GRAPH SEPARATORS AND BOUNDARIES OF RIGHT-ANGLED ARTIN AND COXETER GROUPS

By

Wes Camp

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

MATHEMATICS

May, 2013

Nashville, Tennessee

Approved:

Dr. Michael Mihalik

Dr. Victoria Greene

Dr. John Ratcliffe

Dr. Mark Sapir

Dr. Steven Tschantz

ACKNOWLEDGMENTS

I would first like to give my utmost gratitude to my advisor, Mike Mihalik. This work would not have been possible without his continuous support and constant guidance during our almost daily conversations.

Secondly, I would like to thank the many other faculty members who have helped guide me through both my undergraduate and graduate career at Vanderbilt.

Finally, my thanks to my fellow graduate students and friends, who made both work and play a better experience.

Table of Contents

List of Figures

Figure | Page

Chapter I

INTRODUCTION

In this paper, we examine some connections between topological properties of boundaries of CAT(0) spaces and amalgamated product splittings of groups that act on these spaces. In particular, we study right-angled Artin and Coxeter groups, which are groups that have a natural correspondence with finite graphs in that any finite graph has a unique corresponding right-angled Artin and Coxeter group. A separator of such a graph gives a decomposition of the corresponding group as an amalgamated product; these are the splittings with which we are concerned. The Stallings theorem on ends of groups, a well-known result of group theory, states that a finitely generated group has more than one end (equivalently, has non-connected boundary) if and only if it splits as an amalgamated product or HNN extension over a finite subgroup. The results in this paper, in the same vein, determine connections between amalgamated product splittings of a group (arising from separators of the graph) and local and/or path connectivity of the group's boundaries.

In Chapter II, we introduce the basics of CAT(0) spaces and their boundaries. In Chapter III, we classify the right-angled Coxeter groups with no $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ subgroups that have locally/non-locally connected boundary (this was a joint result with Michael Mihalik). It is known ([15]) that if the presentation graph of any right-angled Coxeter group admits a certain type of separator, then the group has all of its boundaries non-locally connected; we show that the absence of such a separator implies the group has all its boundaries locally connected, given that the group does not contain a visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ subgroup. To do this, we use a construction called a filter, first introduced in [16], that allows us to 'fill in' the space between two rays in ∂X. A filter is a connected, one-ended planar graph with edges labeled by generators of our right-angled Coxeter group, giving a natural map from the filter to the Cayley graph of our group, and so also to any CAT(0) space X on which the group acts. The limit set of a filter always maps to a connected set in ∂X, and so the challenge is to show that if two rays are close in ∂X, then the filter between them can be constructed to have small diameter in ∂X. In [16], conditions are placed on the considered groups that allow the construction of filters with essentially hyperbolic geometry; our hypotheses give no such guarantee.

The main theorem of this chapter is as follows:

Theorem. *Suppose (W, S) is a one-ended right-angled Coxeter system that has no visual subgroup isomorphic to $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$.*

1. *If W visually splits as $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$ then A is word hyperbolic, W has unique boundary homeomorphic to the suspension of the boundary of A, and the boundary of W is non-locally connected iff A is infinite ended.*

2. *Otherwise, W has locally connected boundary iff (W, S) has no virtual factor separator.*

If W has no visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ subgroup but splits as in item (1), then A has no visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^2$ subgroup and is therefore word hyperbolic ([20]), and so ∂A is unique ([10]) and is locally connected iff A is one-ended (iff the presentation graph of A does not split over a complete graph, see Remark 3.2.4). If W splits as $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$ then any boundary of W is the suspension of a boundary of A (see [16]), and is locally connected iff this boundary of A is locally connected. Therefore (1) is easily checked, and the actual content of our theorem lies in (2).

In section 3, we provide examples that show boundary local connectivity cannot be approached through any reasonable graph of groups technique. In our first example, we demonstrate a right-angled Coxeter group that splits as $A *_C B$, and there is $c \in C$ such that c^∞ determines a point of non-local-connectivity in both ∂A and ∂B; however, by our theorem, $A *_C B$ has locally connected boundary. Our second example splits as $A *_C B$ with A and B word hyperbolic (so ∂A and ∂B locally connected) and C virtually a hyperbolic surface group with boundary S^1, but $A *_C B$ contains a virtual factor separator and therefore has non-locally-connected boundary.

In Chapter IV, we turn our attention to boundaries of right-angled Artin groups. In [8] Croke and Kleiner demonstrate a group that acts geometrically on two CAT(0) spaces with non-homeomorphic boundaries, and it was later shown ([23]) that the same group has uncountably many distinct CAT(0) boundaries. The group is the right-angled Artin group whose presentation graph is the path on four vertices P_4, and so has presentation

$$\langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1 \rangle.$$

In [7], it is shown that the boundary of the standard CAT(0) cube complex on which this group acts is non-path-connected. The boundary of such a cube complex is connected if and only if the the presentation graph of the group is connected (and so the group is one-ended). In this chapter, the method in [7] is generalized to a class of right-angled Artin groups whose presentation graphs admit a certain type of splitting. The main theorem here is as follows:

Theorem. *Let* Γ *be a connected graph. Suppose* Γ *contains an induced subgraph* $(\{a,b,c,d\},\{\{a,b\},\{b,c\},\{c,d\}\})$ *(isomorphic to* P_4*), and there are subsets* $B \subset lk(c)$ *and* $C \subset lk(b)$ *with the following properties:*

1. *B separates c from a in Γ, with $d \notin B$;*
2. *C separates b from d in Γ, with $a \notin C$;*
3. *$B \cap C = \emptyset$.*

Then $\partial \mathcal{S}_\Gamma$ *is not path connected.*

Here, $\mathcal{S}_\Gamma$ is the standard CAT(0) cube complex on which the right-angled Artin group A_Γ with presentation graph Γ acts geometrically, and $lk(v)$ is the set of vertices of Γ sharing an edge with v. We in fact show a slightly stronger result, with the hypothesis $B \cap C = \emptyset$ replaced with the statement of Claim 4.2.7. The hypotheses here essentially require a copy of P_4 in Γ that is either not contained in a cycle, or has every cycle containing it separated by chords based at b and c. It is a known fact of graph theory that any graph that does not split as a join contains an induced subgraph isomorphic to P_4, and any graph Γ that splits as a non-trivial join has $\partial \mathcal{S}_\Gamma$ path connected, so the hypothesis that Γ contain a copy of P_4 is satisfied in any interesting case.

If a connected boundary of a CAT(0) space is locally connected, then it is a Peano space (a continuous image of [0,1]) and therefore path connected. The boundaries of some right-angled Coxeter groups are therefore known to be path connected ([16] and [6]), because they are locally connected. However, a consequence of a theorem in [15] is that for right-angled Artin groups, $\partial\mathcal{S}_\Gamma$ is locally connected iff Γ is a complete graph; i.e. $A_\Gamma \cong \mathbb{Z}^n$ and $\partial\mathcal{S}_\Gamma \cong S^{n-1}$. Thus no approach involving local connectivity works for right-angled Artin groups.

In [19], the construction of [8] is generalized to demonstrate a class of groups with non-unique boundary. These groups are of the form

$$G = (G_1 \times \mathbb{Z}^n) *_{\mathbb{Z}^n} (\mathbb{Z}^n \times \mathbb{Z}^m) *_{\mathbb{Z}^m} (\mathbb{Z}^m \times G_2),$$

where G_1 and G_2 are infinite CAT(0) groups. It is easily verified that if G_1 and G_2 are right-angled Artin groups, then G is a right-angled Artin group whose presentation graph satisfies the conditions of the main theorem of this paper; in fact, the method of this paper should work even if G_1 and G_2 are arbitrary infinite CAT(0) groups.

It seems this boundary path connectivity problem may be related to the question of when two right-angled Artin groups are quasi-isometric. In [1], Behrstock and Neumann show that all right-angled Artin groups whose presentation graphs are trees of diameter greater than 2 are quasi-isometric; in [3], Bestvina, Kleiner, and Sageev show that right-angled Artin groups with atomic presentation graphs (no valence 1 vertices, no separating vertex stars, and no cycles of length ≤ 4) have A_Γ quasi-isometric to $A_{\Gamma'}$ iff $\Gamma \cong \Gamma'$. The connection between these results and the result of this paper is that if Γ is a tree of diameter greater than 2, then Γ satisfies the hypotheses of the main theorem here, and therefore $\partial\mathcal{S}_\Gamma$ has non-path-connected boundary; if Γ is atomic, then Γ cannot satisfy the hypotheses of the main theorem here.

Chapter II

CAT(0) PRELIMINARIES

Definition 2.0.1. *A metric space* (X, d) *is* ***proper*** *if each closed ball is compact.*

Definition 2.0.2. *Let* (X, d) *be a complete proper metric space. Given a geodesic triangle* $\triangle abc$ *in* X*, we consider a* ***comparison triangle*** $\triangle \overline{a}\overline{b}\overline{c}$ *in* $\mathbb{R}^2$ *with the same side lengths. We say* X *satisfies the* ***CAT(0) inequality*** *(and is thus a* ***CAT(0) space****) if, given any two points* p, q *on a triangle* $\triangle abc$ *in* X *and two corresponding points* $\overline{p}, \overline{q}$ *on a corresponding comparison triangle* $\triangle \overline{a}\overline{b}\overline{c}$*, we have*

$$d(p, q) \leq d(\overline{p}, \overline{q}).$$

Proposition 2.0.3. *If* (X, d) *is a CAT(0) space, then*

1. *the distance function* $d : X \times X \to \mathbb{R}$ *is convex,*

2. X *has unique geodesic segments between points, and*

3. X *is contractible.*

Definition 2.0.4. *A* ***geodesic ray*** *in a CAT(0) space* X *is an isometry* $[0, \infty) \to X$*.*

Definition 2.0.5. *Let* (X, d) *be a proper CAT(0) space. Two geodesic rays* $c, c' : [0, \infty) \to X$ *are called* ***asymptotic*** *if for some constant* K*,* $d(c(t), c'(t)) \leq K$ *for all* $t \in [0, \infty)$*. Clearly this is an equivalence relation on all geodesic rays in* X*, regardless* *of basepoint. We define the* ***boundary*** *of X (denoted ∂X) to be the set of equivalence classes of geodesic rays in X. We denote the union $X \cup \partial X$ by $\overline{X}$.*

The next proposition guarantees that the topology we wish to put on the boundary is independent of our choice of basepoint in X.

Proposition 2.0.6. *Let (X, d) be a proper CAT(0) space, and let $c : [0, \infty) \to X$ be a geodesic ray. For a given point $x \in X$, there is a unique geodesic ray based at x which is asymptotic to c.*

For a proof of this (and more details on what follows), see [5].

We wish to define a topology on $\overline{X}$ that induces the metric topology on X. Given a point in ∂X, we define a neighborhood basis for the point as follows:
Pick a basepoint $x_0 \in X$. Let c be a geodesic ray starting at x_0, and let $\epsilon > 0$, $r > 0$. Let $S(x_0, r)$ denote the sphere of radius r based at x_0, and let $p_r : X \to S(x_0, r)$ denote the projection onto $S(x_0, r)$. Define

$$U(c, r, \epsilon) = \{x \in \overline{X} : d(x, x_0) > r, d(p_r(x), c(r)) < \epsilon\}.$$

This consists of all points in $\overline{X}$ whose projection onto $S(x_0, r)$ is within ϵ of the point of the sphere through which c passes. These sets together with the metric balls in X form a basis for the **cone topology**. The set ∂X with this topology is sometimes called the **visual boundary**. For our purposes, we will just call it the boundary of X.

Definition 2.0.7. *We say a finitely generated group G* ***acts geometrically*** *on a proper geodesic metric space X if there is an action of G on X such that:*

1. *Each element of G acts by isometries on X,*

2. *The action of G on X is cocompact, and*

3. *The action is properly discontinuous.*

Definition 2.0.8. *We call a group G a* ***CAT(0) group*** *if it acts geometrically on a CAT(0) space.*

The next theorem, due to Milnor [18], will be used in conjunction with Lemmas 3.1.26 and 4.1.17 and to identify geodesic rays in X with certain rays in a right-angled Coxeter/Artin group which acts on X.

Theorem 2.0.9. *If a group G with a finite generating set S acts geometrically on a proper geodesic metric space X, then G with the word metric with respect to S is quasi-isometric to X under the map $g \mapsto g \cdot x_0$, where x_0 is a fixed base point in X.*

Proposition 2.0.10. *If X and Y are proper CAT(0) spaces, then $\partial(X \times Y) \cong \partial X * \partial Y$, where $*$ denotes the spherical join.*

Chapter III

LOCAL CONNECTIVITY OF BOUNDARIES OF RIGHT-ANGLED COXETER GROUPS

3.1 Coxeter group preliminaries

We use [4] and [9] as basic references for the results in this section.

Definition 3.1.1. *A* ***Coxeter system*** *is a pair* (W, S)*, where* W *is a group with* ***Coxeter presentation****:*

$$\langle S : (st)^{m(s,t)} \rangle$$

where $m(s,t) \in \{1, 2, \ldots, \infty\}$*,* $m(s,t) = 1$ *iff* $s = t$*, and* $m(s,t) = m(t,s)$*. The relation* $m(s,s) = 1$ *means each generator is of order 2, and* $m(s,t) = 2$*, iff* s *and* t *commute.*

Definition 3.1.2. *We call a Coxeter group* (W, S) ***right-angled*** *if* $m(s,t) \in \{2, \infty\}$ *for all* $s \neq t$*.*

We are only interested in right-angled Coxeter groups in this chapter but we state many of the lemmas of this section in full generality. In what follows, we will let $\Lambda = \Lambda(W, S)$ denote an abbreviated version of the Cayley graph for W with respect to the generating set S. As usual, the vertices of Λ are the elements of W, and there is an edge between the vertices w and ws for each $s \in S$, but instead of having two edges between adjacent vertices in the graph (since each generator has order 2), we allow only one.

Definition 3.1.3. *For a Coxeter system (W,S), the **presentation graph** $\Gamma(W,S)$ for (W,S) is the graph with vertex set S and an edge labeled $m(s,t)$ connecting distinct $s,t \in S$ when $m(s,t) \neq \infty$.*

Definition 3.1.4. *For a Coxeter system (W,S), a **word** in S is an n-tuple $w = [a_1, a_2, \ldots, a_n]$, with each $a_i \in S$. Let $\overline{w} \equiv a_1 \cdots a_n \in W$. We say the word w is **geodesic** if there is no word $[b_1, b_2, \ldots, b_m]$ such that $m < n$ and $\overline{w} = b_1 \cdots b_m$. Define $lett(w) \equiv \{a_1, \ldots, a_n\}$.*

Definition 3.1.5. *For a Coxeter system (W,S), let $\overline{e} \in S$ be the label of the edge e of $\Lambda(W,S)$. An **edge path** $\alpha \equiv (e_1, e_2, \ldots, e_n)$ in a graph Γ is a map $\alpha : [0,n] \to \Gamma$ such that α maps $[i, i+1]$ isometrically to the edge e_i. For α an edge path in $\Lambda(W,S)$, let $lett(\alpha) \equiv \{\overline{e}_1, \ldots, \overline{e}_n\}$, and let $\overline{\alpha} \equiv \overline{e}_1 \cdots \overline{e}_n$. If β is another geodesic with the same initial and terminal points as α, then call β a **rearrangement** of α.*

Lemma 3.1.6. *Suppose (W,S) is a Coxeter system, and a and b are S-geodesics for $w \in W$ (so $w = \overline{a} = \overline{b}$). Then $lett(a) = lett(b)$.*

Definition 3.1.7. *If (W,S) is a Coxeter system and $A \subset S$, then $lk(A) \equiv \{t \in S : m(a,t) = 2 \text{ for all } a \in A\}$. So when (W,S) is right-angled, $lk(A)$ is the combinatorial link of A in $\Gamma(W,S)$, and the subgroups $\langle A \rangle$ and $\langle lk(A) \rangle$ of W commute.*

Lemma 3.1.8. ***(The Deletion Condition)**. Suppose (W,S) is a Coxeter system. If the S-word $w = [a_1, a_2, \ldots, a_n]$ is not geodesic, then two of the a_i delete; i.e. we have for some $i < j$, $\overline{w} = a_1 a_2 \cdots a_n = a_1 a_2 \cdots a_{i-1} a_{i+1} \cdots a_{j-1} a_{j+1} \cdots a_n$.*

For a Coxeter system (W, S), an edge path $\alpha = (e_1, e_2, \ldots, e_n)$ in $\Lambda(W, S)$ is geodesic iff the word $[\bar{e}_1, \bar{e}_2, \ldots, \bar{e}_n]$ is geodesic. If α is not geodesic and $\bar{e}_i$ deletes with $\bar{e}_j$, for $i < j$, let τ be the the path beginning at the end point of e_{i-1} with edge labels $[\bar{e}_{i+1}, \ldots, \bar{e}_{j-1}]$. Then τ ends at the initial point of e_{j+1}, so that $(e_1, \ldots, e_{i-1}, \tau, e_{j+1}, \ldots, e_n)$ is a path with the same end points as α. We say the edges e_i and e_j **delete** in α.

Definition 3.1.9. *If (W, S) is a Coxeter system and $A \subset S$, then the subgroup of W generated by A is called a* ***special*** *(or* ***visual****) subgroup of W.*

Lemma 3.1.10. *Suppose (W, S) is a Coxeter system, and $A \subset S$. Then the special subgroup $\langle A \rangle$ of W has Coxeter (sub)-presentation*

$$\langle A : (st)^{m(s,t)}; s, t \in A \rangle$$

In particular, distinct $s, t \in S$ determine unique elements of W, and $m(s, t)$ is the order of st for all $s, t \in S$.

Lemma 3.1.11. *Suppose (W, S) is a Coxeter system, and $U, V \subset S$, with $U \cap V = \emptyset$. If u is a geodesic in the letters of U and v is a geodesic in the letters of V, then $[u, v]$ is an S-geodesic.*

Definition 3.1.12. *For (W, S) a Coxeter system and α a geodesic in $\Lambda(W, S)$, let $B(\alpha) \equiv \{\bar{e} \in S : e$ is a Λ-edge at the end of α and (α, e) is not geodesic$\}$.*

Lemma 3.1.13. *Suppose (W, S) is a Coxeter system, and α a geodesic in Λ. Then $B(\alpha)$ generates a finite group.*

Lemma 3.1.14. *If (W, S) is a right-angled Coxeter system, and $s, t \in S$ delete in some S-word. Then $s = t$.*

Lemma 3.1.15. *Suppose (W, S) is a right-angled Coxeter system, $[a_1, a_2, \ldots, a_n]$ is geodesic and $[a_1, a_2, \ldots, a_n, a_{n+1}]$ is not. Then a_{n+1} deletes with some a_m. If $i \neq n+1$ is the largest integer such that $a_i = a_{n+1}$, then a_{n+1} deletes with a_i and a_{n+1} commutes with each letter $a_{i+1}, a_{i+2}, \ldots, a_n$.*

Definition 3.1.16. *Suppose Γ is the presentation graph of a Coxeter system (W, S), and $C \subset S$ separates the vertices of Γ. Let A' be the vertices of a component of $\Gamma - C$ and $B = S - A'$. Let $A = A' \cup C$. Then W splits as $\langle A \rangle *_{\langle C \rangle} \langle B \rangle$ (see [17]) and this splitting is called a* **visual decomposition** *for (W, S).*

Definition 3.1.17. *Let (W, S) be a Coxeter system, and let e be an edge of $\Lambda(W, S)$ with initial vertex $v \in W$. The* ***wall*** *$w(e)$ is the set of edges of $\Lambda(W, S)$ each fixed (setwise) by the action of the conjugate $v\bar{e}v^{-1}$ on Λ.*

Remark 3.1.18. *Certainly $e \in w(e)$ and if d is an edge of $w(e)$, with vertices u and w, then $(v\bar{e}v^{-1})u = w$ and $(v\bar{e}v^{-1})w = u$. Also, $\Lambda(W, S) - w(e)$ has exactly two components and these components are interchanged by the action of $v\bar{e}v^{-1}$ on $\Lambda(W, S)$.*

If (W, S) is right-angled, then given an edge a of $\Lambda(W, S)$ with initial vertex y_1 and terminal vertex y_2, a is in the same wall as e iff there is an edge path $(t_1, \ldots, t_n)$ in $\Lambda(W, S)$ based at w_1 so that $w_1\bar{t}_1 \cdots \bar{t}_n = y_1$ and $w_2\bar{t}_1 \cdots \bar{t}_n = y_2$, where y_1 and y_2 are the vertices of e and $m(\bar{e}, \bar{t}_i) = 2$ for each $1 \leq i \leq n$.

Definition 3.1.19. *Let (W, S) be a right-angled Coxeter system. We say the walls $w(e) \neq w(d)$ of $\Lambda(W, S)$* **cross** *if there is a relation square in $\Lambda(W, S)$ with edges in $w(e)$ and $w(d)$.*

Remark 3.1.20. *We have the following basic properties of walls in a right-angled Coxeter system (W, S):*

1. *If edges a and e of $\Lambda(W, S)$ are in the same wall, then $\overline{a} = \overline{e}$.*

2. *Being in the same wall is an equivalence relation on the set of edges of $\Lambda(W, S)$.*

3. *If $(e_1, e_2, \ldots, e_n)$ is an edge path in $\Lambda(W, S)$ then e_i and e_j are in the same wall iff $\overline{e}_i$ and $\overline{e}_j$ delete in the word $[e_1, e_2, \ldots, e_n]$. Furthermore, the path $(e'_{i+1}, \ldots, e'_{j-1})$ that begins at the initial point of e_i, and has the same labeling as $(e_{i+1}, \ldots, e_{j-1})$, ends at the end point of e_j and $w(e_k) = w(e'_k)$ for all $i < k < j$. If γ is a path in $\Lambda(W, S)$, then γ is geodesic iff no two edges of γ are in the same wall.*

4. *If γ and τ are geodesics in $\Lambda(W, S)$ between the same two points, then the edges of γ and τ define the same set of walls.*

The basics of van Kampen diagrams can be found in Chapter 5 of [13]. Suppose (W, S) is a right-angled Coxeter system. We need only consider relation squares with boundary labels $abab$ in van Kampen diagrams for right-angled Coxeter groups (since those of the type aa are easily removed). Let $(w_1, \ldots, w_n)$ be an edge path loop in $\Lambda(W, S)$, so $\overline{w}_1 \ldots \overline{w}_n = 1$ in W. Consider a van Kampen diagram D for this word. For a given boundary edge d of D (corresponding to say w_i), d can belong to at most one relation square of D and there is an edge d_1 opposite d on this square. Similarly, if d_1 is not a boundary edge, it belongs to a unique relation square adjacent to the one containing d and d_1. Let d_2 be the edge opposite d_1 in the second relation square. These relation squares define a *band* in D starting at d and ending at say d' on the boundary of D and corresponding to some w_j with $j \neq i$. This means that w_i and w_j are in the same wall. However, w_k and w_ℓ being in the same wall does not necessarily mean that they are part of the same band in D; but if $(w_1, \ldots, w_r)$ and $(w_{r+1}, \ldots, w_n)$ are both geodesic, then by (3) in the above remark, bands in D correspond exactly to walls in $\Lambda(W, S)$. This is the situation we will usually consider.

Lemma 3.1.21. *Let (W, S) be a right-angled Coxeter system, and let γ be a geodesic in $\Lambda(W, S)$ with initial vertex x and terminal vertex y. Let A be a set of edges of γ, and τ_A be a shortest path based at x containing an edge in the same wall as a for all $a \in A$. Then τ_A can be extended to a geodesic to y.*

Proof. Let v denote the endpoint of τ_A, and let λ be a geodesic from v to y. Let $\tau_A = (a_1, \ldots, a_n)$ and consider a van Kampen diagram D for $(\tau_A, \lambda, \gamma^{-1})$. If $W(a_j) = W(a)$ for some $a \in A$ and the band for a_j does not end on γ, then it must end on λ, by (3) of Remark 3.1.20. However, then the band for a cannot end on λ, γ, or τ_A (which is impossible). Therefore the band for a_j must end on the edge of D corresponding to the edge a of γ. Now suppose for some $1 \leq i \leq n$, the band for a_i ends on λ. Deleting edges of (τ_A, λ) corresponding to this shared wall gives a path shorter than τ_A with an edge in the same wall as a for all $a \in A$ (see Remark 3.1.20 (3)), a contradiction. Therefore, all bands on λ and τ_a end on γ, so (τ_a, λ) has the same length as γ and is therefore geodesic. □

The following lemma has some of its underlying ideas in Lemma 5.10 of [16]. It is an important tool for measuring the size of (connected) sets in the boundaries of our groups and is used repeatedly in our proof of the main theorem.

Lemma 3.1.22. *Suppose (W, S) is a right-angled Coxeter system, and (α_1, α_2) and (β_1, β_2) are geodesics in $\Gamma(W, S)$ between the same two points. There exist geodesics $(\gamma_1, \tau_1), (\gamma_1, \delta_1), (\delta_2, \gamma_2)$, and (τ_2, γ_2) with the same end points as $\alpha_1, \beta_1, \alpha_2, \beta_2$ respectively, such that:*

1. *τ_1 and τ_2 have the same edge labeling,*
2. *δ_1 and δ_2 have the same edge labeling, and*
3. *$lett(\tau_1)$ and $lett(\delta_1)$ are disjoint and commute.*

Furthermore, the paths (τ_1^{-1}, δ_1) and (δ_2, τ_2^{-1}) are geodesic.

Proof. Consider a van Kampen diagram for the loop $(\alpha_1, \alpha_2, \beta_2^{-1}, \beta_1^{-1})$ (Figure 3.1), and recall that since (α_1, α_2) and (β_1, β_2) are geodesic, bands in this van Kampen diagram correspond exactly to walls in $\Lambda(W, S)$. Let $a_1, \ldots, a_n$ be the edges of α_1 (in the order they appear on α_1) that are in the same wall as an edge of β_1. Notice that if e is an edge of α_1 occurring before a_1, then $w(e)$ crosses $w(a_1)$. Therefore α_1 can be rearranged to begin with an edge in $w(a_1)$, since $\overline{a}_1$ commutes with every edge label of α_1 before it. Similarly, $w(a_2)$ must cross $w(e)$ for any edge $e \neq a_1$ of α_1 occurring before a_2, so α_1 can be rearranged to begin with an edge in $w(a_1)$ followed by an edge in $w(a_2)$. Continuing for each a_i gives us a rearrangement (γ_1, τ_1) of α_1 where the

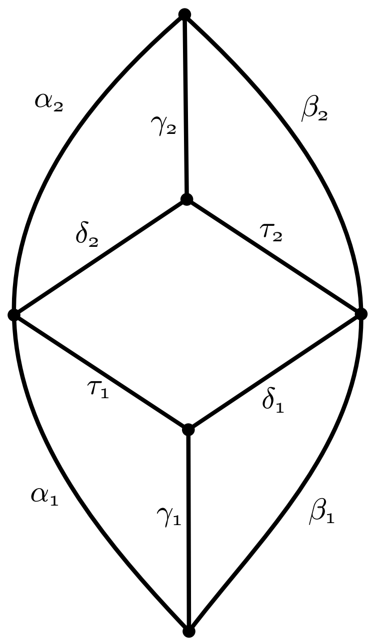

Figure 3.1: Lemma 3.1.22

walls of γ_1 are exactly $w(a_1), \ldots, w(a_n)$. If $b_1, \ldots, b_m$ are the edges of β_1 in the same wall as an edge of α_1, then the same process gives us a rearrangement (γ_1', δ_1) of β_1 where the walls of γ_1' are exactly $w(b_1), \ldots, w(b_m)$. However, $\{w(a_1), \ldots, w(a_n)\} = \{w(b_1), \ldots, w(b_m)\}$, so $m = n$ and γ_1 and γ_1' are geodesics between the same points, so (γ_1, δ_1) is a rearrangement of β_1. Construct rearrangements (δ_2, γ_2) and (τ_2, γ_2) of α_2 and β_2 respectively in the same way, and note that τ_1 and τ_2 have the same walls, δ_1 and δ_2 have the same walls, and every wall of τ_1 crosses every wall of δ_1. In particular, (see Remark 3.1.20 (3)) (τ_1^{-1}, δ_1) is geodesic. □

Remark 3.1.23. *For the entirety of this chapter, we will only consider the case of Lemma 3.1.22 where* $|\alpha_1| = |\beta_1|$*. In this case,* $|\tau_1| = |\tau_2| = |\delta_1| = |\delta_2|$*, so the* ***diamond*** *formed by the loop* $\tau_1^{-1}\delta_1\tau_2\delta_2^{-1}$ *is actually a product square. If* y *is the*

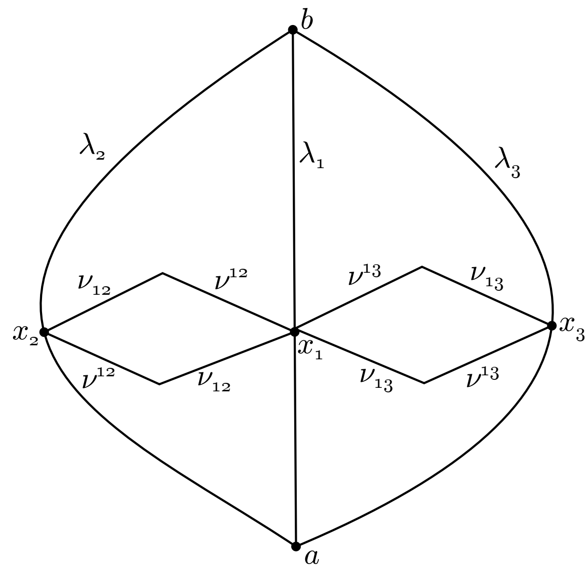

Figure 3.2: Lemma 3.1.24

endpoint of α_1 *and* μ *is any other geodesic between the same points as* (α_1, α_2)*, the diamond between* (α_1, α_2) *and* μ *at* y *is therefore well defined. We call* τ_1^{-1} *the* ***down edge path*** *at* y *and* δ_2 *the* **up edge path** *at* y *of the diamond for* (α_1, α_2) *and* (β_1, β_2)*.*

Lemma 3.1.24. *Suppose* (W, S) *is a right-angled Coxeter system with no visual subgroup isomorphic to* $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$*. Let* λ_1*,* λ_2*,* λ_3 *be* $\Lambda(W, S)$*-geodesics between two points* a *and* b*, and let* x_1*,* x_2*,* x_3 *be points on* λ_1*,* λ_2*,* λ_3 *respectively, such that the* x_i *are all equidistant from* a*. Let* ν_{12} *and* ν_{13} *be the down edge paths respectively of the diamonds at* x_1 *between* λ_1 *and* λ_2 *and between* λ_1 *and* λ_3*, as in Lemma 3.1.22, and suppose* $|\nu_{12}| \geq |\nu_{13}| \geq 2|S|$*. If* $\{a, b\} \subset lett(\nu_{12}) \cap lett(\nu_{13})$ *and* $m(a, b) = \infty$*, then* $d(x_2, x_3) < 2(|\nu_{12}| - |\nu_{13}|) + 4|S|$*.*

Proof. To simplify notation we use the same label for two paths with the same edge labeling. Let ν^{12} and ν^{13} be the up edge paths respectively of the diamonds at x_1 between λ_1 and λ_2 and between λ_1 and λ_3. Note that at x_2, $\nu^{12}\nu_{12}\nu_{13}\nu^{13}$ is a path from x_2 to x_3. By Lemma 3.1.22, $\{a,b\}$ is disjoint from and commutes with $lett(\nu^{12}) \cup lett(\nu^{13})$. Thus, ν^{13} cannot have a pair of walls with unrelated labels cross a pair of walls with unrelated labels from ν^{12}, since that would give a visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ in W. Rearrange ν^{12} and ν^{13} so they have a longest common initial segment (see definition 3.1.5). As ν^{12} and ν^{13} are initial segments of a geodesic from x_1 to b, the walls of the unshared edges of ν^{13} cross those of ν^{12}. In particular, the unshared part of ν^{13} has length $\leq |S|-1$, and ν^{12} and ν^{13} share two walls with unrelated labels. By symmetry, this last part implies ν_{13} and ν_{12} at x_1 can be rearranged to have a shared initial segment so the unshared part of ν_{13} has length $\leq |S|-1$. Deleting edges of the path $\nu^{12}\nu_{12}\nu_{13}\nu^{13}$ (from x_2 to x_3) corresponding to these shared walls leaves us with a geodesic from x_2 to x_3 of length less than $2(|\nu_{12}| - |\nu_{13}|) + 4|S|$. □

Let (W,S) be a right-angled Coxeter group acting geometrically on a CAT(0) space X. Pick a base point $* \in X$ and identify a copy of the Cayley graph for (W,S) inside X as in Theorem 2.0.9. If vertices u,v of $\Lambda(W,S)$ are adjacent, then we connect $u*$ and $v*$ with a CAT(0) geodesic in X. This defines a map $C : \Lambda \to X$ respecting the action of W. If α is a Λ-geodesic, we call $C(\alpha)$ a Λ-geodesic in X.

Definition 3.1.25. *Let $r : [a,b] \to X$ be a geodesic segment in X with $r(a) = x$ and $r(b) = y$. For $\delta > 0$, we say that a Cayley graph geodesic α δ-**tracks** r if every point of $C(\alpha)$ is within δ of a point of the image of r and the endpoints of r and $C(\alpha)$ are within δ of each other.*

Proofs of the next two lemmas can be found in section 4 of [16].

Lemma 3.1.26. *There exists some* $\delta_1 > 0$ *such that for any geodesic ray* $r : [0, \infty) \to X$ *based at* x_0*, there is a geodesic ray* α_r *in* $\Lambda(W, S)$ *that* δ_1*-tracks* r*.*

Lemma 3.1.27. *There exist* $c, d > 0$ *such that, for any infinite geodesic rays* r *and* s *and* X *based at* x_0 *that are within* ϵ *of each other at distance* M *from* x_0*, there are Cayley graph geodesic rays* α *and* β *which* $(c\epsilon + d)$*-track* r *and* s *respectively, and which share a common initial segment of length* $M - c\epsilon - d$*.*

3.2 Local connectivity and filter construction

Definition 3.2.1. *We say a CAT(0) group* G *has* ***(non-)locally connected boundary*** *if for every CAT(0) space* X *on which* G *acts geometrically,* ∂X *is (non-)locally connected.*

Definition 3.2.2. *Let* (W, S) *be a right-angled Coxeter system, and let* Γ *be the presentation graph for* (W, S)*. A* ***virtual factor separator*** *for* (W, S) *(or* Γ*) is a pair* (C, D) *where* $D \subset C \subset S$*,* C *separates vertices of* Γ*,* $\langle C - D \rangle$ *is finite and commutes with* $\langle D \rangle$*, and there exist* $s, t \in S - D$ *such that* $m(s, t) = \infty$ *and* $\{s, t\}$ *commutes with* D*.*

In this section we prove the following theorem:

Theorem 3.2.3. *Suppose* (W, S) *is a one-ended right-angled Coxeter system that has no visual subgroup isomorphic to* $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$*.*

1. *If* W *visually splits as* $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$ *then* A *is word hyperbolic,* W *has unique boundary homeomorphic to the suspension of the boundary of* A*, and the bound-*

ary of W is non-locally connected iff A is infinite ended.

2. *Otherwise, W has locally connected boundary iff (W, S) has no virtual factor separator.*

Part (1) of this result is clear; if the right angled Coxeter system (W, S) does not visually split as a direct product $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$ and has a virtual factor separator, then W has non-locally connected boundary (see [16]). It remains to show local connectivity of the boundaries of CAT(0) spaces acted upon geometrically by one-ended right-angled Coxeter groups with no virtual factor separators. To do this, we pick two rays whose end points are "close" in ∂X, and use Lemma 3.1.27 to find two tracking Cayley geodesics which share a long initial segment. We then construct a filter of geodesics (a way of "filling in" the space) between the branches of these Cayley geodesics such that its limit set gives a small connected set in ∂X containing our original rays.

In what follows, let (W, S) be a right-angled, one-ended Coxeter system with no virtual factor separator and containing no visual subgroup isomorphic to $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$. Set $N = |S|$. We will show that if W acts geometrically on a CAT(0) space X, then given $\epsilon > 0$, there exists δ such that if two points $x, y \in \partial X$ satisfy $d(x, y) < \delta$, then there is a connected set in ∂X of diameter $\leq \epsilon$ containing x and y.

Remark 3.2.4. *The right-angled Coxeter group W is one-ended iff $\Gamma(W, S)$ contains no complete separating subgraph (i.e., a subgraph whose vertices generate a finite group in W). For a proof of this, see [17].*

Remark 3.2.5. *If e is an edge in $\Lambda(W, S)$, we let $\overline{e} \in S$ denote the label of e. Recall*

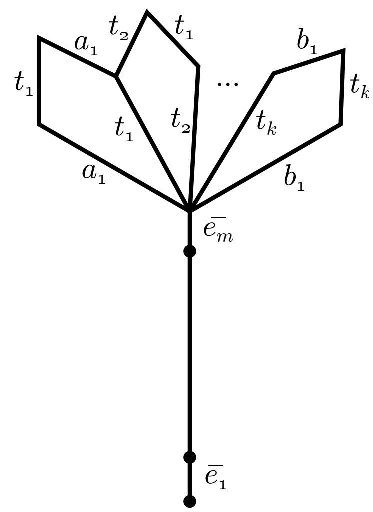

Figure 3.3: A fan for the geodesics $(e_1, \ldots, e_m, e_{m+1})$ and $(e_1, \ldots, e_m, d_{m+1})$

that for $g \in W$, $B(g)$ is the set of $s \in S$ such that gs is shorter than g, and that $\langle B(g) \rangle$ is finite (Lemma 3.1.13).

Remark 3.2.6. *If α is a geodesic in $\Lambda(W, S)$ from a vertex a to another vertex b, then for any other geodesic γ from a to b, we have $B(\alpha) = B(\gamma)$. Since this set depends only on a and b, we may use the notation $B(b \to a)$ to denote $B(\alpha)$, where it is more convenient to do so.*

We begin with an example that demonstrates one important idea behind our proof. Let (W, S) be a right-angled Coxeter system where W is one-ended and acts geometrically on a CAT(0) space X. Suppose that $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ and $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$ are Λ-geodesics in X, based at a vertex $*$, that $(c+d)$-track two CAT(0) geodesics r and s in X (as in Lemma 3.1.27), and let x_m denote the endpoint of $(e_1, \ldots, e_m)$. Set $a_1 = \bar{e}_{m+1}$ and $b_1 = \bar{d}_{m+1}$. By the previous remarks, $B(x_m \to *)$ does not separate the presentation graph $\Gamma(W, S)$, and $a_1, b_1 \notin B(x_m \to$

$*$). Let $a_1, t_1, \ldots, t_k, b_1$ be the vertices of a path from a_1 to b_1 in $\Gamma(W, S)$ where each $t_i \notin B(x_m \to *)$. We can construct a (labeled) planar diagram (Figure 3.3) that maps naturally into Λ (respecting labels) and then to X. As in [16], we call Figure 3.3 a **fan** for the geodesics $(e_1, \ldots, e_m, e_{m+1})$ and $(e_1, \ldots, e_m, d_{m+1})$. Each loop corresponds to the relation given by t_i and t_{i+1} commuting. Since each t_i commutes with t_{i+1} and $t_i, t_{i+1} \notin B(x_m \to *)$, the path $(e_1, \ldots, e_m, t_i, t_{i+1})$ is geodesic for each i (this is an easy consequence of Lemma 3.1.15). Now, let $a_2 = \overline{e}_{m+2}$, $b_2 = \overline{d}_{m+2}$, and continue. We overlap our original fan with fans for the pairs of geodesics $(e_1, \ldots, e_m, e_{m+1}, e_{m+2})$ and $(e_1, \ldots, e_m, e_{m+1}, t_1)$, $(e_1, \ldots, e_m, t_1, a_1)$ and $(e_1, \ldots, e_m, t_1, t_2)$, and so on, ending with a fan for $(e_1, \ldots, e_m, d_{m+1}, t_k)$ and $(e_1, \ldots, e_m, d_{m+1}, d_{m+2})$.

By continuing to build fans in this manner, we construct (Figure 3.4) a connected, one-ended, planar graph (with edge labels in S) called a **filter** for the geodesics $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ and $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$. Note that if v is a vertex of the filter, then the obvious edge paths in the filter from $*$ to v define Λ-geodesics. The limit set determined by this filter in ∂X is a connected set containing our original rays r and s. However, this connected set may not be small. We refer to the image of a filter, in Λ or in X, again as a filter.

If we wish for the limit set of our filter to be small in ∂X, we need to ensure that the CAT(0) geodesics between $*$ and points in our filter are not far from the base point x_m of our filter. Using Lemma 3.1.22, we know what a wide bigon between two geodesics in Λ must look like. Our first goal is to classify the "down edge paths", from x_m towards $*$, of any potential diamond given by a wide bigon in Λ, and show there are only two "types" of such paths.

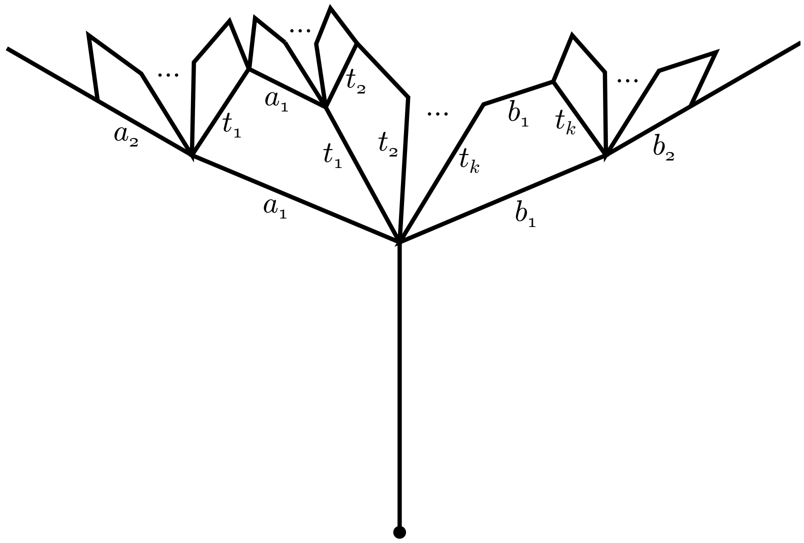

Figure 3.4: A filter for a pair of geodesics

Remark 3.2.7. *For the rest of this section, we assume that* Γ *has no virtual factor separators and* (W, S) *contains no visual subgroup isomorphic to* $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$.

Definition 3.2.8. *Construct a geodesic from* x_m *to* $*$ *in* Λ *as follows: let* α_1 *be a longest geodesic with edge labels in the finite group* $\langle B(x_m \to *)\rangle$, *and let* y_1 *be the endpoint of* α_1 *based at* x_m. *Let* α_2 *be a longest geodesic in the finite group* $\langle B(y_1 \to *)\rangle$. *Continuing in this way, we obtain a geodesic* $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ *from* x_m *to* $*$. *We call this a* ***back combing*** *geodesic from* x_m *to* $*$.

Remark 3.2.9. *We have the following properties of a back combing geodesic* $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ *from* x_m *to* $*$:

1. *Every edge label of* α_i *commutes with every other edge label of* α_i.

2. *No edge label of* α_{i+1} *commutes with every edge label of* α_i.

3. *Let* (γ_1, γ_2) *be a* Λ*-geodesic from* x_m *to* $*$ *and let* v *be the endpoint of* γ_1*. If* $(\beta_1, \beta_2, \ldots, \beta_s)$ *is a back combing geodesic from* x_m *to* v*, then the set of walls of* β_i *is a subset of the set of walls of* α_i*, for* $1 \leq i \leq s$*. In particular:*

4. *Let* (γ_1, γ_2) *be a* Λ*-geodesic from* x_m *to* $*$*. If* γ_1 *has an edge in the same wall as an edge of* α_j *for some* $1 \leq j \leq r$*, then* γ_1 *contains an edge in the same wall as an edge of* α_i *for all* $1 \leq i \leq j$*.*

5. *Let* (γ_1, γ_2) *and* (τ_1, τ_2) *be* Λ*-geodesics from* x_m *to* $*$*. If each of* τ_1*,* γ_1*, and* α_j *(for some* $1 \leq j \leq r$*) has an edge of the wall* $w(e)$*, then for each* $1 \leq i \leq j$*, each of* α_i*,* τ_1*, and* γ_1 *has an edge of the wall* $w(e_i)$*.*

We will always assume that x_m and $*$ are sufficiently far apart, so for now suppose $d(x_m, *) > 7N^2$. Let $\alpha_{7N+1} = (u_1, u_2, \ldots, u_k)$ (note $k < N$), and for $1 \leq i \leq k$, let U_i be a shortest Λ-geodesic based at x_m such that last edge of U_i is in the same wall as u_i (so by Lemma 3.1.21, U_i extends to a geodesic from x_m to $*$). There may be several such geodesics, but they all have the same set of walls.

Lemma 3.2.10. *If* (γ_1, γ_2) *is a* Λ*-geodesic from* x_m *to* $*$ *with* $|\gamma_1| \geq 7N^2$*, then* γ_1 *can be rearranged to begin with exactly* U_i*, for some* $1 \leq i \leq k$*.*

Proof. Consider a van Kampen diagram (Figure 3.2.10) for the geodesic bigon determined by (γ_1, γ_2) and a Λ-geodesic from x_m to $*$ that begins with U_i. Let $\gamma_1 = (t_1, t_2, \ldots, t_s)$, where $s \geq 7N^2$. Let j be the smallest number such that the edge t_j shares a wall with an edge u_i of α_{7N+1}, for some $1 \leq i \leq k$ (such a j exists from Remark 3.2.9 (3) and because the lengths of $\alpha_1, \ldots, \alpha_{7N}$ are each less than N). Now,

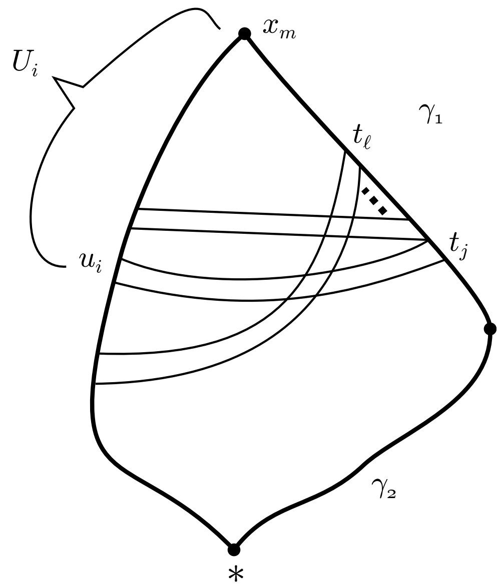

Figure 3.5: Lemma 3.2.10

choose ℓ maximal with $1 < \ell < j$ where the wall of t_ℓ is not on U_i. Clearly the wall of t_ℓ crosses the walls of $t_{\ell+1}, \dots, t_j$, so $\bar{t}_\ell$ commutes with $\bar{t}_{\ell+1}, \dots, \bar{t}_j$, and so $\bar{t}_1 \cdots \bar{t}_j$ can be rewritten $\bar{t}_1 \cdots \bar{t}_{\ell-1}\bar{t}_{\ell+1} \cdots \bar{t}_j\bar{t}_\ell$. Repeating this process, we obtain a rearrangement of γ_1 that begins with a rearrangement of U_i, which can be replaced by U_i. $\square$

We now have a finite number $k < N$ of "directions", given by our U_i, in which a bigon can be wide at x_m. The next lemma (3.2.11) is a fundamental combinatorial consequence of our no $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ hypothesis which allows us to refine this collection to at most two directions.

We will say that U_i and U_j **R-overlap** if there is an edge a of α_R that shares a wall with an edge of U_i and an edge of U_j. Let τ_a be a shortest Λ-geodesic based at x_m that can be extended to a geodesic ending at $*$ and whose last edge is in the same wall as a. Then U_i and U_j can be rearranged to begin with τ_a. We will now refine our list of U_i through the following five step process which at each application either terminates the process, or removes at least one of the U_i from our list and replaces all those that remain by geodesics with last edge in a wall of α_R where R begins at $7N$ and is reduced by one at each successive application:

1. Choose i minimal so that for some $j > i$, U_i and U_j R-overlap (by sharing some wall with an edge a of α_R). If no such i exists, our process stops.

2. Replace U_i with a shortest geodesic based at x_m and ending with an edge in the wall of a (which can be extended to a geodesic to $*$ by Lemma 3.1.21), and redefine u_i to be a.

3. Eliminate U_j from the list of U_ℓ.

4. For each remaining U_ℓ with $\ell \neq i$, choose an edge of U_ℓ in the same wall as an edge b_ℓ of α_R, replace U_ℓ with a shortest geodesic based at x_m and ending with an edge in the wall of b_ℓ, and redefine u_ℓ to be b_ℓ.

5. At this point each U_ℓ ends with an edge sharing a wall with an edge of α_R. If two U_ℓ end with edges in the same wall, remove one of them from the list. Now, relabel the remaining U_ℓ to form a list $U_1, \ldots, U_p$. Reduce R to $R - 1$.

When this process stops, no two U_i R-overlap, and each u_i shares a wall with an edge of α_{R+1}. Since U_i is a shortest geodesic with last edge in the wall of u_i, every geodesic from x_m to the end point of U_i ends with u_i. By the minimality of U_i and Remark 3.2.9 (3), if c is an edge of U_i in a wall of α_R, then $\overline{u}_i$ and $\overline{c}$ do not commute. Note that when this process stops, $6N < R \leq 7N$.

Lemma 3.2.11. *At most two U_i survive this reduction process.*

Proof. Suppose none of U_1, U_2, and U_3 R-overlap. Let a_1, a_2, a_3 be edges of U_1, U_2, U_3 respectively such that each a_i shares a wall with an edge of α_R. Since the process terminated, $\overline{a}_1, \overline{a}_2, \overline{a}_3$ are distinct and commute. But a_i does not commute with u_i for $i = 1, 2, 3$, and the pairs (a_i, u_i) all commute, so this gives a visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ in (W, S), a contradiction. □

We now have at most two directions U_1 and U_2 remaining. If there is no U_2, then to simplify notation for now, define $U_2 = U_1$.

If there is no geodesic extension of $\beta = (e_1, \ldots, e_m)$ that can be rearranged to form a bigon of width $16N^2$ with the down edge path of the diamond at x_m (Lemma 3.1.22) containing every wall of U_2, then we redefine $U_2 = U_1$, and similarly for U_1. If no geodesic extension of β can lead to a wide bigon in either direction, then an arbitrary filter (built as in the example in the beginning of this section) has "small" connected set limit set in ∂X.

Note that U_1 and U_2 have length at least $6N$. Now, if U_1 and U_2 share two walls with unrelated labels, then let $(\alpha_1, \alpha_2, \ldots)$ be a back combing from x_m to the endpoint of U_1, and choose an edge a in α_2 so that U_1 and U_2 both have edges in the same wall as a (such an edge exists by (5) of Remark 3.2.9). Let $U_1 = U_2$ be a shortest geodesic at x_m containing an edge in the same wall as a.

Remark 3.2.12. *If $U_1 \neq U_2$, then U_1 and U_2 share less than N walls, and the sets $lett(U_1) - (lett(U_1) \cap lett(U_2))$ and $lett(U_2) - (lett(U_1) \cap lett(U_2))$ commute.*

For this next remark, note that x_m is the $(m+1)^{st}$ vertex of β (since $*$ is the first).

Remark 3.2.13. *If $U_1 \neq U_2$, (β, γ) is a Λ-geodesic and γ' is some rearrangement of (β, γ) whose $(m+1)^{st}$ vertex is of distance at least $14N^2$ from x_m, then the down edge path τ at x_m of the diamond (Lemma 3.1.22) for these two geodesics can be rearranged to begin with either U_1 or U_2, by Lemma 3.2.10. Both cannot initiate rearrangements of τ, since otherwise there is a $(\mathbb{Z}_2 * \mathbb{Z}_2)^2$ in lett(τ), and the diamond at x_m containing τ determines a $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ in (W, S).*

Recall that $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ and $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$ are geodesics in Λ $(c+d)$-tracking two CAT(0) geodesics in X, and x_m is the endpoint of $(e_1, \ldots, e_m)$. Let x_i denote the endpoint of $(e_1, \ldots, e_i)$ where $i > m$, and y_i denote the endpoint of d_i where $i > m$. Set $U_1^{x_m} = U_1$ and construct $U_1^{x_i}$, $U_2^{x_i}$, $U_1^{y_i}$, $U_2^{y_i}$ exactly as above, by replacing x_m with x_i or y_i.

Let $\lambda = (\ell_1, \ell_2, \ldots, \ell_n)$ be a geodesic based at some x_i extending $(\beta, e_{m+1}, \ldots, e_i)$ (or based at y_i and extending $(\beta, d_{m+1}, \ldots, d_i)$), but not passing through e_{i+1} (d_{i+1}). Our goal is to classify the directions back toward $*$ at the endpoint of λ in a way that gives us some correspondence between our direction(s) at x_i (y_i) and the direction(s) at the endpoint of λ. We'll do this inductively, by corresponding directions at the endpoint of each edge of λ to the directions at the endpoint of the previous edge of λ. For what follows, let v denote the endpoint of ℓ_1.

1. If $U_1^{x_i} = U_2^{x_i}$ and $\overline{\ell}_1$ commutes with *lett*(U_1), then let $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1)$ be the edge path at v with the same labeling as $U_1^{x_i}$. Note that if $\overline{\ell}_1$ commutes with *lett*(U_1), then $\overline{\ell}_1 \notin$ *lett*(U_1), since (ℓ_1^{-1}, U_1) is geodesic.

2. If $U_1^{x_i} = U_2^{x_i}$ and $\overline{\ell}_1$ does not commute with $lett(U_1^{x_i})$, then set $U_1^{x_i}((\ell_1)) = U_2^{x_i}((\ell_1)) = (\ell_1{}^{-1}, U_1^{x_i})$.

3. If $U_1^{x_i} \neq U_2^{x_i}$, we construct directions from v back toward $*$ just as we've done from x_m back toward $*$. If there is only one direction V_1, set $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1) = V_1$. If there are two directions V_1 and V_2, but there is no geodesic extension of $(\beta, e_{m+1}, \ldots, e_i, \ell_1)$ that can lead to a $16N^2$ wide bigon in the V_2 direction at v, then set $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1) = V_1$ (and similarly for V_1). If there is no geodesic extension that can lead to a wide bigon in either direction, then building arbitrary fans, as in the example at the beginning of this section, fills in this section of the filter with rays in X that are sufficiently close to our original two rays in X. Otherwise, take a geodesic extension γ of $(\beta, e_{m+1}, \ldots, e_i, \ell_1)$ so that a rearrangement of $(\beta, e_{m+1}, \ldots, e_i, \ell_1, \gamma)$ gives a $16N^2$ wide bigon at v whose down edge path of the diamond at v (Lemma 3.1.22) begins with V_1. By Remark 3.2.13, the down edge path of the diamond at x_i for this bigon can be rearranged to begin with either $U_1^{x_i}$ or $U_2^{x_i}$ (but not both). If it's $U_1^{x_i}$ set $U_1^{x_i}(\ell_1) = V_1$ and $U_2^{x_i}(\ell_1) = V_2$, else set $U_1^{x_i}(\ell_1) = V_2$ and $U_2^{x_i}(\ell_1) = V_1$. It will be made clear by Lemma 3.2.17 that this choice does not depend on the choice of γ.

We now define $U_1^{x_i}((\ell_1, \ell_2))$ and $U_2^{x_i}((\ell_1, \ell_2))$ by replacing $U_1^{x_i}$ by $U_1^{x_i}(\ell_1)$ and $U_2^{x_i}$ by $U_2^{x_i}(\ell_1)$ in the above process, and continue repeating this process to define $U_1^{x_i}(\lambda)$ and $U_2^{x_i}(\lambda)$. Note that for any geodesic extension (λ_1, λ_2) of $(\beta, e_{m+1}, \ldots, e_i)$ that does not pass through e_{i+1}, if $U_1^{x_i}(\lambda_1) = U_2^{x_i}(\lambda_1)$, then $U_1^{x_i}((\lambda_1, \lambda_2)) = U_2^{x_i}((\lambda_1, \lambda_2))$.

Remark 3.2.14. *From here on, when we mention a geodesic extension* λ *of* $(\beta, e_{m+1}, \ldots, e_i)$ *(or* $(\beta, d_{m+1}, \ldots, d_i)$*), we assume* λ *does not pass through* e_{i+1} *(*d_{i+1}*).*

Lemma 3.2.15. *Let* λ *be a geodesic extension of* $(\beta, e_{m+1}, \ldots, e_i)$ *(or* $(\beta, d_{m+1}, \ldots, d_i)$*) with* $U_1^{x_i}(\lambda) \neq U_2^{x_i}(\lambda)$ *(*$U_1^{y_i}(\lambda) \neq U_2^{y_i}(\lambda)$*), and let* (γ_1, γ_2) *be any geodesic from the endpoint of* λ *to* $*$*. If* $|\gamma_1| \geq 7N^2$*, then* (γ_1, γ_2) *can be rearranged to begin with either* $U_1^{x_i}(\lambda)$ *or* $U_2^{x_i}(\lambda)$*.*

Proof. This follows from the proof of Lemma 3.2.10 and the construction of the $U_i^{x_i}(\lambda)$. □

Remark 3.2.16. *Remarks 3.2.12 and 3.2.13 apply to* $U_1^{x_i}(\lambda)$ *and* $U_2^{x_i}(\lambda)$*, whenever they are not equal.*

Lemma 3.2.17. *Suppose* λ *geodesically extends* $(\beta, e_{m+1}, \ldots, e_i)$*,* e *is an edge with* $(\beta, e_{m+1}, \ldots, e_i, \lambda, e)$ *geodesic, and* $U_1^{x_i}((\lambda, e)) \neq U_2^{x_i}((\lambda, e))$*, then* $U_j^{x_i}((\lambda, e))$ *and* $U_j^{x_i}(\lambda)$ *have at least* $6N - 3$ *walls in common.*

Proof. It suffices to show this for U_1 $(= U_1^{x_m})$ and $U_1(\ell_1)$ $(= U_1^{x_m}(\ell_1))$, as in the first step of our $U_i(\lambda)$ construction. Let γ be the geodesic extension of (β, ℓ_1) used in the construction of the $U_i(\ell_1)$, so that there is a rearrangement γ' of (β, ℓ_1, γ) whose $(m+2)^{nd}$ vertex is at least $16N^2$ from the endpoint of (β, ℓ_1). Let τ be the down edge path at the endpoint of ℓ_1 for the diamond for these two geodesics, as in Lemma 3.1.22. Note $|\tau| \geq 8N^2$. By Lemma 3.2.15 (and without loss of generality), τ can be rearranged to begin with $U_1(\ell_1)$. However, if τ has an edge in the same wall as ℓ_1 then τ can be rearranged to begin with ℓ_1, and so (ℓ_1, U_1). Otherwise, τ can be rearranged to begin with U_1, so either way every edge of U_1 shares a wall with an edge of τ. Let $(\alpha_1, \ldots, \alpha_{6N}, \ldots)$ be a back combing from x_m to $*$, choose an edge a_1 of α_{6N-1} that shares a wall with an edge of $U_1(\ell_1)$, and pick an edge a_2 of α_{6N-2} whose label does not commute with $\overline{a}_1$ (so a_2 also shares a wall with an edge of $U_1(\ell_1)$). Pick an edge b_1 of α_{6N-2} that shares a wall with an edge of U_1, and pick an edge b_2 of α_{6N-3} whose label doesn't commute with $\overline{b}_1$. If neither b_1 nor b_2 have their walls on $U_1(\ell_1)$, then the pair $\overline{a}_1, \overline{a}_2$ commutes with the pair $\overline{b}_1, \overline{b}_2$, and the up edge path at x_m for this diamond gives a third pair of unrelated elements that commute with the pairs $\overline{a}_1$, $\overline{a}_2$ and $\overline{b}_1$, $\overline{b}_2$, which is a contradiction. Thus the wall of b_2 must cross $U_1(\ell_1)$, and so $U_1(\ell_1)$ and U_1 have at least $6N-3$ walls in common. □

We claimed in the construction of the $U_j^{x_i}(\lambda)$ that Lemma 3.2.17 shows the association between $U_j^{x_i}$ and $U_j^{x_i}(\ell_1)$ is independent of the choice of γ. If the association depended on the choice of γ, then by the above proof, $U_1^{x_i}(\ell_1)$ would have $6N-3$ walls in common with both $U_1^{x_i}$ and $U_2^{x_i}$. By Remark 3.2.12, $lett(U_1^{x_i}(\ell_1))$ must then contain a $(\mathbb{Z}_2 * \mathbb{Z}_2)^2$, meaning the walls of $U_1^{x_i}(\ell_1)$ cannot all appear on the down edge path at x_m of the diamond for a wide bigon, which is a contradiction.

This next lemma gives an important correspondence between the directions $U_j^{x_i}(\lambda_1)$ and $U_j^{x_i}((\lambda_1, \lambda_2))$.

Lemma 3.2.18. *Let $(\lambda_1, \lambda_2, \lambda_3)$ be a geodesic extending $(\beta, e_{m+1}, \ldots, e_i)$ (not passing through x_{i+1}) with endpoint v, let τ be another Λ-geodesic from $*$ and v, let z_J and z_M denote the endpoints of λ_1 and λ_2, respectively, and suppose $U_1^{x_i}((\lambda_1, \lambda_2)) \neq U_2^{x_i}((\lambda_1, \lambda_2))$. Suppose $R \geq 14N^2$ and every vertex $z_J, z_{J+1}, \ldots, z_M$ of λ_2 is of Λ-*

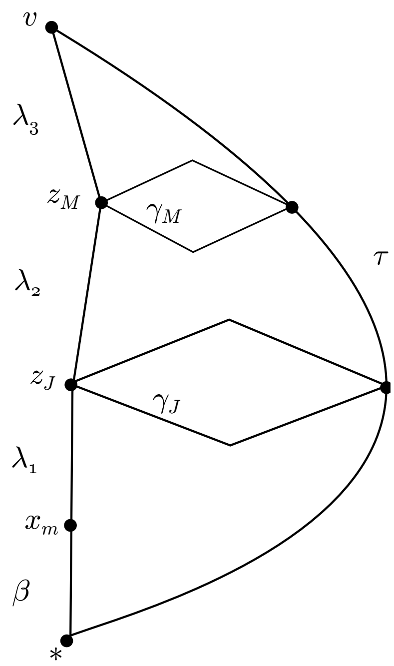

Figure 3.6: Lemma 3.2.18

distance at least R from τ. If the down edge path of the diamond at z_J for τ and $(\beta, e_{m+1}, \dots, e_i, \lambda_1, \lambda_2, \lambda_3)$ can be rearranged to begin with $U_1^{x_i}(\lambda_1)$, then the down edge path of the diamond at z_M for these geodesics can be rearranged to begin with $U_1^{x_i}((\lambda_1, \lambda_2))$ (and similarly for U_2).

Proof. It suffices to show this for $U_1^{x_m}((\lambda_1, \lambda_2)) = U_1((\lambda_1, \lambda_2))$ when $(\lambda_1, \lambda_2, \lambda_3)$ is a geodesic based at x_m, since the constructions are identical for each x_i. Let γ_J and γ_M be the down edge paths at z_J and z_M respectively of the diamonds for $(\beta, \lambda_1, \lambda_2, \lambda_3)$ and τ, as given by Lemma 3.1.22. (See Figure 3.6).

For each K with $J < K < M$, let λ_K denote the initial segment of (λ_1, λ_2) ending at z_K. Suppose γ_J can be rearranged to begin with $U_1(\lambda_1)$ but γ_M cannot be arranged to begin with $U_1((\lambda_1, \lambda_2))$. There is then K with $J < K < M$ where the down edge path γ_K at z_K of the diamond for these geodesics can be rearranged to begin with $U_1(\lambda_K)$ and the down edge path γ_{K+1} at z_{K+1} can be rearranged to begin with $U_2(\lambda_{K+1})$, by Lemma 3.2.15. By Lemma 3.2.17 and since $U_1(\lambda_{K+1}) \neq U_2(\lambda_{K+1})$, there is a pair of unrelated edge labels a_1, b_1 of $U_1(\lambda_K)$ that commute with some unrelated pair of labels a_2, b_2 from $U_2(\lambda_{K+1})$. Let ν^K and ν^{K+1} be the up edge paths of the diamonds at z_K and z_{K+1} respectively. From Lemma 3.1.22, these paths differ by at most two walls, and so they have two unrelated edge labels a_3 and b_3 in common. But then the pairs (a_i, b_i) must all commute, giving a visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ in W, a contradiction. □

The proof of the next lemma basically follows that of Lemma 5.5 of [16].

Lemma 3.2.19. *Let λ be a geodesic based at x_i extending $(\beta, e_{m+1}, \ldots, e_i)$ with endpoint v, and let s and t be vertices of Γ not in $B(v \to *)$. If (γ_1, γ_2) is any rearrangement of $(\beta, e_{m+1}, \ldots, e_i, \lambda)$ where $\langle lett(\gamma_2)\rangle$ is infinite, then there is a path from s to t of length at least two in Γ, none of whose vertices (except possibly s and t) are in $lk(lett(\gamma_2)) \cup B(v \to *)$.*

Proof. Since $(\beta, e_{m+1}, \ldots, e_i, \lambda)$ can be rearranged to end with γ_2, for $e \in B(v \to *)$, either $\bar{e} \in lett(\gamma_2)$ or $\bar{e} \in \mathrm{lk}(lett(\gamma_2))$. To see that $\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ does not separate $\Gamma(W, S)$, observe that otherwise G is not one-ended if $\langle \mathrm{lk}(lett(\gamma_2))\rangle$ is finite or $(\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *), \mathrm{lk}(lett(\gamma_2)))$ is a virtual factor separator for Γ if $\langle \mathrm{lk}(\gamma_2)\rangle$ is infinite.

If $s = t$ and $s \in \mathrm{lk}(lett(\gamma_2))$, then there is a vertex $a \in \Gamma$ adjacent to s with $a \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$, since $lett(\gamma_2)$ generates an infinite group and $B(v \to *)$

does not. If e is the edge between s and a, we use the path e followed by e^{-1}.

If $s = t$ and $s \notin \mathrm{lk}(lett(\gamma_2))$, then there is a vertex $a \in \Gamma$ adjacent to s with $a \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$, else $(\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *), \mathrm{lk}(lett(\gamma_2)))$ is a virtual factor separator for Γ. We construct the path as before.

If $s \neq t$, $s, t \notin \mathrm{lk}(lett(\gamma_2))$ and no such path exists, then $(\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *), \mathrm{lk}(lett(\gamma_2)))$ is a virtual factor separator for Γ. Note that if there is an edge e between s and t, we use the path e, e^{-1}, e to satisfy the length two requirement.

If $s \neq t$ and $s \in \mathrm{lk}(lett(\gamma_2))$, then there is a vertex $a \in \Gamma$ adjacent to s with $a \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$, since $lett(\gamma_2)$ generates an infinite group and $B(v \to *)$ does not. Now if $t \in \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ we obtain a b adjacent to t with $b \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ and we have a path between a and b as above (or, if $a = b$, we already had the path), or else we connect a and t as above. □

Remark 3.2.20. *Edge paths in Γ of the form (e, e^{-1}) and (e, e^{-1}, e) may seem unorthodox, but as in [16], they are combinatorially useful in the filter construction.*

Remark 3.2.21. *Note that $U_1^{x_i}(\lambda)^{-1}$ and $U_2^{x_i}(\lambda)^{-1}$ satisfy the hypotheses of γ_2 in the previous lemma.*

Recall the filter construction presented near the beginning of this section, and notice that Lemma 3.2.19 gives us more control during the fan construction process: instead of avoiding only $B(v \to *)$ when choosing paths in $\Gamma(W, S)$ to construct a fan based at v, we can avoid $B(v \to *)$ together with $\mathrm{lk}(lett(\gamma))$, where γ could potentially begin the down edge path of a diamond based at v. This is the key idea that allows us to keep the Cayley geodesics in our filter "straight" (in the CAT(0) sense), which makes the limit set of the filter small in ∂X. We'll now specify our choice of γ at each vertex v in the filter.

Recall that W acts geometrically on a CAT(0) space X giving a map $C : \Lambda \to X$ (respecting the action of W). The Γ geodesics $(\beta, e_{m+1}, e_{m+2}, \dots)$ and $(\beta, d_{m+1}, d_{m+2}, \dots)$ $(c+d)$-track two CAT(0) geodesics in X as in Lemma 3.1.27, and x_i denotes the endpoint of $(\beta, e_{m+1}, \dots, e_i)$, for $i \geq m$.

Definition 3.2.22. *For each vertex v of Λ, let ρ_v be a Λ-geodesic from $*$ to v such that $C(\rho_v)$ δ_1-tracks the X-geodesic from $C(*)$ to $C(v)$ (Lemma 3.1.26).*

Definition 3.2.23. *Suppose λ is a geodesic extending $(\beta, e_{m+1}, \dots, e_i)$ for some $i \geq m$, and y and z are vertices of λ with $d(z,*) > d(y,*) = k$. We say z is $R-$**wide in the τ direction** at y if the Λ-distance from y to $\rho_z(k)$ is at least R, and the down edge path at y of the diamond for $(\beta, e_{m+1}, \dots, e_i, \lambda)$ and ρ_z can be rearranged to begin with τ. If z is the endpoint of λ, we say λ is $R-$wide in the τ direction at y.*

Remark 3.2.24. *Using the notation in the definition, if $y = x_i$ and $d(\rho_z(i), x_i) \geq 14N^2$, then z is $14N^2$-wide in either the $U_1^{x_i}$ or $U_2^{x_i}$ direction at x_i, by Lemma 3.2.15.*

Let $\delta_0 = \max\{1, \delta_1, c+d\}$, where δ_1 is the tracking constant from Lemma 3.1.26, and c, d are the tracking constants from Lemma 3.1.27.

Let λ be a geodesic extending $(\beta, e_{m+1}, \dots, e_i)$ for some $i \geq m$. Set $A^i = U_1^{x_i}$, and define $A^i(\lambda)$ as follows:

1. If $U_1^{x_i}(\lambda) = U_2^{x_i}(\lambda)$, then set $A^i(\lambda) = U_1^{x_i}(\lambda)$.

2. If $U_1^{x_i}(\lambda) \neq U_2^{x_i}(\lambda)$ and λ is not at least $20N^2\delta_0$ wide in the $U_1^{x_i}$ or $U_2^{x_i}$ direction at x_i, then set $A^i(\lambda) = U_1^{x_i}(\lambda)$.

3. If $U_1^{x_i}(\lambda) \neq U_2^{x_i}(\lambda)$ and λ is at least $20N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i but less than $21N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i, then set $A^i(\lambda) = U_1^{x_i}(\lambda)$ (and similarly for $U_2^{x_i}$).

4. If $U_1^{x_i}(\lambda) \neq U_2^{x_i}(\lambda)$ and λ is at least $21N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i, then let λ_0 be the longest initial segment of λ such that λ_0 is at least $20N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i but not $21N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i. Then set $A^i(\lambda)$ to be a shortest geodesic based at the endpoint of λ containing an edge in each wall of $U_1^{x_i}(\lambda_0)$ (and similarly for $U_2^{x_i}$). By Lemma 3.1.21, $A^i(\lambda)$ geodesically extends to $*$.

At the endpoint of each such λ, we will construct fans avoiding $\mathrm{lk}(lett(A^i(\lambda))) \cup B((\beta, e_{m+1}, \ldots, e_i, \lambda))$ as in Lemma 3.2.19.

The next lemma explains why the last step in the above process is significant.

Lemma 3.2.25. *Let (λ_1, λ_2) be a geodesic extension of $(\beta, e_{m+1}, \ldots, e_i)$. Let τ be a shortest geodesic based at the endpoint of λ_2 containing an edge in each wall of $U_1^{x_i}(\lambda_1)$. If e is an edge that geodesically extends $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2)$ with $\overline{e} \notin lk(lett(\tau))$, then for any geodesic extension γ of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e)$ and any rearrangement γ' of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e, \gamma)$, no edge in $w(e)$ can appear on the up edge path at the endpoint of λ_1 of the diamond for $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e, \gamma)$ and γ' if the down edge path at the endpoint of λ_1 contains edges in all the walls of $U_1^{x_i}(\lambda_1)$.*

Proof. Suppose not; i.e. there is a geodesic extension γ of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e)$ and a rearrangement γ' of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e, \gamma)$ such that an edge e' of $w(e)$

appears on the up edge path at the endpoint of λ_1 of the diamond for these geodesics, and the down edge path at the endpoint of λ_1 contains edges in all the walls of $U_1^{x_i}(\lambda_1)$. Then $w(e') = w(e)$ crosses all walls of $U_1^{x_i}(\lambda_1)$. Let c_1 be an edge of τ such that $\overline{e}$ does not commute with $\overline{c}_1$. In particular, $w(c_1)$ is not a wall of $U_1^{x_i}(\lambda_1)$. By the definition of τ, there is an edge c_2 of τ, following c_1, such that $\overline{c}_1$ does not commute with $\overline{c}_2$. The walls $w(c_2)$ and $w(e)$ are on opposite sides of $w(c_1)$ (see Remark 3.1.18), so they do not cross. In particular, $w(c_2)$ is not a wall of $U_1^{x_i}(\lambda_1)$. Clearly we can continue picking c_i in such a way, but since the length of τ is finite, this process must stop. This gives the desired contradiction. □

Remark 3.2.26. *Note that Lemma 3.2.25 does not require that $U_1^{x_i}(\lambda_1) \neq U_2^{x_i}(\lambda_1)$ or $U_1^{x_i}((\lambda_1, \lambda_2)) \neq U_2^{x_i}((\lambda_1, \lambda_2))$. It is easy to show from our construction that if $U_1^{x_i}(\lambda_1) = U_2^{x_i}(\lambda_1)$, then τ (as defined in Lemma 3.2.25) has the same walls as $U_1^{x_i}((\lambda_1, \lambda_2)) = U_2^{x_i}((\lambda_1, \lambda_2))$, and so avoiding $lk(lett(A^i(\lambda)))$ has the effect that no wall of λ_2 can contain an edge of an up edge path at the end point of λ_1 for a diamond as described in Lemma 3.2.25.*

For a geodesic extension λ of $(\beta, d_{m+1}, \ldots, d_i)$, we define $A_d^i(\lambda)$ in the analagous way. To simplify notation, we will only deal with geodesic extensions λ of $(\beta, e_{m+1}, \ldots, e_i)$, except where necessary.

We now return to the filter construction. Set $a_1 = \overline{e}_{m+1}$ and $b_1 = \overline{d}_{m+1}$. We have $a_1, b_1 \notin B(x_m \to *)$, so let $a_1, t_1, \ldots, t_k, b_1$ be the vertices of a path of length at least 2 (Lemma 3.2.19) from a_1 to b_1 in $\Gamma(W, S)$, where each $t_i \notin \mathrm{lk}(lett(A^m)) \cup B(x_m \to *)$. We construct a fan in Λ as before (Figure 3.7).

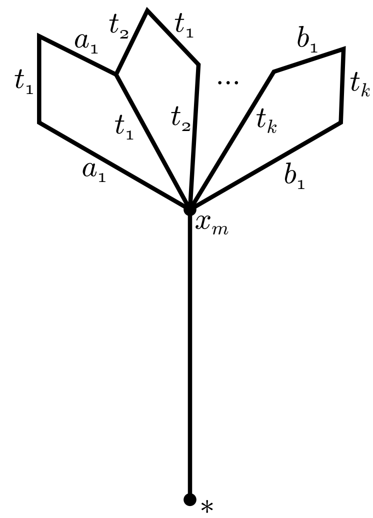

Figure 3.7: A fan, again

Definition 3.2.27. *The edges labeled a_1 and b_1 at x_m in the fan are called (respectively) the **left** and **right fan edges** at x_m. The edges labeled $t_1, \ldots, t_k$ at x_m are called **interior fan edges**. This fan is called the **first-level** fan, and the vertices at the endpoints of the edges based at x_m and labeled $x_{m+1}, t_1, \ldots, t_k, y_{m+1}$ are called **first-level** vertices.*

Now, let $a_2 = \bar{e}_{m+2}$, $b_2 = \bar{d}_{m+2}$ and let w_i be the edge at x_m labeled t_i for $1 \le i \le k$. Continue constructing the filter (Figure 3.8) by constructing fans avoiding $\mathrm{lk}(lett(A^m((w_i)))) \cup B((\beta, w_i))$ at the endpoint of each w_i, avoiding $\mathrm{lk}(lett(A^{m+1})) \cup B(x_{m+1} \to *)$ at x_{m+1}, and avoiding $\mathrm{lk}(lett(A_d^{m+1})) \cup B(y_{m+1} \to *)$ at y_{m+1}. Each of these fans is called a **second-level** fan, and each vertex of distance 2 from x_m (that will be the base vertex of a third-level fan) is called a **second-level** vertex.

It could occur that two edges of this graph share a vertex and are labeled the same; for example, we could have $t_1 = a_2$ in Figure 3.8. We do not identify these

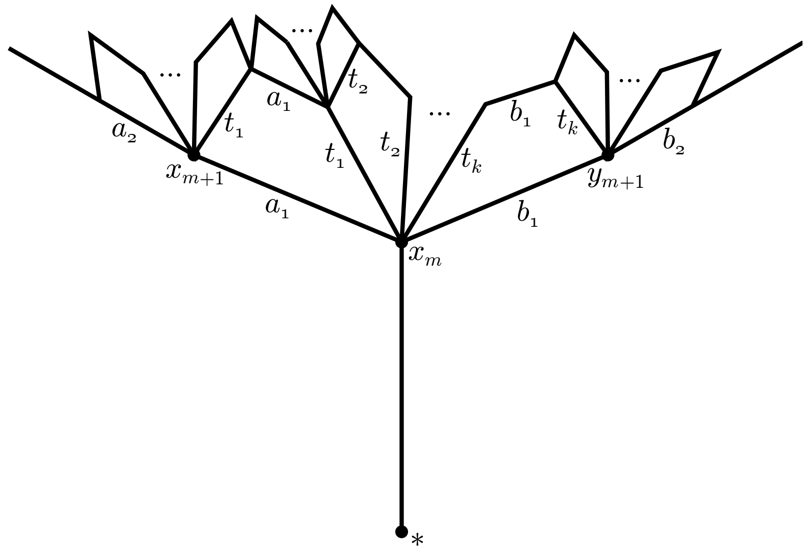

Figure 3.8: A filter, again

edges; instead, we will construct an edge path between them as described in Lemma 3.2.19 and extend the graph between them.

In order to build the third-level fans, we must specify geodesics from x_m to each vertex defined so far, so that $A^i(\lambda)$ is well-defined at each second-level vertex. We'll do this by picking the upper left edge from each first-level fan-loop to be a **non-tree** edge. This specifies a geodesic from x_m to each second-level vertex. We designate the upper right edge from each second-level fan as a non-tree edge, and continue alternating at each level, so the upper right edge of a n-th level fan is a non-tree edge if n is even, and the upper left edge of a n-th level fan is a non-tree edge if n is odd. By continuing to construct fans and designate non-tree edges, we construct a filter for our Λ-geodesics $(\beta, e_{m+1}, e_{m+2}, \dots)$ and $(\beta, d_{m+1}, d_{m+2}, \dots)$.

Recall that for an edge a of $\Lambda(W, S)$ with initial vertex y_1 and terminal vertex y_2, an edge e with initial vertex w_1 and terminal vertex w_2 is in the same wall as a if there is an edge path $(t_1, \ldots, t_n)$ in $\Lambda(W, S)$ based at w_1 so that $w_1\bar{t}_1 \cdots \bar{t}_n = y_1$ and $w_2\bar{t}_1 \cdots \bar{t}_n = y_2$, and $m(\bar{e}, \bar{t}_i) = 2$ for each $1 \leq i \leq n$. For two edges a and e of F, we say a and e are in the same **filter wall** if there is such a path $(t_1, \ldots, t_n)$ in F.

Remark 3.2.28. *The following are useful facts about a filter F for two such geodesics ((1)-(5) from [16]):*

1. *Each vertex v of F has exactly one or two edges beneath it, and there is a unique fan containing all edges (a left and right fan edge, and at least one interior edge) above v. We would not have this fact if we allowed association of same-labeled edges at a given vertex.*

2. *If a vertex of F has exactly one edge below it, then the edge is either e_i (for some i), d_i (for some i), or an interior fan edge.*

3. *If a vertex of F has exactly two edges below it, then one is a right fan edge (the one to the left), and one is a left fan edge, and both belong to a single fan loop.*

4. *F minus all non-tree edges is a tree containing $(\beta, e_{m+1}, e_{m+2}, \ldots)$ and $(\beta, d_{m+1}, d_{m+2}, \ldots)$ and all interior edges of all fans.*

5. *If T is the tree obtained from F by removing all non-tree edges, then there are no dead ends in T; i.e. for every vertex v of T, there is an interior edge extending from v.*

6. *No two consecutive edges of T not on $(\beta, e_{m+1}, e_{m+2}, \ldots)$ or $(\beta, d_{m+1}, d_{m+2}, \ldots)$ are right (left) fan edges.*

7. *If λ is a geodesic in F extending $(\beta, e_{m+1}, \dots, e_i)$ (and not passing through x_{i+1}), then λ shares at most one filter wall with $(e_{i+1}, e_{i+2}, \dots)$, and it is the wall of e_{i+1}.*

By rescaling, we may assume the image of each edge of Λ under C is of length at most 1 in X. Then for vertices v and w of Λ, $d_\Lambda(v, w) \geq d_X(C(v), C(w))$.

Lemma 3.2.29. *If $(\beta, e_{m+1}, \dots, e_i, \lambda)$ is geodesic in the tree T with endpoint v and $U_1^{x_i}(\lambda) \neq U_2^{x_i}(\lambda)$, then some point on the CAT(0) geodesic between $C(v)$ and $C(*)$ is within X-distance $101N^2\delta_0$ of $C(x_i)$.*

Proof. Suppose otherwise; then the endpoint v of λ is at least $100N^2\delta_0$ wide at x_i, and so suppose v is wide in the $U_1^{x_i}$ direction at x_i. Choose the last vertex w on λ such that w is between $20N^2\delta_0$ and $21N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i, so that every vertex between v and w on λ is at least $21N^2\delta_0$ wide in the $U_1^{x_i}$ direction at x_i. Let λ_w be the segment of λ starting at x_i and ending at w. We will show that v is wide in the $U_1^{x_i}(\lambda_w)$ direction at w and that v cannot be wide in the $U_1^{x_i}(\lambda_w)$ direction at w, obtaining a contradiction.

Claim 1: The vertex v is wide in the $U_1^{x_i}(\lambda_w)$ direction at w.

Recall that ρ_w and ρ_v are Λ-geodesics δ_1-tracking the X-geodesics from $C(*)$ to $C(w)$ and $C(v)$ respectively. By CAT(0) geometry, ρ_v is at least $75N^2\delta_0$ wide at w, since w is less than $21N^2\delta_0$ wide at x_i. Consider Figure 3.9, with diamonds for these geodesics as in Lemma 3.1.22. Let y be the endpoint of the up edge path of the diamond at x_i

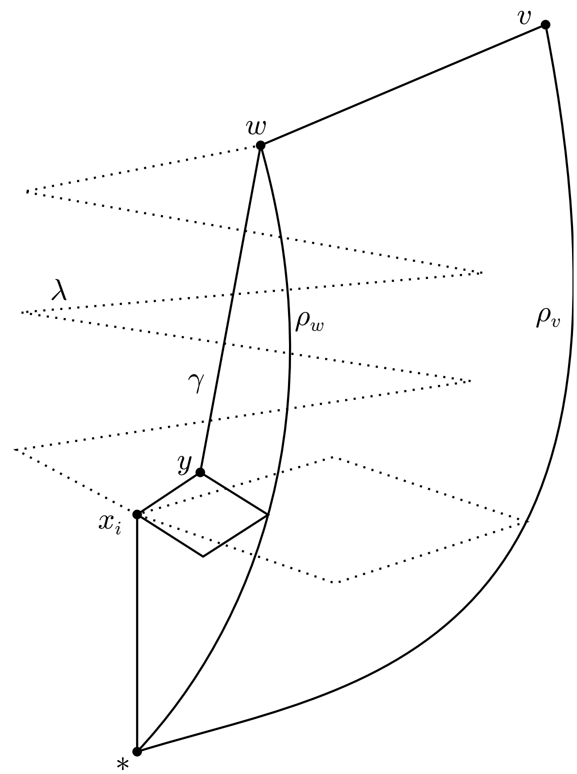

Figure 3.9: Lemma 3.2.29, Claim 1

for ρ_w and $(\beta, e_{m+1}, \ldots, e_i, \lambda_w)$, and let γ_0 be any geodesic from y to w. A simple van Kampen diagram argument shows that there is a rearrangement γ_1 of γ_0 such that if $w(c_1), w(c_2), \ldots, w(c_n)$ are the walls of the edges of γ_0, then γ_1 crosses these walls in the same order as ρ_w. Let γ be any geodesic from x_i to y followed by γ_1. By Lemma 3.1.22, it is clear that each vertex x of γ is of Λ-distance less than $21N^2\delta_0$ from the corresponding vertex x' of ρ_w (satisfying $d(x,*) = d(x',*)$ in Λ). Therefore γ is of Λ-distance at least $54N^2\delta_0$ from ρ_v. Now, if no vertex of λ_w is within Λ-distance $14N^2$ of the corresponding vertex of ρ_v, then by Lemma 3.2.18 (with λ_1 trivial), v is $75N^2\delta_0$ wide in the $U_1^{x_i}(\lambda_w)$ direction at w, as claimed. Suppose there are vertices of λ_w within Λ-distance $14N^2$ of the corresponding vertices on ρ_v, and list the consecutive points

$z_1, \ldots, z_\ell$ of λ_w (with z_1 closest to x_i) where each z_j has the property that if g_j and m_j are the points on γ and ρ_v respectively with $d(z_j, *) = d(g_j, *) = d(m_j, *)$, then $|d(z_j, g_j) - d(z_j, m_j)| < N$ (so each z_j is almost Λ-equidistant from its corresponding points on γ and ρ_v). Let λ_{z_j} denote the initial segment of λ_w ending at each z_j. Now, ρ_v (equivalently v) is wide in the $U_1^{x_i}(\lambda_{z_1})$ direction at z_1, since λ_w has not yet passed close to ρ_v. Now consider the down edge path of the diamond at z_1 for λ_w and γ; this path is of length more than $7N^2$ and must have edges in all the walls of $U_2^{x_i}(\lambda_{z_1})$ (Lemma 3.2.15), else by Lemma 3.1.24, γ and ρ_v would be close. Now, if ρ_v is wide in the $U_2^{x_i}(\lambda_{z_2})$ direction at z_2, then the down edge path at z_2 for the diamond for λ_w and γ must have edges in all the walls of $U_1^{x_i}(\lambda_{z_2})$; however, by Lemma 3.2.18, at most one of these directions could have switched, since λ does not pass close to one of ρ_v or γ between z_1 and z_2. Continuing this argument along the z_i shows that v is wide in the $U_1^{x_i}(\lambda_w)$ direction at w, as claimed.

Claim 2: The vertex v cannot be wide in the $U_1^{x_i}(\lambda_w)$ direction at w.

Note that no interior fan edges on λ between v and w can have walls appearing on the up edge path of a $U_1^{x_i}(\lambda_w)$ diamond at w by Lemma 3.2.25, since all of these edges have labels chosen to avoid $\text{lk}(lett(U_1^{x_i}((\lambda_w, ...))))$. Also note that if the first edges of λ after λ_w are a right fan edge followed by a left fan edge, the left fan edge shares a wall with an interior fan edge adjacent to λ, and so it was also chosen to avoid $\text{lk}(lett(U_1^{x_i}((\lambda_w, ...))))$, and so no edge in its wall can appear on a $U_1^{x_i}(\lambda_w)$ diamond at w (and similarly for a left fan edge followed by right fan edge). The same analysis holds for any right or left fan edge appearing after an interior fan edge (except for at most one edge of λ, which could share a wall with a right/left fan edge based at w). Thus the only way λ can have enough edges in the same walls as edges on the up edge path of a $U_1^{x_i}(\lambda_w)$ diamond is if a large sequence of the edges of λ immediately after λ_w are all right fan edges or all left fan edges, which cannot happen by (6) of Remark 3.2.28. Thus v is not wide in the $U_1^{x_i}(\lambda_w)$ direction at x_i, which gives the desired contradiction. □

Lemma 3.2.30. *If λ is a geodesic in the tree T with endpoint v that extends $(\beta, e_{m+1}, \ldots, e_i)$ and $U_1^{x_i}(\lambda) = U_2^{x_i}(\lambda)$, then some point on the CAT(0) geodesic between $C(v)$ and $C(*)$ is within X-distance $118N^2\delta_0$ of $C(x_i)$.*

Proof. Let λ_w be the shortest initial segment of λ such that $U_1^{x_i}(\lambda_w) = U_2^{x_i}(\lambda_w)$, and let w be the endpoint of λ_w. By Lemma 3.2.29, the CAT(0) geodesic between $C(w)$ and $C(*)$ comes within X-distance $101N^2\delta_0$ of $C(x_i)$. Note that if the CAT(0) geodesic between $C(v)$ and $C(*)$ is more than $17N^2\delta_0$ from $C(w)$, then v (equivalently λ) is at least $16N^2\delta_0$ wide in the $U_1^{x_i}(\lambda_w)$ direction at w. When $U_1^{x_i}(\lambda_w) = U_2^{x_i}(\lambda_w)$ we have the following cases:

Case 1: No geodesic extension of $(\beta, e_{m+1}, \ldots, e_i, \lambda_w)$ leads to a bigon $16N^2$ wide at w.

In this case, λ is not $16N^2$ wide in any direction at w, so by CAT(0) geometry, some point on the CAT(0) geodesic between $C(v)$ and $C(*)$ is within X-distance $118N^2\delta_0$ of $C(x_i)$..

Case 2: For any geodesic μ from $*$ to the endpoint of $(\beta, e_{m+1}, \ldots, e_i, \lambda)$, if the bigon determined by μ and $(\beta, e_{m+1}, \ldots, e_i, \lambda)$ is $16N^2$ wide at w, then it is wide in the $U_1^{x_i}(\lambda_w)$ direction at w.

From Lemma 3.2.25, Remark 3.2.26, and our filter construction, we know that any interior edge on λ after w cannot have its wall on the up edge path of a $U_1(\lambda_w)$ diamond at w. If the first edges of λ after w are a right fan edge followed by a left fan edge, the left fan edge shares a wall with an interior fan edge adjacent to λ, and so the left fan edge also cannot have an edge in its wall on the up edge path of a $U_1(\lambda_w)$ diamond at w. The same analysis holds for any left or right fan edge following an interior fan edge (except for at most one edge of λ, which could share a wall with a right/left fan edge based at w). Thus by (6) of Remark 3.2.28, λ cannot be $16N^2$ wide in the $U_1(\lambda_w)$ direction at w, so some point on the CAT(0) geodesic between $C(v)$ and $C(*)$ is within X-distance $118N^2\delta_0$ of $C(x_i)$. $\square$

Suppose X is a CAT(0) space, $* \in X$ a base point and $B_n(*)$ the open n-ball about $*$. Let $\overline{X}$ be the compact metric space $X \cup \partial X$. If F is a filter in X, let $\overline{F}$ be the closure of F in $\overline{X}$. Since F is connected, $\overline{F}$ is connected. Since F is one-ended, $\overline{F} - F$ (the limit set of F) is contained in C_n, a component of $\overline{F} - B_n(*)$, for each $n > 0$. Then $\overline{F} - F = \cap_{n=1}^{\infty} C_n$ is the intersection of compact connected subsets of a metric space and so is connected.

Theorem 3.2.31. *Suppose (W, S) is a one-ended right-angled Coxeter system, $\Gamma(W, S)$ contains no visual subgroup isomorphic to $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$, and W does not visually split as $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$. Then W has locally connected boundary iff $\Gamma(W, S)$ does not contain a virtual factor separator.*

Proof. If (W, S) has a virtual factor separator, then by [16], W has non-locally connected boundary. Suppose W acts geometrically on a CAT(0) space X, and let r be a CAT(0) geodesic ray based at a point $*$ of X. Let $\epsilon > 0$ be given. We find δ such that if s is a geodesic ray within δ of r in ∂X, then our filter for r and s has (connected) limit set of diameter less than ϵ in ∂X. In what follows, the constants c and d are the tracking constants from Lemma 3.1.27, δ_1 is the tracking constant from Lemma 3.1.26, and $\delta_0 = max\{1, \delta_1, c+d\}$. Recall $C : \Lambda(W, S) \to X$ W-equivariantly, and suppose for simplicity $C(*) = *$. Choose M large enough so that for all $m \geq M - c - d$, if s is an X-geodesic ray based at $*$ within $120N^2\delta_0$ of $C(\beta(m))$ for any Cayley geodesic β that δ_0-tracks r, then r and s are within $\epsilon/2$ in ∂X. Choose δ so that if r and s are within δ in ∂X, then r and s satisfy $d(r(M), s(M)) < 1$. Now, if r and s are within δ in ∂X, by Lemma 3.1.27, r and s can be δ_0-tracked by Cayley geodesics α_r and α_s sharing an initial segment of length at least $M - c - d$. Let $m = M - c - d$ and denote the "split point" of α_r and α_s by x_m, as in the filter construction. Similarly, let $\alpha_r(i) = x_i$ and $\alpha_s(i) = y_i$ for $i \geq m$. By the previous two lemmas, for any vertex v in the filter F for α_r and α_s, the X-geodesic from $C(v)$ to $*$ passes within $118N^2\delta_0$ of $C(x_i)$ (or $C(y_i)$), where $i \geq m$. By CAT(0) geometry, this geodesic must also pass within $119N^2\delta_0$ of $C(x_m)$. Thus every geodesic ray in the limit set of $C(F)$ is within $\epsilon/2$ of r in ∂X, so this set has diameter less than ϵ in ∂X. □

3.3 Two interesting examples

Let (W, S) be the (one-ended) right-angled Coxeter system with presentation graph Γ given by Figure 3.10. For what follows, let $A = \{a_1, a_2\}$, $B = \{b_1, b_2\}$,

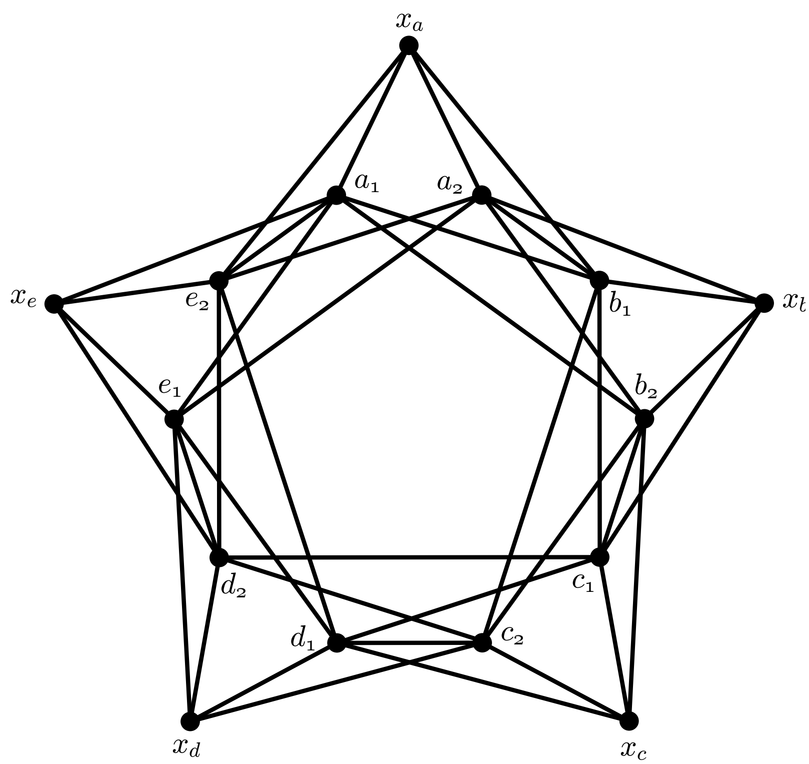

Figure 3.10: Example 1

$C = \{c_1, c_2\}$, $D = \{d_1, d_2\}$ and $E = \{e_1, e_2\}$.

It is not hard to check that Γ has no virtual factor separator, (W, S) does not visually split as a direct product and that (W, S) has no visual $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$. However, Γ contains product separators: for example, $A \cup D$ commutes with E, and $A \cup D \cup E$ separates x_e from the rest of Γ.

Corollary 5.7 of [15] gives specific conditions for when the boundary of a right-angled Coxeter group is non-locally connected:

Corollary 3.3.1. *Suppose (W, S) is a right-angled Coxeter system. Then W has non-locally connected boundary if there exist $v, w \in S$ with the following properties:*

1. *v and s are unrelated in W, and*
2. *$lk(v) \cap lk(w)$ separates $\Gamma(W, S)$, with at least one vertex in $S - lk(v) \cap lk(w)$ other than v and w.*

In particular, they show that if such v, w exist, then $(vw)^\infty$ is a point of non-local connectivity in any CAT(0) space acted on geometrically by W. Note that if v, w exist as in this corollary, then $(\mathrm{lk}(v) \cap \mathrm{lk}(w), \mathrm{lk}(v) \cap \mathrm{lk}(w))$ is a virtual factor separator for $\Gamma(W, S)$.

Let $G_1 = \langle S - x_a \rangle$. Note that $\mathrm{lk}(e_1) \cap \mathrm{lk}(e_2) = A \cup D \cup \{x_e\}$ separates e_2 from the rest of $\Gamma(G_1, S - \{x_a\})$, so G_1 has non-locally connected boundary, with $(e_1 e_2)^\infty$ a point of non-local connectivity for G_1. Similarly, let $Q = A \cup B \cup E$ and let $G_2 = \langle Q \cup \{x_a\} \rangle$. Then $\mathrm{lk}(e_1) \cap \mathrm{lk}(e_2) = A \cup D \cup \{x_e\}$ separates e_1 from the rest of $\Gamma(G_2, Q \cup \{x_a\})$, and so G_2 also has non-locally connected boundary, also with $(e_1 e_2)^\infty$ a point of non-local connectivity. Note that we now have $W = G_1 *_Q G_2$,

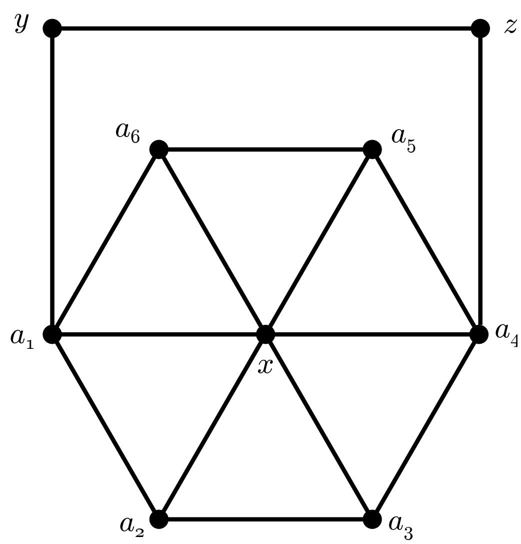

Figure 3.11: Example 2

where ∂G_1 and ∂G_2 have $(e_1e_2)^\infty$ as a point of non-local connectivity and Q contains e_1 and e_2, so it would seem that ∂W should also have $(e_1e_2)^\infty$ as a point of non-local connectivity. However, our theorem implies W has locally connected boundary.

For our second example consider the right-angled Coxeter group (G, S) with presentation graph given by Figure 3.11. Let $A = \{a_1, \ldots, a_6\}$ and (G', S') have the same presentation graph as (G, S) but with each vertex v labeled v'. Let (W, S) be the right-angled Coxeter group of the amalgamated product $G *_{A=A'} G'$ (where $S = \{x, x', y, y'z, z', A\}$, and $\{x, x'\}$ commutes with A). Both G and G' are word hyperbolic and one-ended so they have locally connected boundary. The subgroup $\langle A \rangle$ of G is virtually a hyperbolic surface group and so determines a circle boundary in the boundary of G. Still, W has non-locally connected boundary since (A, A) is a virtual factor separator for (W, S).

Aside from being rather paradoxical, these examples show that boundary local connectivity of right-angled Coxeter groups is not accessible through graphs of groups techniques.

Chapter IV

PATH CONNECTIVITY OF BOUNDARIES OF RIGHT-ANGLED ARTIN GROUPS

4.1 Artin group and cube comblex preliminaries

Definition 4.1.1. *Given a (undirected) graph Γ with vertex set $S = a_1, \ldots, a_n$, the corresponding* ***right-angled Artin group*** *A_Γ is the group with presentation*

$$\langle a_1, \ldots, a_n \mid [a_i, a_j] \text{ if } i < j \text{ and } \{a_i, a_j\} \text{ is an edge of } \Gamma \rangle.$$

We call Γ the ***presentation graph*** *for A_Γ.*

Definition 4.1.2. *If A_Γ is a right-angled Artin group with Cayley graph Λ_Γ, let $\overline{e} \in S$ be the label of the edge e of Λ_Γ. An* **edge path** *$\alpha \equiv (e_1, e_2, \ldots, e_n)$ in Λ_Γ is a map $\alpha : [0, n] \to \Lambda_\Gamma$ such that α maps $[i, i+1]$ isometrically to the edge e_i. For α an edge path in Λ_Γ, let lett$(\alpha) \equiv \{\overline{e}_1, \ldots, \overline{e}_n\}$, and let $\overline{\alpha} \equiv \overline{e}_1 \cdots \overline{e}_n$. If β is another geodesic with the same initial and terminal points as α, then call β a* **rearrangement** *of α.*

Lemma 4.1.3. *If $w = g_1 \ldots g_k$ is a word in A_Γ (with each $g_i \in S^{\pm}$) that is not of minimal length, then two letters of $g_1 \ldots g_k$* ***delete****; that is, for some $i < j$, $g_i = g_j^{-1}$, the sets $\{g_i, g_j\}$ and $\{g_{i+1}, \ldots, g_{j-1}\}$ commute, and $w = g_1 \ldots g_{i-1} g_{i+1} \ldots g_{j-1} g_{j+1} \ldots g_k$.*

Proof. Let $w = h_1 \ldots h_m$ be a minimal length word representing w, and draw a van Kampen diagram D for the loop $g_1 \ldots g_k h_m^{-1} \ldots h_1^{-1}$. For each boundary edge e_i corresponding to a g_i, trace a band across the diagram by picking the opposite edge of e_i in the relation square containing e_i, and continuing to pick opposite edges (without going backwards). Note that such a band cannot cross itself, and so this band must end on another boundary edge of D. Since $k > m$, there is some boundary edge e_i corresponding to some g_i that has its band B end on a boundary edge e_j corresponding to g_j, with $i < j$. Note this implies $g_i = g_j^{-1}$. Now, either all the bands corresponding to $g_{i+1}, \ldots, g_{j-1}$ cross B (implying each of $g_{i+1}, \ldots, g_{j-1}$ commutes with g_i and g_j), or some band corresponding to one of $g_{i+1}, \ldots, g_{j-1}$ ends on a boundary edge corresponding to another of $g_{i+1}, \ldots, g_{j-1}$. Picking an "innermost" such band and repeating the above argument gives the desired result. □

Remark 4.1.4. *Note that the bands in the van Kampen diagram D share the same labels along their 'sides'. This means that deleting the band B from the diagram and matching up the separate parts of what remains (along paths with the same labels) gives a van Kampen diagram D' for the loop*

$w = g_1 \ldots g_{i-1} g_{i+1} \ldots g_{j-1} g_{j+1} \ldots g_k h_m^{-1} \ldots h_1^{-1}$.

Remark 4.1.5. *Given a non-geodesic edge path $(e_1, \ldots, e_k)$ in the Cayley graph Λ_Γ for A_Γ, we say edges e_i and e_j delete if their corresponding labels delete in the word $\overline{e_1} \ldots \overline{e_k}$.*

Lemma 4.1.6. *Suppose A_Γ is a right-angled Artin group, and (α_1, α_2) and (β_1, β_2) are geodesics between the same two points in in the Cayley graph Λ_Γ for A_Γ. There exist geodesics $(\gamma_1, \tau_1), (\gamma_1, \delta_1), (\delta_2, \gamma_2)$, and (τ_2, γ_2) with the same end points as $\alpha_1, \beta_1, \alpha_2, \beta_2$ respectively, such that (see Figure 4.1):*

1. *τ_1 and τ_2 have the same labels,*

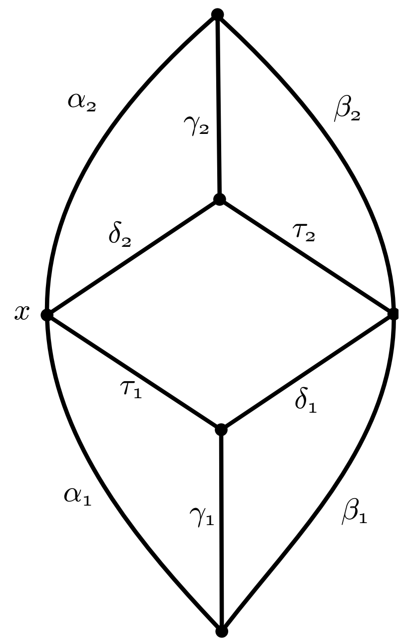

Figure 4.1: Lemma 4.1.6

2. δ_1 *and* δ_2 *have the same labels, and*

3. $lett(\tau_1)$ *and* $lett(\delta_1)$ *are disjoint and commute.*

Furthermore, the paths (τ_1^{-1}, δ_1) *and* (δ_2, τ_2^{-1}) *are geodesic.*

Proof. Let D be a van Kampen diagram for the loop $(\alpha_1, \alpha_2, \beta_2^{-1}, \beta_1^{-1})$, and let $\alpha_1 = (a_1, \ldots, a_k)$, $\beta_1 = (b_1, \ldots, b_m)$. Let $a_{i_1}, \ldots, a_{i_j}$ be (in order) the edges of α_1 whose bands in D end on β_1. Note that by Lemma 4.1.3, β_1 can be rearranged to begin with an edge labeled $\overline{a_{i_1}}$, since a_{i_1} and b_{ℓ_1} delete in (α_1^{-1}, β_1) for some ℓ_1 and all the bands based at $b_1, \ldots, b_\ell, a_1, \ldots, a_{i_1-1}$ cross the band based at a_{i_1} and ending at b_{ℓ_1}. Similarly, β_1 can be rearranged to begin with an edge labeled $\overline{a_{i_1}}$ followed by an edge labeled $\overline{a_{i_2}}$, and continuing in this manner, we obtain a rearrangement of β_1 that begins with $\gamma_1 = (\overline{a_{i_1}}, \ldots, \overline{a_{i_j}})$, and we let δ_1 be the remainder of this rearrangment. This argument also implies α_1 can be rearranged to begin with γ_1, and we let τ_1 be the remainder of this rearrangement. Note that if e is an edge of τ_1, no edge of δ_1 is labeled $\overline{e}$ or $\overline{e}^{-1}$, since bands with those labels must have crossed in D. We obtain γ_2, τ_2 and δ_2 in the analogous way from α_2 and β_2, and note that in a van Kampen diagram B' for $(\tau_1, \delta_2, \tau_2^{-1}, \delta_1^{-1})$, no band based on τ_1 can end on δ_2, since (τ_1, δ_2) is geodesic, and no band based on τ_1 ends on δ_1, since τ_1 and δ_1 share no labels or inverse labels. Therefore all bands on τ_1 end on τ_2, so τ_1 and τ_2 have the same labels, as do δ_1 and δ_2. □

Definition 4.1.7. *Under the hypotheses of the previous lemma, we call τ_1 the **down edge path** at x, and we call δ_2 the **up edge path** at x. If α_1 and β_1 have the same length, we call Figure 4.1 the **diamond** at x for (α_1, α_2) and (β_1, β_2).*

Definition 4.1.8. *P_4 is the (undirected) graph on four vertices a, b, c, d, with edge set $\{\{a, b\}, \{b, c\}, \{c, d\}\}$.*

Definition 4.1.9. *The **union** of two graphs (V_1, E_1) and (V_2, E_2) is the graph $(V_1 \cup V_2, E_1 \cup E_2)$.*

Definition 4.1.10. *The **join** of two graphs (V_1, E_1) and (V_2, E_2) is the graph $(V_1 \cup V_2, E_1 \cup E_2 \cup (V_1 \times V_2))$.*

Definition 4.1.11. *A graph is **decomposable** if it can be expressed as joins and unions of isolated vertices.*

The following is Theorem 9.2 in [14].

Theorem 4.1.12. *A finite graph G is decomposable iff it does not contain P_4 as an induced subgraph.*

In particular, if a connected graph G does not contain P_4 as an induced subgraph, then it must split as the join $G_1 \vee G_2$, for some subgraphs G_1, G_2 of G.

Definition 4.1.13. *For a graph Γ and a vertex a of Γ,*

$$lk(a) = \{b \in \Gamma \mid \{a, b\} \text{ is an edge of } \Gamma\}.$$

Let Λ_Γ be the Cayley graph for the group A_Γ.

Definition 4.1.14. *The **standard complex** $\mathcal{S}_\Gamma$ for the group A_Γ is the CAT(0) cube complex whose one-skeleton is Λ_Γ, with each cube given the geometry of $[0, 1]^n$ for the appropriate n.*

If the graph Γ splits as a non-trivial join $\Gamma_1 \vee \Gamma_2$, then the group A_Γ splits as the direct product $A_{\Gamma_1} \times A_{\Gamma_2}$, and so we have $\mathcal{S}_\Gamma \cong \mathcal{S}_{\Gamma_1} \times \mathcal{S}_{\Gamma_2}$. Proposition 2.0.10 then gives that $\partial\mathcal{S}_\Gamma \cong \partial\mathcal{S}_{\Gamma_1} * \partial\mathcal{S}_{\Gamma_2}$. Any non-trivial spherical join is path connected, and so $\partial\mathcal{S}_\Gamma$ is path connected.

For more on cube complexes and the definitions below, see [21].

Definition 4.1.15. *A **midcube** in a cube complex C is the codimension 1 subspace of an n-cube $[0, 1]^n$ obtained by restricting exactly one coordinate to $\frac{1}{2}$. A **hyperplane** is a connected nonempty subspace of C whose intersection with each cube is either empty or consists of one of its midcubes.*

Lemma 4.1.16. *If D is a hyperplane of the cube complex C, then $C - D$ has exactly two components.*

Given a graph Γ, a vertex v of Γ, and the corresponding standard complex $\mathcal{S}_\Gamma$, note that if a hyperplane of $\mathcal{S}_\Gamma$ intersects an edge of $\mathcal{S}_\Gamma$ with label v, then every edge intersected by this hyperplane is also labeled v. Thus we can refer to hyperplanes in $\mathcal{S}_\Gamma$ as v-hyperplanes, for v a vertex of Γ. If x is a vertex of $\mathcal{S}_\Gamma$, then xv and x are separated by a v-hyperplane D. Let $x\mathcal{S}_{lk(v)}$ denote the cube complex generated by the coset $x\langle lk(v)\rangle$; then D and $x\mathcal{S}_{lk(v)}$ are isometric and parallel, of distance $\frac{1}{2}$ apart.

A proof of the following can be found in Section 3 of [11]

Lemma 4.1.17. *There is a bound $\delta > 0$ such that if α is a CAT(0) geodesic path in $\mathcal{S}_\Gamma$, then there is a Cayley graph geodesic path β in Λ_Γ (contained naturally in $\mathcal{S}_\Gamma$) such that each vertex of β is within distance δ of α, and each point of α is within δ of a vertex of β.*

4.2 Non-path-connectivity of some right-angled Artin boundaries

The goal of this section is to prove the following theorem:

Theorem 4.2.1. *Let Γ be a connected graph. Suppose Γ contains an induced subgraph $(\{a, b, c, d\}, \{\{a, b\}, \{b, c\}, \{c, d\}\})$ (isomorphic to P_4), and there are subsets $B \subset lk(c)$ and $C \subset lk(b)$ with the following properties:*

1. *B separates c from a in Γ, with $d \notin B$;*
2. *C separates b from d in Γ, with $a \notin C$;*

3. $B \cap C = \emptyset$.

Then $\partial \mathcal{S}_\Gamma$ *is not path connected.*

In fact, we prove a stronger result, with the hypothesis $B \cap C = \emptyset$ replaced by the statement of Claim 4.2.7. For the remainder of this section, suppose $a, b, c, d \in \Gamma$, $B \subset lk(c)$, and $C \subset lk(b)$ are as in Theorem 4.2.1. Note that $b \in B$, $c \in C$. We wish to consider the following rays in Λ_Γ (equivalently in $\mathcal{S}_\Gamma$), based at the identity vertex $*$:

$$r = cdab(cb)^2cdab(cb)^6 \cdots = \prod_{i=1}^{\infty} (cb)^{k_i} cdab$$

and

$$s = dbcb^2adbc(b^2c)^2b^2adbc(b^2c)^6b^2a \cdots = \prod_{i=1}^{\infty} dbc(b^2c)^{k_i}b^2a$$

where the k_i are defined recursively with $k_0 = -1$, $k_{i+1} = 2k_i + 2$.

Define the following vertices of r, for $n \geq 0$:

$$v_n = \left(\prod_{i=1}^{n} (cb)^{k_i} cdab \right) (cb)^{k_{n+1}} cd$$

$$v'_n = v_n a$$

Define the following vertices of s, for $n \geq 0$:

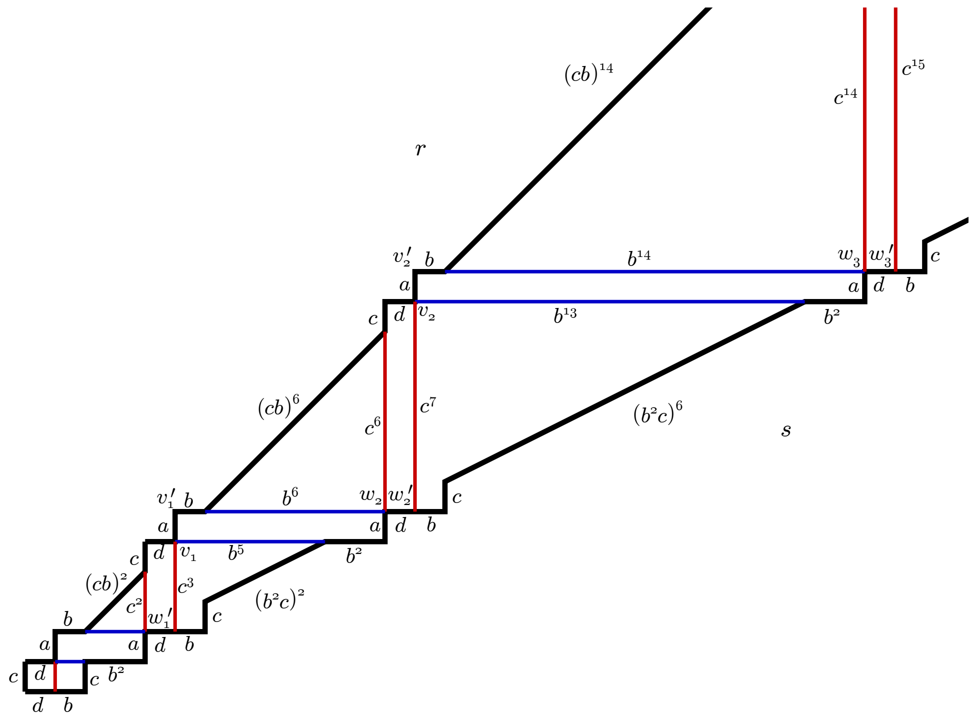

Figure 4.2: The rays r and s

$$w_n = \left(\prod_{i=1}^{n} dbc(b^2c)^{k_i} b^2 a \right)$$

$$w_n' = w_n d$$

We have $v_0 = cd$, $v_0' = cda$, $v_1 = cdab(cb)^2cd$, $w_0 = *$, $w_0' = d$, $w_1 = dbcb^2a$. It will be helpful to refer to Figure 4.2 for many of the claims that follow.

The following is proved in [7].

Claim 4.2.2. *For $n \geq 0$, $v_n = w_n' c^{k_{n+1}+1}$ and $v_n' b^{k_{n+2}+1} = w_{n+1}$.*

Since $b \in B$ and $c \in C$, we then have $v_n\langle C\rangle = w_n'\langle C\rangle$ and $w_n\langle B\rangle = v_{n-1}'\langle B\rangle$.

If Q_c denotes the component of c in $\Gamma - B$, and Q_b denotes the component of b in $\Gamma - C$, then A_Γ can be represented as $\langle Q_c \cup B\rangle *_B \langle \Gamma - Q_c\rangle$ or $\langle Q_b \cup C\rangle *_C \langle \Gamma - Q_b\rangle$, and so at each vertex x of Λ_Γ, the cosets $x\langle B\rangle$ and $x\langle C\rangle$ separate Λ_Γ. Therefore, if $x\mathcal{S}_B$ and $x\mathcal{S}_C$ denote the cube complexes generated by $\langle B\rangle$ and $\langle C\rangle$ respectively at a vertex x of $\mathcal{S}_\Gamma$, then $x\mathcal{S}_B$ and $x\mathcal{S}_C$ separate $\mathcal{S}_\Gamma$. Note that $\mathcal{S}_\Gamma - x\mathcal{S}_B$ has at least two components: one containing xc^{-1}, and one containing xa. Similarly, $\mathcal{S}_\Gamma - x\mathcal{S}_C$ has at least two components: one containing xb^{-1}, and one containing xd.

For each i, define the following components of $\mathcal{S}_\Gamma$:

1. V_i^+ is the component of $\mathcal{S}_\Gamma - v_i\mathcal{S}_B$ containing $v_i a$;
2. V_i^- is the component of $\mathcal{S}_\Gamma - v_i\mathcal{S}_B$ containing $v_i c^{-1}$;
3. W_i^+ is the component of $\mathcal{S}_\Gamma - w_i\mathcal{S}_C$ containing $w_i d$;
4. W_i^- is the component of $\mathcal{S}_\Gamma - w_i\mathcal{S}_C$ containing $w_i b^{-1}$.

Note V_i^+ contains the vertices of r after v_i, and W_i^+ contains the vertices of s after w_i. For each $V_i^\pm$, (respectively $W_i^\pm$), let $\overline{V_i^\pm}$ denote the closure of $V_i^\pm$ in $\mathcal{S}_\Gamma$, so $\overline{V_i^\pm} = V_i^\pm \cup v_i\mathcal{S}_B$ ($\overline{W_i^\pm} = W_i^\pm \cup w_i\mathcal{S}_C$). For a subset S of $\mathcal{S}_\Gamma$, let $L(S)$ denote the limit set of S in $\partial\mathcal{S}_\Gamma$.

Claim 4.2.3. *1. The sets $\overline{V_i^\pm}$, $\overline{W_i^\pm}$ are convex.*

2. $L(\overline{V_i^+}) \cap L(\overline{V_i^-}) = L(v_i\mathcal{S}_B)$ and $L(\overline{W_i^+}) \cap L(\overline{W_i^-}) = L(w_i\mathcal{S}_C)$.

3. The set $L(v_i\mathcal{S}_B)$ (respectively $L(w_i\mathcal{S}_C)$) separates $L(\overline{V_i^+})$ and $L(\overline{V_i^-})$ (respectively $L(\overline{W_i^+})$ and $L(\overline{W_i^-})$) in ∂X.

Proof. For (1), the only way out of the set $\overline{V_i^+}$ is through the convex subcomplex $v_i\mathcal{S}_B$.

For (2), if q is a ray in $L(\overline{V_i^+}) \cap L(\overline{V_i^-})$, then there are geodesic rays $q_1 \in \overline{V_i^+}$, $q_2 \in \overline{V_i^-}$ that are a bounded distance from q, and therefore from one another. Thus both q_1 and q_2 remain a bounded distance from $v_i\mathcal{S}_B$, as required.

For (3), suppose $\alpha : [0,1] \to \partial\mathcal{S}_\Gamma$ is a path connecting $x \in L(V_i^+)$ and $y \in L(V_i^-)$. Choose $w \in v_i\mathcal{S}_B$, and for each $t \in [0,1]$, let $\beta_t : [0,\infty) \to \mathcal{S}_\Gamma$ be the geodesic ray from w to $\alpha(t) \in \partial\mathcal{S}_\Gamma$. This gives a continuous map $H : [0,1] \times [0,\infty) \to \mathcal{S}_\Gamma$ where $H(t,s) = \beta_t(s)$. Note $H(0,s) \subset V_i^+$, $H(1,s) \subset V_i^-$. For each $n \geq 0$, let z_n be a point of $H([0,1] \times \{n\})$ in $v_i\mathcal{S}_B$; then $L(\cup_{n=1}^{\infty}\{z_n\}) \subset Im(\alpha) \cap L(v_i\mathcal{S}_B)$ as required. □

In [7], it is shown that r and s track distinct CAT(0) geodesics in $\mathcal{S}_\Gamma$, so $L(r)$ and $L(s)$ are distinct one-element sets.

Claim 4.2.4. *For $n \geq 1$, the sets $L(w_{2n-1}\mathcal{S}_C)$ and $L(r)$ are separated in $\partial\mathcal{S}_\Gamma$ by $L(v_{2n+1}\mathcal{S}_B)$.*

Proof. First note that $L(r) \in L(V_i^+)$ for each $i \geq 1$. Let D_{2n} be the d-hyperplane that separates w_{2n} from w'_{2n} (and also separates v_{2n} from the previous vertex of r), and let A_{2n} be the a-hyperplane that separates v_{2n} from v'_{2n} (and also separates w_{2n+1} from the previous vertex of s). Note that $w_{2n-1}\mathcal{S}_C$ is contained in the same component of $\mathcal{S}_\Gamma - D_{2n}$ as $*$ since $d \notin C$ and therefore no path in $\langle C \rangle$ based at w_{2n-1} crosses D_{2n}. Also note $A_{2n} \subset V_{2n+1}^-$. Since D_{2n} and A_{2n} cannot cross (since d does not commute with a), and D_{2n} is not in the same component as $v_{2n+1}\mathcal{S}_B$ in $\mathcal{S}_\Gamma - A_{2n}$, we have that $w_{2n-1}\mathcal{S}_C \subset V_{2n+1}^-$. The previous claim gives the result. □

Claim 4.2.5. *For $n \geq 1$, the sets $L(v_{2n-1}\mathcal{S}_B)$ and $L(r)$ are separated in $\partial\mathcal{S}_\Gamma$ by $L(w_{2n+1}\mathcal{S}_C)$.*

Proof. The proof is analagous to the proof of the previous claim, replacing the hyperplanes D_{2n} and A_{2n} with the hyperplanes A_{2n-1} and D_{2n} respectively. □

Remark 4.2.6. *The previous two claims imply that if there is a path in $\partial\mathcal{S}_\Gamma$ between a point of $L(w_1\mathcal{S}_C)$ and $L(r)$, the path must pass through (in order) $L(v_3\mathcal{S}_B)$, $L(w_5\mathcal{S}_C)$, $L(v_7\mathcal{S}_B)$, $L(w_9\mathcal{S}_C)$, and so on.*

We will now show that the sets $L(v_i\mathcal{S}_B)$ (resp. $L(w_i\mathcal{S}_C)$) are eventually 'close' to $L(s)$ (resp. $L(r)$), implying the path described in Remark 4.2.6 cannot exist.

Claim 4.2.7. *$C \cap lk(a) \cap lk(d) = C \cap lk(a) \cap lk(c) = \emptyset$, and $B \cap lk(a) \cap lk(d) = B \cap lk(d) \cap lk(b) = \emptyset$.*

Proof. If $e \in C \cap lk(a) \cap lk(d)$, then (a, e, d, c) is a path from a to c in Γ. Since B separates a from c and $d \notin B$, we must have $e \in B$, but $B \cap C = \emptyset$. Similarly, if $e \in C \cap lk(a) \cap lk(c)$, then (a, e, c) is a path from a to c in Γ, and so $e \in B$, contradiction. The remaining statements are proved identically. □

For $i \geq 1$, let r_i (respectively s_i) be the segment of r (respectively s) between $*$ and v'_i (respectively $*$ and w'_i). Let β_i be a Cayley graph geodesic ray based at w'_i with labels in B, and let γ_i be a Cayley graph geodesic ray based at v'_i with labels in C.

Claim 4.2.8. *Any Cayley graph geodesic from $*$ to a point of γ_i must pass within 4 units of v'_i. Any Cayley graph geodesic from $*$ to a point of β_i must pass within 4* *units of* w_i'.

Proof. First observe that if (r_i, γ_i) is not Λ_Γ-geodesic, then an edge of γ_i must delete with an edge of r_i. Since $a, b, d \notin C$, the labels of these deleting edges must be c and c^{-1}. However, the labels of these edges must also be in $lk(a) \cap lk(d)$, by Lemma 4.1.3 (see Figure 4.2). Therefore (r_i, γ_i) is a Cayley geodesic.

Now, suppose there is a Λ_Γ-geodesic ρ between $*$ and a point of γ_i with $d(\rho, v_i') > 4$. Let α denote the segment of (r_i, γ_i) between $*$ and the endpoint of ρ. Consider a diamond based at v_i' for ρ and α as in Lemma 3.1.22. Let τ and δ be the down edge path and up edge path respectively at v_i', and note τ and δ have length at least 3. Every Λ_Γ-geodesic from $*$ to v_i' must end with an edge labeled a, so every label of δ is in $lk(a)$. If an edge of τ has label d, then every label of δ is in $C \cap lk(a) \cap lk(d)$, but this set is empty by Claim 4.2.7. By Lemma 4.1.3 every other edge of τ has its label in $lk(d) \cap \{a, b, c, d\}$, so the remaining edges of τ must be labeled c, but $C \cap lk(a) \cap lk(c)$ is also empty. Thus $d(\rho, v_i') \leq 4$. The proof of the second statement is identical. □

Claim 4.2.9. $\partial\mathcal{S}_\Gamma$ *is not path connected.*

Proof. Observe that since $v_{n-1}' b^{k_{n+1}+1} = w_n$ by Claim 4.2.2 and $C \subset lk(b)$, any ray α based at w_n with labels in C stays a bounded distance from the ray based at v_{n-1}' with the same labels. Combining Claim 4.2.8 and Lemma 4.1.17, we have that a CAT(0) geodesic from $*$ to a point of $L(\alpha)$ must pass within $\delta + 4$ of v_{n-1}', where δ is the tracking constant given by Lemma 4.1.17. We therefore have that any sequence of points $\{p_i\}_{i=1}^\infty$ with each $p_i \in L(w_i\mathcal{S}_C) \subset \partial\mathcal{S}_\Gamma$ must converge to $L(r) \in \partial\mathcal{S}_\Gamma$. Similarly, any sequence of points $\{q_i\}_{i=1}^\infty$ with each $q_i \in L(v_i\mathcal{S}_B) \subset \partial\mathcal{S}_\Gamma$

must converge to $L(s) \in \partial\mathcal{S}_\Gamma$. Therefore, by Remark 4.2.6, given any ϵ, any path from a point of $L(w_1\mathcal{S}_C)$ to $L(r)$ eventually bounces back and forth infinitely between the ϵ-neighborhood of $L(s)$ and the ϵ-neighborhood of $L(r)$, which is impossible; therefore, no such path exists. $\square$

References

[1] J. Behrstock and W. Neumann. *Quasi-isometric classification of graph manifold groups.* Duke Math. J. 141 (2008), 217-240.

[2] M. Bestvina and G. Mess. *The boundary of negatively curved groups.* J. Amer. Math. Soc. 4 (1991), 469-481.

[3] M. Bestvina, B. Kleiner, and M. Sageev. *The asymptotic geometry of right-angled Artin groups, I.* Geom. Topol. 12 (2008), 1653-1699.

[4] N. Bourbaki. *Groupes et algèbres de Lie.* Chapitres 4, 5, et 6, Hermann, Paris, 1968.

[5] M.R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature.* Grundlehren Math. Wissensch. 319. Springer-Verlag, 1999.

[6] W. Camp and M. Mihalik. *A classification of right-angled Coxeter groups with no 3-flats and locally connected boundary.* Submitted for publication. arXiv:1206.5234 [math.GR], 2012.

[7] G. Conner, M. Mihalik and S. Tschantz. *Homotopy of Ends and Boundaries of CAT(0) Groups.* Geom. Dedicata 120 (1) (2006), 1-17.

[8] C. Croke and B. Kleiner. *Spaces with non-positive curvature and their ideal boundaries.* Topology 39 (2000), 549-556.

[9] M.W. Davis. *The geometry and topology of Coxeter groups.* London Mathematical Society Monographs Series Vol. 32, Princeton University Press, Princeton, NJ, 2008.

[10] M. Gromov. *Hyperbolic groups.* Essays in Group Theory, S. Gersten ed., MSRI Publications 8 (1987), 75-265.

[11] T. Hsu and D. Wise. *Separating quasiconvex subgroups of right-angled Artin groups.* Math. Z. 240 (2002), 521-548.

[12] T. Januszkiewicz, J. Swiatkowski. *Hyperbolic Coxeter groups of large dimension.* Comment. Math. Helv. 78 (2003), 555-583.

[13] R. Lyndon and P. Schupp. *Combinatorial Group Theory.* Springer-Verlag, New York, 2001.

[14] R. Meris. *Graph Theory.* Wiley, New York, 2001.

[15] M. Mihalik and K. Ruane. *CAT(0) groups with non-locally connected boundary.* J. London Math. Soc. (2) 60 (1999), 757-770.

[16] M. Mihalik, K. Ruane, and S. Tschantz. *Local connectivity of right-angled Coxeter group boundaries*. J. Group Theory 10 (2007), 531-560.

[17] M. Mihalik and S. Tschantz. *Visual Decompositions of Coxeter Groups*. Groups, Geometry, and Dynamics 3 (2009), 173-198.

[18] J. Milnor. *A note on curvature and the fundamental group*. J. Differential Geometry 2 (1968), 1-7.

[19] C. Mooney. *Generalizing the Croke-Kleiner construction*. Topology Appl. 157 (7) (2010), 1168-1181.

[20] G. Moussong. *Hyperbolic Coxeter groups*. Ph.D. Thesis. Ohio State University (1988).

[21] Michah Sageev. *Ends of group pairs and non-positively curved cube complexes*. Proc. London Math. Soc. (3) 71 (1995), 585-617.

[22] G. Swarup. *On the cut point conjecture*. Electron. Res. Announc. Amer. Math. Soc. 2 (1996), 98-100 (electronic).

[23] J. Wilson. *A CAT(0) group with uncountably many distinct boundaries*. J. Group Theory 8 (2005), 229-238.