, t, and y. This triangle is δ-thin and the geodesic between s and

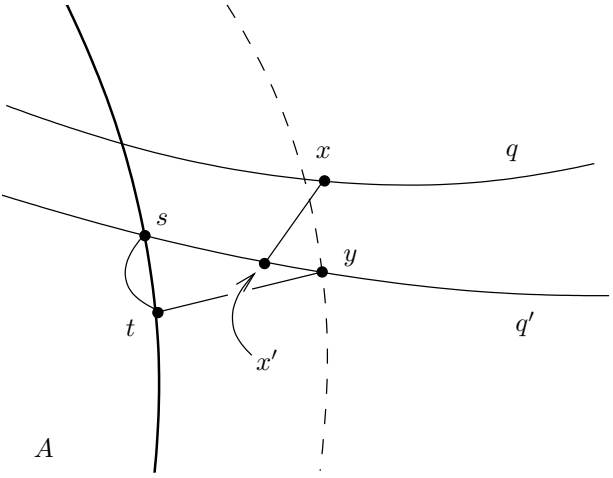

Figure 7. Tube lemma.

t is contained in $N_\delta(A)$. Since $d(x,A) \geq R$, we have $d(x',A) \geq R - 2\delta > 2\delta$, so x' cannot be as close as δ to the geodesic between s and t; thus there is a point z on the geodesic between y and t so that $d(x',z) \leq \delta$ and $d(z,y) = d(x',y)$. We have

$$d(z,A) \geq d(x',A) - \delta \geq R - 3\delta,$$

and so $d(z,y) \leq 3\delta$; this implies $d(x',y) = d(z,y) \leq 3\delta$. Finally, $d(x,y) \leq d(x,x') + d(x',y) \leq 5\delta$, and the lemma is proved. ∎

Remark 4.6: The assumptions of Lemma 4.5 can be weakened. At the cost of a slightly more complicated proof, the assumptions of 'closed' and 'proper' can be removed.

However we will only ever use Lemma 4.5 exactly as it is stated.

Definition 4.7: For $a, b \in \Upsilon$, let $\alpha(a,b)$ be the set of all geodesics between a and b.

Definition 4.8: Suppose that $a, b \in \Upsilon$. Let $\mathcal{C}^0(a,b)$ denote the set of globules which intersect some element of $\alpha(a,b)$ nontrivially.

For any R let $\mathcal{C}^R(a,b)$ denote the set of globules P so that for every $\gamma \in \alpha(a,b)$

$$N_R(P) \cap \gamma \neq \emptyset.$$

Since bigons are δ-thin, $\mathcal{C}^0(a,b) \subseteq \mathcal{C}^K(a,b)$ for all $a, b \in \Upsilon$. (Recall $K = 10\delta$.)

Remark 4.9: For a geodesic γ in Υ, the set of globules $P \in \mathcal{G}$ so that $N_{2K}(P) \cap \gamma \neq \emptyset$ inherits a natural order by projection to the geodesic. This is because globules lie at least $50\delta = 5K$ apart from each other. This induces an order on $\mathcal{C}^{2K}(a,b)$ for each $a, b \in \Upsilon$ and each $\gamma \in \alpha(a,b)$; this order is independent of γ.

Note that if $\mathcal{C}^{2K}(a,b)$ intersects $\mathcal{C}^{2K}(c,d)$ nontrivially, then the order on the intersection inherited from $\mathcal{C}^{2K}(c,d)$ will either coincide with the order inherited from $\mathcal{C}^{2K}(a,b)$ or with its reverse.

For any $a, b \in \Upsilon$, there are only finitely many elements of $\mathcal{C}^K(a,b)$: the size is bounded by $\frac{1}{K}d(a,b) + 1$.

Let $O_{\mathcal{G}}$ be the set of totally ordered finite subsets of $\mathcal{G}$

Thus we have functions

$$\mathcal{C}^0 \colon \Upsilon \times \Upsilon \to O_{\mathcal{G}},$$

and

$$\mathcal{C}^K \colon \Upsilon \times \Upsilon \to O_{\mathcal{G}}.$$

Note that the action of $\mathrm{Isom}(\Upsilon; \mathcal{G})$ on $\mathcal{G}$ induces an action on $O_{\mathcal{G}}$ and that $\mathcal{C}^0$ and $\mathcal{C}^K$ are equivariant with respect to this action. (Note also that the same subset appears many times in $O_{\mathcal{G}}$, once for each possible total order on the subset. By Remark 4.9, only two of these orders actually appear in the image of $\mathcal{C}^0$ or $\mathcal{C}^K$.)

For a function

$$\mathcal{C} \colon \Upsilon \times \Upsilon \to O_{\mathcal{G}},$$

we will denote $\mathcal{C}(a,b)$ by $\mathcal{C}_{a,b}$. In the properties (A1)–(A6) defined below, the symbol "$\subset$" is used to denote "ordered subset", not just "subset".

We will use the following notation for subintervals: Let $a, b \in \Upsilon$. If A and B are contained in $\mathcal{C}_{a,b}$ then

$$\begin{aligned}\mathcal{C}_{a,b}[A,B] &= \{Y \in \mathcal{C}_{a,b} : A \leq Y \leq B\},\\ \mathcal{C}_{a,b}(a,A] &= \{Y \in \mathcal{C}_{a,b} : Y \leq A\}, \text{ and}\\ \mathcal{C}_{a,b}[A,b) &= \{Y \in \mathcal{C}_{a,b} : A \leq Y\}.\end{aligned}$$

Here are some useful conditions $\mathcal{C}$ might satisfy:

(A1) For all $a, b \in \Upsilon$, $\mathcal{C}^0_{a,b} \subseteq \mathcal{C}_{a,b}$;
(A2) For all $a, b \in \Upsilon$, $\mathcal{C}_{a,b} \subseteq \mathcal{C}^K_{a,b}$;
(A3) For all $a, b \in \Upsilon$, $\mathcal{C}_{a,b} = (\mathcal{C}_{b,a})^{op}$;
(A4) $\mathcal{C}$ is $\mathrm{Isom}(\Upsilon; \mathcal{G})$-equivariant.
(A5) If $A, B \in \mathcal{C}_{a,b} \cap \mathcal{C}_{c,d}$ for some $a, b, c, d \in \Upsilon$ and $A, B \in \mathcal{G}$ then $\mathcal{C}_{a,b}[A,B] = \mathcal{C}_{c,d}[A,B]$;
(A6) If $A \in \mathcal{C}_{a,b} \cap \mathcal{C}_{a,c}$ for some $a, b, c \in \Upsilon$ and $A \in \mathcal{G}$ then $\mathcal{C}_{a,b}(a,A] = \mathcal{C}_{a,c}(a,A]$.
(A7) If $A \in \mathcal{C}_{a,b} \cap \mathcal{C}_{c,b}$ for some $a, b, c \in \Upsilon$ and $A \in \mathcal{G}$ then $\mathcal{C}_{a,b}[A,b) = \mathcal{C}_{c,b}[A,b)$.

Note that Axiom (A7) follows from (A6) and (A3).

The purpose of this section is to find a function $\mathcal{D}^\infty$ which satisfies all seven of the axioms (A1)–(A7). For the application in the next section, (A5)–(A7) are the most important. They are also the most difficult to ensure, although our definition of $\mathcal{D}^\infty$ in Definition 4.11 below is designed to make (A5)–(A7) as apparent as possible.

The approach to constructing $\mathcal{D}^\infty$ is as follows. We start with $\mathcal{C}^0$ with the order as in Remark 4.9 and observe:

LEMMA 4.10: *The function $\mathcal{C}^0$ satisfies (A1)–(A4)*

For a pair $a, b \in \Upsilon$, the final $\mathcal{D}^\infty_{a,b}$ will come from the preliminary $\mathcal{C}^0_{a,b}$ by adding new elements in order to make (A5)–(A7) hold. It is not at all obvious that enforcing (A5)–(A7) whilst retaining (A1)–(A4) is possible. This is the content of the proof of Theorem 4.12 below.

By (A2), all of the elements of $\mathcal{D}^\infty_{a,b}$ must lie in $\mathcal{C}^K_{a,b}$.

Below we define a filtration of $\mathcal{D}^\infty$, which will be the minimal (in the obvious

sense of the word minimal) possible function satisfying (A1)–(A7). Before giving the definition, consider the possible ways that (A5)–(A7) might fail to hold for some $\mathcal{C}_{a,b}$. Either (A5) fails (in the "middle" of $\mathcal{C}_{a,b}$; see Figure 8), (A6) fails (on the "left"; see Figure 9) or (A7) fails.

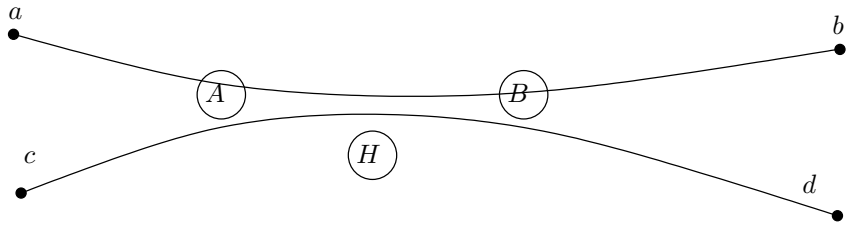

Figure 8. Failure of (A5).

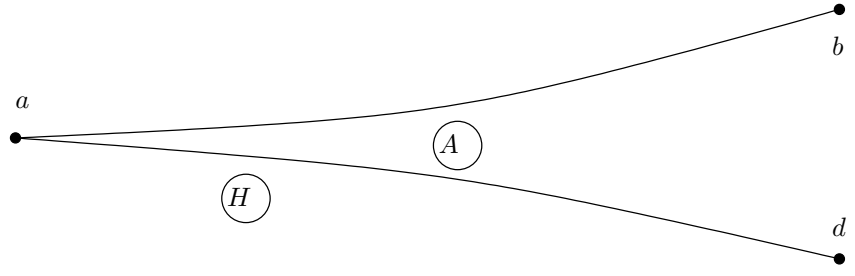

Figure 9. Failure of (A6).

Definition 4.11 below can be thought of in the following way: Start with $\mathcal{C}^0$, and "fix" every failure of axiom (A5)–(A7); after doing these repairs you will have obtained $\mathcal{D}^1$, which still does not satisfy (A5)–(A7). Repairing $\mathcal{D}^1$ yields $\mathcal{D}^2$, and so on.

Definition 4.11: For any pair $a, b \in \Upsilon$, we make the following definitions: Let $\mathcal{D}^0_{a,b} = \mathcal{C}^0_{a,b}$, and define inductively, for $i \geq 0$:

$$\mathcal{M}^i_{a,b} = \bigcup_{c,d \in \Upsilon} \left(\bigcup_{A < B \in \mathcal{D}^i_{a,b} \cap \mathcal{D}^i_{c,d}} \mathcal{D}^i_{c,d}[A, B] \right),$$

$$\mathcal{L}^i_{a,b} = \bigcup_{d \in \Upsilon} \left(\bigcup_{A \in \mathcal{D}^i_{a,b} \cap \mathcal{D}^i_{a,d}} \mathcal{D}^i_{a,d}(a, A] \right),$$

$$\mathcal{R}^i_{a,b} = \bigcup_{c \in \Upsilon} \left(\bigcup_{A \in \mathcal{D}^i_{c,b} \cap \mathcal{D}^i_{a,b}} \mathcal{D}^i_{c,b}[A, b) \right), \text{ and}$$

$$\mathcal{D}^{i+1}_{a,b} = \mathcal{D}^i_{a,b} \cup \mathcal{M}^i_{a,b} \cup \mathcal{L}^i_{a,b} \cup \mathcal{R}^i_{a,b}.$$

Finally, we define

$$\mathcal{D}^\infty_{a,b} = \bigcup_{i=1}^{\infty} \mathcal{D}^i_{a,b}.$$

Notice that the set $\mathcal{M}^i_{a,b}$ fixes the failure of (A5) for $\mathcal{D}^i_{a,b}$, and so on.

The following is the main result of this section, which will allow us to define preferred paths in the next section.

THEOREM 4.12: *Suppose that $\mathcal{G}$ is a 50δ-separated collection of convex subsets of a δ-hyperbolic space Υ. Then the function $\mathcal{D}^\infty$ as defined in Definition 4.11 satisfies the axioms (A1)–(A7).*

Definition 4.11 is tailored so as to make (A5)–(A7) as apparent as possible (once it is known that $\mathcal{D}^\infty$ satisfies (A2), (A5)–(A7) are immediate). As explained below, the hard part of proving Theorem 4.12 is Axiom (A2). We will proceed by induction. On their own, Axioms (A1)–(A4) for $\mathcal{D}^i$ do not seem strong enough to imply Axioms (A1)–(A4) for $\mathcal{D}^{i+1}$. In order for the inductive proof to work, we need to impose further conditions, which are encapsulated in the following definitions.

Definition 4.13: Let $a, b \in \Upsilon$ and $C \in \mathcal{D}^n_{a,b}$ with $a, b \notin C$. A **pair of** (n, a, b)**-guards** of C is a pair (Z, W) (each of which may be either a point or a globule) so that there exists $x, y \in \Upsilon$ for which:

(1) $C \in \mathcal{C}^0_{x,y}$;
(2) Either $Z = a = x$ or $Z \in \mathcal{C}^K_{x,y} \cap \mathcal{D}^n_{a,b}$ and $Z < C$; and
(3) Either $W = b = y$ or $W \in \mathcal{C}^K_{x,y} \cap \mathcal{D}^n_{a,b}$ and $W > C$.

We inductively define (n, a, b)**-sentinels** of C by saying that (n, a, b)-guards of C are (n, a, b)-sentinels of C and that if D is an (n, a, b)-sentinel of C then any (n, a, b)-guards of D are (n, a, b)-sentinels of C.

Remark 4.14: For any n, a, b and C as in Definition 4.13, (n, a, b)-guards of C come in pairs (one to the left of C and one to the right). A single (n, a, b)-guard may occur in many pairs.

We also remark that for any $n \geq 0$, any $a, b \in \Upsilon$ and any $C \in \mathcal{G}$, a pair of (n, a, b)-guards of C is also a pair of $(n + 1, a, b)$-guards of C.

We now introduce the property which will form the inductive hypothesis in the proof of Theorem 4.12.

Definition 4.15: Given an integer $n \geq 0$, a pair $a, b \in \Upsilon$ and a globule $C \in \mathcal{G}$ we let $\mathcal{B}^n(a, b; C)$ be the conjunction of the following three statements:

(1) $C \in \mathcal{D}^n_{a,b}$;
(2) $C \in \mathcal{C}^K_{a,b}$; and
(3) either (i) C has a pair of (n, a, b)-guards; (ii) $a \in C$; or (iii) $b \in C$.

LEMMA 4.16: *For any $n \geq 0$, any $a, b \in \Upsilon$ and any $C \in \mathcal{G}$, if $\mathcal{B}^n(a, b; C)$ holds then $\mathcal{B}^{n+1}(a, b; C)$ also holds.*

Definition 4.17: Let $a, b \in \Upsilon$ and $C, D \in \mathcal{G}$. Suppose that $\gamma \in \alpha(a, b)$ and that γ intersects both $N_\beta(C)$ and $N_\beta(D)$ nontrivially for some $\beta \leq 2K$. Suppose also that $C < D$ in the order on $\mathcal{C}^0_{a,b}$ as in Remark 4.9.

Let y be the last point on γ in $N_\beta(C)$ and z the first point on γ in $N_\beta(D)$. Define

$$\gamma^\beta(C, D)$$

to be the subsegment of γ between y and z.

Proof of Theorem 4.12. Axiom (A4) is obvious. The (unordered) set theoretic parts of (A1) and (A3) are also obvious. The set theoretic part of (A2) is the key: By Remark 4.9, this gives sense to (and implies) all the statements about ordered sets. Axioms (A5)–(A7) will then follow from the construction. Axiom (A2) follows from the following inductive statement:

CLAIM 4.18: For $i \geq 0$, all $a, b \in \Upsilon$ and all $C \in \mathcal{G}$ the following are equivalent:

(1) $C \in \mathcal{D}^i_{a,b}$;

(2) $\mathcal{B}^i(a,b;C)$ holds, and $\mathcal{B}^i(a,b;D)$ holds whenever D is an (i,a,b)-sentinel of C.

Note that for any $i \geq 0$, any $a,b \in \Upsilon$ and any $C \in \mathcal{G}$, the statement "(2) implies (1)" holds tautologically (by Statement 4.15.(1)).

We will prove Claim 4.18 by induction on i.

Base Case: For any $a,b \in \Upsilon$ and $C \in \mathcal{G}$, if $C \in \mathcal{D}^0_{a,b}$, then $\mathcal{B}^0(a,b;C)$. Moreover, whenever D is an $(0,a,b)$-sentinel of C, $\mathcal{B}^0(a,b;D)$ holds.

Proof (Base case). Suppose $C \in \mathcal{D}^0_{a,b}$. If $a \in C$ or $b \in C$, then the result is immediate. Thus we may suppose that $a,b \notin C$.

Statement 4.15.(1) holds by assumption.

Statement 4.15.(2) holds because $\mathcal{D}^0_{a,b} = \mathcal{C}^0_{a,b} \subseteq \mathcal{C}^K_{a,b}$.

Statement 4.15.(3) holds because (a,b) form a pair of $(0,a,b)$-guards of C.

Now suppose that D is a $(0,a,b)$-sentinel of C. This implies that $D \in \mathcal{D}^0_{a,b}$ and the same argument applies. ∎

Inductive Hypothesis: Fix $n \geq 1$. For any $0 \leq i < n$, any $a,b \in \Upsilon$ and any $C \in \mathcal{G}$ if $C \in \mathcal{D}^i_{a,b}$, then $\mathcal{B}^i(a,b;C)$ is true and $\mathcal{B}^i(a,b;D)$ holds for any (i,a,b)-sentinel D of C.

Inductive step: Consider $a,b \in \Upsilon$ and $C \in \mathcal{G}$ and suppose that $C \in \mathcal{D}^n_{a,b}$. We wish to prove that $\mathcal{B}^n(a,b;C)$ holds and that $\mathcal{B}^n(a,b;D)$ holds for any (n,a,b)-sentinel D of C.

Observe that if D is an (n,a,b)-sentinel of C, then it is in particular in $\mathcal{D}^n_{a,b}$. Therefore since $C \in \mathcal{D}^n_{a,b}$ was arbitrary, it suffices to prove that $\mathcal{B}^n(a,b;C)$ holds.

By Definition 4.11, one of four situations must occur:

(1) $C \in \mathcal{D}^{n-1}_{a,b}$;
(2) $C \in \mathcal{M}^{n-1}_{a,b}$;
(3) $C \in \mathcal{L}^{n-1}_{a,b}$; or
(4) $C \in \mathcal{R}^{n-1}_{a,b}$.

We deal with each of these situations in turn.

Case 1: In this case $\mathcal{B}^{n-1}(a,b;C)$ holds by induction, and by Lemma 4.16 $\mathcal{B}^n(a,b;C)$ also holds.

Case 2: Suppose that $C \notin \mathcal{D}^{n-1}_{a,b}$ but that $C \in \mathcal{M}^{n-1}_{a,b}$.

Therefore, there are $c,d \in \Upsilon$ and $A,B \in \mathcal{G}$ so that

(1) $A,B \in \mathcal{D}^{n-1}_{a,b} \cap \mathcal{D}^{n-1}_{c,d}$, and

(2) $C \in \mathcal{D}^{n-1}_{c,d}[A, B]$.

We wish to show that $\mathcal{B}^n(a, b; C)$ holds. Statement 4.15.(1) is clear.

We now prove Statement 4.15.(2) of $\mathcal{B}^n(a, b; C)$.

Since $C \in \mathcal{D}^{n-1}_{c,d}$, the statement $\mathcal{B}^{n-1}(c, d; C)$ holds by the inductive hypothesis, as does $\mathcal{B}^{n-1}(c, d; D)$ for any $(n-1, c, d)$-sentinel D of C. Note that $A < C < B$, so $c, d \notin C$ and C has a pair of $(n-1, c, d)$ guards Z and W, by property $\mathcal{B}^{n-1}(c, d; C)$.(3). Suppose that $x, y \in \Upsilon$ are the points associated to C, Z and W from Definition 4.13.

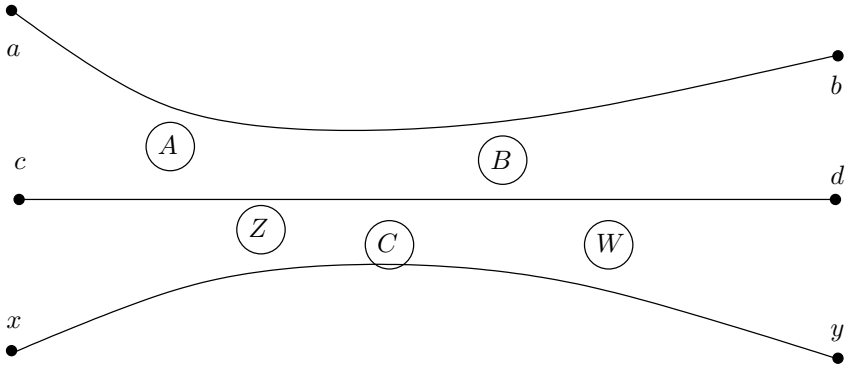

Figure 10. Case (2).

We now define globules P and Q. In case $Z = x = c$, let $P = A$. Otherwise, Z and A are both globules in $\mathcal{D}^{n-1}_{c,d}$ and we let $P = \max\{A, Z\}$ with respect to the order on $\mathcal{D}^{n-1}_{c,d}$. Similarly, if $W = y = d$, then $Q = B$ and otherwise $Q = \min\{B, W\}$, see Figure 10. In the figure, $P = Z$ and $B = Q$, and they are all globules.

Claim 1: $P, Q \in \mathcal{C}^{K+5\delta}_{x,y}$.

We only consider P, as the argument for Q is identical. If $P = Z$, then $Z \in \mathcal{C}^{K}_{x,y} \subseteq \mathcal{C}^{K+5\delta}_{x,y}$.

Thus suppose that $P = A$ and that $A \neq Z$. There are now two cases, depending on whether $Z = x = c$ or not.

If $Z = x = c$, let $\beta^{5\delta, K+5\delta}_{x,y}(Z, C)$ denote that portion of $\beta_{x,y}$ between the last point in $N_{5\delta}(Z)$ and the first point in $N_{K+5\delta}(C)$, and define $\beta^{5\delta, K+5\delta}_{x,y}(Z, C)$ similarly. Note that if $N_{K+5\delta}(A)$ and $N_{5\delta}(x)$ are not disjoint, then $A = P$

is certainly contained in $\mathcal{C}^{K+5\delta}_{x,y}$. Otherwise, since the singleton $\{Z\} = \{x\}$ is convex, Lemma 4.5 implies that the Hausdorff distance between $\beta^{5\delta,K+5\delta}_{x,y}(Z,C)$ and $\beta^{5\delta,K+5\delta}_{x,y}(Z,C)$ is at most 5δ. But $A \in \mathcal{D}^{n-1}_{c,d}$, so by the inductive hypothesis $A \in \mathcal{C}^{K}_{c,d}$. Therefore, $\beta_{x,y}$ passes within $K+5\delta$ of $A = P$, as required, see Figure 11.

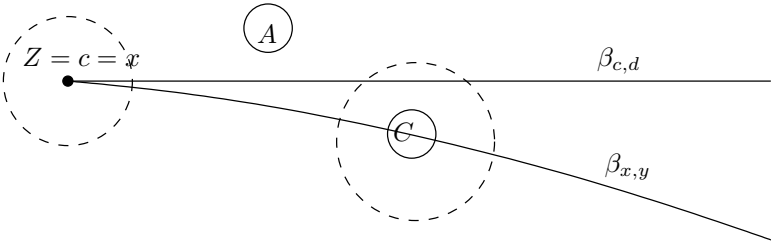

Figure 11. Claim 1, in case $Z = x = c$. Dotted lines indicate $N_K(Z)$ and $N_K(C)$.

Suppose then that $Z \neq x$ or $Z \neq c$. Then Z is a globule and is separated from A by at least L_1. By Lemma 4.5 again, the Hausdorff distance between $\beta^{K+5\delta}_{c,d}(Z,C)$ and $\beta^{K+5\delta}_{x,y}(Z,C)$ is at most 5δ. Once again we know by induction that $A \in \mathcal{C}^{K}_{c,d}$, which proves that $\beta_{x,y}$ passes within $K+5\delta$ of A, see Figure 12. This proves Claim 1.

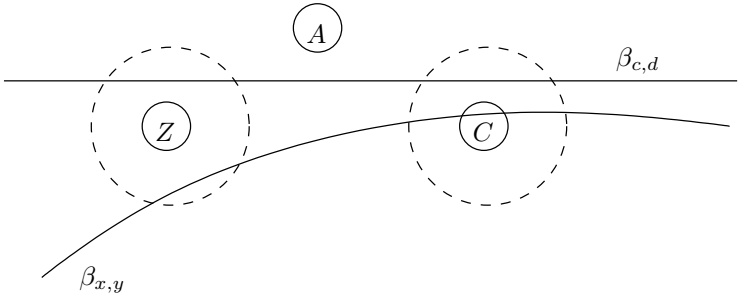

Figure 12. Claim 1, in case Z is not x and c. Dotted lines indicate $N_K(Z)$ and $N_K(C)$.

Claim 2: $P, Q \in \mathcal{C}^{K+5\delta}_{a,b}$.

Once again, we only consider P, as the argument for Q is identical.

If $P = A$, then $P \in \mathcal{D}^{n-1}_{a,b}$ so by induction we have $\mathcal{B}^{n-1}(a, b; P)$ and $P \in \mathcal{C}^K_{a,b} \subseteq \mathcal{C}^{K+5\delta}_{a,b}$.

Suppose then that $P = Z$ and that $Z \neq A$ (this is depicted in Figure 10). Certainly we also have $Z \neq B$.

We know that $A, B \in \mathcal{D}^{n-1}_{a,b} \cap \mathcal{D}^{n-1}_{c,d}$ so by the inductive hypothesis $A, B \in \mathcal{C}^K_{a,b} \cap \mathcal{C}^K_{c,d}$.

Let $\beta_{a,b} \in \alpha(a, b)$ and $\beta_{c,d} \in \alpha(c, d)$ be arbitrary. Then $\beta_{a,b}$ and $\beta_{c,d}$ pass within K of both A and B. Since $P = Z$ is an $(n-1, c, d)$-guard of C, and $c \neq Z$ we know $P \in \mathcal{D}^{n-1}_{c,d}$ and so by another application of the inductive hypothesis $P \in \mathcal{C}^K_{c,d}$. Therefore $\beta_{c,d}$ passes within K of P.

By Lemma 4.5, the Hausdorff distance between $\beta^{K+5\delta}_{c,d}(A, B)$ and $\beta^{K+5\delta}_{a,b}(A, B)$ is at most 5δ. Thus, $\beta_{a,b}$ passes within $K + 5\delta$ of P, proving Claim 2 (see Figure 13).

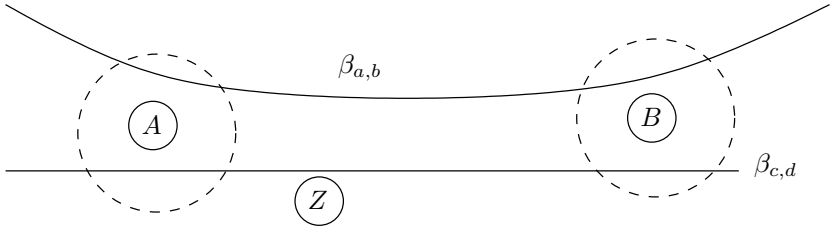

Figure 13. Claim 2, in case $A < Z$. Dotted lines indicate $N_K(A)$ and $N_K(B)$.

Note that whatever P and Q are, we always have $P < C < Q$. Since $P, Q \in \mathcal{C}^{K+5\delta}_{a,b} \cap \mathcal{C}^{K+5\delta}_{x,y}$, Lemma 4.5 implies that the Hausdorff distance between $\beta^{K+10\delta}_{a,b}(P, Q)$ and $\beta^{K+10\delta}_{x,y}(P, Q)$ is at most 5δ. Since $\beta^{K+10\delta}_{x,y}(P, Q)$ intersects C, the geodesic $\beta_{a,b}$ passes within 5δ of C (Figure 14). This proves Statement 4.15.(2) of $\mathcal{B}^n(a, b; C)$.

We now prove Statement 4.15.(3) of $\mathcal{B}^n(a, b; C)$. We claim that P and Q form a pair of (n, a, b)-guards of C. It is certainly true that $P < C < Q$, and that keeping the same x and y we still have $C \in \mathcal{C}^0_{x,y}$. Note that P and Q are globules. Thus we have to prove

(1) $P, Q \in \mathcal{D}^n_{a,b}$; and
(2) $P, Q \in \mathcal{C}^K_{x,y}$.

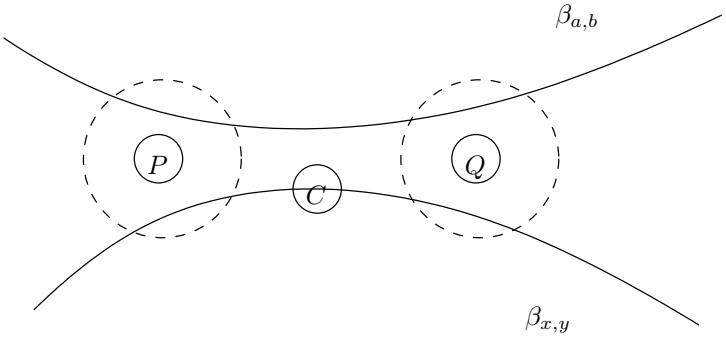

Figure 14. Case (2), Statement 4.15.(2). Dotted lines indicate $N_{K+5\delta}(P)$ and $N_{K+5\delta}(Q)$.

Once again, we prove these statements for P, the arguments for Q being analogous.

Suppose that $P = A$. Then $P \in \mathcal{D}^{n-1}(a,b) \subseteq \mathcal{D}^n_{a,b}$. Otherwise, $P = Z$ and $A < P < B$. Therefore $P \in \mathcal{D}^{n-1}_{c,d}[A,B]$, which implies, by the definition of $\mathcal{D}^n$, that $P \in \mathcal{M}^{n-1}_{a,b} \subseteq \mathcal{D}^n_{a,b}$. This proves Statement 1 above.

We now prove Statement 2. If $P = Z$, then $P \in \mathcal{C}^K_{x,y}$, so we are done. Thus suppose that $P = A$, and $Z < A$.

We know that $A \in \mathcal{D}^{n-1}(c,d)$. Certainly, $A < C$ so $d \notin A$. Suppose that $c \in A$. This forces $Z = x = c$, by Definition 4.13, so $x \in A$ and $A \in \mathcal{C}^0_{x,y} \subseteq \mathcal{C}^K_{x,y}$.

Suppose now that $c \notin A$. Then the inductive hypothesis and property $\mathcal{B}^{n-1}(c,d;A)$ imply that A has a pair of $(n-1,c,d)$-guards, which we denote Z_A and W_A, and an associated pair of points $x_A, y_A \in \Upsilon$, as in Definition 4.13.

We know that $A \in \mathcal{C}^0_{x_A,y_A}$ and that

(1) either $c = x_A = Z_A$ or $Z_A \in \mathcal{C}^K_{x_A,y_A} \cap \mathcal{D}^{n-1}_{c,d}$; and
(2) either $d = y_A = W_A$ or $W_A \in \mathcal{C}^K_{x_A,y_A} \cap \mathcal{D}^{n-1}_{c,d}$.

In case Z_A is a globule, we have $Z_A < A$, and in case W_A is a globule we have $A < W_A$ (see Figure 15).

We now define P_A, which may be either a point or a globule. In case $Z = x = c$, define $P_A = Z_A$ (which may be either a point or a globule). In case Z is a globule, we consider whether Z_A is a point or a globule. If Z_A is a point, then $P_A = Z$. Otherwise $P_A = \max\{Z, Z_A\}$. Note that if Z is a globule, so is P_A.

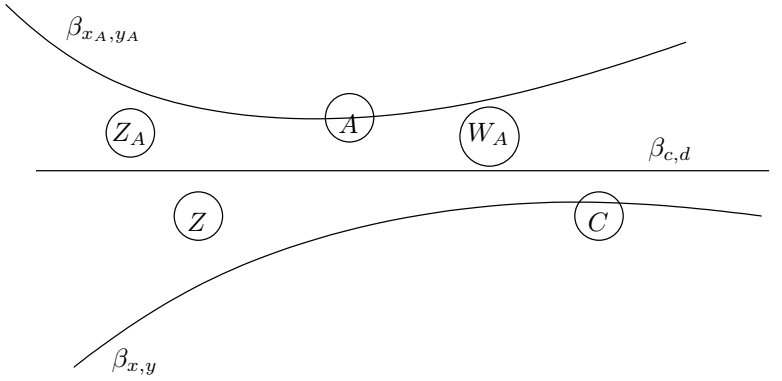

Figure 15. Case (2), Statement 4.15.(3) (one of many possible arrangements).

We also define a globule Q_A. If $W_A = y_A = d$, then $Q_A = C$. If W_A is a globule, then $Q_A = \min\{W_A, C\}$.

Claim 3: If P_A is a globule, then $P_A \in \mathcal{C}^{K+5\delta}_{x_A,y_A} \cap \mathcal{C}^{K+5\delta}_{x,y}$. In any case $Q_A \in \mathcal{C}^{K+5\delta}_{x_A,y_A} \cap \mathcal{C}^{K+5\delta}_{x,y}$.

The proof of this is similar to those of Claims 1 and 2.

Let β_{x_A,y_A} be any geodesic between x_A and y_A. In case P_A is a globule, Claim 3 and Lemma 4.5 imply that the Hausdorff distance between $\beta^{K+10\delta}_{x_A,y_A}(P_A, Q_A)$ and $\beta^{K+10\delta}_{x,y}(P_A, Q_A)$ is at most 5δ. Since $P_A < A < Q_A$ and β_{x_A,y_A} intersects A, the path $\beta_{x,y}$ passes within 5δ of A. Thus in this case $P = A \in \mathcal{C}^K_{x,y}$,as required.

Suppose that P_A is a point (see Figure 16). In this case, $Z = x = c$ and $Z_A = x = x_A$, so $P_A = x = x_A$. In case $N_{5\delta}(P_A)$ and A are not disjoint, x lies within 5δ of A, and certainly $A \in \mathcal{C}^K_{x,y}$. Otherwise, the singleton $\{x\}$ is convex so Lemma 4.5 implies that the Hausdorff distance between those parts of $\beta_{x,y}$ and β_{x_A,y_A} between $N_{5\delta}(x)$ and $N_{K+5\delta}(C)$ are at Hausdorff distance at most 5δ from each other. Since $A \in \mathcal{C}^0_{x_A,y_A}$ this implies once again that $P = A \in \mathcal{C}^K_{x,y}$ as required.

This finally proves that P and Q form a pair of (n, a, b)-guards of C, and finishes the proof of Case (2).

Case 3: Suppose that $C \notin \mathcal{D}^{n-1}_{a,b} \cup \mathcal{M}^{n-1}_{a,b}$, but that $C \in \mathcal{L}^{n-1}_{a,b}$.

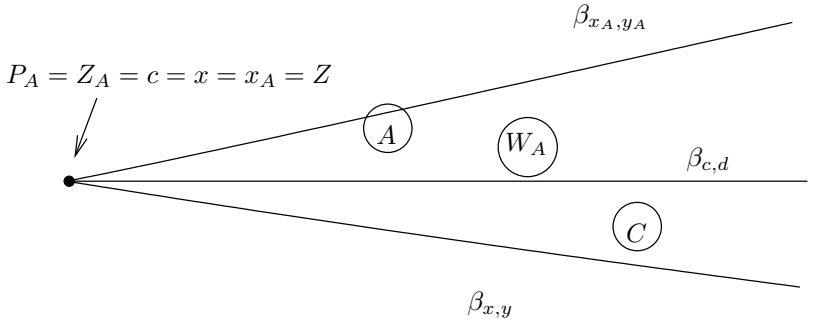

Figure 16. Case (2), Statement 4.15.(3) (another of many possible arrangements).

Therefore, there is $d \in \Upsilon$ and $A \in \mathcal{D}^{n-1}_{a,b} \cap \mathcal{D}^{n-1}_{a,d}$ so that $C \in \mathcal{D}^{n-1}_{a,d}(a, A]$.

We are required to show that $\mathcal{B}^n(a, b; C)$ holds. As usual, statement 4.15.(1) is immediate.

The inductive hypothesis implies that $\mathcal{B}^{n-1}(a, d; C)$ holds.

Since $C \notin \mathcal{D}^{n-1}_{a,b}$ and $A \in \mathcal{D}^{n-1}_{a,b}$ are globules, they are different and $C < A$. Therefore, $d \notin C$, and so either $a \in C$ or C has a pair of $(n-1, a, d)$-guards. If $a \in C$, then statements 4.15.(2) and 4.15.(3) are clear. Therefore, we suppose that $a \notin C$. We now prove 4.15.(2).

Let Z and W be a pair of $(n-1, a, d)$-guards for C, with associated points $x, y \in \Upsilon$. Note that either $Z = x = a$ or $Z \in \mathcal{C}^K_{x,y} \cap \mathcal{D}^{n-1}_{a,d}$. Also, either $W = y = d$ or $W \in \mathcal{C}^K_{x,y} \cap \mathcal{D}^{n-1}_{a,d}$.

Let $P = Z$ and $Q = \min\{A, W\}$. Note that P may be a point or a globule whereas Q is definitely a globule.

Then following is analogous to Claims 1, 2 and 3:

Claim 4: If P is a globule, then $P \in \mathcal{C}^{K+5\delta}_{a,b} \cap \mathcal{C}^{K+5\delta}_{x,y}$. In any case $Q \in \mathcal{C}^{K+5\delta}_{a,b} \cap \mathcal{C}^{K+5\delta}_{x,y}$.

An identical argument to that which followed Claim 3 implies that $C \in \mathcal{C}^K_{a,b}$, which proves 4.15.(2).

We now prove 4.15.(3). Note that $b \notin C$ and we are assuming that $a \notin C$. Therefore we are required to show that C has a pair of (n, a, b)-guards. We will show that P and Q form such a pair (with associated points x and y).

If $P = a$, then $P = x$ also, so 4.13.(2) holds in this case. Otherwise, $P = Z$ is a globule and $P < C < A$ (See Figure 17). Since $P, A \in \mathcal{D}^{n-1}_{a,d}$, and $A \in \mathcal{D}^{n-1}_{a,b}$ this implies that $P \in \mathcal{L}^{n-1}_{a,b} \subseteq \mathcal{D}^n_{a,b}$. Thus we have proved 4.13.(2).

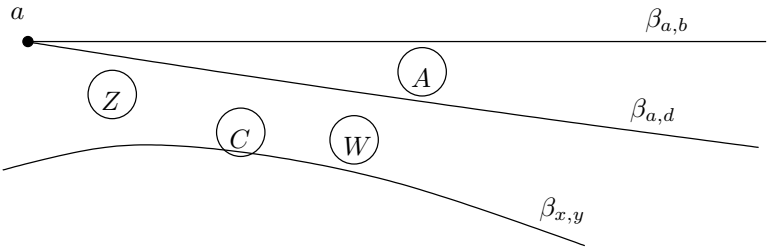

Figure 17. The proof of Case (3), Statement $\mathcal{B}^n(a, b; C)$, in case C has $(n-1, a, d)$-guards.

We now show 4.13.(3). Note that Q is a globule and that $C < Q$. If $Q = W$, then W is a globule and $W \in \mathcal{C}^K_{x,y} \cap \mathcal{D}^{n-1}_{a,d}$. In this case also $W \in \mathcal{D}^{n-1}_{a,d}(a, A] \subseteq \mathcal{L}^{n-1}_{a,b} \subseteq \mathcal{D}^n_{a,b}$. This proves 4.13.(3) in this case.

Suppose then that $Q = A \neq W$. Then $Q \in \mathcal{D}^n_{a,b}$ and we need only prove that $Q \in \mathcal{C}^K_{x,y}$.

By the inductive hypothesis, since A is an $(n-1, a, d)$-sentinel of C and $C \in \mathcal{D}^{n-1}_{a,d}$, property $\mathcal{B}^{n-1}(a, d; A)$ holds. Certainly $C < A$, and both are globules, so $a \notin A$. If $d \in A$ then since $Q = A$ and $A \neq W$, it cannot be that W is a globule. Therefore, in this case $d = W = y$, so $y \in A$ and $A \in \mathcal{C}^0_{x,y} \subseteq \mathcal{C}^K_{x,y}$.

Suppose then that $d \notin A$. In this case, 4.15.(3) implies that A has a pair of $(n-1, a, d)$-guards, which we denote by Z_A and W_A, with associated points $x_A, y_A \in \Upsilon$ (See Figure 18). As before, we define $P_A = \max\{Z_A, C\}$ and $Q_A = \min\{W_A, W\}$. Note that P_A is always a globule but Q_A might be a point.

We now make the customary

Claim 5: If Q_A is a globule, then $Q_A \in \mathcal{C}^{K+5\delta}_{x_A,y_A} \cap \mathcal{C}^{K+5\delta}_{x,y}$. In any case $P_A \in \mathcal{C}^{K+5\delta}_{x_A,y_A} \cap \mathcal{C}^{K+5\delta}_{x,y}$.

Since $A \in \mathcal{C}^0_{x_A,y_A}$, the argument following Claim 3 now implies that $A \in \mathcal{C}^K_{x,y}$ as required. This establishes 4.13.(3). This proves that property $\mathcal{B}^n(a, b; C)$, which finally completes the proof of Case 3.

Case 4: Finally, suppose that $C \notin \mathcal{D}^{n-1}_{a,b}$ but $C \in \mathcal{R}^{n-1}_{a,b}$.

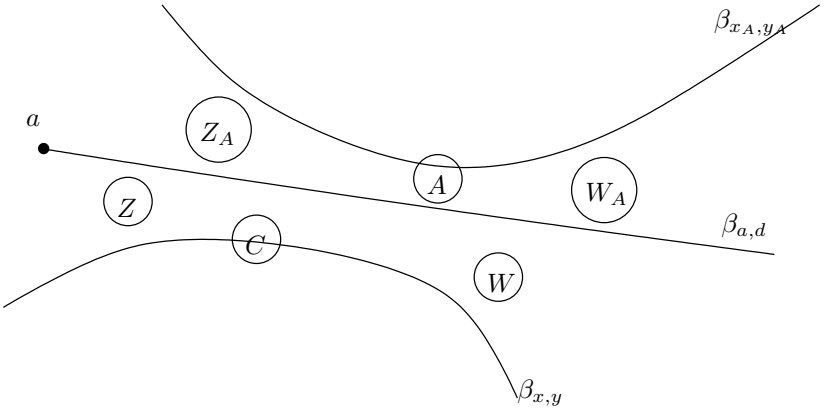

Figure 18. Case (3), proving that P and Q form a pair of (n, a, b)-guards for C (when $Q = A$).

The proof that $\mathcal{B}^n(a, b; C)$ holds is identical (with a left-right reflection) to that of Case 3.

This completes the proof of Claim 4.18, which as we noted above implies that axiom (A2) holds for the collections $\mathcal{D}^\infty_{A,B}$. In turn, by Remark 4.9, this gives a coherent order on each of the sets. The function $\mathcal{D}^\infty$ was designed to make axioms (A5)–(A7) immediately apparent, once axiom (A2) holds, and axioms (A1), (A3) and (A4) are now clear also.

This completes the proof of Theorem 4.12. ∎

COROLLARY 4.19: *Suppose that $A, B \in \mathcal{G}$, and that $a_1, a_2 \in A$ and $b_1, b_2 \in B$. Then $\mathcal{D}^\infty_{a_1,b_1} = \mathcal{D}^\infty_{a_2,b_2}$.*

Also, if $x \in \Upsilon$, then $\mathcal{D}^\infty_{a_1,x} = \mathcal{D}^\infty_{a_2,x}$.

Proof. Axiom (A1) implies that $A, B \in \mathcal{D}^\infty_{a_1,b_1} \cap \mathcal{D}^\infty_{a_2,b_2}$. An application of (A5) implies that $\mathcal{D}^\infty_{a_1,b_1} \subset \mathcal{D}^\infty_{a_2,b_2}$, and, symmetrically, $\mathcal{D}^\infty_{a_2,b_2} \subset \mathcal{D}^\infty_{a_1,b_1}$ (note that we are implicitly using the order and separation properties of globules).

The proof of the second assertion is similar (using axiom (A7) instead). ∎

Notation 4.20: Suppose that $A, B \in \mathcal{G}$. Then we denote the set $\mathcal{D}^\infty_{a,b}$ for $a \in A$ and $b \in B$ by $\mathcal{D}^\infty_{A,B}$. This is well-defined by the above corollary.

Similarly, if $x \in \Upsilon$ and $A \in \mathcal{G}$, then the sets $\mathcal{D}^\infty_{A,x}$ and $\mathcal{D}^\infty_{x,A}$ are well-defined.

5. Preferred paths, preferred triangles, and skeletal fillings

Throughout this section, and the remainder of the paper, G will denote a torsion-free group which is hyperbolic relative to a (finite) collection of (finitely generated) subgroups P. We suppose that S is a finite compatible generating set for G, and that δ is the constant of hyperbolicity for $X(G, P, S)$.

In this section we define **preferred paths**. Preferred paths give an equivariant, symmetric quasi-geodesic bicombing of the space $X = X(G, P, S)$, and will be a key to our construction of the bicombing in Section 6 below. Having defined preferred paths, we then use them to understand the combinatorial types of triangles whose sides are preferred paths.

We fix some notation and terminology: Recall that δ gives rise to the constants $K = 10\delta$, $L_1 = 100K$, and $L_2 = 3L_1$. Denote by $\mathcal{H}$ the collection of L_1-horoballs in X (Definition 3.21). In this section, the word 'horoball' (without a prefix) will refer to an element of $\mathcal{H}$.

Suppose that P is a horoball and $N > 0$. There is a unique N-horoball which intersects P nontrivially. We denote this N-horoball by P^N.

Any horoball P has a single accumulation point in ∂X, which we denote e_P. Denote by $\partial_{\mathcal{H}} X$ the collection of such accumulation points.

Definition 5.1: Recall that in Lemma 3.27 we chose a G-equivariant antisymmetric geodesic bicombing γ. Let $\Omega(X)$ denote the set of unit-speed geodesic paths in X, up to orientation preserving reparametrization. We will extend γ to an antisymmetric map from most of $(X \cup \mathcal{H} \cup \partial_{\mathcal{H}} X)^2$ to $\Omega(X)$. (If A is a horoball, we leave $\gamma(A, A)$, $\gamma(e_A, A)$, $\gamma(A, e_A)$ and $\gamma(e_A, e_A)$ undefined.)

First, for each pair of horoballs A, B, we choose (in an antisymmetric and G-equivariant way), a geodesic $\gamma(A, B)$, which realizes the distance between A and B. Second, for each point $a \in X$ and each horoball $A \in \mathcal{H}$, we (equivariantly) choose a geodesic path $\gamma(a, A)$ which realizes the distance from a to A. The path $\gamma(A, a)$ is the time-reverse of $\gamma(a, A)$.

Third, we extend γ to points in $\partial_{\mathcal{H}} X$. If $x \in X$ and $e_A \in \partial_{\mathcal{H}} X$, then we define $\gamma(x, e_A)$ to be the concatenation of $\gamma(x, A)$ with the vertical geodesic ray in A from the endpoint of $\gamma(x, A)$ to e_A.

Finally, if $e_A, e_B \in \partial_{\mathcal{H}} X$ correspond to distinct horoballs A and B in $\mathcal{H}$, then we define $\gamma(e_A, e_B)$ to be the path $\gamma(A, B)$, together with vertical paths on either end. Note that $\gamma(e_A, e_B)$ is always a geodesic line.

Remark (about torsion) 5.2: As noted in Remark 3.29, 2-torsion may prevent a choice of paths as in Definition 5.1 above.

Also, if there is torsion, then parabolic subgroups of G could intersect non-trivially, which makes the G-equivariance problematic.

For this section, these problems can be solved by considering the union of all shortest paths between A and B (of which there are only finitely many, for a given A and B). In future sections, this is a more subtle problem: in fact we should also use the "average" as well as the union.

Similarly, we should take the union (or average) of all geodesics from a point a to a horoball A which realize the distance from a to A.

5.1. Definition and basic properties of preferred paths. In this subsection we apply the construction of Section 4. The family $\mathcal{H}$ satisfies the hypotheses of Theorem 4.12, therefore there is a function $\mathcal{D}^\infty\colon X \times X \to \mathcal{O}_\mathcal{H}$ satisfying Axioms (A1)–(A7).

Suppose that $a, b \in X \cup \partial_\mathcal{H} X$. We now associate a collection of horoballs $\mathcal{H}_{a,b}$ to the pair $\{a, b\}$. If $a, b \in X$, then $\mathcal{H}_{a,b} = \mathcal{D}^\infty_{a,b}$. If $a \in X$ and $b = e_B \in \partial_\mathcal{H} X$, then $\mathcal{H}_{a,b} = \mathcal{D}^\infty_{a,B}$ as defined in Notation 4.20. Similarly, if $a = e_A \in \partial_\mathcal{H} X$ and $b \in X$, then $\mathcal{H}_{a,b} = \mathcal{D}^\infty_{A,b}$ and if $a = e_A, b = e_B \in \partial_\mathcal{H} X$, then $\mathcal{H}_{a,b} = \mathcal{D}^\infty_{A,B}$.

Remark 5.3: For any $a, b \in X \cup \partial_\mathcal{H} X$, the set $\mathcal{H}_{a,b}$ is a linearly ordered collection of horoballs, and this order is compatible with the order obtained from projection to $\gamma(a, b)$.

Definition 5.4: For each $A \in \mathcal{H}$, and each pair $x, y \in A \cup \{e_A\}$, we define $\sigma(x, y) = \sigma(y, x)$ in the following manner: If $x = e_A$, then $\sigma(x, y)$ is the vertical ray from y to e_A. Thus suppose that $x, y \in A$. Then for some $t \in T$ we have $x = (t, p_1, k_1)$ and $y = (t, p_2, k_2)$. In case $p_1 = p_2$, then $\sigma(x, y)$ is the vertical geodesic between x and y. Otherwise, suppose that $d_P(p_1, p_2)$ satisfies

$$2^{N-1} < d_P(p_1, p_2) \leq 2^N,$$

for $N \in \mathbb{N}$, where $d_P(p_1, p_2)$ is the distance between p_1 and p_2 in P, with respect to the generating set $S \cap P$. Then define $R' = \max\{N, k_1, k_2\}$ and $R = R'$, unless $R = L_2$ in which case $R = L_2 + 1$. Then define $\sigma(x, y)$ to be the path which consists of vertical paths from x and y to depth R, and then joins the endpoints of these vertical paths with the (unique) edge of length 1 at depth R.

Let m be the midpoint of the unique horizontal edge of $\sigma(x, y)$. The point m cuts $\sigma(x, y)$ into two geodesic segments. The path $\sigma(x, y)$ is not usually itself a geodesic, but is very close to one.

Remark 5.5: We insist that the paths $\sigma(x, y)$ do not have a horizontal edge at depth L_2 in order to simplify some future arguments. In particular, this means that our preferred paths are 'transverse' to $D^{-1}(L_2)$.

LEMMA 5.6: *Let x and y lie in the same horoball A. The path $\sigma(x, y)$ is Hausdorff distance at most 5 from $\gamma(x, y)$.*

Proof. The geodesic $\gamma(x, y)$ coincides with $\sigma(x, y)$ to depth S, where S is the maximum depth of $\gamma(x, y)$. Moreover, the sub-path $\sigma(x, y) \cap A^S$ has length at most 7, while the (horizontal) sub-path $\gamma(x, y) \cap A^S$ has length at most 3 (see Lemma 3.10). Thus $\sigma(x, y)$ and $\gamma(x, y)$ are Hausdorff distance at most 5 from one another. ∎

We now want to use the collections of horoballs $\mathcal{H}$ to define **preferred paths** between any two points in $X \cup \partial_{\mathcal{H}} X$.

Let $\alpha, \beta \in X \cup \partial_{\mathcal{H}} X$, and let $\mathcal{H}_{\alpha,\beta} = \{A_1, \ldots, A_k\}$. Note that $\mathcal{H}_{\alpha,\beta}$ is a finite linearly ordered set.

Definition 5.7: **Preferred paths.** If $\mathcal{H}_{\alpha,\beta}$ is empty, then we define

$$p_{\alpha,\beta} = \gamma(\alpha, \beta).$$

If $\mathcal{H}_{\alpha,\beta}$ contains a single horoball A, then we set

$$p_{\alpha,\beta} = \gamma(\alpha, A) \cup \sigma(a_1, a_2) \cup \gamma(A, \beta),$$

where a_1 is the terminal point of $\gamma(\alpha, A)$, and a_2 is the terminal point of $\gamma(\beta, A)$. Otherwise, we define the preferred path $p_{\alpha,\beta}$ to be

$$\gamma(\alpha, A_1) \cup r_1 \cup \gamma(A_1, A_2) \cup r_2 \cup \cdots \cup \gamma(A_{k-1}, A_k) \cup r_k \cup \gamma(A_k, \beta),$$

where the r_i are paths $\sigma(a_{i,1}, a_{i,2})$, where $a_{1,1}$ is the terminal point of $\gamma(\alpha, A_1)$ and otherwise $a_{i,1}$ is the terminal point of $\gamma(A_{i-1}, A_i)$; the point $a_{2,k}$ is the terminal point of $\gamma(A_k, \beta)$ and otherwise $a_{i,2}$ is the initial point of $\gamma(A_i, A_{i+1})$.

This describes $p_{\alpha,\beta}$ as a set, but we will often consider it as a map, and parametrize by arc length.

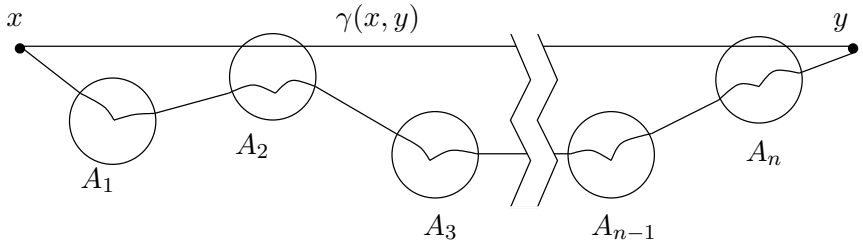

Figure 19. A preferred path.

Remark (about torsion) 5.8: We have indicated in Remarks 3.29 and 5.2 how the paths $\gamma(x,y)$ and $\gamma(A,B)$ might be defined using averages in order to ensure antisymmetry and G-equivariance.

In this framework, preferred paths will consist of all of the "possible" preferred paths, superimposed like the states of a quantum system.

These more complicated preferred paths will continue to have many of the properties that honest preferred paths do, though these properties become a deal more cumbersome to state, let alone prove.

We now make a generalization of Definition 4.17.

Definition 5.9: Suppose that the path $\gamma(\cdot,\cdot)$ passes within $\beta \leq L_1 - 2\delta$ of the (distinct) horoballs A and B. Define

$$\gamma_{AB}^{\beta}(\cdot,\cdot)$$

to be that portion of $\gamma(\cdot,\cdot)$ which connects the β-neighborhoods of A and B, where the symbol '$\cdot$' is used to mean a point or a horoball.

Definition 5.10: $\overline{p}_{x,y}$ denotes the path obtained from $p_{x,y}$ by replacing each subsegment of the form $\sigma(a_{i,1},a_{i,2})$ by the corresponding geodesic $\gamma(a_{i,1},a_{i,2})$.

LEMMA 5.11: *If $A_i \in \mathcal{H}_{x,y}$, then either $\overline{p}_{x,y} \cap A^L$ is geodesic for any $L \geq 2\delta$, or $d(a_{i,1},a_{i,2}) < 3$, where $a_{i,1}$ and $a_{i,2}$ are the first and last points on $p_{x,y} \cap A_i$.*

Proof. Let $\overline{p}_{x,y} \cap A \cap D^{-1}(2\delta) = \{b_1,b_2\}$ and let $\gamma = \gamma_{b_1,b_2}$. This consists of two vertical segments, and a single horizontal segment of length at most 3.

If γ intersects $D^{-1}(L_1)$ nontrivially, then the vertical segments in γ and the vertical segments in $\overline{p}_{x,y} \cap A^L$ coincide, and the horizontal segments have the same length. Therefore, in this case, $\overline{p}_{x,y} \cap A^L$ is a geodesic.

Otherwise let c_1 and c_2 be the deepest points of the vertical subsegments of γ, with c_1 lying directly above $a_{i,1}$ and c_2 directly above $a_{i,2}$. Then the distance between c_1 and c_2 is at most 3, which implies that the distance between $a_{i,1}$ and $a_{i,2}$ is less than 3. ∎

5.1.1. *Parametrizations.* We pause briefly to discuss the parametrizations of the paths $\gamma(x,y)$, $p_{x,y}$ and $\overline{p}_{x,y}$. These parametrizations are always by arc length. For each $x, y \in X \cup \partial_{\mathcal{H}} X$ choose $\mathcal{I}_{x,y}$, $\overline{\mathcal{I}}_{x,y}$ and $\mathcal{J}_{x,y}$ as follows:

If $x, y \in X$, then $\mathcal{I}_{x,y} = [0, \text{length}(p_{x,y})]$, and similarly $\overline{\mathcal{I}}_{x,y} = [0, \text{length}(\overline{p}_{x,y})]$ and $\mathcal{J}_{x,y} = [0, d(x,y)]$.

If $x \in X$ and $y \in \partial_{\mathcal{H}} X$, then all of the intervals are $[0, \infty)$.

If $x \in \partial_{\mathcal{H}} X$, then each of the paths $\gamma(x,y)$, $p_{x,y}$ and $\overline{p}_{x,y}$ begins with an infinite vertical ray. If y is contained in the same L_1-horoball as x, then the three paths coincide and we define $\mathcal{I}_{x,y} = \overline{\mathcal{I}}_{x,y} = \mathcal{J}_{x,y} = (-\infty, L_1 - D(y)]$.

Suppose then that $x \in \partial_{\mathcal{H}} X$ and that y is not contained in the same L_1-horoball as x. Let x' be the first point on $p_{x,y}$ contained in $D^{-1}(L_1)$. Then the interval $\mathcal{I}_{x',y}$ is already defined, and we define $\mathcal{I}_{x,y} = (\infty, 0] \cup \mathcal{I}_{x',y}$, and suppose that $p_{x,y}(0) = x'$.

We make similar definitions for $\overline{\mathcal{I}}_{x,y}$ and $\mathcal{J}_{x,y}$.

PROPOSITION 5.12: *For any two points x and y in $X \cup \partial_{\mathcal{H}} X$, there are nondecreasing functions*

$$f = f_{xy} \colon \mathcal{I}_{x,y} \to \mathcal{J}_{x,y}$$

and

$$g = g_{xy} \colon \mathcal{J}_{x,y} \to \mathcal{I}_{x,y}$$

satisfying:

(1) *For all but finitely many points $t \in \mathcal{I}_{x,y}$ and $s \in \mathcal{J}_{x,y}$, $f'(t)$ and $g'(s)$ exist and are either 0 or 1.*

(2) *For any $t \in \mathcal{I}_{x,y}$, and $s \in \mathcal{J}_{x,y}$ we have*

$$d(p_{x,y}(t), \gamma(x,y)(f(t)) \leq K + 12\delta + 9, \text{ and}$$

$$d(p_{x,y}(g(s)), \gamma(x,y)(s) \leq K + 12\delta + 9.$$

(3) *If $p_{x,y}(t)$ is outside of $A^{L_1-(K+2\delta)}$ for every $A \in \mathcal{H}_{x,y}$, then $f'(t) = 1$.*
Similarly, if $\gamma(x,y)(t)$ is outside of $A^{L_1-(K+2\delta)}$ for every $A \in \mathcal{H}_{x,y}$, then $g'(t) = 1$.

(4) *Let $A \in \mathcal{H}_{x,y}$ and $\hat{A} = A^{L_1-(K+2\delta)}$. Let $[t_1, t_2] = p_{x,y}^{-1}(\hat{A})$ and $[s_1, s_2] = \gamma(x,y)^{-1}(\hat{A})$. The total waiting times in A, i.e., the maximum of the measures of*

$$\{t_1 \le t \le t_2 \colon f'(t) = 0\} \quad \text{and} \quad \{s_1 \le s \le s_2 \colon g'(s) = 0\}$$

is at most $4K + 18\delta + 9$. Moreover, the total jumping times in A, i.e. the maximum of the measures of

$$\{s_1 \le s \le s_2 \colon \nexists t,\, f(t) = s\} \quad \text{and} \quad \{t_1 \le t \le t_2 \colon \nexists s,\, g(s) = t\}$$

is at most $4K + 18\delta + 9$.

Proof. We suppose for the moment that $x, y \in D^{-1}[0, L_1]$. We will deal with the other cases at the end of the proof. In fact, we only deal with the cases that $x, y \in D^{-1}[0, L_1] \cup \partial_{\mathcal{H}} X$. The other cases are similar, but are not required for later applications.

We use the same symbol for a path and its image, throughout. Note that a map from the domain of one path to the domain of another induces a map from the image of one path to the image of the second. Conversely, if both paths are embeddings, then a map between their images induces a map between their domains.

Recall that the path $p_{x,y}$ is defined to be a concatenation of geodesic segments outside of horoballs, and paths of the form $\sigma(u,v)$ through horoballs.

Consider the difference between the paths $\sigma(u,v)$ and $\gamma(u,v)$. This consists of a path of length at most 9 in $\sigma(u,v)$ and a path of length at most 5 in $\gamma(u,v)$. Define functions

$$\mathcal{I}_{x,y} \xrightarrow{f_1} \overline{\mathcal{I}}_{x,y}, \quad \text{and} \quad \overline{\mathcal{I}}_{x,y} \xrightarrow{g_1} \mathcal{I}_{x,y},$$

so that the induced maps on the intersection of the images are the identity, and the induced maps on the differences send a component to its initial point.

The functions f_1 and g_1 have derivatives at all but finitely many points, the derivatives are always 0 or 1, and the waiting times and jumping distances are at most 9 per horoball in $\mathcal{H}_{x,y}$.

It is not difficult to see that it suffices now to define appropriate functions

$$\overline{\mathcal{I}}_{x,y} \xrightarrow{f_2} \mathcal{J}_{x,y} \quad \text{and} \quad \mathcal{J}_{x,y} \xrightarrow{g_2} \overline{\mathcal{I}}_{x,y}.$$

There are three cases, depending on whether $\mathcal{H}_{x,y}$ is empty, contains a single horoball, or contains at least two horoballs.

CASE 1: If $\mathcal{H}_{x,y}$ is empty, then $p_{x,y} = \overline{p}_{x,y} = \gamma(x,y)$, and we can set both f and g equal to the identity map.

CASE 2: Suppose $\mathcal{H}_{x,y}$ is a single horoball. This case is strictly easier than Case 3, and we leave it as an exercise.

CASE 3: There is more than one horoball in $\mathcal{H}_{x,y}$. By Axiom (A2), the intersection $\gamma(x,y) \cap A_i^{L_1-K}$ is non-empty for each $1 \leq i \leq n$.

Let u be the first point on $\gamma(x,y)$ which intersects $A_1^{L_1-K}$, and let v be the first point on $\overline{p}_{x,y}$ which intersects $A_1^{L_1-K}$. We define f_2 and g_2 to be the identity on $[0,(u,v)_x]$. Note that the segments $\overline{p}_{x,y}[0,(u,v)_x]$ and $\gamma(x,y)[0,(u,v)_x]$ have the same length and are identified under the map to the comparison tripod $Y_{u,v,x}$. Thus, for example, the distance between $\overline{p}_{x,y}(t)$ and $\gamma(x,y)(f_2(t))$ is at most δ for all $t \in [0,(u,v)_x]$. There is clearly no waiting or jumping in this interval.

Similarly, let w be the last point on $\gamma(x,y)$ intersecting A_n, and let z be the last point on $\overline{p}_{x,y}$ intersecting A_n. We can define the map

$$f_2\colon \Big[|\overline{p}_{x,y}| - (y_1,y_2)_y, |\overline{p}_{x,y}|\Big] \to \Big[d(x,y) - (w,z)_y, d(x,y)\Big],$$

to be the unique (orientation preserving) isometry, and $g_2 = (f_2)^{-1}$ on this interval. Once again, the images of these intervals are identified under the map to the comparison tripod $Y_{w,z,y}$. We also define $y_{1,n} = \gamma(x,y)(d(x,y) - (w,z)_y)$ and $y_{2,n} = \overline{p}_{x,y}(d(x,y) - (w,z)_y)$.

We now define the maps f_2 and g_2 on the region between adjacent horoballs A_i and A_{i+1}.

Let $1 \leq i \leq n-1$. Let a be the last point on $\gamma(x,y) \cap A_i^{L_1-K}$, and let b be the first point on $\gamma(x,y) \cap A_{i+1}^{L_1-K}$. Similarly, let c be the last point on $\overline{p}_{x,y} \cap A_i^{L_1-K}$ and let d be the first point on $\overline{p}_{x,y} \cap A_{i+1}^{L_1-K}$. Suppose that $a = \gamma(x,y)(t_a)$, and define times t_b, t_c and t_d analogously. By drawing two comparison tripods (one for the triangle $\Delta(a,b,c)$ and one for $\Delta(b,c,d)$) we see that

$$I_{a,b} = \Big[t_a + (b,c)_a, t_b - (c,d)_b\Big] \quad \text{and} \quad I_{c,d} = \Big[t_c + (a,b)_c, t_d - (b,c)_d\Big],$$

have the same length (note that between $A_i^{L_1-K}$ and $A_{i+1}^{L_1-K}$, both $\overline{p}_{x,y}$ and $\gamma(x,y)$ are geodesics parametrized by arc length).

Define $f_2\colon I_{c,d} \to I_{a,b}$ and $g_2\colon I_{a,b} \to I_{c,d}$ to be orientation preserving isometries. Note that the comparison tripods imply that for $t \in I_{a,b}$,

$$d(\gamma(x,y)(t), \overline{p}_{x,y}(g_2(t))) \leq 2\delta,$$

and for $s \in I_{c,d}$,

$$d(\overline{p}_{x,y}(s), \gamma(x,y)(f_2(s))) \leq 2\delta.$$

We claim now that those parts of $\overline{\mathcal{I}}_{x,y}$ on which f_2 is not yet defined are segments $I_1, \ldots, I_n$ so that $\overline{p}_{x,y}(I_j) \subseteq N_{2\delta}(A_j^{L_1-K})$, and similarly that those parts of $\mathcal{J}_{x,y}$ on which g_2 is not yet defined are segments $I_1', \ldots, I_n'$ so that $\gamma(x,y)(I_j') \subseteq N_{2\delta}(A_j^{L_1-K})$. To see this consider, for example, the points $\gamma(x,y)(t_a + (b,c)_a)$ and $\overline{p}_{x,y}(t_c + (a,b)_c)$, for the points a, b, c and d as above. These are typical endpoints of the intervals on which f_2 and g_2 are not yet defined. Any geodesic between a and c lies entirely within the convex set $A_i^{L_1-K}$. Therefore, the point $\gamma(x,y)(t_a + (b,c)_a)$ lies within δ of $A_i^{L_1-K}$, because it is in the preimage of the central point of the comparison tripod $Y_{a,b,c}$. Now, $\overline{p}_{x,y}(t_c + (a,b)_c)$ lies within δ of the point on the geodesic $[c,b]$ in the preimage of the central point of $Y_{a,b,c}$, and in particular, $\overline{p}_{x,y}(t_c + (a,b)_c)$ lies within 2δ of $A_i^{L_1-K}$. Similar arguments for all other points prove the claim.

We now define f_2 and g_2 on the remaining subsegments, I_j and I_j'. We fix $1 \leq j \leq n$. Let x_1 and y_1 be the images of the endpoints of I_j' under $\gamma(x,y)$, and let x_2 and y_2 be the images of the endpoints of I_j under $\overline{p}_{x,y}$. The path $\overline{p}_{x,y}(I_j)$ is a concatenation of three geodesics, one traveling towards A_j as quickly as possible, one through A_j and one traveling away from A_j as quickly as possible. Let x_3 be the first point on $\overline{p}_{x,y}(I_j) \cap A_j$ and let y_3 be the last point on $\overline{p}_{x,y}(I_j) \cap A_j$ (see Figure 20.)

It is clear from the previous arguments that $d(x_1,x_2), d(y_1,y_2) \leq 2\delta$. Also, x_2 is within $K + 2\delta$ of A_j. However, between x_2 and x_3, the path $\overline{p}_{x,y}$ is traveling as quickly as possible towards A_j, so $d(x_2,x_3) \leq K + 2\delta$. Similarly, $d(y_2,y_3) \leq K + 2\delta$.

Suppose that $x_1 = \gamma(x,y)(s_1), y_1 = \gamma(x,y)(s_2), x_2 = \overline{p}_{x,y}(t_1), x_3 = \overline{p}_{x,y}(t_2)$, $y_3 = \overline{p}_{x,y}(t_3)$ and $y_2 = \overline{p}_{x,y}(t_4)$.

For $t \in [t_1,t_2]$, define $f_2(t) = s_1$, and for $t \in [t_3,t_4]$ define $f_2(t) = s_2$.

Drawing comparison tripods as in Figure 21 allows us to see that

$$J_1 = \Big[s_1 + (x_3,y_1)_{x_1}, s_2 - (x_3,y_3)_{y_1}\Big], \quad \text{and}$$

$$J_2 = \Big[t_2 + (x_1, y_1)_{x_3}, t_3 - (x_3, y_1)_{y_3}\Big]$$

have the same length.

Assume first that J_1 and J_2 are nonempty, and define $f_2\colon J_2 \to J_1$ and $g_2\colon J_1 \to J_2$ to be orientation preserving isometries.

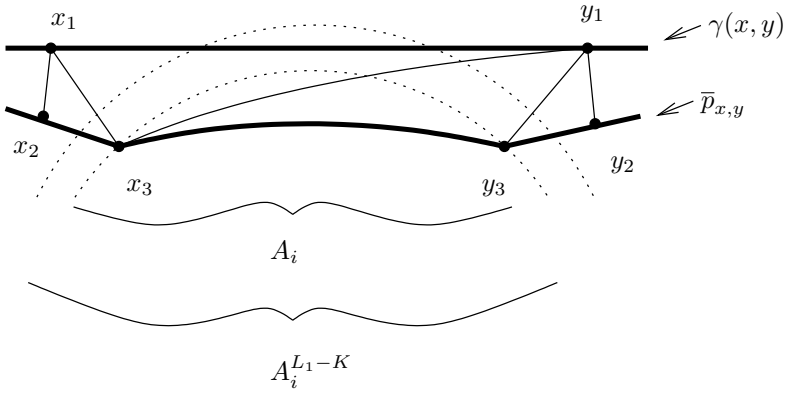

Figure 20. $\gamma(x,y)$ and $\overline{p}_{x,y}$ near $A_i \in \mathcal{H}_{x,y}$.

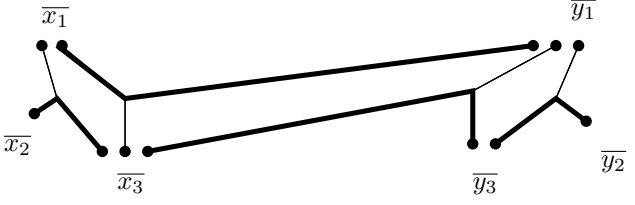

Figure 21. Comparison tripods for the triangles in Figure 20.

It remains to define f_2 on the intervals $K_1 = \Big[t_2, t_2 + (x_1, y_1)_{x_3}\Big]$ and $K_2 = \Big[t_3 - (x_3, y_1)_{y_3}, t_3\Big]$, and g_2 on the intervals $K_3 = \Big[s_1, s_1 + (x_3, y_1)_{x_1}\Big]$ and $K_4 = \Big[s_2 - (x_3, y_3)_{y_1}, s_2\Big]$. Since $d(x_1, x_2) \le 2\delta$ and $d(x_2, x_3) \le K + 2\delta$, we have $d(x_1, x_3) \le K + 4\delta$. This implies that the lengths of K_1 and K_3 are at most $K + 4\delta$. The same is true of the lengths of K_2 and K_4.

If $t \in K_1$, define $f_2(t) = s_1 + (x_3, y_1)_{x_1}$. If $t \in K_2$, define $f_2(t) = s_2 - (x_3, y_3)_{y_1}$. If $s \in K_3$, define $g_2(s) = t_2 + (x_1, y_1)_{x_3}$. If $s \in K_4$, define $g_2(s) = t_3 - (x_3, y_1)_{y_3}$.

If J_1 and J_2 are empty, we must define f_2 and g_2 slightly differently:

$$f_2([t_2, t_2 + (y_1, y_3)_{x_3}]) = s_1$$
$$f_2([t_2 + (y_1, y_3)_{x_3}, t_4]) = s_2$$
$$g_2([s_1, s_1 + (y_1, x_3)_{x_1}]) = t_2$$
$$g_2([s_1 + (y_1, x_3)_{x_1}, s_2]) = t_3$$

In this case, the intersections of $\gamma(x, y)$ and $\overline{p}_{x,y}$ with $A_i^{L_1 - K}$ are quite small.

We have now defined f_2 and g_2 in the case that $x, y \in D^{-1}[0, L_2]$. It is not difficult to see that the functions $f = f_2 \circ f_1$ and $g = g_1 \circ g_2$ satisfy the requirements of the proposition. In fact (and this is used below), we have shown that they satisfy the requirements of the proposition with the constants reduced by 6δ.

Finally, we now turn to the case that one or both of x, y lie in $\partial_{\mathcal{H}} X$.

Suppose $x \in \partial_{\mathcal{H}} X$ and $y \in D^{-1}[0, L_1]$. As in the discussion above the proposition, let x' be the first point on $p_{x,y}$ so that $x' \in D^{-1}(L_1)$. Let x'' be the first point on $\gamma_{x,y}$ so that $x'' \in D^{-1}(L_1)$. Then $\mathcal{I}_{x,y} = (-\infty, 0] \cup \mathcal{I}_{x',y}$ and $\mathcal{J}_{x,y} = (-\infty, 0] \cup \mathcal{J}_{x'',y}$. Note also that $p_{x',y} = p_{x,y}|_{\mathcal{I}_{x',y}}$, by the definition of preferred paths.

We define f_{xy} to be the identity on $(-\infty, 0]$. If we can prove that $d(x', x'') \leq 6\delta$, then there are obvious maps $\phi\colon \mathcal{J}_{x',y} \to \mathcal{J}_{x'',y}$ and $\psi\colon \mathcal{J}_{x'',y} \to \mathcal{J}_{x',y}$ which identify the long side of the comparison tripod $Y_{x,x',y}$, and collapse the short sides to the center of the tripod. If we have such functions ϕ and ψ, then we define

$$f_{x,y}|_{\mathcal{I}_{x',y}} = \phi \circ f_{x',y}\colon \mathcal{I}_{x',y} \to \mathcal{J}_{x'',y},$$

and

$$g_{x,y}|_{\mathcal{J}_{x'',y}} = g_{x',y} \circ \psi\colon \mathcal{J}_{x'',y} \to \mathcal{I}_{x',y}.$$

These will satisfy the requirements of the proposition so long as $d(x', x'') \leq 6\delta$.

To see this, we argue as follows: Let σ be that part of $p_{x,y}$ from x until the second horoball in $\mathcal{H}_{x,y}$ (or to be all of $p_{x,y}$ if $\mathcal{H}_{x,y}$ is a singleton). It is not difficult to see that σ must be a geodesic, because that part after x' travels away from the horoball as quickly as possible (by the definition of preferred paths). Let z be the endpoint of σ which is not x'. By Axiom (A2) of Section 4, and

Lemma 4.5, there is a point $z' \in \gamma(x,y)$ which lies within K of the horoball, A, which contain z. The geodesic between z and z' lies entirely within the convex set A^{L_1-K}. The (partially ideal) triangle with vertices x, z, z' is 3δ-slim, which proves that $d(x', x'') \leq 6\delta$ (because both $p_{x,y}$ and $\gamma(x,y)$ are vertical on the interval $(-\infty, 3\delta)$.

We now suppose that $x, y \in \partial_{\mathcal{H}} X$. Then there are points $x', y' \in p_{x,y}$ and $x'', y'' \in \gamma(x,y)$ so that

$$\mathcal{I}_{x,y} = (-\infty, 0] \cup \mathcal{I}_{x',y'} \cup [|p_{x',y'}|, \infty),$$

and

$$\mathcal{J}_{x,y} = (-\infty, 0] \cup \mathcal{J}_{x'',y''} \cup [d(x'', y''), \infty).$$

As before, $d(x', x''), d(y'y,'') \leq 6\delta$. Thus, there are obvious functions $\phi' \colon \mathcal{J}_{x',y'} \to \mathcal{J}_{x'',y''}$ and $\psi' \colon \mathcal{J}_{x'',y''} \to \mathcal{J}_{x',y'}$ and we make the following definitions:

$$f_{xy}(t) = \begin{cases} t, & \text{if } t \in (-\infty, 0) \\ \phi'(f_{x',y'}(t)), & \text{if } t \in \mathcal{I}_{x',y'}, \text{ and} \\ t + (d(x'', y'') - |p_{x',y'}|) & \text{otherwise,} \end{cases}$$

and

$$g_{xy}(t) = \begin{cases} t, & \text{if } t \in (-\infty, 0) \\ \psi'(f_{x',y'}(t)), & \text{if } t \in \mathcal{I}_{x',y'}, \text{ and} \\ t + (|p_{x',y'}| - d(x'', y'')) & \text{otherwise.} \end{cases}$$ ∎

We now state a couple of corollaries of Proposition 5.12

COROLLARY 5.13: *For any $a, b \in X \cup \partial_{\mathcal{H}} X$, the Hausdorff distance between $p_{a,b}$ and $\gamma(a,b)$ is at most $K + 12\delta + 9$. The Hausdorff distance between $p_{a,b}[t_1, t_2]$ and $\gamma(a,b)[f(t_1), f(t_2)]$ is at most $3K + 21\delta + 14$.*

Proof. The first assertion is obvious from Proposition 5.12.(2). Suppose that $t \in [t_1, t_2]$. Since the function f from Proposition 5.12 is monotone, there is some $s \in [f(t_1), f(t_2)]$ so that $f(t) = s$. By Proposition 5.12.(2), $d(\gamma(a,b)(s), p_{a,b}(t) \leq K + 12\delta + 9$. Conversely, suppose that $s \in [f(t_1), f(t_2)]$. It follows from monotonicity and Proposition 5.12.(4) that there is a point $s' \in f([t_1, t_2]$ with $|s - s'| \leq \frac{4K+18\delta+9}{2}$. It follows that $\gamma(a,b)(s)$ is no further than $K + 12\delta + 9 + \frac{4K+18\delta+9}{2} \leq 3K + 21\delta + 14$ from $p_{a,b}[t_1, t_2]$. ∎

COROLLARY 5.14: *The path $p_{x,y}$ is a (λ,ϵ)-quasi-geodesic, for*

$$\lambda = 2 \quad \textit{and} \quad \epsilon = 20K + 120\delta + 72.$$

Proof. Let t_1 and t_2 be points in $[0, \text{length}(p_{x,y})]$, and let $x_1 = p_{x,y}(t_1)$ and $x_2 = p_{x,y}(t_2)$. Since $p_{x,y}$ is parametrized by arc length, $|t_2 - t_1| \geq d(x_1, x_2)$ is automatic.

Let $x_i' = \gamma(x,y)(f(t_i))$, for $i \in \{1,2\}$ and f as described in Proposition 5.12. Then $d(x_1', x_2') \leq d(x_1, x_2) + 2(K + 12\delta + 9)$. The arc length $|t_2 - t_1|$ is at most $|f(t_2) - f(t_1)| = d(x_1', x_2')$ plus the amount of time in $[t_1, t_2]$ that $f'(t) = 0$. The number of horoballs in $\mathcal{H}_{x,y}$ that $p_{x,y}$ passes within $K + 2\delta$ of between time t_1 and t_2 is at most $2 + \frac{|t_2 - t_1|}{L_1}$; the amount of time $f'(t) = 0$ when $p_{x,y}(t)$ is near a horoball is at most $4K + 18\delta + 9 < \frac{1}{2}L_1$, by part (4) of Proposition 5.12. Thus

$$|t_2 - t_1| \leq d(x_1', x_2') + \Big(2 + \frac{|t_2 - t_1|}{L_1}\Big)(4K + 18\delta + 9),$$

which implies

$$\frac{1}{2}|t_2 - t_1| \leq \Big(1 - \frac{4K + 18\delta + 9}{L_1}\Big)|t_2 - t_1| \leq d(x_1', x_2') + 2(4K + 18\delta + 9),$$

and so

$$\begin{aligned} |t_2 - t_1| &\leq 2d(x_1', x_2') + 4(4K + 18\delta + 9) \\ &\leq 2d(x_1, x_2) + 4(K + 12\delta + 9) + 4(4K + 18\delta + 9) \\ &= 2d(x_1, x_2) + 20K + 120\delta + 72 \qquad \blacksquare \end{aligned}$$

We now make some some other useful observations about preferred paths.

LEMMA 5.15: *Preferred paths are G-equivariant.*

Remark 5.16: It follows from Corollary 5.14 that sub-paths of preferred paths are quasi-geodesic, and hence form δ''-slim triangles for **some** δ''. However, the δ'' coming from (λ,ϵ)-quasi-geodesicity is **much** worse than the δ' below, and our previous choice of L_1 secretly takes into account the particular bound given in the next proposition.

PROPOSITION 5.17: *There exists a constant*

$$\delta' = 6K + 48\delta + 28 \leq \frac{1}{4}L_1$$

so that any triangle or bigon whose sides are sub-paths of preferred paths is δ'-slim.

This includes triangles any or all of whose vertices are in $\partial_{\mathcal{H}} X$.

Proof. Let P, Q, and R be three points in $X \cup \partial_{\mathcal{H}} X$, joined by parts of preferred paths p_{ab}, p_{cd} and p_{ef}. Specifically, suppose that $p_{ab}[t_1, t_2]$ joins P to Q, $p_{cd}[t_3, t_4]$ joins Q to R, and $p_{ef}[t_5, t_6]$ passes from R to P, where $t_i \in [-\infty, \infty)$ for i odd and $t_i \in (-\infty, \infty]$ for i even.

By Corollary 5.13, the side $p_{ab}[t_1, t_2]$ is Hausdorff distance at most $3K + 21\delta + 14$ from the geodesic segment $\gamma(a, b)[f(t_1), f(t_2)]$, where $f = f_{ab}$ is the function from Proposition 5.12. Furthermore, the endpoints $\gamma(a, b)(f(t_1))$ and $\gamma(a, b)(f(t_2))$ of this segment are (if not ideal) at most $K + 12\delta + 9$ from P and Q respectively, by Proposition 5.12.(2). Similar statements can be made about the other two sides. If none of the endpoints of the geodesic segments are ideal, they may be joined together by segments of length at most $2(K+12\delta+9)$, to form a geodesic hexagon, which must be 4δ-slim. Otherwise, they may be partially joined up to form a pentagon with one ideal vertex, a quadrilateral with two ideal vertices, or simply an ideal triangle. Call the (possibly partially ideal) polygon so obtained H. Subdividing and using Lemma 2.11, it is easy to see that H is 6δ-slim.

Suppose that $x \in p_{ab}[t_1, t_2]$. There is some point x' on $\gamma(a, b)[f(t_1), f(t_2)]$ at most $3K + 21\delta + 14$ away from x (see Figure 22.) Since H is 6δ-slim,

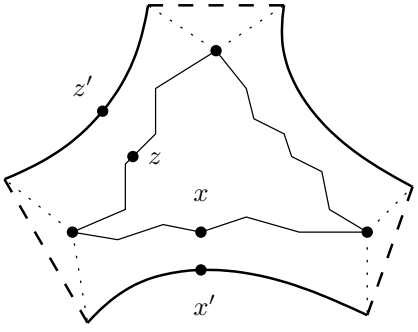

Figure 22. The generic case is pictured. It is also possible for z' to lie on one of the dashed segments, in which case it is even closer to the other two sides.

there is some point z' on another side of H so that $d(z', x') \leq 6\delta$. If z' is on

$\gamma(c,d)[f(t_3),f(t_4)]$ or $\gamma(e,f)[f(t_5),f(t_6)]$, then there is a point z on $p_{cd}[t_3,t_4]$ or $p_{ef}[t_5,t_6]$ with $d(z,z') \leq 3K+21\delta+14$. Otherwise, z' is at most $K+12\delta+9+\delta$ from one of the vertices of our original triangle. In either case x is within $2(3K+21\delta+14)+6\delta = 6K+48\delta+28$ of one of the other two sides of the triangle. ∎

Remark (about torsion) 5.18: Any of the 'possible' preferred paths referred to in Remark 5.8 will satisfy all of the above properties, and an identical proof suffices to establish this.

5.2. Preferred triangles.

Definition 5.19: Let Δ be a 2-simplex. A **preferred triangle** is a continuous map

$$\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$$

so that if e is one of the three sides of Δ, then $\psi|_e$ is an embedding, whose image is a preferred path.

Remark 5.20: If $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ is a preferred triangle, then $\psi|_{\Delta^{(0)}}$ is injective.

Assumption 5.21: For ease of exposition, we will always assume that if $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ is a preferred triangle, then ψ does not send the vertices of Δ to $D^{-1}(L_2)$. In our applications, this will always be the case. The following definitions can be extended to the case that vertices are sent by ψ to $D^{-1}(L_2)$, but at the expense of making the statements a bit more cumbersome.

For the rest of this section, we fix a preferred triangle

$$\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$$

so that $\psi(v) \neq L_2$ for $v \in \Delta^{(0)}$. We refer to the elements of $\Delta^{(0)}$ as **corners**.

We next define a **skeletal filling** of a preferred triangle. This filling will take the form of a 1-complex (with 3 different kinds of edges) inscribed on the 2-simplex Δ.

Definition 5.22: The points in $(D\circ\psi)^{-1}(L_2)$ will be referred to as L_2**-vertices**.

We now describe the different types of preimages of L_2-horoballs which can occur in a preferred triangle.

Definition 5.23: Suppose that $P \in \mathcal{H}$ is such that $P^{L_2} \cap \psi(\partial\Delta) \neq \emptyset$. Recall that the unique point in ∂X which is a limit of points in P is called e_P.

(1) If $\psi^{-1}(P^{L_2} \cup e_P)$ contains a single corner of $\partial\Delta$, then we say that $\psi^{-1}(P^{L_2} \cup e_P)$ is a **bite**.
(2) If $\psi^{-1}(P^{L_2} \cup e_P)$ intersects exactly one side of the triangle $\partial\Delta$, then $\psi^{-1}(P^{L_2} \cup e_P)$ is a **nibble**.
(3) If $\psi^{-1}(P^{L_2} \cup e_P)$ intersects exactly two sides of the triangle $\partial\Delta$ and is not a bite, then $\psi^{-1}(P^{L_2} \cup e_P)$ is a **dip**.
(4) If $\psi^{-1}(P^{L_2} \cup e_P)$ intersects all three sides of the triangle $\partial\Delta$, then it is a **plunge**.

See Figure 23.

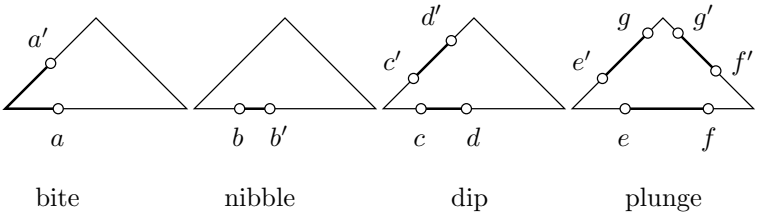

Figure 23. Possibilities for the preimage of a single L_2-horoball are shown in bold. For any letter x, the vertices x and x' form a **pair** (Definition 5.26).

LEMMA 5.24: *Any L_2-vertex is in the boundary of a unique bite, dip, nibble, or plunge.*

LEMMA 5.25: *There is a partition of the L_2-vertices into sets, $\{S_i\}$, of cardinality 2 , so that if $S_i = \{x, y\}$, then $d(\psi(x), \psi(y)) \leq 1$.*

Proof. Suppose that x is an L_2-vertex. Then x is in the boundary of a unique bite, dip, nibble, or plunge. We consider each of them in turn, in order to define the set S_i of which x is a member. Let P^{L_2} be the L_2-horoball containing $\phi(x)$. We delay the proof that if $S_i = \{x, y\}$, then $d(\psi(x), \psi(y)) \leq 1$ until later in the proof.

Suppose that x is contained in the boundary of a bite. Then there is a unique L_2-vertex $y \neq x$ in $\partial\Delta$ so that $\psi(y) \in P^{L_2}$. Define $S_i = \{x, y\}$ in this case.

Suppose that x is contained in the boundary of a dip. Then the set of L_2-vertices on the boundary of this dip has cardinality 4. We partition this into two sets by pairing the L_2-vertices on different sides which are closest to the common vertex of $\partial\Delta$.

Suppose that x is contained in a nibble. Then the boundary of the nibble has cardinality 2, and we define the set containing x to be this boundary.

Finally, suppose that x is contained in a plunge. The boundary of the plunge has cardinality at most 6; we partition the boundary into sets of size 2 by pairing L_2-vertices which are closest to a given vertex of $\partial\Delta$.

The procedure described above (and also in Figure 23) defines a partition of the L_2-vertices into sets $\{S_i\}$ of cardinality 2. It remains to investigate the distance between $\psi(x)$ and $\psi(y)$ where $\{x, y\} = S_i$.

We consider first the case that $S_i = \{x, y\}$, and that x and y lie in different sides of $\partial\Delta$. This covers the cases when x and y are contained in the boundary of a bite, a dip, or a plunge.

Now, $\psi(x) \in D^{-1}(L_2) \cap p_{a,b}$, say. Suppose that $\psi(y) \in D^{-1}(L_2) \cap p_{a,c}$ (we can arrange both of these by relabeling a, b and c if necessary). Preferred paths are chosen to be vertical from depth L_1 to at least L_2+1 (see Definition 5.4). In fact, because preferred paths travel as quickly as possible towards L_1-horoballs, they must be vertical from depth 2δ to at least $L_2 + 1$.

Let z_x be the point on $p_{a,b}$ which is at depth L_1 and lies closest to $\psi(x)$, and let z_y be defined similarly in relation to $\psi(y)$.

Preferred paths are $\frac{1}{4}L_1$-slim by Proposition 5.17. Thus, there is a point $w \in p_{a,c} \cup p_{b,c}$ which is within $\frac{1}{4}L_1$ of z_x.

Suppose first that $w \in p_{a,c}$. Then, since preferred paths are vertical between depths 2δ and $L_2 + 1$, it is not difficult to see that $d(z_x, z_y) \leq \frac{1}{2}L_1$. Because of the geometry of combinatorial horoballs, distance between points on vertical paths strictly decreases with increasing depth. Thus, by depth $1\frac{1}{2}L_1$ in P, the paths $p_{a,b}$ and $p_{a,c}$ must be within 1 of each other. It is now easy to see that $d(\psi(x), \psi(y)) \leq 1$.

Suppose then that $w \in p_{b,c}$. Then, arguing as above, there is a point $w_x \in D^{-1}(L_1) \cap P \cap p_{b,c}$ so that $d(z_x, w_x) \leq \frac{1}{4}L_1$.

Similarly, there is a point $w' \in p_{a,b} \cup p_{b,c}$ within $\frac{1}{4}L_1$ of z_y. If $w' \in p_{a,b}$, then we see that $d(\psi(x), \psi(y)) \leq 1$, just as above. Therefore, suppose that $w' \in p_{b,c}$, in which case we have some $w_y \in D^{-1}(L_1) \cap P \cap p_{b,c}$ so that $d(z_y, w_y) \leq \frac{1}{4}L_1$.

Note that the part of $p_{b,c}$ between w_x and w_y is $\sigma(w_x, w_y)$. We consider the maximum depth $\sigma(w_x, w_y)$. First note that because of the choice of x and y, it must be that $p_{b,c} \cap P^{L_2} = \emptyset$. Therefore, the maximum depth of $\sigma(w_x, w_y)$ is less than L_2. If this depth is at most $L_2 - 2$, then at depth $L_2 - 2$ the preferred paths $p_{a,b}$ and $p_{a,c}$ are at most 3 apart, which implies that at depth L_2 they are at most 1 apart, as required. Therefore, we are left with the case that the maximum depth of $\sigma(w_x, w_y)$ is exactly $L_2 - 1$. In this case, at depth $L_2 - 2$, the preferred paths $p_{a,b}$ and $p_{b,c}$ are distance at most 1 apart, as are the preferred paths $p_{a,c}$ and $p_{b,c}$. Because $\sigma(w_x, w_y)$ has maximum depth $L_2 - 1$, at depth $L_2 - 2$ the two vertices in $p_{b,c}$ are distance at most 2 apart. This implies that the distance (at depth $L_2 - 2$) between the vertices in $p_{a,b}$ and $p_{a,c}$ at depth $L_2 - 2$ is at most 4. This implies that $d(\psi(x), \psi(y)) \leq 1$.

We are now left with the case that x and y are contained in the same side of $\partial\Delta$, so in the boundary of a nibble. Suppose that $\psi(x), \psi(y) \in p_{a,b}$. Define the points $z_x, z_y \in D^{-1}(L_1) \cap P \cap p_{a,b}$. There are points $w_x, w_y \in D^{-1}(L_1) \cap P \cap (p_{b,c} \cup p_{a,c})$ so that $d(z_x, w_x), d(z_y, w_y) \leq \frac{1}{2}L_1$. Here there are two cases to consider, depending on whether w_x and w_y are contained in the same preferred path, or not. In any case, The preferred paths $p_{a,c}$ and $p_{b,c}$ do not intersect P^{L_2} (because we are considering a nibble). However, by Proposition 5.17, one of $p_{a,c}$ or $p_{b,c}$ must intersect $P^{L_2 - \frac{1}{4}L_1}$ nontrivially.

There are a number of cases to consider, depending on the maximum depth in P of $p_{a,c}$ and $p_{b,c}$. We consider the most complicated, and leave the remainder as exercises for the reader.

Suppose that $w_x \in p_{a,c}$, $w_y \in p_{b,c}$ and that both $p_{a,c}$ and $p_{b,c}$ intersect $P^{L_2 - 1}$ nontrivially. Consider the points in $p_{a,c} \cap P \cap D^{-1}(L_1)$. There are two of these points, w_x, and u_x, say. Similarly, let $p_{b,c} \cap P \cap D^{-1}(L_1) = \{w_y, u_y\}$.

It is not difficult to see (using $\frac{1}{4}L_1$-slim triangles and the properties of the paths $\sigma(r, s)$) that $d(u_x, u_y) \leq \frac{1}{2}L_1$.

Therefore, at depth $L_2 - 3$, we have points $z'_x, z'_y \in p_{a,b}$, $w'_x, u'_x \in p_{a,c}$ and $w'_y, u'_y \in p_{b,c}$ so that a primed point is directly beneath its unprimed counterpart. Now, $d(z'_x, w'_x), d(u'_x, u'_y), d(w'_y, z'_y) \leq 1$. Also, the $(L_2 - 3)$-distance between w'_x and u'_x is at most 4 (because the maximum depth of $p_{a,c}$ in P is $L_2 - 1$). Similarly, the $L_2 - 3$ distance between w'_y and u'_y is at most 4. This shows that the $L_2 - 3$ distance between z'_x and z'_y is at most 7. This implies (because $7 < 2^3$) that $d(\psi(x), \psi(y)) \leq 1$, as required.

This completes the proof of Lemma 5.25. ∎

Definition 5.26: An element of the partition in Lemma 5.25 will be called a **pair**.

Definition 5.27: We define a 1-complex $\mathrm{Skel}(\psi)$ which we call the **skeletal filling** of ψ. The vertex set of $\mathrm{Skel}(\psi)$ is equal to $\Delta^0 \cup (D \circ \psi)^{-1}(L_2)$. There are three kinds of edges:

(1) The first kind come from the subdivision of $\partial\Delta$ by the vertex set.
(2) If $\{x, y\}$ are a pair of L_2-vertices, in the sense of Definition 5.26, and $\psi(x) = \psi(y)$, then we connect x and y by an edge which we call a **ligament**.
(3) If $\{x, y\}$ are a pair of L_2-vertices so that $\psi(x) \neq \psi(y)$, then we connect x and y by an edge which we call a **rib**.

LEMMA 5.28: *The identity map on $\partial\Delta$ extends to an embedding of* $\mathrm{Skel}(\psi)$ *into* Δ.

Remark 5.29: By Lemma 5.25, the map ψ also extends to a map

$$\ddot{\psi}\colon \mathrm{Skel}(\psi) \to X \cup \partial_{\mathcal{H}} X$$

which sends each ligament to a point and each rib to a single horizontal edge.

Definition 5.30: Suppose there are (possibly degenerate) subintervals σ_1 and σ_2 of sides of $\partial\Delta$ so that $\psi(\sigma_1) = \psi(\sigma_2)$, and suppose that these are chosen to be maximal such intervals with endpoints in the vertex set of $\mathrm{Skel}(\psi)$. Then the minimal subcomplex of $\mathrm{Skel}(\psi)$ containing all edges with both endpoints in $\sigma_1 \cup \sigma_2$ is called a **leg**. A leg contains one or two **distinguished** ligaments, joining the endpoints of σ_1 to the endpoints of σ_2.

We note that ψ collapses each leg to a subsegment of a preferred path.

LEMMA 5.31: *Each leg of* $\mathrm{Skel}(\psi)$ *contains exactly on corner of* Δ. *In particular,* $\mathrm{Skel}(\psi)$ *has at most* 3 *legs.*

Proof. Consider a corner v of $\partial\Delta$, and the two sides e_1 and e_2 containing v. Let $x_1 \in e_1$ and $x_2 \in e_2$ be the endpoints of the ligament joining e_1 and e_2 which is furthest from v.

Let $\psi(e_1) = p_{a,b}$ and $\psi(e_2) = p_{a,c}$, so $\psi(v) = a$. Suppose that $\psi(x_1) = \psi(x_2) \in D^{-1}(L_2) \cap P$. Then $P \in \mathcal{H}_{a,b} \cap \mathcal{H}_{a,c}$. By Axiom (A6) and the definition of preferred paths those parts of $p_{a,b}$ and $p_{a,c}$ between a and $\psi(x_1) = \psi(x_2)$ are

identical. This implies that the segments between v and x_1 and v and x_2, and the ligaments joining them, form a leg.

If there is no ligament joining the edges e_1 and e_2 then there is no leg intersecting both of these sides. ∎

Definition 5.32: Let Γ be a graph and let e be an edge of Γ which separates Γ into two components. We say Γ_1 and Γ_2 are obtained from Γ by **surgery along** e if there are edges e_1 in Γ_1 and e_2 in Γ_2 which can be identified to give a graph isomorphic to Γ.

Definition 5.33: Successive surgery along the distinguished ligaments of $\mathrm{Skel}(\psi)$ yields a collection of at most 4 graphs, at most three of which come from legs, and exactly one of which contains pieces of all three edges of $\partial\Delta$. We call the graph which contains pieces of all three edges of $\partial\Delta$ the **middle** of $\mathrm{Skel}(\psi)$. [4] (See Figure 24.)

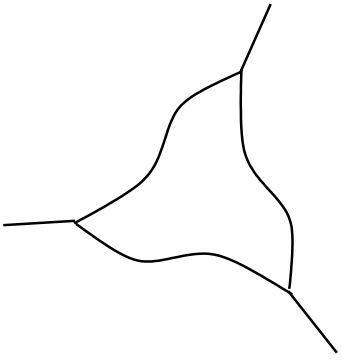

Figure 24. A somewhat stylized picture of the image of a preferred triangle.

LEMMA 5.34: *Let $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ be a preferred triangle, with edges e_1, e_2, e_3, and suppose that $\psi(e_1) = p_{a,b}$, $\psi(e_1) = p_{a,c}$ and $\psi(e_3) = p_{b,c}$. Suppose further that P is a horoball which is such that $\psi(e_1) \cap P^{L_2} \neq \emptyset$. Then $P \in \mathcal{H}_{a,b}$ and either $P \in \mathcal{H}_{a,c}$ or $P \in \mathcal{H}_{b,c}$.*

[4] In an earlier version of this paper, preferred triangles could have **feet**, as well as legs and a middle. The improved construction in Section 4 does away with the need for feet.

Proof. Follows from the fact that the image $\psi(\partial\Delta)$ is a quasi-geodesic triangle which is $\frac{1}{4}L_1$-slim and that $L_2 - \frac{1}{4}L_1 = 3L_1 - \frac{1}{4}L_1 > L_1$. ∎

Definition 5.35: Suppose that $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}}X$ is a preferred triangle, with corners v_1, v_2, v_3 so that $\psi(v_1) = a, \psi(v_2) = b$ and $\psi(v_3) = c$. Let e_1 be the edge of $\partial\Delta$ joining v_1 to v_2. If $\psi^{-1}(P^{L_2}) \neq \emptyset$ and $P \in \mathcal{H}_{a,b} \cap \mathcal{H}_{a,c}$ then we say that P is **associated to** v_1. We define horoballs associated to v_2 and v_3 analogously.

If P is a horoball so that $\psi^{-1}(P^{L_2})$ is nonempty and contained in a leg, l, then we say that P is **associated to** l.

Remark 5.36: By Lemma 5.34, if $P^{L_2} \cap \psi(e_1) \neq 0$, then P is contained in two of $\mathcal{H}_{a,b}, \mathcal{H}_{a,c}$ and $\mathcal{H}_{b,c}$.

Therefore, for any horoball P, if $\psi^{-1}(P^{L_2})$ is nonempty, then P is associated to some corner of Δ.

It could be that a horoball P as in Definition 5.35 is contained in all three of $\mathcal{H}_{a,b}, \mathcal{H}_{a,c}$ and $\mathcal{H}_{b,c}$, in which case P is associated to all three corners of $\partial\Delta$. However

LEMMA 5.37: *There is at most one horoball associated to all three corners of* $\partial\Delta$.

PROPOSITION 5.38: *Let* $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}}X$ *be a preferred triangle and let* v *be a vertex of* $\partial\Delta$.

There is at most one horoball which is associated to v *and is not also associated to a leg in* $\mathrm{Skel}(\psi)$.

Proof. Suppose that the horoballs A and B are associated to v.

Let a, b, c be the images of the vertices of $\partial\Delta$, with $a = \psi(v)$. Then $A, B \in \mathcal{H}_{a,b} \cap \mathcal{H}_{a,c}$ and suppose that $A < B$ in the order coming from $\mathcal{H}_{a,b}$.We will prove that A is associated to a leg in $\partial\Delta$.

The preferred paths $p_{a,b}$ and $p_{a,c}$ coincide at least until they reach the horoball B and in particular, the intersections of these paths with A is identical. Since one of $A^{L_2} \cap p_{a,b}$ and $A^{L_2} \cap p_{a,c}$ is nonempty (since A is associated to v), both of these sets are non-empty.

It is not difficult to see that $\psi^{-1}(A^{L_2})$ is contained in a leg. ∎

COROLLARY 5.39: *Let* $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ *be a preferred triangle. The total number of ribs and ligaments in the feet and middle of* $\mathrm{Skel}(\psi)$ *is at most* 6. *In particular, the number of ribs in* $\mathrm{Skel}(\psi)$ *is at most* 6.

Proof. Each L_2-vertex in $\mathrm{Skel}(\psi)$ maps to a horoball associated to some corner of $\partial\Delta$. If u is an L_2-vertex in a foot or middle of $\mathrm{Skel}(\psi)$ and P is the horoball containing $\psi(u)$, then P is not associated to a leg.

For each horoball, A, the set of L_2-vertices in $\mathrm{Skel}(\psi)$ which map to A^{L_2} has cardinality at most 6, and at most 4 unless $\psi^{-1}(A^{L_2} \cup e_A)$ is a plunge. Thus there are 3 ribs or ligaments for a plunge, and at most 2 for other horoballs.

Lemma 5.37 implies that there is at most one plunge and Proposition 5.38 implies that there are at most 3 horoballs not associated to legs. The result now follows easily. ∎

COROLLARY 5.40: *Let* $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ *be a preferred triangle and let* M *be the union of the feet and middles of* $\mathrm{Skel}(\psi)$. *Then*

$$|M^{(0)}| \leq 15.$$

Proof. There are 3 vertices for the corners of $\partial\Delta$.

Also, there are at most 6 ribs or ligaments in the feet and middle of $\mathrm{Skel}(\psi)$, and each of these contributes two vertices. ∎

We now proceed to decompose preferred triangles. A key feature of preferred triangles is Proposition 5.44 below.

Definition 5.41 (Sub-pictures): Suppose that $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ is a preferred triangle. Successive surgery along all of the ribs and ligaments of $\mathrm{Skel}(\psi)$ yields a collection of graphs. We call these graphs **sub-pictures**.

The following result follows immediately from the fact that the depth function D is continuous.

LEMMA 5.42: *Suppose that* **Pic** *is a sub-picture of* $\mathrm{Skel}(\psi)$. *Then either* $\psi(\mathbf{Pic}) \subset D^{-1}([0, L_2])$ *or* $\psi(\mathbf{Pic}) \subset D^{-1}([L_2, \infty))$.

Definition 5.43 (Thick-thin decomposition of preferred triangles): Suppose that $\psi\colon \partial\Delta$ *to*$X \cup \partial_{\mathcal{H}} X$ is a preferred triangle, and that **Pic** is a sub-picture of $\mathrm{Skel}(\psi)$.

If $\psi(\mathbf{Pic}) \subset D^{-1}([0, L_2])$, then we call **Pic** a **thick sub-picture**. If $\psi(\mathbf{Pic}) \subset D^{-1}([L_2, \infty))$ then we call **Pic** a **thin sub-picture**.

PROPOSITION 5.44: *If $\psi\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$ is a preferred triangle, and u and v are vertices in the same thick sub-picture of* Skel(ψ) *then any geodesic $\gamma(\psi(u), \psi(v))$ joining $\psi(u)$ to $\psi(v)$ does not penetrate any $(L_2 + L_1)$-horoball.*

Proof. Let **Pic** be a thick sub-picture of Skel(ψ). Then $\psi(\mathbf{Pic}) \subset D^{-1}([0, L_2])$. Furthermore, the image of **Pic** under ψ consists of subsegments of preferred paths $p_{a,b}, p_{a,c}$ and $p_{b,c}$, say, together with single edges from $D^{-1}(L_2)$, corresponding to the images of ribs in **Pic**.

If $p_{a,b}$ and $p_{a,b}$ both intersect $\psi(\mathbf{Pic})$ nontrivially, then they either intersect in $\psi(\mathbf{Pic})$ (in case their preimages are joined by a ligament) or else are joined by an edge of length 1 in $D^{-1}(L_2)$ (in case their preimages are joined by a rib).

Thus, any point in $\psi(\mathbf{Pic})$ lies within distance at most $\frac{1}{2}$ from $p_{a,b} \cap p_{a,c} \cap p_{b,c}$.

Suppose $p_{a,b} \cap \psi(\mathbf{Pic}) \neq \emptyset$. Let $m_{a,b}$ be the point in $p_{a,b} \cap \psi(\mathbf{Pic})$ closest to a, and let $n_{a,b}$ be the point closest to b. Define points $m_{a,c}, n_{a,c}, m_{b,c}$ and $n_{b,c}$ analogously. Not all of these points need be defined, if the intersection of some preferred path with $\psi(\mathbf{Pic})$ is empty. Also, it is certainly possible for some of these points to be the same, if they lie in the image of a ligament, or if they are the image of a corner of $\partial\Delta$.

Suppose that $m_{a,b} = p_{a,b}(t)$. Let $\hat{m}_{a,b} = \gamma(a, b)(f(t))$, where f is the function from Proposition 5.12.

Use the function f to define points $\hat{n}_{a,b}$, $\hat{m}_{a,c}$, etc.

Note that $d(m_{a,b}, \hat{m}_{a,b}) \leq K + 12\delta + 9$, and similarly for the other points and their hatted counterparts.

The points $m_{a,b}, n_{a,b}, \ldots$ are partitioned into some of the pairs defined earlier in this section, and, in particular, each pair consists of one or two points which are distance at most 1 apart.

Thus, the geodesics between the points $\hat{m}_{a,b}, \hat{n}_{a,b}, \ldots$, consisting of the subsegments of $\gamma(a, b), \gamma(a, c), \gamma(b, c)$, and paths joining the points whose un-hatted counterparts correspond to a 'pair'.

Therefore, depending on whether there are two or three preferred paths which intersect **Pic** nontrivially, the points $\hat{m}_{a,b}$, $\hat{n}_{a,b}$, etc. form either a hexagon or a quadrilateral. If it is a hexagon, then three of the sides have length at most $2(3K + 21\delta + 14) + 1$, and if it is a quadrilateral two of the sides of length at most $2(3K + 21\delta + 14) + 1$. Call the sides $[\hat{m}_{a,b}, \hat{n}_{a,b}]$ etc., the 'long' sides

of this polygon. Any point on a long side of this polygon is distance at most $2(3K + 21\delta + 14) + 1 + 3\delta$ from one of the other long sides.

Now consider any vertices $u, v \in \mathbf{Pic}$. As already noted, there are points u', v' so that u' and v' are contained in preferred paths and $d(u, u'), d(v, v') \leq 1/2$.

By Proposition 5.17 there are points u_1, v_1 on different preferred paths to u' and v', respectively so that $d(u', u_1), d(v', v_1) \leq \delta' + 1$. Since there are only three preferred paths, at least two of u', u_1, v', v_1 must lie on the same preferred path. Since u' and u_1 lie on different preferred paths, as do v' and v_1, one of the two points lies within δ' of u' and one within δ' of v'. Let this pair of points on the same preferred path be $\overline{u}$ and $\overline{v}$, so that $d(\overline{u}, u'), d(\overline{v}, v') \leq \delta'$. That part of the preferred path which lies between $\overline{u}$ and $\overline{v}$ lies in a thick sub-picture, and so does not penetrate any $(L_2 + 2)$-horoball. Let μ be this part of the preferred path.

By Proposition 5.12, there are points $\hat{u}, \hat{v}$ on a long side of the polygon described above, so that $d(\overline{u}, \hat{u}), d(\overline{u}, \hat{u}) \leq 3K + 21\delta + 14$. The geodesic between $\hat{u}$ and $\hat{v}$ (which lies on a long side of the polygon) lies at Hausdorff distance at most $3K + 21\delta + 14$ from μ, and therefore does not penetrate any $(L_2 + (3K + 21\delta + 14) + 2)$-horoball.

Now, consider the geodesic quadrilateral formed by $u, v, \hat{u}, \hat{v}$. We have

$$d(u, \hat{u}) \leq d(u, u') + d(u', \overline{u}) + d(\overline{u}, \hat{u}) \leq \delta' + (3K + 21\delta + 14) + 1\frac{1}{2},$$

and, similarly, $d(v, \hat{v}) \leq \delta' + (3K + 21\delta + 14) + 1\frac{1}{2}$.

Therefore, any point on the geodesic between u and v lies in at most $\delta' + (3K + 21\delta + 14) + 1\frac{1}{2} + 2\delta$ of some point on the geodesic between $\hat{u}$ and $\hat{v}$.

This implies that the geodesic between u and v does not penetrate any $(L_2 + 2(3K + 21\delta + 14) + 2\delta + 3\frac{1}{2})$-horoball. Since $L_2 + 2(3K + 21\delta + 14) + 2\delta + 3\frac{1}{2} < L_2 + L_1$, the proof is finished. ∎

Remark (about torsion) 5.45: In the context of the "quantum" preferred paths, when G is not torsion-free, not all choices of triples of preferred paths will have the controlled properties of preferred triangles in this section. In applications, it is possible to make consistent choices. We will say more about this in Part 2.

6. A homological bicombing

In this section we describe how the results from previous sections in this paper

may be combined with results of Mineyev from [29] to construct a quasi-geodesic homological bicombing of $X \cup \partial_{\mathcal{H}} X$ (in the sense of Definitions 2.35 and 2.36).

6.1. MINEYEV'S BICOMBING. In this subsection, we briefly recall a construction of Mineyev from [29]. We need a slightly more general statement than appears in [29], and we explain how Mineyev's proof implies Theorem 6.2 below.

Suppose that Γ is a locally finite graph which is δ-hyperbolic for some integer $\delta \geq 1$.

Remark 6.1: In his construction, Mineyev further assumes that Γ has bounded valence. This is important for the area bounds, but not for making the definitions.

Suppose that the group G acts freely on Γ. Let γ be a G-equivariant geodesic bicombing on Γ (see Definition 2.31). Let P be the homological bicombing induced by γ (see Remark 2.34).

For each vertex a in Γ, define

$$pr_a : \Gamma^{(0)} \to \Gamma^{(0)},$$

as follows:

- $pr_a(a) = a$; and
- if $b \neq a$, then $pr_a(b) = \gamma(a,b)(r)$, where r is the largest (integral) multiple of 10δ which is strictly less than $d(a,b)$.

For vertices a, b in Γ, the **flower at b with respect to a** is the set

$$Fl(a,b) = S(a, d(a,b)) \cap B(b, \delta) \subset \Gamma^{(0)}.$$

Now, for each pair of vertices $a, b \in \Gamma$ define a 0-chain $f(a,b)$ in Γ inductively on the distance $d(a,b)$ as follows:

- if $d(a,b) \leq 10\delta$ then $f(a,b) = b$;
- if $d(a,b) > 10\delta$ and $d(a,b)$ is not an integral multiple of 10δ then $f(a,b) = f(a, pr_a(b))$; and
- if $d(a,b) > 10\delta$ and $d(a,b)$ is an integral multiple of 10δ then

$$f(a,b) = \frac{1}{\#Fl(a,b)} \sum_{x \in Fl(a,b)} f(a, pr_a(x)).$$

For each vertex $a \in \Gamma^{(0)}$, define a 0-chain, $star(a)$ by

$$star(a) = \frac{1}{\#B(a, 7\delta)} \sum_{x \in B(a,7\delta)} x.$$

By linearity, $star(a)$ is defined for any 0-chain a.

Now define, for $a, b \in \Gamma^{(0)}$,

$$\overline{f}(a, b) = star(f(a, b)).$$

We now define a homological bicombing Q' on Γ. First note that, by linearity, $P_{a,b}$ makes sense when a is any 0-chain. The 1-chain $Q'_{a,b}$ is defined inductively on $d(a, b)$, as follows: if $d(a, b) \leq 10\delta$, then $Q'_{a,b} = P_{a,b}$. Suppose now that $d(a, b) > 10\delta$. By [29, Proposition 7(2)]

$$\operatorname{supp}(\overline{f}(a, b)) \subseteq B(P_{a,b}(10\delta), 8\delta).$$

Note that Mineyev's proof of this does not use the bounded valence assumption. Therefore, for each $x \in \operatorname{supp}(\overline{f}(b, a))$ we have $d(a, x) < d(a, b)$, so $Q'_{a,x}$ is defined by induction. Define $Q'_{a,\overline{f}(b,a)}$ by linearity over the second variable and define

$$Q'_{a,b} = Q'_{a,\overline{f}(b,a)} + P_{\overline{f}(b,a),b}.$$

Finally, we define

$$Q_{a,b} = \frac{1}{2}(Q'_{a,b} - Q'_{b,a}),$$

so that Q is anti-symmetric.

Mineyev proves that when Γ has bounded valence, the bicombing Q has **bounded area**, in the sense of Theorem 6.2 below. When Γ does not have bounded valence, these area bounds break down completely. However, we are only going to use the bicombing from [29] on a subset of X of uniformly bounded depth, and on such a subspace the valence is uniformly bounded, and Mineyev's techniques apply.

The proof of the area bound for a given triangle occurs entirely in the 60δ-neighborhood of the union of the three geodesic sides. Moreover, if we have a bound on the valence in some part of a graph, then we can calculate a bound on the number of vertices in any given ball which lies entirely within the chosen part. The proof from [29] now applies directly to prove the following theorem.

THEOREM 6.2 (Mineyev): *There is a function $T = T(\delta, v)$ so that: For any finite valence δ-hyperbolic graph Γ, and group G acting freely on Γ, there is a*

(18δ)-quasi-geodesic G-equivariant antisymmetric homological bicombing Q so that

$$|Q_{ab} + Q_{bc} + Q_{ca}|_1 \leq T(\delta, v)$$

whenever a, b, c are vertices of Γ spanning a geodesic triangle $\mathcal{T}_{abc}$ so that every vertex in the 60δ-neighborhood of $\mathcal{T}_{abc}$ has valence less than v.

Theorem 6.2 is a key ingredient in the construction in this section. The other key ingredients are the construction of preferred paths and the analysis of preferred triangles from Section 5

6.2. The bicombing q. In the remainder of this section we define our homological bicombing q of $(X \cup \partial_{\mathcal{H}} X) \times (X \cup \partial_{\mathcal{H}} X)$, which uses preferred paths and Mineyev's bicombing Q from Theorem 6.2. (See Definition 2.35 for the definition of homological bicombing which allows some points to be ideal.)

Suppose that $a, b \in X \cup \partial_{\mathcal{H}} X$. Let $p_{a,b}$ be the preferred path between a and b.

Decompose $p_{a,b}$ into subintervals, oriented consistently with $p_{a,b}$, which lie either entirely within $D^{-1}([0, L_2])$ or entirely within $D^{-1}([L_2, \infty))$, and so that the endpoints of these subintervals lie in $\{a, b\} \cup D^{-1}(L_2)$.

By the way that preferred paths were defined, there is a unique way of performing this decomposition.

Suppose that μ is a subinterval in the decomposition of $p_{a,b}$ so that μ lies in $D^{-1}([0, L_2])$, and let x and y be the endpoints of μ.

Lemma 6.3: *The geodesic between x and y does not intersect any*

$$(L_2 + 2(3K + 21\delta + 14) + 2\delta)\text{-horoball.}$$

Proof. By Proposition 5.12, there exist $x', y' \in \gamma(a, b)$ so that $d(x, x'), d(y, y') \leq 3K + 21\delta + 14$, and so that the Hausdorff distance between $p_{a,b}[x, y]$ and $\gamma(a, b)[x', y']$ is at most $3K + 21\delta + 14$. Thus, $\gamma(a, b)[x', y']$ does not penetrate any $(L_2 + (3K + 21\delta + 14))$-horoball.

The geodesic quadrilateral with vertices x, x', y, y' has two sides of length at most $3K + 21\delta + 14$. Therefore, the geodesic between x and y lies in a $(3K+21\delta+14)+2\delta$ neighborhood of $\gamma(a, b)[x', y']$, and hence does not penetrate any $(L_2 + 2(3K + 21\delta + 14) + 2\delta)$-horoball, as required. ∎

Definition 6.4 (Definition of q): Suppose that $a, b \in X \cup \partial_{\mathcal{H}} X$ are distinct. Let $p_{a,b}$ be the preferred path between a and b, and let $\underline{p_{a,b}}$ be the induced 1-chain. The decomposition of $p_{a,b}$ described above induces a decomposition of $\underline{p_{a,b}}$. Let μ be an element of the decomposition of $\underline{p_{a,b}}$ for which $\operatorname{supp}(\mu) \subset D^{-1}[0, L_2]$. Let $\partial\mu = \mu_+ - \mu_-$.

Taking the sum over all such μ we define

$$q_{a,b} = \underline{p_{a,b}} - \sum_{\mu} \big(Q(\mu_-, \mu_+) - \mu\big).$$

Because $\partial\mu = \partial Q(\mu_-, \mu_+) = \mu_+ - \mu_-$, we have $\partial \underline{p_{a,b}} = \partial q_{a,b}$. We claim that q is a homological bicombing on $X \cup \partial_{\mathcal{H}} X$ in the sense of Definition 2.35 and that furthermore it is $6000\delta^2$-quasi-geodesic in the sense of Definition 2.36. It also has nice properties analogous to those in Theorem 6.2 above. See Theorem 6.10 below for the precise statements about the bicombing q.

PROPOSITION 6.5: *For any $a, b \in X \cup \partial_{\mathcal{H}} X$, $\operatorname{supp}(q_{a,b})$ lies in a $\big(K + 25\delta + 9\big)$-neighborhood of any geodesic between a and b.*

Proof. By Corollary 5.13, $p_{a,b}$ lies in a $(K + 7\delta + 9)$-neighborhood of $\gamma(a, b)$.

By construction and Theorem 6.2, $\operatorname{supp}(q_{a,b})$ lies in an 18δ-neighborhood of $p_{a,b}$. ∎

Since $K + 25\delta + 9 < 6000\delta^2$, Proposition 6.5 proves the first of the two required statements for q to be $6000\delta^2$-quasi-geodesic. We now prove the remaining requirement.

PROPOSITION 6.6: *q is a $6000\delta^2$-quasi-geodesic homological bicombing on $X \cup \partial_{\mathcal{H}} X$.*

Proof. By Proposition 6.5, it remains to prove statement (2) of Definition 2.36. Let $a, b \in X^{(0)}$ be distinct. By Theorem 6.2 and the definition of q, we have

$$|q(a, b)|_1 \leq 18\delta |p_{a,b}|_1.$$

By Corollary 5.14, the length of $p_{a,b}$ is at most $2d(a, b) + 20K + 120\delta + 72 < 325\delta d(a, b)$. ∎

Remark 6.7: It is also possible to prove a suitable refinement of statement (2) as alluded to in Remark 2.37.

6.3. Bounded thick area. The main result of Section 6 is Theorem 6.10.

Definition 6.8: Suppose $a, b, c \in X^{(0)} \cup \partial_{\mathcal{H}} X$. We define a 1-cycle c_{abc} as follows:

Let $\phi\colon \Delta \to X \cup \partial_{\mathcal{H}} X$ be the preferred triangle associated to the triple (a, b, c). Associated to ϕ is the graph Skel(ϕ) (see Definition 5.27), which has associated thick sub-pictures (see Definition 5.43). Let **Pic** be a thick sub-picture of Skel(ϕ). The vertices of **Pic** inherit a circular order $(v_1, \ldots, v_n)$ from the order (a, b, c).

Define

$$c_{\mathbf{Pic}} = \sum_{i \in \mathbb{Z}/m} Q(\ddot{\phi}(v_i), \ddot{\phi}(v_{i+1})).$$

Finally, define

$$c_{abc} = \sum_{\mathbf{Pic}} c_{\mathbf{Pic}},$$

where the sum is over all thick sub-pictures of Skel(ϕ).

Observation 6.9: If **Pic** is a thick sub-picture lying in a leg of Skel(ϕ), then $c_{\mathbf{Pic}} = 0$.

The following is a key theorem for our proof of Theorem 11.11, one of the major steps in proving Theorem 7.2:

Theorem 6.10: *There exists a constant T_1, depending only on X, so that for all $a, b, c \in X \cup \partial_{\mathcal{H}} X$ there is a 1-cycle c_{abc}, as described in Definition 6.8 above then*

$$|c_{abc}|_1 \leq T_1.$$

Also,

$$supp\big(q(a, b) + q(b, c) + q(c, a) - c_{abc}\big) \subset D^{-1}[L_2, \infty).$$

Proof. We have already noted that the thick sub-pictures which lie inside a leg do not contribute anything to c_{abc}. Thus we may concentrate on thick sub-pictures lying in the feet or middle of Skel(ϕ). By Corollary 5.40, there are at most 15 vertices in total in all such sub-pictures. Therefore, we can triangulate all of these sub-pictures with at most 13 triangles, whose vertices all appear as a vertex in one of the thick sub-pictures.

Let u and v be such vertices. By Proposition 5.44, the geodesic between u and v does not intersect any $(L_2 + L_1)$-horoball.

Let v be the maximum valence of any vertex in $D^{-1}[0, L_2 + L_1 + 18\delta]$, and let $T(\delta, v)$ be as in Theorem 6.2. Define $T_1 = 13T(\delta, v)$.

We can express c_{abc} as the sum of at most 13 1-chains of the form $Q_{u_1,u_2} + Q_{u_2,u_3} + Q_{u_3,u_1}$, where u_1, u_2 and u_3 are the images under $\ddot{\phi}$ of vertices in thick sub-pictures of Skel(ϕ) which lie in the feet or middle.

The result now follows from Theorem 6.2. ∎

The following is immediate from Theorem 6.10 and Theorem 3.25.

COROLLARY 6.11: *For all $a, b, c \in X \cup \partial_{\mathcal{H}} X$ there exists a 2-chain ω_{abc} so that*

(1) $\partial\omega_{abc} = c_{abc}$; *and*
(2) $|\omega_{abc}|_1 \leq M_X T_1$,

where M_X is the constant for the linear homological isoperimetric function for X.

Remark (about torsion) 6.12: In the presence of torsion, the bicombing q, and the 2-chains ω_{abc}, can be defined without much difficulty using "averaged" preferred paths, and the ideas already contained in this section.

Part 2. Dehn filling in relatively hyperbolic groups

7. Dehn filling in groups

In Part 2 of this paper, we provide an application of the constructions from Part 1.

Definition 7.1: Let G be a group, and P a subgroup. Suppose that G is generated by S and P is generated by $P \cap S$. If $K \trianglelefteq P$ is a normal subgroup of P, then we define

$$|K|_P = \inf_{k \in K \smallsetminus \{1\}} |k|_{P \cap S},$$

where $|k|_{P\cap S}$ is the distance from k to the identity in the Cayley graph $\Gamma(P, P \cap S)$. By convention $|\{1\}|_P = \infty$.

In the special case that P is free abelian and K the cyclic group generated by $\kappa \in K$, $|K|_P$ is just the length of κ in P.

The main result of Part 2 is the following theorem.

THEOREM 7.2: *Let G be a torsion-free group, which is hyperbolic relative to a collection $\mathcal{P} = \{P_1, \dots, P_n\}$ of finitely generated subgroups. Suppose that S is a generating set for G so that for each $1 \leq i \leq n$ we have $P_i = \langle P_i \cap S \rangle$.*

There exists a constant B depending only on $(G, \mathcal{P})$ so that for any collection $\{K_i\}_{i=1}^n$ of subgroups satisfying

- *$K_i \trianglelefteq P_i$; and*
- *$|K_i|_{P_i} \geq B$,*

then the following hold, where K is the normal closure in G of $K_1 \cup \cdots \cup K_n$.

(1) *The map $P_i/K_i \xrightarrow{\iota_i} G/K$ given by $pK_i \mapsto pK$ is injective for each i.*
(2) *G/K is hyperbolic relative to the collection $\mathcal{Q} = \{\iota_i(P_i/K_i) \colon 1 \leq i \leq n\}$.*

In fact, much more than this is true. For example, for $i \neq j$ we have

$$\iota_i(P_i/K_i) \cap \iota_j(P_j/K_j) = \{1\},$$

(see Corollary 9.4) and each $\iota_i(P_i/K_i)$ is malnormal in G/K (see Corollary 9.5). Also, if G is non-elementarily hyperbolic relative to $\{P_i\}$, then G/K is non-elementarily hyperbolic relative to $\{\iota_i(P_i/K_i)\}$ (see Theorem 11.12).

The remainder of the paper is devoted to the proof of Theorem 7.2, and the subsidiary assertions mentioned above.

Theorem 7.2 clearly holds if $n = 1$ and $G = P_1$, so we henceforth assume (without mention) that this is not the case.

8. Equations involving parabolics and skeletal fillings of surfaces

In this section we suppose that G is hyperbolic relative to $\mathcal{P} = \{P_1, \dots, P_n\}$, and that $X = X(G, \mathcal{P}, S, \mathcal{R})$ is the cusped space associated to some compatible set of generators S, and some collection of relators $\mathcal{R}$, as described in Section 3. Finally $H \trianglelefteq G$ is an arbitrary normal subgroup of G.

In order to use the geometry of X to study the quotients of G, we will need to turn equations in G/H or (represented as maps of compact planar surfaces into X/H) into "pleated surfaces" in X/H. Exactly what this means depends on the context, and will become clear as we proceed.

8.1. LIFTING AND STRAIGHTENING. Let $\Gamma_H = \Gamma(G, S)/H \subset X/H$. Notice that Γ_H is a Cayley graph for G/H. We first show how to extend maps of

compact surfaces with boundary components in Γ_H to proper maps of non-compact surfaces. We then say what we mean by "lifting and straightening" such proper maps.

Let Σ be a compact planar surface, and let

$$\phi\colon \Sigma \to X/H$$

be a cellular map so that $\phi|_{\partial\Sigma} \subset \Gamma_H$.

Definition 8.1: If $P \in \mathcal{P}$, let $\partial_P\Sigma$ be the union of those boundary components c so that $\phi|_c$ lifts to an arc in X which lies in a single left coset of P. Let

$$\partial_{\mathcal{P}}\Sigma = \bigcup_{P\in\mathcal{P}} \partial_P\Sigma.$$

be the union of those boundary components of Σ, each of which is sent into the image of a single 0-horoball of X (see Definition 3.21). We refer to $\partial_{\mathcal{P}}$ as the **parabolic boundary** of Σ.

Definition 8.2: Let

$$\check{\Sigma} = \Sigma \cup_{\partial_{\mathcal{P}}\Sigma} \partial_{\mathcal{P}}\Sigma \times [0,\infty)$$

be the surface obtained from Σ by attaching a half-open annulus to each component of $\partial_{\mathcal{P}}\Sigma$. We extend ϕ to a proper map

$$\check{\phi}\colon \check{\Sigma} \to X/H$$

as follows: Let $c \cong S^1$ be a component of $\partial_{\mathcal{P}}\Sigma$. The map ϕ sends $c \times \{0\}$ to a loop γ_0 which is contained in $\Gamma(P, P \cup S)/P$ for some $P \in \mathcal{P}$. This loop lifts to some path $\widetilde{\gamma_0}$ in X, each edge of which is the top edge of some vertical square. Let γ_1 be the loop which is the projection of the path obtained by traversing the bottom edges of those squares, and define $\check{\phi}|_{c\times[0,1]}$ to be a homotopy across the images of those squares in X/G so that $\check{\phi}|_{c\times\{1\}} = \gamma_1$. Similarly, define $\check{\phi}|_{c\times[k,k+1]}$ for each $k \geq 1$ so that $D(\check{\phi}(c \times \{t\})) = t$ for each $t \geq 0$.

A **peripheral path** in X/H is a path in the 1-skeleton which lies in the image of the Cayley graph of some $P_i \in \mathcal{P}$.

Definition 8.3: Let Σ be a compact planar surface, and let $\phi\colon \Sigma \to X/H$ be as above. A **reducing arc** for ϕ is an essential, properly embedded interval $\sigma\colon I \to \Sigma$ so that $\sigma(\partial I) \subset \partial_P\Sigma$ for some $P \in \mathcal{P}$ and $\phi \circ \sigma$ is homotopic rel endpoints to a peripheral path.

Definition 8.4: Analogously we may define reducing arcs for proper maps into X/G: Suppose that Ξ is obtained from a compact planar surface by removing finitely many points. Let $f : \Xi \to X/H$ be a proper map, and suppose $\sigma : \mathbb{R} \to \Xi$ is an essential proper arc. We say that σ is a reducing arc if $f \circ \sigma$ is properly homotopic into $D^{-1}[L, \infty)$ for some (and hence for any) $L > 0$.

LEMMA 8.5: *The map $\phi\colon \Sigma \to X/H$ has a reducing arc if and only if $\check{\phi}\colon \check{\Sigma} \to X/H$ has a reducing arc.*

LEMMA 8.6: *Let Ξ be a surface of finite type (possibly with boundary), and let $\mathcal{T}$ be a triangulation of Ξ (partially ideal if appropriate). Let $\psi\colon \Xi \to X/H$ be a proper map satisfying:* (i) *ψ has no reducing arcs; and* (ii) *if e is an edge of $\mathcal{T}$ which limits on a puncture p of Ξ then there is a neighborhood U of p so that $\psi|_{e\cap U}$ is vertical.*

Then there is a proper homotopy from ψ to a map

$$\psi_{\mathcal{T}}\colon \Xi \to X/H$$

so that if e is any edge of $\mathcal{T}$, then $\psi_{\mathcal{T}}|_e$ lifts to a preferred path in X.

Proof. Let e be an edge of $\mathcal{T}$. It suffices to show that $\psi|_e$ is homotopic to a preferred path, and that this homotopy is level-preserving near the ends of e.

Choose a lift $\widetilde{\psi|_e}$ of $\psi|_e$ to X. The map $\widetilde{\psi|_e}$ extends to a map from I to $X \cup \partial_{\mathcal{H}} X$. Let a be the image of 0 and b the image of 1. Since e is not a reducing arc, $a \neq b$ and there is a preferred path $p_{a,b}$ between these points. Consider $p_{a,b}$ as a map from e to X.

Suppose that a is ideal. Then both $\widetilde{\psi|_e}$ and $p_{a,b}$ are vertical on some initial segment, so we may reparametrize so that $D \circ \widetilde{\psi|_e}$ and $D \circ p_{a,b}$ agree on this initial segment. Therefore, there is an obvious horizontal homotopy from $\widetilde{\psi|_e}$ to $p_{a,b}$ on this initial segment. We consider the projection of this homotopy to X/H.

In case b is ideal, we may similarly perform a homotopy on a terminal subsegment of e.

We are now left to deal with a compact loop, formed by the paths $\widetilde{\psi|_e}$ and $p_{a,b}$ (or the sub-paths with which we have not yet dealt). The space X is simply-connected. ∎

Remark 8.7: Suppose that $\phi\colon \Sigma \to X/H$ is as at the beginning of this subsection. The surface $\check{\Sigma}$ and the map $\check{\phi}\colon \check{\Sigma} \to X/H$ satisfy condition (ii) of the

hypothesis of Lemma 8.6.

Remark (about torsion) 8.8: In the presence of torsion it is not, in general, possible to find a map $\psi_{\mathcal{T}}$ as in Lemma 8.6 so that each triangle is mapped to something which is combinatorially controlled the way that preferred triangles are. Having no reducing arcs is not sufficient; an additional hypothesis is required.

8.2. The skeleton of a map. In this subsection, we define the **skeleton** of the map of a surface into X/H, assuming the surface has been triangulated by edges which are sent to preferred paths.

Definition 8.9: Suppose that Ξ is a surface with a (possibly partially ideal) triangulation $\mathcal{T}$, and that $\theta\colon \Xi \to X/H$ sends each edge of $\mathcal{T}$ to a non-degenerate path which lifts to a preferred path between points in $X \cup \partial_{\mathcal{H}} X$. Further assume that θ sends no vertex of $\mathcal{T}$ to a point in an L_2-horoball. Let $\overline{\Xi}$ be the compact surface obtained from Ξ by filling in the punctures; $\mathcal{T}$ induces a triangulation $\overline{\mathcal{T}}$ of $\overline{\Xi}$. Then θ extends to a map

$$\overline{\theta}\colon \overline{\Xi} \to X/H \cup (\partial_{\mathcal{H}} X)/H.$$

If Δ is a triangle of $\overline{\mathcal{T}}$, then $\overline{\theta}|_{\partial\Delta}$ lifts to a preferred triangle in X (Definition 5.19). The skeletal filling $\mathrm{Skel}\theta|_{\partial\Delta}$ can then be inscribed on $\overline{\Xi}$ (Lemma 5.28). The **skeleton of** θ, $\mathrm{Skel}\theta$, is the 1-complex in $\overline{\Xi}$ which is the union of the $\mathrm{Skel}\theta|_{\partial\Delta}$ for $\Delta \in \overline{\mathcal{T}}$.

Remark 8.10: A typical application of the above definition and the lemmas below is the situation where $\phi\colon \Sigma \to X/H$ is as at the beginning of the section (with no reducing arcs). In this situation, we will take Ξ to be $\check{\Sigma}$ and θ to be $\check{\phi}_{\mathcal{T}}$.

Remark 8.11: Let $\theta\colon \Xi \to X/H$ and $\mathcal{T}$ be as described in Definition 8.9, and write $\Xi^{(1)}$ for the union of the edges of $\mathcal{T}$. By Remark 5.29, we can extend $\theta|_{\Xi^{(1)}}$ to a map

$$\ddot{\theta}\colon \mathrm{Skel}\theta \to X/H \cup (\partial_{\mathcal{H}} X)/H$$

which collapses each ligament to a point, and sends each rib to a horizontal edge. Observe that $D(\ddot{\theta}(x)) = L_2$ for any vertex v of $\mathrm{Skel}\theta$ not coming from a vertex of $\overline{\mathcal{T}}$.

Definition 8.12: Let $\theta\colon \Xi \to X/H$ and $\mathcal{T}$ be as described in Definition 8.9, and let $\ddot{\theta}\colon \mathrm{Skel}\theta \to X/H \cup (\partial_{\mathcal{H}} X)/H$ be as in Remark 8.11. Let $\overline{D}\colon \mathrm{Skel}\theta \to [0,\infty]$ be as follows:

$$\overline{D}(x) = \begin{cases} \infty & \text{if } x \text{ is a vertex coming from a puncture.} \\ D(\ddot{\theta}(x)) & \text{otherwise.} \end{cases}$$

Let v be a vertex of $\mathrm{Skel}\theta$ coming from a puncture of Ξ. Let $E(v) \subset \overline{\Xi}$ be the smallest closed disk containing the component of $\overline{D}^{-1}[L_2,\infty]$ which also contains v. The **link of** v, $\mathrm{Lk}(v)$, is the boundary of $E(v)$.

Remark 8.13: Note that $\mathrm{Lk}(v)$ is contained in the skeleton of θ. If Ξ is a punctured sphere, then $\mathrm{Lk}(v)$ is a circle made up entirely of ribs and ligaments; otherwise it may contain parts of edges of $\mathcal{T}$. Also notice that $E(v)$ contains no vertex of $\overline{\mathcal{T}}$ other than v. In Figure 30, one can see an example showing both kinds of links.

The next two lemmas follow easily from the fact that θ restricted to a triangle of $\overline{\mathcal{T}}$, lifts to a preferred triangle in X, and that both X and L_2-horoballs in X are simply connected.

LEMMA 8.14: *Let $\theta\colon \Xi \to X/H$ and $\mathcal{T}$ be as described in Definition 8.9, and let $\ddot{\theta}\colon \mathrm{Skel}\theta \to X/H \cup (\partial_{\mathcal{H}} X)/H$ be as in Remark 8.11. Let $\iota\colon \mathrm{Skel}\theta \to \overline{\Xi}$ be an inclusion which is the identity on edges in $\overline{\mathcal{T}}$. The map θ is properly homotopic to a map θ' so that $\theta' \circ \iota = \ddot{\theta}$.*

LEMMA 8.15: *Let $\theta\colon \Xi \to X/H$ and $\mathcal{T}$ be as described in Definition 8.9, and let $\ddot{\theta}\colon \mathrm{Skel}\theta \to X/H \cup (\partial_{\mathcal{H}} X)/H$ be as in Remark 8.11. Let $\sigma\colon S^1 \to \Xi$ be a loop surrounding a puncture x, so that $\theta \circ \sigma(S^1)$ lies entirely in the component of $(\theta \circ D)^{-1}[L_2,\infty)$ surrounding x. Let $\gamma\colon S^1 \to \mathrm{Lk}(x)$ be a homeomorphism. Then $\theta \circ \sigma$ is homotopic to $\ddot{\theta} \circ \gamma$ inside $D^{-1}[L_2,\infty)$.*

9. Punctured spheres and disks

In this section we investigate relations in G/K amongst the images of the parabolic elements of G. In particular, we prove assertion (1) of Theorem 7.2. We also prove Theorem 9.3.

For this section, we make the following standing assumptions:

(1) G is a finitely generated, torsion-free group, which is hyperbolic relative to a collection $\mathcal{P} = \{P_1, \ldots, P_n\}$ of finitely generated subgroups;
(2) S is a finite compatible generating set for G with respect to $\mathcal{P}$ (in the sense of Definition 2.15);
(3) $X(G, \mathcal{P}, S)$ is δ-hyperbolic.

9.1. From parabolic equations to surfaces and punctured spheres. In this paragraph, we explain how any equation amongst parabolic words (and their conjugates) may be turned into a map of a compact planar surface with boundary into X/G; we can then extend this to a proper map of a punctured sphere into X/G using Definition 8.2.

Suppose that in G there is an equality of the form:

$$1 = \prod_{i=1}^{m} g_i p_i g_i^{-1}, \tag{1}$$

where, for each i, $p_i \in P_{j_i}$ and $g_i \in G$. Choosing words for each g_i and words in $S \cap P_{j_i}$ for each p_i, we find a map $\tilde{\phi}\colon \tilde{\Sigma} \to X$ of a disk, sending the boundary to a loop representing the equation.

Projection gives a map $\pi \circ \tilde{\phi}\colon \tilde{\Sigma} \to X/G$.

Let Σ be the surface obtained from the disk $\tilde{\Sigma}$ by identifying those parts of the boundary corresponding to the g_j in pairs (Figure 25).

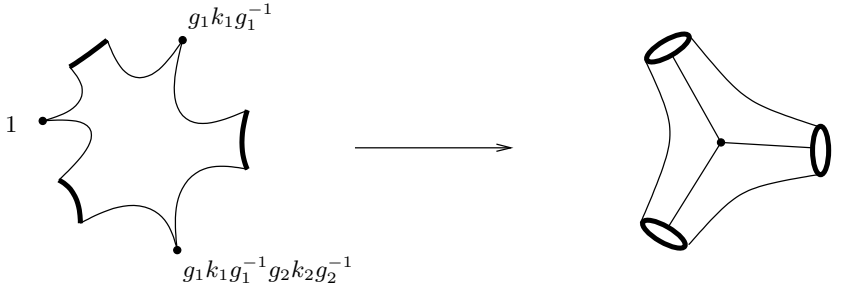

Figure 25. Identifying those parts of the disk $\tilde{\Sigma}$ labelled by the conjugating elements yields a compact planar surface which maps into X/G.

The map $\pi \circ \tilde{\phi}$ factors through the quotient map from $\tilde{\Sigma}$ to Σ. Let $\phi\colon \Sigma \to X/G$ be the resulting map.

The m boundary components of Σ are sent by ϕ to peripheral loops in X/G whose labels are the words chosen for the p_i above.

Conversely, given a compact planar surface with all but (possibly) one boundary component labelled by elements of the K_i, we can reconstruct an equality like (1) for the word representing the other boundary component, by choosing the g_j via some paths through the surface.

9.2. The groups P_i/K_i inject into G/K.

THEOREM 9.1: *There is a constant $R = R(\delta) \leq 12 \cdot 2^{3000\delta}$ so that the following holds: If $\{K_1, \ldots, K_n\}$ are so that $K_i \trianglelefteq P_i$ and $|K_i|_{P_i} > R$ then the natural map $\iota_i \colon P_i/K_i \to G/K$ is an injection, where K is the normal closure in G of $K_1 \cup \cdots \cup K_n$.*

Proof. (cf. [26, Proof of Theorem 3.1])

We will show that if, for some i, the map ι_i is not injective, then $|K_l|_{P_l}$ must be small for some l.

Let $\alpha \in P_i \setminus K_i$ be an element of K. Thus there is some equation

$$\alpha = \prod_{j=1}^{m} g_j k_j g_j^{-1} \tag{2}$$

for some finite sequence of $g_j \in G$ and $k_j \in K_{i_j}$. We say that such an equation **represents the death of α in G/K.** We may suppose that the product in (2) is minimal in the following sense: If α' is any element of $\bigcup_i P_i \setminus \bigcup_i K_i$, and

$$\alpha' = \prod_{r=1}^{m'} g_r k_r g_r^{-1}$$

for some collection of $g_r \in G$ and $k_r \in K_{i_r}$, then $m' \geq m$. (In other words, the expression in (2) is minimal in length not only for α, but over all such equations with the left hand side an element of $\bigcup_i P_i \setminus \bigcup_i K_i$.)

The equality (2) is realized by a map $\phi \colon \Sigma \to X/G$, as in Subsection 9.1.

CLAIM 9.2: The map ϕ has no reducing arcs.

Proof. Let σ be a reducing arc. Each of the boundary components of Σ has a corresponding word, and thus a corresponding base point (the starting point of this word). We may assume that σ starts and finishes at one of these base points. The path σ determines some element p of G. Because σ is homotopic to

a peripheral path, p is contained in some P_k. We suppose ϕ has been homotoped so that $\phi \circ \sigma$ is a peripheral path.

There are five cases to consider, depending on the endpoints of σ.

CASE 1: Suppose that the endpoints of σ are on different boundary components, neither of which corresponds to α. Cutting along σ yields a new surface Σ' with fewer boundary components than Σ (see Figure 26); two have been removed,

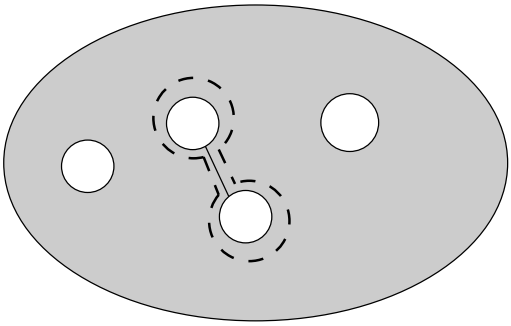

Figure 26. Reducing the number of boundary components labelled with elements of K_i.

and the new one is sent to a word representing $k_i p k_j p^{-1}$ for some i, j. This contradicts the minimality assumption.

CASE 2: Suppose next that the initial point of σ lies on the boundary component corresponding to α, and that the terminal point lies on the boundary component corresponding to k_i for some i. Cutting along σ again yields a surface Σ' with one fewer boundary component than Σ. The boundary components corresponding to α and k_i have been removed. The new boundary component has label $\alpha' = \alpha p k_i p^{-1}$. Since $K_i \trianglelefteq P_i$, and $\alpha \notin K_i$, we also have $\alpha' \notin K_i$. This again contradicts the minimality of equation (2).

CASE 3: Suppose that both endpoints of σ lie on the boundary component corresponding to α. Cutting Σ along σ yields two compact planar surfaces, each with fewer boundary components than Σ. One of them represents the death of p in G/K, whilst the other represents the death of $p\alpha$.

In case $p \notin K_i$, the equation for p contradicts minimality. Thus we may suppose that $p \in K_i$, in which case $p\alpha \notin K_i$, and again we get a contradiction to minimality.

Case 4: Finally, suppose that σ is a loop whose endpoints are on the boundary component corresponding to k_i (as in Figure 27). Cutting along σ yields a pair

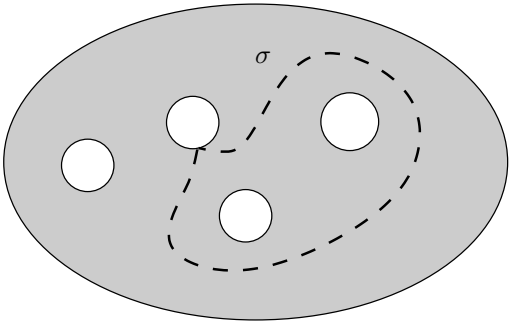

Figure 27. Case 4 of Claim 9.2.

of surfaces, both of which have fewer boundary components than Σ. If $p \notin K_{j_i}$ then the surface Σ' not containing α corresponds to an equation representing the death of p or pk_i (depending on whether k_i is contained in Σ'). On the other hand, if $p \in K_{j_i}$, then we can replace the surface not containing α by a single puncture, to yield a surface Σ'' with fewer punctures than Σ, representing the death of α.

In either case, this contradicts the minimality of m.

This proves Claim 9.2 ∎

As in Section 8.1, let $\check{\Sigma} = \Sigma \cup \partial\Sigma \times [0, \infty)$, and let

$$\check{\phi}\colon \check{\Sigma} \to X/K$$

be the extension from Definition 8.2.

We now choose an ideal triangulation $\mathcal{T}$ of the punctured sphere $\check{\Sigma}$. As $\check{\Sigma}$ is an $m+1$-times punctured sphere, $\mathcal{T}$ contains $2m-2$ triangles. Let

$$\check{\phi}_{\mathcal{T}}\colon \check{\Sigma} \to X/K$$

be the map from Lemma 8.6, which sends each edge of $\mathcal{T}$ to a preferred path, and let

$$\ddot{\phi}_{\mathcal{T}}\colon \operatorname{Skel}(\check{\phi}_{\mathcal{T}}) \to X/G \cup (\partial_{\mathcal{H}} X)/G,$$

be as in Remark 8.11.

If $T \in \mathcal{T}$, then $\phi|_{\partial T}$ lifts to a preferred triangle $\widetilde{\phi_T}\colon \partial T \to X$. Let $R(T)$ be the number of ribs in $\operatorname{Skel}(\widetilde{\phi_T})$, and note that this number does not depend on the lift chosen. Corollary 5.39 implies that $R(T) \leq 6$.

Let

$$A(\phi) = \sum_{T \in \mathcal{T}} R(T).$$

Corollary 5.39 immediately implies

$$A(\phi) \leq 6(2m-2) < 12m. \tag{3}$$

Let x be one of the punctures of $\check{\Sigma}$ not corresponding to α. By Lemma 8.15, $\ddot{\phi}|_{Lk(x)}$ is a loop at the L_2-level which, considering the L_2-level to be the image of a Cayley graph for P_i, represents an element k of K_i.

It follows that $Lk(x)$ must contain at least $2^{-L_2}|k|_{P_i}$ ribs.

Thus

$$A(\phi) \geq 2^{-L_2}\Big(\min_s |K_{i_s}|_{P_{i_s}}\Big)m.$$

Therefore, by (3),

$$\min_s |K_{i_s}|_{P_{i_s}} \leq 2^{L_2}\frac{A(\phi)}{m} < 12 \cdot 2^{L_2}. \qquad \blacksquare$$

9.3. On the structure of the quotients G/K. The theme in this subsection is that, by choosing large enough algebraic slope lengths, we can preserve much of the structure of G in its quotients. See [20] for more results along these lines.

Theorem 9.3: *Suppose G is torsion-free and that $|K_i|_{P_i} > 12 \cdot 2^{L_2}$ for each i. Let P and P' be conjugate into $\mathcal{P}$ and suppose that whenever $P^g = P'$ then $g \notin K$. Then the images of P and P' in G/K intersect trivially.*

Proof. Suppose that P is conjugate to $P_j \in \mathcal{P}$, and let K_P be the conjugate of K_j in P. Define $K_{P'}$ similarly. If the theorem is false, then there is an equation in G of the form

$$q = q' \prod_{i=1}^{m} g_i k_i g_i^{-1}, \tag{4}$$

where $q \in P \smallsetminus K_P$ and $q' \in P' \smallsetminus K_{P'}$. Suppose we have chosen such an equation with m minimal over all such equations (over all choices of P and P').

This gives rise to a map ϕ of a compact planar surface with $m+2$ boundary components into X/G. Once again, we claim that this map has no reducing arcs.

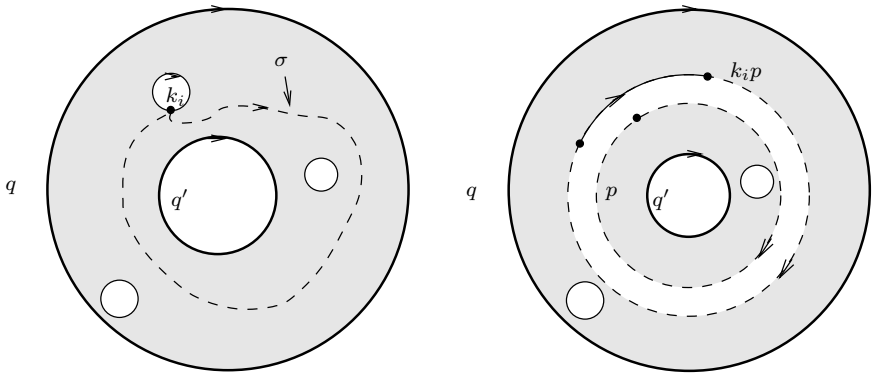

Figure 28. This kind of reducing arc gives rise to at least one new equation of type (4), if $p \notin K_{j_i}$.

Let σ be a reducing arc which is a peripheral arc, and let p be the corresponding element of G, contained in K_i, say. Most of the cases are entirely analogous to those in the proof of Theorem 9.1; they all lead to a contradiction to minimality, or to injectivity. We deal with the most interesting case, when σ is a loop whose endpoints lie on a puncture corresponding to some k_i, and so that cutting along σ separates the puncture corresponding to q from the puncture corresponding to q' (see Figure 28).

By Theorem 9.1, q and q' are nontrivial in G/K. Therefore, $p \notin K_i$. Now, the hypothesis of the theorem implies that it cannot be that there are $k, k' \in K$ so that $P = P_i^k$ and $P' = P_i^k$. Therefore, one of the two diagrams obtained by cutting along σ yields a contradiction to the minimality of m. This, and the omitted cases, show that there are no reducing arcs.

We now proceed as in the proof of Theorem 9.1. The only difference is that we now have $2m$ triangles, rather than $2m-2$. However, it is still the case that

$\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$ has at most $12m$ ribs, and therefore

$$\min_s |K_{i_s}|_{P_{i_s}} \leq 12 \cdot 2^{L_2},$$

as required. ∎

COROLLARY 9.4: *Under the hypotheses of Theorem 9.3, if $i \neq j$ then*

$$\iota_i(P_i/K_i) \cap \iota_j(P_j/K_j) = \{1\}.$$

COROLLARY 9.5: *Under the hypotheses of Theorem 9.3, $\iota_i(P_i/K_i)$ is malnormal in G/K, for each i.*

PROPOSITION 9.6: *Suppose that G is torsion-free and that $|K_i|_{P_i} > 12 \cdot 2^{L_2}$ for each i. Suppose that $g \in G \smallsetminus \{1\}$ is such that there exists $x \in X$ so that the preferred path $p_{x,g.x}$ lies entirely within $D^{-1}([0, L_2 - 1])$. Then $g \notin K$.*

Proof. As usual, we suppose that the theorem is false, build a surface in X/K, and use its geometry to derive a contradiction.

If $g \in K$ then there is an equality in G of the form

$$1 = g^{-1} \prod_{i=1}^{m} g_i k_i g_i^{-1}, \tag{5}$$

where $k_i \in K_{j_i}$ and $g_i \in G$. We choose such an expression for g which minimizes m. Since $g \neq 1$ in G, we have $m \geq 1$.

Choosing words for each g_i, each k_i and g, the expression for g in (5) gives a loop in X, beginning at 1. In turn, this induces a map $\tilde{\phi}\colon \tilde{\Sigma} \to X$ of a disk, as described in Subsection 9.1. Also as in Subsection 9.1, we may glue $\tilde{\Sigma}$ along parts of the boundary corresponding to the g_i, to obtain a compact planar surface, Σ, together with a map $\phi'\colon \Sigma \to X/K$. The surface Σ has one distinguished boundary component which is labelled by a word representing g, and we call this component the **g-boundary** of Σ.

We claim that there are no reducing arcs for ϕ whose endpoints do not lie on the g-boundary. This follows as before: Any such reducing arc either contradicts Theorem 9.1 or else the minimality of m.

Let ξ be a simple path in X from 1 to x, and consider the loop in X which is the concatenation $\xi \cdot p_{x,gx} \cdot \mu \cdot \eta$, where μ is the path $g\xi$ traversed backwards, and η is a lift to X of the image under ϕ' of the g-boundary of Σ.

This loop may be filled with a disk in X, which projects to an annulus in X/K, which has one boundary component the image of the g-boundary of Σ.

Homotoping ϕ' over the annulus gives a new map $\phi\colon\ \Sigma \to X$ which maps the g-boundary to the image of the preferred path $p_{x,g.x}$ in X/K, see Figure 29.

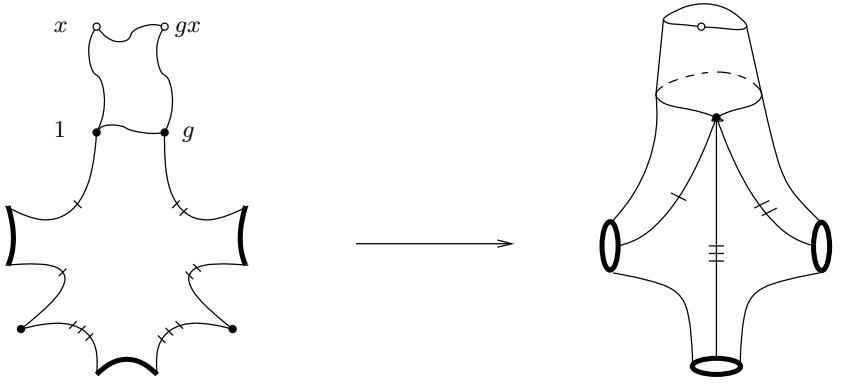

Figure 29. The surface Σ.

We now define a surface $\check{\Sigma}$ almost as in Definition 8.2 by attaching a half-open annulus to each component of $\partial\Sigma$ *except* for the g-boundary. This gives a map

$$\check{\phi}\colon\ \check{\Sigma} \to X/K.$$

The g-boundary of $\check{\Sigma}$ is defined in the obvious way.

Choose a (partially ideal) triangulation $\mathcal{T}$ of $\check{\Sigma}$ which has one vertex which is the preimage of x, and all other vertices ideal, and one edge the g-boundary of $\check{\Sigma}$.

Homotope $\check{\phi}$ to

$$\check{\phi}_{\mathcal{T}}\colon\ \check{\Sigma} \to X/K.$$

as in Lemma 8.6.

Because no edge in $\mathcal{T}$ is a reducing arc for $\check{\phi}_{\mathcal{T}}$, the image of each triangle $T \in \mathcal{T}$ lifts to a preferred triangle T_{abc} in X. Therefore, we can define the map

$$\ddot{\phi}_{\mathcal{T}}\colon\ \mathrm{Skel}(\check{\phi}_{\mathcal{T}}) \to X/G \cup (\partial_{\mathcal{H}}X)/G,$$

as in Remark 8.11.

We now proceed as in the proof of Theorem 9.1. By Euler characteristic, there are $2m-1$ triangles. Therefore, $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$ has less than $12m$ ribs.

Because $p_{x,g.x}$ lies entirely within $D^{-1}([0, L_2 - 1])$, none of the ribs intersect the g-boundary of $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$ (where the g-boundary of $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$ is defined in the obvious way). By Lemma 8.15 there must be at least

$$\min_s |K_{i_s}|_{P_{i_s}} m,$$

ribs in $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$. Thus we have

$$\min_s |K_{i_s}|_{P_{i_s}} \leq 12 \cdot 2^{L_2},$$

which is the required contradiction. ∎

This proposition has the following corollary:

COROLLARY 9.7: *Let $F \subset G$ be a finite set. Then there is a constant $C = C(F)$ so that if each $|K_i|_{P_i} > C$, then the quotient map $G \to G/K$ is injective on F.*

Proof. We prove the equivalent statement that there exists some $C' = C'(F)$ so that $G \to G/K$ sends no element of $F \smallsetminus \{1\}$ to 1. (To see the equivalence notice that $C(F) \leq C'((F \cup F^{-1})^2)$.)

Without loss of generality, we may assume that any element of F which is conjugate into some P_i actually lies in P_i. Write $F = F_1 \cup F_2$, where $F_1 = F \cap (\cup \mathcal{P})$ is the set of parabolic elements of F, and $F_2 = F \smallsetminus F_1$ is the set of non-parabolic elements of F. Let $S' = S \cup F_2$. Note that since S is a compatible generating set, so is S'. We may thus form the space $X' = X(G, \mathcal{P}, S')$. By Theorem 3.25, X' is δ'-hyperbolic for some $\delta' > 0$.

Since $S' \cap P_i = S \cap P_i$, the meaning of $|K_i|_{P_i}$ does not change in moving from X to X'. We may thus apply Proposition 9.6 in the context of X' rather than X. The preferred path joining 1 to f in X' is a single edge for any $f \in F_2$. Let $C'' = 24 \cdot 2^{3000\delta'}$. Proposition 9.6 implies that if $|K_i|_{P_i} > C''$ for each i, then no element of F_2 is sent to $1 \in G/K$. Each $f \in F_1$ is contained in some P_{i_f}. Let

$$C''' = \max\{|f|_{P_{i_f} \cap S} \colon f \in F_1\}$$

and let $C' = \max\{C'', C'''\}$. ∎

Using exactly analogous arguments to those in Proposition 9.6 we can obtain the following result.

PROPOSITION 9.8: *Suppose that G is torsion-free, and that $|K_i|_{P_i} > 18 \cdot 2^{L_2}$ for each i. Suppose that $g \in G$ is not conjugate into any $P_i \in \mathcal{P}$, and that there*

is some $x \in X$ with $p_{x,gx} \subset D^{-1}[0, L_2 - 1]$. Then there is no $h \in G$, $p \in P \in \mathcal{P}$ so that $hgh^{-1}p^{-1} \in K$.

Proof. The constant for $|K_i|_{P_i}$ changes because we consider a disk with several punctures. However, we have no control of the length of one of the punctures (that corresponding to p).

Therefore, if there are m punctures corresponding to elements of the K_i, then there are $2(m+1) - 1 = 2m + 1 \leq 3m$ triangles.

Otherwise, the proof is just as before. ∎

Remark 9.9: Once we have proved that G/K is hyperbolic relative to the images of the P_i in $\mathcal{P}$, this will imply that g does not project to a parabolic element of G/K.

Remark (about torsion) 9.10: In the presence of torsion, many of the arguments in this section (and later sections) become more difficult to implement. In particular, the notion of minimality for surfaces needs to be refined. Also, some of the results in this section need to be reformulated in the presence of torsion. The main issue here is that parabolics in G may already intersect nontrivially, and this causes a number of problems.

10. The surgered space

In this section we make the following assumptions: G is a torsion-free group which is hyperbolic relative to a collection $\mathcal{P} = \{P_1, \ldots, P_n\}$ of finitely generated subgroups. The finite set S is a compatible generating set for G with respect to $\mathcal{P}$. Finally, $\langle S, \mathcal{P} \mid \mathcal{R} \rangle$ is a finite relative presentation for G.

Recall from Theorem 3.25, that under these assumptions, the space $X(G, \mathcal{P}, S)$ is δ-hyperbolic for some δ. Moreover, given the finite relative presentation $\langle S, \mathcal{P} \mid \mathcal{R} \rangle$, we may adjoin 2-cells to $X(G, \mathcal{P}, S)$ to form a two-complex $X = X(G, \mathcal{P}, S, \mathcal{R})$ which satisfies a linear combinatorial isoperimetric inequality.

We further suppose that, for $1 \leq i \leq n$, we have $K_i \trianglelefteq P_i$. Let $K \trianglelefteq G$ be the normal closure of the union of the K_i. In order to construct a nice model space for G/K it is essential that the P_i/K_i inject into G/K. We thus make the standing assumption that:

Assumption 10.1: For each $1 \leq i \leq n$, $|K_i|_{P_i} > 12 \cdot 2^{L_2}$.

The space X contains a copy of the Cayley graph $\Gamma(G,S)$ of G, and has an associated depth function $D\colon X \to \mathbb{R}_+$ so that $\Gamma(G,S) \subset D^{-1}(0)$.

We now describe a "neutered" version of X, and how to modify it to get a model for G/K.

Definition 10.2: Let $Y = D^{-1}([0,L_2])$. The **boundary of** Y is $\partial Y = D^{-1}(L_2)$. If $H < G$ then the boundary of Y/H is the image of ∂Y in Y/H.

Remark 10.3: If the parabolics are not finitely presented, then Y will not be simply connected. Its fundamental group is generated by those loops in $D^{-1}(L_2)$ which cannot be filled in $D^{-1}(L_2)$.

The boundary of Y/G has n connected components which correspond to the subgroups $P_1,\dots,P_n$, as described in the next few paragraphs.

Let $\overline{1} \in X/G$ be the image of the vertices of $\Gamma(G,S)$ in X/G. For each $1 \le i \le n$ there is a unique L_2-horoball in X stabilized by P_i, which we denote H_i. There is a unique vertical path $\tilde{\gamma}_i$ joining 1 to $(i,1,1,L_2) \in H_i$.

Let γ_i be the image of $\tilde{\gamma}_i$ in $Y/G \subset X/G$, and let γ_i^{-1} be γ_i in the opposite direction. Let $T_i \subset Y/G$ be the component of $\partial(Y/G)$ containing the image y_i of $(i,1,1,L_2)$. The vertical path γ_i induces an inclusion of $\pi_1(T_i,y_i)$ into $\pi_1(Y/G,\overline{1})$.

Any loop in $Y/G \subset X/G$ based at $\overline{1}$ lifts to a unique path in X starting at 1 and ending at some group element. This gives a well-defined homomorphism from $\pi_1(Y/G,\overline{1})$ onto G, which maps $\pi_1(T_i)$ onto P_i.

The next two lemmas follow from Assumption 10.1, together with Theorem 9.1.

LEMMA 10.4: *If c is any loop in a boundary component of Y/K of length less than $12 \le |K_i|_{P_i}/2^{L_2}$, then c lifts to a loop in ∂Y.*

In particular, there is no loop in $\partial(Y/K)$ consisting of a single edge.

Definition 10.5: Let $Z = Z(K)$ be the 2-complex obtained from Y/K by gluing a combinatorial horoball onto each component of the boundary of Y/K.

Remark 10.6: The depth function on Y/K naturally extends to a depth function D on Z.

LEMMA 10.7: *If each $|K_i|_{P_i} > 2^{L_2}$, the complex Z is simply connected.*

Proof. The fundamental group of Y is generated by peripheral loops. Note that we have an exact sequence:

$$1 \longrightarrow \pi_1(Y) \longrightarrow \pi_1(Y/K) \longrightarrow K \longrightarrow 1$$

Thus, the fundamental group of Y/K is generated by the fundamental group of Y together with peripheral loops in Y/K which represent elements of $\pi_1(Y/K)$ which map to normal generators of K.

Any peripheral loop eventually dies in a horoball, by Lemma 10.4 and Proposition 3.7. ∎

LEMMA 10.8: *G/K acts freely and properly on Z.*

We thus have the following diagram of spaces:

(6)
$$\begin{array}{ccccc} X & \longleftarrow & Y & & \\ \downarrow & & \downarrow & & \\ X/K & \longleftarrow & Y/K & \longrightarrow & Z \\ \downarrow & & \downarrow & & \downarrow \\ X/G & \longleftarrow & Y/G & \longrightarrow & Z/(G/K) \end{array}$$

Each horizontal arrow in (6) is an inclusion, and each vertical arrow is a covering map.

We will show in Section 11 that Z satisfies a linear isoperimetric inequality if the $|K_i|_{P_i}$ are sufficiently large. Together with Proposition 10.12 below, this will imply that G/K is hyperbolic relative to $\mathcal{P}' = \{P_1/K_1, \ldots, P_n/K_n\}$, and hence complete the proof of Theorem 7.2.

LEMMA 10.9: *There is a G/K-equivariant embedding $\rho\colon Z^{(1)} \to (X/K)^{(1)}$.*

Proof. The spaces Z and X/K are identical at depths less than or equal to L_2. Thus there is an obvious map at these depths. It is also obvious that vertical edges in Z can be uniquely associated to vertical edges in X/K.

A horizontal edge e in Z at depth $L_2 + L$ corresponds to a path p at depth L_2 in Z of length at most 2^L. This path has already been mapped to a path p' in X/K at depth L_2 (still of length at most 2^L). The path p' lifts to a path $\tilde{p'}$ in X (which is still of length at most 2^L). The path $\tilde{p'}$ lies above an edge $\tilde{e'}$ in X at depth $L_2 + L$, which projects to an edge e' in X/K. This is $\rho(e)$. ∎

Let $\mathcal{P}' = \{\iota_i(P_i/K_i) \colon 1 \leq i \leq n\}$.

Observation 10.10: $D^{-1}(0) \subset Y/K \subset X/K$ is a relative Cayley complex for G/K with respect to the finite relative presentation $\langle S, \mathcal{P}' \mid \mathcal{R}\rangle$.

Recall the notation $\mathcal{H}_N$ from Definition 3.5, where $\mathcal{H}$ is a combinatorial horoball.

LEMMA 10.11: *Let $\tilde{A}$ be a 0-horoball in X, and A its projection in X/K. Then the intersection $A \cap Y/K$ is isomorphic to $\mathcal{H}_{L_2}$, where $\mathcal{H} = \mathcal{H}(A \cap D^{-1}(0))$.*

Proof. We define maps $\eta_1 \colon A \cap Y/K \to \mathcal{H}_{L_2}$ and $\eta_2 \colon \mathcal{H}_{L_2} \to A \cap Y/K$.

There are obvious bijections on the 0-skeleta, which extend to isomorphisms at the 0-level, and the vertical edges.

The first thing to note is that horizontal edges in $A \cap Y/K$ are not loops, by Assumption 10.1 and Theorem 9.1. It is also true that horizontal edges in $\mathcal{H}_{L_2}$ are not loops.

Consider a horizontal edge $e \in A \cap Y/K$, at depth L. This lifts to a horizontal edge $\tilde{e}$ in X at depth L. This can be pushed up to a path $\tilde{p}$ of length at most 2^L in $D^{-1}(0) \cap \tilde{A}$, which project to a path p in $D^{-1}(0) \cap A$ above e. We have already defined $\eta_1(p)$, and this path lies above an edge in $\mathcal{H}_{L_2}$ at depth L. This edge is $\eta_1(e)$. Thus we have defined η_1 on the 1-skeleton of $A \cap Y/K$.

We now define η_2 on the 1-skeleton of $\mathcal{H}_{L_2}$. Let e' be an edge at depth L in $\mathcal{H}_{L_2}$. There is a path p' above e' at the 0-level of $\mathcal{H}_{L_2}$. The path $\eta_2(p')$ is already defined. The path $\eta_2(p')$ lifts to a path in $D^{-1}(0) \cap \tilde{A}$, and lies above an edge $\tilde{e}$ in $D^{-1}(L) \cap \tilde{A}$, which in turn projects to an edge e in $A \cap Y/K$. Set $\eta_2(e') = e$.

We have now defined η_1 and η_2 on the 1-skeleta, and we leave it as an exercise to prove that they are mutually inverse.

The map η_1 obviously extends over the 2-skeleta. Consider then a 2-cell c in $\mathcal{H}_{L_2}$. Then $\eta_2(\partial c)$ is a loop in $A \cap Y/K$ of length at most 5. Suppose $\eta_2(\partial c)$ lifts to a path σ which is not a loop. Let $k \in K_i$ be the element which sends one endpoint of σ to the other. Then $|k|_{P_i \cap S} \leq 5.2^{L_2}$, a contradiction. Therefore, $\eta_2(\partial c)$ does lift to a loop in Y, so there is a 2-cell filling $\eta_2(\partial c)$. This 2-cell is $\eta_2(c)$. ∎

The following proposition follows easily from Observations 3.6 and 10.10 and Lemma 10.11.

PROPOSITION 10.12: *If $|K_i|_{P_i} \geq 12 \cdot 2^{L_2}$ for all $i \in \{1, \ldots, n\}$, then Z is equivariantly isomorphic to $X(G/K, \mathcal{P}', S, \mathcal{R})$.*

11. A linear isoperimetric inequality

We have reduced the proof of Theorem 7.2 to proving that the space Z satisfies a homological isoperimetric inequality. This is proved in Theorem 11.11 below. The proof of this result breaks neatly into two pieces: a combinatorial piece (Proposition 11.10), and a homological piece (which becomes Theorem 11.11 below).

We continue to assume that G is a group which is hyperbolic relative to $\mathcal{P} = \{P_1, \ldots, P_n\}$, and that $K_i \lhd P_i$ for each i. We let $X = X(G, \mathcal{P}, S, \mathcal{R})$, Y, and Z be as described in section 10. In order to slightly simplify the proof of Proposition 11.10, we replace Assumption 10.1 with the slightly stricter:

Assumption 11.1:

$$|K_i|_{P_i} \geq 24 \cdot 2^{L_2} = 24 \cdot 2^{3000\delta}$$

for each i.

We show (Theorem 11.11) that under these conditions, the surgered space Z described in section 10 satisfies a linear homological isoperimetric inequality. It follows via Theorem 3.25 and Proposition 10.12 that G/K is hyperbolic relative to the images of the subgroups in $\mathcal{P}$.

Let

$$\rho\colon\ Z^{(1)} \to X/K^{(1)}$$

be the map from Lemma 10.9.

LEMMA 11.2: *Let w be a 2-cell in Z so that $\rho(\partial w)$ does not bound a 2-cell in X/K. Then $w \subset D^{-1}[L_2 + 1, \infty)$.*

Proof. $\rho(\partial w)$ must lie in $D^{-1}([L_2, \infty))$ in X/K, since X/K and Z are identical between depth 0 and L_2.

If $\rho(\partial w)$ does not surround a 2-cell, then it does not lift to a loop in X. Thus, since the length of $\rho(\partial w)$ is at most 5, it lifts to a path in X of length 5, whose endpoints are in the same orbit under the action of K. Theorem 9.1 and Assumption 11.1 now imply that these points must lie in the same orbit under the action of the stabilizer of the horoball in which they lie. Thus there is some $k \in K_i$ so that gkg^{-1} sends one endpoint to the other for some g.

Let L be the minimal depth of $\rho(\partial w)$. There is a lift of $\rho(\partial w)$ starting and ending at depth L in X, and the endpoints of this lift are joined by a horizontal path of length at most 4 (as the diligent reader may readily verify). Therefore, the points at depth 0 above these endpoints may be joined by a horizontal path of length at most $4 \cdot 2^L$. The length of such a path is an upper bound for $|k|_{P_i \cap S}$, so by Assumption 11.1 we have $4 \cdot 2^L \geq 24 \cdot 2^{L_2}$. In particular, $L > L_2 + 1$ as required. ∎

Definition 11.3: We refer to 2-cells in Z as described in the above lemma as **missing** 2-cells. Note that the map ρ extends to those 2-cells of Z which are not missing.

Definition 11.4: Suppose that E is a cellulated disk, and that $\phi\colon E \to Z$ is a combinatorial map. A **partly missing piece** of ϕ is a component C of $(D \circ \phi)^{-1}(L_2, \infty)$ so that $\phi(C)$ contains some missing 2-cell. Let $E_* \subseteq E$ be the complement of the partly missing pieces of E.

Let $\overline{E}$ be a closed regular neighborhood of $E_* \cup \partial E$ whose image under ϕ contains no missing 2-cells (we may need to adjust ϕ by a small homotopy rel ∂E, and re-cellulate, to ensure that $\overline{E}$ exists). A **reducing arc** for ϕ is an map $\sigma\colon [0,1] \to E$ satisfying the following:

(1) σ is an essential arc in $\overline{E}$.
(2) The endpoints of σ lie in partly missing pieces of E.
(3) σ is homotopic rel endpoints into $D^{-1}(L_2, \infty)$, and this homotopy does not pass over any missing 2-cell.

Remark 11.5: This is slightly different from the way reducing arcs were defined in Section 8 for a couple of reasons. First, we need to deal with the possibility that, for instance, two components of $E \setminus E_*$ have intersecting closures. Second, we do not want to reduce along arcs with an endpoint in ∂E, because this would change the loop being filled.

LEMMA 11.6: *Let E be a disk. Suppose that γ is a combinatorial loop in Z, and $\phi\colon E \to Z$ is a filling of γ. Then there is some $\phi'\colon E \to Z$ so that each partly missing piece of ϕ' is simply connected, and the number of partly missing pieces of ϕ' is no more than the number of partly missing pieces of ϕ.*

Proof. Fix a partly missing piece C, and let ξ be the component of the boundary of $\overline{E}$ which lies in C and contains all other components of $\partial\overline{E} \cap C$. The loop

$\phi(\xi)$ lies entirely in $Z \smallsetminus Y$, and can thus be filled there. We set $\phi' = \phi$ outside of ξ, and set ϕ' equal to this new filling inside ξ. ∎

LEMMA 11.7: *Let E be a disk. Suppose that γ is a combinatorial loop in Z, and $\phi\colon E \to Z$ is a filling of γ chosen so that the number of partly missing pieces of E is minimized. Then there are no reducing arcs for ϕ.*

Proof. If there is a reducing arc σ, we may reduce the number of partly missing pieces as follows:

Suppose first that the arc joins distinct partly missing pieces. Then we may cut open E along the arc and add in two copies of the homotopy into $D^{-1}[L_2, \infty)$, thus combining the two partly missing pieces into one. Since the homotopy passes over no missing 2-cells, we have decreased the number of partly missing pieces.

Next suppose that the arc σ joins some partly missing piece C to itself. Again we can cut open along σ and insert the homotopy, thus creating a non-simply connected partly missing pieces. By Lemma 11.6, this non-simply connected component can be replaced with a simply connected one. Since σ was essential, it enclosed at least one partly missing piece other than C, and so the modified map has fewer partly missing pieces than it did before. ∎

Definition 11.8: If ϕ is a map of a disk E into Z, then $\rho \circ \phi|_{E_* \cup \partial E}$ can be extended to a proper map of a punctured disk

$$\check{\phi}\colon \check{E}_* \to X/K.$$

The surface $\check{E}_*$ can be obtained from E by removing a point from the interior of each partly missing piece. The map $\check{\phi}$ is then defined to be equal to $\rho \circ \phi$ on $E_* \cup \partial E$. The complement of $E_* \cup \partial E$ in $\check{E}_*$ is a union of annuli; the map $\check{\phi}$ is defined so that the image of these annuli consists entirely of vertical squares, just as in Definition 8.2.

The proof of Lemma 8.6 easily adapts to a proof of the following:

LEMMA 11.9: *Let E be a disk, and let $\phi\colon E \to Z$ be a map with no reducing arcs, so that ϕ has at least one partly missing piece and $\phi(\partial E) \cap \Gamma_K$ is nonempty. Let $\mathcal{T}$ be a partially ideal triangulation of $\check{E}_*$ with a single vertex $v_0 \in \partial E = \partial(\check{E}_*)$ satisfying $\phi(v_0) \in \Gamma_K$. Then $\check{\phi}$ is properly homotopic to a map*

$$\check{\phi}_{\mathcal{T}}\colon \check{E}_* \to X/K$$

which sends each edge of $\mathcal{T}$ to the image in X/K of a preferred path in X.

The key combinatorial step in proving Theorem 11.11 is the following:

PROPOSITION 11.10: *There is a constant $C = C(\delta) > 0$ so that the following holds: Let $w\colon S^1 \to Z$ be a combinatorial loop, and suppose that*

$$\phi\colon E \to Z$$

is a filling of this loop by a disk E, and that this filling has no reducing arcs. If ϕ has $m \geq 1$ partly missing pieces, then $|w|_1 \geq Cm$.

Proof. We will show that C can be chosen equal to $\min\{1, \frac{L_2}{2(\lambda+\epsilon)}\}$, where λ and ϵ are the constants of quasi-geodesicity from Corollary 5.14. If w does not intersect $\Gamma_K = \Gamma(G,S)/K$, then w can be filled with a disk containing at most one partly missing piece. We may therefore assume that $w(1) \in \Gamma_K \subset Z$.

Let γ be a loop homotopic (rel 1) to the loop $\rho \circ w$ in X/K, so that γ lifts to a preferred path $\widetilde{\gamma}$ in X. Let $g \in G$ be the unique group element which sends the initial point of $\widetilde{\gamma}$ to the terminal point.

By Corollary 5.14, $|\gamma|_1 = |\widetilde{\gamma}|_1 \leq \lambda d(1,g) + \epsilon$, where $d(1,g)$ is the distance from 1 to g in X. As $|w|_1$ is bounded below by $d(1,g)$, we have

$$|w|_1 \geq \frac{1}{\lambda+\epsilon}|\gamma|_1.$$

It thus suffices to bound $|\gamma|_1$ below linearly in terms of m. We remark that $\gamma(1) = (\rho \circ w)(1) \in \Gamma_K \subset X/K$.

By assumption, we have a map $\phi\colon E \to Z$ of a disk into Z with no reducing arcs. Let $\check{\phi}\colon \check{E}_* \to X/K$ be the proper map from Definition 11.8 of an m-times punctured disk into X/K. Note that $\check{\phi}|_{\partial(\check{E}_*)} = \rho \circ w$.

Let $\mathcal{T}$ be a (partially ideal) triangulation of $\check{E}_*$ with one vertex the preimage of $(\rho \circ w)(1)$ and all other vertices ideal, and one edge on the boundary. This induces an obvious triangulation $\mathcal{T}'$ of the disk E.

By Lemma 11.9, $\check{\phi}$ is homotopic to $\check{\phi}_{\mathcal{T}}\colon \check{E}_* \to X/K$ so that $\check{\phi}_{\mathcal{T}}|_{\partial \check{E}_*} = \gamma$, and all edges of $\mathcal{T}$ map to paths in X/K which lift to preferred paths in X.

We now consider the map

$$\ddot{\phi}_{\mathcal{T}}\colon \mathrm{Skel}(\check{\phi}_{\mathcal{T}}) \to X/G \cup (\partial_{\mathcal{H}} X)/G,$$

as in Remark 8.11.

Let p be a puncture on $\check{E}_*$. The link $Lk(p)$ in $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$ is a loop and we have

$$\ddot{\check{\phi}}_{\mathcal{T}}|_{Lk(p)} \colon S_1 \to D^{-1}([L_2, \infty)) \subset X/G.$$

This loop represents a conjugacy class in P_i for some i, and this class is contained in K_i, by Lemma 8.15.

A puncture p is called **interior** if $Lk(p)$ is composed entirely of ribs and ligaments, and **exterior** otherwise. The puncture p corresponds to a vertex of the triangulation $\mathcal{T}'$ of E, which we also describe as **interior** or **exterior** (see Figure 30).

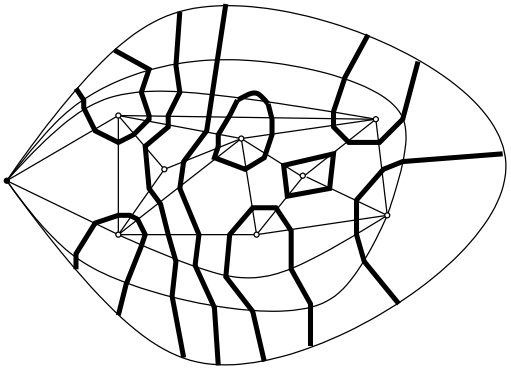

Figure 30. This picture shows a possible (actually, impossible) picture of $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$, with two interior punctures, and the rest exterior. Ribs and ligaments are shown in bold. Each of the exterior punctures contributes at least 2 points to $(D \circ \widetilde{\gamma})^{-1}(L_2)$.

Let V_I be the number of interior vertices of $\mathcal{T}'$, and let $V_\partial = m+1-V_I$ be the number of exterior vertices. It is clear from the definitions that the cardinality of $(D \circ \widetilde{\gamma})^{-1}(L_2)$ is at least $2(V_\partial - 1)$. The set $D^{-1}(L_2)$ partitions $\widetilde{\gamma}$ into subsegments, which alternate between lying in $D^{-1}[0, L_2]$ and lying in $D^{-1}[L_2, \infty)$. As the initial and terminal points of $\widetilde{\gamma}$ lie in $D^{-1}(0)$, those subsegments of $\widetilde{\gamma}$ with image in $D^{-1}[0, L_2]$ must have length at least L_2 (In fact, all but the initial and terminal subsegments must have length at least $2L_2$, as they must pass between distinct L_2-horoballs.). There are at least V_∂ such subsegments, and

so

$$|\gamma|_1 = |\widetilde{\gamma}|_1 \geq V_\partial L_2 = (m + 1 - V_I)L_2 > (m - V_I)L_2. \tag{7}$$

In order to bound $|\gamma|_1$ below by a linear function of m, it therefore suffices to show that V_I is at most some definite proportion of m (bounded away from 1).

Let R be the number of ribs in $\mathrm{Skel}(\check{\phi}_{\mathcal{T}})$. Then $R \leq 6(2m-1)$, by Corollary 5.39. However, $R \geq 24V_I$, by Assumption 11.1 and Lemma 8.15.

This implies that $V_I < \frac{1}{2}m$. By (7), we deduce that

$$|\gamma|_1 \geq \frac{L_2}{2}m,$$

and hence

$$|w|_1 \geq \frac{L_2}{2(\lambda+\epsilon)}m. \qquad \blacksquare$$

11.1. Proof of Theorem 7.2. We make the same assumptions about G, $\mathcal{P}$, S, $\mathcal{R}$ and X stated at the beginning of the last section.

Theorem 11.11: *If* $\min_j\{|K_j|_{P_j}\} > 24 \cdot 2^{L_2} = 24 \cdot 2^{3000\delta}$, *then* Z *satisfies a homological linear isoperimetric inequality in the sense of Definition 2.28.*

Proof. The idea behind the proof is as follows. Take a loop in Z, which is filled by some disk. Attempt to move this disk to X/K, fail, and find a punctured disk mapping into X/K. Triangulate, lift and straighten the triangles, project back to X/K, attempt to transfer back to Z, fail and fill the failures.

By Theorem 2.26, it suffices to show that there is a constant $M > 0$ so that any simple loop c bounds a rational 2-chain w with

$$|w|_1 \leq M|c|_1.$$

Let c be a simple (combinatorial) loop in Z. If c lies entirely inside a single horoball, we may fill with a disk of area at most $3|c|_1$ by Proposition 3.7. We may thus suppose that $c(1) \in D^{-1}(0) \subset Z$. Thus $\rho \circ c$ lifts to a path $\widetilde{c}$ in X between two (not necessarily distinct) vertices g and h of the Cayley graph $\Gamma(G,S) \subset X$. Consider the 1-chain $\widetilde{\underline{c}}$ corresponding to the path $\widetilde{c}$. Then $c - q_{g,h}$ is a 1-cycle. By Proposition 6.6, $|q_{g,h}|_1 \leq 6000\delta^2 d(g,h) \leq 6000\delta^2|c|_1$. Hence, if M_X is the homological isoperimetric constant for X, there is a 2-chain $\widetilde{\omega}$ so that $\partial\widetilde{\omega} = \widetilde{\underline{c}} - q_{g,h}$, and $|\widetilde{\omega}|_1 \leq M_X\big(6000\delta^2 + 1\big)|c|_1$.

As Z is simply connected, there is a combinatorial map of a cellulated disk E

$$\phi\colon E \to Z$$

so that $\phi|_{\partial E}$ is c. We may suppose that the map ϕ has the minimal possible number of partly missing pieces in the sense of Definition 11.4. By Lemma 11.7, ϕ has no reducing arcs. By Proposition 11.10, the number of partly missing pieces is at most $C^{-1}|c|$, for a $C > 0$ which depends only on δ. Thus we can triangulate the punctured surface $\check{E}_*$ defined in Definition 11.8 with a triangulation $\mathcal{T}$ consisting of fewer than $2C^{-1}|c|$ triangles. Let $\check{\phi}_{\mathcal{T}}|_e$ be the map from Lemma 11.9. This induces, for each triangle $T \in \mathcal{T}$, a preferred triangle $\psi_T\colon \partial\Delta \to X \cup \partial_{\mathcal{H}} X$.

Suppose that $T \in \mathcal{T}$, and that the image of the vertices of T are a, b, and c (with order coming from the orientation of T). Let $c_T = c_{abc}$, the 1-chain defined in Definition 6.8, and let $\omega_T = \omega_{abc}$ be the 2-chain as in Corollary 6.11. Let $\widetilde{\xi} = \sum_{T\in\mathcal{T}} \omega_T$. Each ω_T satisfies $|\omega_T|_1 \le M_X T_1$ by Corollary 6.11.

Now let $\mu = \pi_\sharp(\widetilde{\omega} + \widetilde{\xi})$, and let μ^{thick} be the 2-chain which comes from including only those 2-cells which lie entirely in $Y = D^{-1}[0, L_2]$. Note that

$$\begin{aligned} |\mu^{\text{thick}}|_1 &\le |\mu|_1 \le |\widetilde{\omega}|_1 + |\widetilde{\xi}|_1 \\ &\le \left[(M_X(6000\delta^2 + 1)) + 2C^{-1}T_1 M_X\right]|c|_1 \end{aligned}$$

Let $M_\mu = (M_X(6000\delta^2 + 1) + 2C^{-1}T_1 M_X$.

Since μ^{thick} is supported entirely in Y, it determines a 2-chain μ_Z in Z. Furthermore, $c^{\text{thin}} := c - \partial\mu_Z$ satisfies the following:

(1) The support of c^{thin} lies entirely in $D^{-1}[L_2, \infty) \subset Z$, and
(2) $|c^{\text{thin}}|_1 \le |c|_1 + |\partial\mu_Z| \le (1 + MM_\mu)|c|_1$, where M is the maximum length of the boundary of a 2-cell in Z (which is the same as that maximum length in X).

Thus by Proposition 3.7 and Theorem 2.26 there is a 2-chain ζ satisfying

(1) $\partial\zeta = c^{\text{thin}}$, and
(2) $|\zeta|_1 \le 3|c^{\text{thin}}|_1 \le 3(1 + MM_\mu)|c|_1$.

Finally, we note that $\partial(\mu_Z + \zeta) = c$ and $|\mu_Z + \zeta|_1 \le |\mu_Z|_1 + |\zeta|_1 \le (M_\mu + 3 + 3MM_\mu)|c|_1$. This completes the proof of Theorem 11.11, and hence also of Theorem 7.2. ∎

We close by proving that G/K is nonelementary.

THEOREM 11.12: *Suppose that G is torsion-free and that $|K_i|_{P_i} > 24 \cdot 2^{L_2}$ for each i. Then G/K is non-elementary relatively hyperbolic (relative to $\{\iota_i(P_i/K_i)\}_i$.*

Proof. If all parabolics are finite, then G/K is hyperbolic.

By Theorem 3.33, there is a hyperbolic element $g \in G$ with an axis which is entirely contained in $D^{-1}[0, 19\delta]$. If $x \in X$ is contained in this axis, then for all j, the preferred path $p_{x,g^j.x}$ is the geodesic between these points (since $19\delta + \delta \leq L_1$). Therefore, by Proposition 9.6, none of these elements die in G/K. Therefore, G/K is infinite. Furthermore, by Proposition 9.8, g does not project to a parabolic element of G/K. Thus G/K is not equal to any of the $\iota_i(P_i/K_i)$. By Theorem 9.3, the intersection of two distinct parabolic subgroups of G/K is trivial.

Suppose G/K is virtually cyclic. Then G/K has a finite normal subgroup N with quotient either infinite cyclic or infinite dihedral. In fact, N must be contained in every parabolic; and so it is trivial. However, the parabolic subgroups of G/K have size (much) greater than 2.

We may now suppose that some $\iota_i(P_i/K_i)$ is infinite. We have already seen that G/K does not equal $\iota_i(P_i/K_i)$ for any i, so Theorem 3.34 implies that G/K is non-elementary relatively hyperbolic. ∎

References

[1] I. Agol, *Bounds on exceptional Dehn filling*, Geometry and Topology **4** (2000), 431–449 (electronic).

[2] E. Alibegović, *A combination theorem for relatively hyperbolic groups*, Bulletin of the London Mathematical Society **37** (2005), 459–466.

[3] D. J. Allcock and S. M. Gersten, *A homological characterization of hyperbolic groups*, Inventiones Mathematicae **135** (1999), 723–742.

[4] B. Bowditch, *Relatively hyperbolic groups*, preprint, available at www.maths.soton.ac.uk/staff/Bowditch/preprints.html, 1999.

[5] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319 Springer-Verlag, Berlin, 1999.

[6] I. Bumagin, *On definitions of relatively hyperbolic groups*, in *Geometric Methods in Group Theory*, Contemp. Math., vol. 372, Amer. Math. Soc., Providence, RI, 2005, pp. 189–196.

[7] J. W. Cannon and D. Cooper, *A characterization of cocompact hyperbolic and finite-volume hyperbolic groups in dimension three*, Transactions of the American Mathematical Society **330** (1992), 419–431.

[8] F. Dahmani, *Classifying spaces and boundaries for relatively hyperbolic groups*, Proceedings of the London Mathematical Society (3), **86** (2003), 666–684.

[9] F. Dahmani, *Combination of convergence groups*, Geometry and Topology **7** (2003), 933–963 (electronic).

[10] F. Dahmani, *Les Groupes Relativement Hyperboliqes et Leurs Bords*, PhD thesis, Strasbourg, 2003.

[11] C. Druţu and M. Sapir, *Relatively hyperbolic groups with rapid decay property*, International Mathematics Research Notices **19** (2005), 1181–1194.

[12] C. Druţu and M. Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), 959–1058. With an appendix by Denis Osin and Sapir.

[13] B. Farb, Relatively hyperbolic groups, Geometric and Functional Analysis **8** (1998), 810–840.

[14] K. Fujiwara and J. F. Manning, *CAT*(0) *and CAT*(−1) *fillings of hyperbolic manifolds*, in preparation.

[15] S. Gersten, *A cohomological characterisation of hyperbolic groups*, preprint, 1996.

[16] M. Gromov, *Hyperbolic groups*, in *Essays in Group Theory*, Mathematical Sciences Research Institute Publications, vol. 8. Springer, New York, 1987, pp. 75–263.

[17] D. Groves, *Actions of relatively hyperbolic groups on strongly bolic spaces*, in preparation.

[18] D. Groves, *Limit groups for relatively hyperbolic groups, I: The basic tools*, preprint, `arXiv:math.GR/0412492`, 2004.

[19] D. Groves and J. Manning, *Relative bounded cohomology and relatively hyperbolic groups*, in preparation.

[20] D. Groves, J. Manning, and D. Osin, in preparation.

[21] D. Groves and J. F. Manning, *Fillings, finite generation and direct limits of relatively hyperbolic groups*, Groups, Geometry, and Dynamics **1** (2007), 329–342.

[22] C. D. Hodgson and S. Kerckhoff, *Universal bounds for hyperbolic Dehn surgery*, Annals of Mathematics **162** (2005), 67–421.

[23] G. Hruska, Relative hyperbolicity and relative quasi-convexity for countable groups, preprint, `arXiv:0801.4596`, 2008.

[24] G. C. Hruska and B. Kleiner, *Hadamard spaces with isolated flats*, Geometry and Topology **9** (2005), 1501–1538 (electronic). With an appendix by the authors and Mohamad Hindawi.

[25] C. Hummel and V. Schroeder, *Cusp closing in rank one symmetric spaces*, Inventiones Mathematicae **123** (1996), 283–307.

[26] M. Lackenby, *Word hyperbolic Dehn surgery*, Inventiones Mathematicae **140** (2000), 243–282.

[27] M. Lackenby, *Attaching handlebodies to 3-manifolds*, Geometry and Topology **6** (2002), 889–904 (electronic).

[28] V. Lafforgue, *K-théorie bivariante pour les algèbres de Banach et conjecture de Baum-Connes*, Inventiones Mathematicae **149** (2002), 1–95.

[29] I. Mineyev, *Straightening and bounded cohomology of hyperbolic groups*, Geometric and Functional Analysis **11** (2001), 807–839.

[30] I. Mineyev, *Bounded cohomology characterizes hyperbolic groups*, The Quarterly Journal of Mathematics **53** (2002), 59–73.

[31] I. Mineyev, *Flows and joins of metric spaces*, Geometry and Topology **9** (2005), 403–482 (electronic).

[32] I. Mineyev, N. Monod, and Y. Shalom, *Ideal bicombings for hyperbolic groups and applications*, Topology **43** (2004), 1319–1344.

[33] I. Mineyev and G. Yu, *The Baum-Connes conjecture for hyperbolic groups*, Inventiones Mathematicae **149** (2002), 97–122.

[34] L. Mosher and M. Sageev, *Nonmanifold hyperbolic groups of high cohomological dimension*, preprint, available at `http://newark.rutgers.edu/~mosher/`.

[35] D. V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Memoirs of the American Mathematical Society **179**(843) (2006), vi+100.

[36] D. V. Osin, *Peripheral fillings of relatively hyperbolic groups*, Inventiones Mathematicae **167** (2007), 295–326.

[37] P. Papasoglu, *Strongly geodesically automatic groups are hyperbolic*, Inventiones Mathematicae **121** (1995), 323–334.

[38] D. Rebbechi, *Algorithmic Properties of Relatively Hyperbolic Groups*. PhD thesis, Rutgers, Newark, 2001.

[39] V. Schroeder, *A cusp closing theorem* Proceedings of the American Mathematical Society **106** (1989), 797–802.

[40] W. P. Thurston, *Geometry and Topology of Three-Manifolds*, Princeton lecture notes, Princeton, 1979.

[41] A. Yaman, *A topological characterisation of relatively hyperbolic groups* Journal für Reine und Angewandte Mathematik **566** (2004), 41–89.

[42] G. Yu, *Hyperbolic groups admit proper affine isometric actions on l^p-spaces*, Geometric and Functional Analysis **15** (2005), 1144–1151.