A CLASSIFICATION OF RIGHT-ANGLED COXETER GROUPS WITH NO 3-FLATS AND LOCALLY CONNECTED BOUNDARY

WES CAMP AND MICHAEL MIHALIK

Abstract. If (W, S) is a right-angled Coxeter system and W has no $\mathbb{Z}^3$ subgroups, then it is shown that the absence of an elementary separation property in the presentation diagram for (W, S) implies all CAT(0) spaces acted on geometrically by W have locally connected CAT(0) boundary. It was previously known that if the presentation diagram of a general right-angled Coxeter system satisfied the separation property, then all CAT(0) spaces acted on geometrically by W have non-locally connected boundary. In particular, this gives a complete classification of the right-angled Coxeter groups with no 3-flats and with locally connected boundary.

1. Introduction

In this paper, we classify the right-angled Coxeter groups with no $\mathbb{Z}^3$ subgroups that have locally connected CAT(0) boundary. We say a CAT(0) group has locally (respectively, non-locally) connected boundary if all CAT(0) boundaries of the group are locally (respectively non-locally) connected. Our main theorem states that if the Coxeter presentation of the group satisfies an elementary combinatorial condition, then this group has locally connected boundary and otherwise has non-locally connected boundary. This condition was first considered in [10], and the results there make it natural to conjecture that any right-angled Coxeter group has locally connected boundary if and only if the group presentation satisfies this condition. The primary working tool for both this paper and [10] is the notion of a *filter* for CAT(0) geodesics r and s in a CAT(0) space X on which the Coxeter group W acts geometrically. A filter is a connected, one-ended planar graph whose edges are labeled by the Coxeter generators S of W. Hence there is a natural (proper) map of the filter into the Cayley graph of (W, S), which in turn maps properly and W-equivariantly into the CAT(0) space X. The two sides of the filter track the geodesics r and s and the limit set of the filter is a connected set in ∂X (the boundary of X), containing the limit points of r and s. The idea is to construct a filter in such a way so that if r and s are "close" in ∂X, then the filter has "small" limit set containing the limit points of r and s, and local connectivity of the boundary of X follows.

In [10], two types of separators are defined for the Coxeter presentation graph Γ of the group, the first of which is a *virtual factor separator*: a virtual factor separator for (W, S) (or for 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. The main theorem of [10] states: if Γ has a virtual factor separator, then the Coxeter group W has non-locally connected boundary, and if Γ has neither type of separator, then the Coxeter group has locally connected boundary. In fact, when Γ has neither type of separator, the filters constructed in [10] basically have hyperbolic geometry; i.e. any geodesic path in the Cayley graph from the base point of the filter to another point of the filter must track the filter geodesic connecting these two points (just as in a word hyperbolic group). In this paper, the geometry of our filters is necessarily more complex. The no $\mathbb{Z}^3$ subgroup hypothesis does restrict the pathology of the geometry of the filter, but our results are the natural next step towards a full classification of right-angled Coxeter groups with locally connected boundary, and provide hard evidence that the following conjecture is sound:

Conjecture. *Let (W, S) be a directly indecomposable one-ended right-angled Coxeter group with presentation graph Γ. Then W has locally connected boundary if and only if Γ has no virtual factor separator.*

If a Coxeter group has no $\mathbb{Z}^2$ subgroup, then it is word hyperbolic [13], and all one-ended word hyperbolic groups have (unique) locally connected boundary [14]. Januszkiewicz and Swiatkowski ([7]) produce word hyperbolic, right-angled Coxeter groups of virtual cohomological dimension n for all positive integers n, so our no $\mathbb{Z}^3$ hypothesis does not restrict the virtual cohomological dimension of the groups under consideration. In [5], Croke and Kleiner exhibit a one-ended CAT(0) group with non-homeomorphic boundaries. Each of these boundaries is in fact connected but not path connected (see [4]). In particular, (by classical point set topology) these boundaries are not locally connected. It seems that many of the serious pathologies one sees in boundaries of CAT(0) groups, but not in boundaries of word hyperbolic groups, happen in the presence of non-local connectivity. At the time of this writing, no CAT(0) group has been shown to have non-homeomorphic boundaries, one of which is locally connected. There are numerous questions about how or even if the homology and homotopy of two boundaries of a CAT(0) group can differ. These questions may be more tractable if the boundaries considered are locally connected. If our results extend to all right-angled Coxeter groups, then those with locally connected boundary should provide an interesting testing ground for such questions.

The paper is laid out as follows. In Section 2, basic definitions and background results are listed, including a lemma (2.21) that provides the fundamental combinatorial tool for measuring how large the limit set of a filter might be. In Section 3, the basics of CAT(0) spaces and groups are outlined, and we list two tracking results (developed in [10]) that connect the CAT(0) geometry and algebraic combinatorics of right-angled Coxeter groups. In Section 4, we construct a basic filter, and show that while the limit set of such a filter is always a connected subset of the CAT(0) boundary, this limit set may not be small. In Section 5, we use our no $\mathbb{Z}^3$ hypothesis to find at most two 'directions' in which a geodesic could lead to a filter having a large limit set. In Section 6, we construct a filter with 'small' limit set, and prove our main theorem:

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 if and only if A is infinite ended.

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

Corollary. *Suppose (W,S) is a one-ended right-angled Coxeter system that has no visual subgroup isomorphic to $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$. Then all CAT(0) boundaries of W are locally connected or all are non-locally connected.*

The group W visually splits as in item (1) of the theorem precisely when there are $s,t \in S$ such that st has infinite order in W and $\{s,t\}$ commutes with $S-\{s,t\}$, so this condition is easily checked. If a CAT(0) group splits as $G = (\mathbb{Z}_2 * \mathbb{Z}_2) \times A$, then any boundary of G is the suspension of a boundary of A (see [10]) and this suspension is locally connected if and only if the boundary of A is locally connected. If (W,S) is a one-ended right-angled Coxeter system with no visual subgroup isomorphic to $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ and W visually splits as $(\mathbb{Z}_2 * \mathbb{Z}_2) \times A$, then A is word hyperbolic (see [13]). It is straightforward to check if a Coxeter group is infinite-ended (see Remark 4.3). Thus item (1) of the theorem is easily verified and the real content of the theorem is contained in item (2). If (C,D) is a virtual factor separator for (W,S), then W visually splits nontrivially as an amalgamated product $\langle A\rangle *_{\langle C\rangle} \langle B\rangle$ (here A and B are subsets of S with $A\cup B = S$ and $A \cap B = C$). Therefore local connectivity of boundaries of W is directly tied to visual splittings.

In Section 7, we give examples to show there are no combination or splitting results for right-angled Coxeter groups that respect local connectivity of boundaries. One example describes a right-angled Coxeter group as the (visual) amalgamated product $W = A *_C B$ where A and B are one-ended and word hyperbolic (so both have locally connected boundary) and C is virtually a surface group (with boundary a circle), but W has non-locally connected boundary. The second example describes a right-angled Coxeter group W that visually splits as $A *_C B$, and a single element of infinite order in C determines a boundary point of non-local connectivity in both A and B. Nevertheless, our main theorem implies W has locally connected boundary. These examples indicate there are no reasonable graph of groups approaches to this problem. Morse theory also seems unhelpful, but we do not expand here.

2. Preliminaries

We use [2] and [6] as basic references for the results in this section.

Definition 2.1. A *Coxeter system* is a pair (W,S), where W is a group with *Coxeter presentation*:

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

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

Definition 2.2. 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 paper 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 2.3. 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 2.4. 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 *S-geodesic* (or simply *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 2.5. 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 $[i-1, i]$ 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 α and β are geodesic edge paths with the same initial and terminal points, we call β a *rearrangement* of α.

Lemma 2.6. *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 2.7. If (W, S) is a Coxeter system and $A \subset S$, then $lk(A) \equiv \{t \in S : m(a, t) = 2$ 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 2.8. *(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 if and only if the word $[\overline{e}_1, \overline{e}_2, \ldots, \overline{e}_n]$ is geodesic. If α is not geodesic and $\overline{e}_i$ deletes with $\overline{e}_j$, for $i < j$, let τ be the the path beginning at the end point of e_{i-1} with edge labels $[\overline{e}_{i+1}, \ldots, \overline{e}_{j-1}]$. 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 2.9. If (W, S) is a Coxeter system and $A \subset S$, then the subgroup of W generated by A is called a *visual* (or *special*) subgroup of W.

Lemma 2.10. *Suppose (W, S) is a Coxeter system, and $A \subset S$. Then the visual 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 2.11. *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 2.12. For (W, S) a Coxeter system and α a geodesic in $\Lambda(W, S)$, let $B(\alpha) \equiv \{\overline{e} \in S : e$ is a Λ-edge based at the terminal vertex of α and (α, e) is not geodesic$\}$.

Lemma 2.13. *Suppose (W,S) is a Coxeter system, and α a geodesic in Λ. Then $B(\alpha)$ generates a finite group.*

Lemma 2.14. *If (W,S) is a right-angled Coxeter system, and $s,t \in S$ delete in some S-word. Then $s=t$.*

Lemma 2.15. *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 2.16. 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 [11]) and this splitting is called a *visual decomposition* for (W,S).

Definition 2.17. 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 2.18. 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 if and only 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 $w_2\bar{t}_1\cdots\bar{t}_n = y_2$, where y_1 and y_2 are the vertices of e and $m(\bar{e},\bar{t}_i) = 2$ for each $1 \le i \le n$.

Definition 2.19. 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 2.20. 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 if and only if $\overline{e}_i$ 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 if and only if 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 [8]. 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.

The following lemma has some of its underlying ideas in Lemma 5.10 of [10]. 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 2.21. *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.

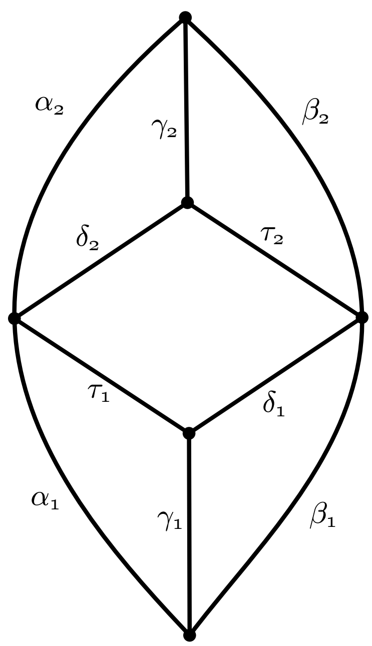

FIGURE 1

Proof. Consider a van Kampen diagram for the loop $(\alpha_1, \alpha_2, \beta_2^{-1}, \beta_1^{-1})$ (Figure 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 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 2.20 (3)) (τ_1^{-1}, δ_1) is geodesic. □

Remark 2.22. Using the notation of Lemma 2.21 (and Figure 1), we have the following:

(1) The walls of γ_1 are exactly the walls shared by α_1 and β_1;
(2) The walls of γ_2 are exactly the walls shared by α_2 and β_2;
(3) The walls of δ_1 are the same as the walls of δ_2, and these are exactly the walls shared by β_1 and α_2;
(4) The walls of τ_1 are the same as the walls of τ_2, and these are exactly the walls shared by α_1 and β_2;
(5) All the walls of τ_1 cross all the walls of δ_1.
(6) If α_1', α_2', β_1' and β_2' are rearrangements of α_1, α_2, β_1 and β_2 respectively, and γ_1', γ_2', τ_1', τ_2', δ_1' and δ_2' are the paths given by Lemma 2.21 for (α_1', α_2') and (β_1', β_2'), then γ_1' γ_2', τ_1', τ_2', δ_1' and δ_2' are rearrangements of γ_1, γ_2, τ_1, τ_2, δ_1 and δ_2 respectively. I.e. up to rearrangements the paths γ_1, γ_2, τ_1, τ_2, δ_1 and δ_2 are uniquely determined by the (four) end points of α_1, α_2, β_1 and β_2.

Remark 2.23. For the entirety of this paper, we will only consider the case of Lemma 2.21 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. In this case, if y is the endpoint of α_1 and μ is any other geodesic with the same initial and terminal vertices as (α_1, α_2), the diamond between (α_1, α_2) and μ at y is uniquely defined (up to rearrangements within the subpaths). 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 2.24. *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. Furthermore, if τ_A' is another such shortest path at x, then τ_A and τ_A' have the same end point (and so cross the same set of walls).*

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 2.20. 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 2.20 (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.

Now suppose τ'_A is another such shortest path, but with end point different than that of τ_A. Extend both τ_A and τ'_A to geodesics ending at y. Applying Lemma 2.21 to the resulting bigon gives a diagram as in Figure 1, with $\alpha_1 = \tau_A$ and $\beta_1 = \tau'_A$. As τ_A has minimal length, the last edge of τ_1 (in Figure 1) must belong to a wall of A. Then τ_2 would also contain an edge of that wall. That is impossible since $\beta_1 \equiv \tau'_A$ crosses all walls of A and the geodesic (β_1, τ_2) would cross a wall of A twice. □

Lemma 2.25. *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 2.21, and suppose $|\nu_{12}| \geq |\nu_{13}| \geq 2|S|$. If $\{c, d\} \subset lett(\nu_{12}) \cap lett(\nu_{13})$ and $m(c,d) = \infty$, then $d(x_2, x_3) < 2(|\nu_{12}| - |\nu_{13}|) + 4|S|$.*

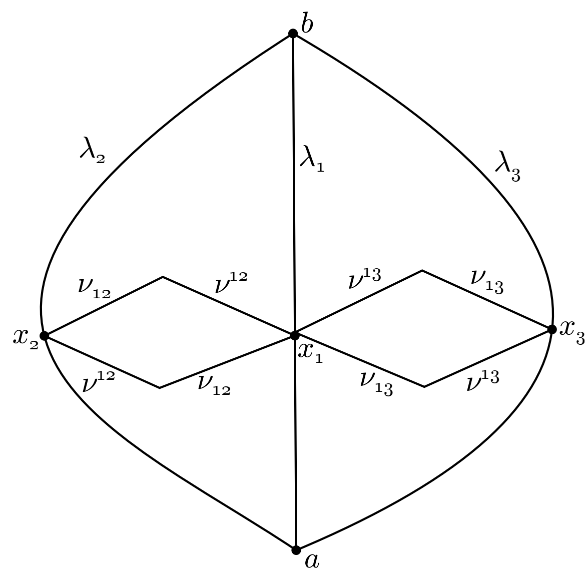

FIGURE 2

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. Since the x_i are all equidistant from a and equidistant from b, we have $|\nu^{12}| = |\nu_{12}|$ and $|\nu^{13}| = |\nu_{13}|$. Note that at x_2, $\nu^{12}\nu_{12}\nu_{13}\nu^{13}$ is a path from x_2 to x_3. By Lemma 2.21, $\{c, d\}$ 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 2.5). 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 so the shared initial segment of ν^{12} and ν^{13} has length at least $|S|+1$ (since $|\nu_{12}| \geq |\nu_{13}| \geq 2|S|$), and so 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 the shared walls of ν_{12} and ν_{13} and the shared walls of ν^{12} and ν^{13} leaves us with a geodesic from x_2 to x_3 of length less than $2(|\nu_{12}| - |\nu_{13}|) + 4|S|$. □

3. CAT(0) Spaces and Actions by Coxeter Groups

Definition 3.1. A metric space (X, d) is *proper* if each closed ball is compact.

Definition 3.2. 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 3.3. *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 3.4. A *geodesic ray* in a CAT(0) space X is an isometry $[0, \infty) \to X$.

Definition 3.5. 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 3.6. *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 [3].

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 3.7. 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 3.8. We call a group G a *CAT(0) group* if it acts geometrically on a CAT(0) space.

The next theorem, due to Milnor [12], will be used in conjunction with the next two technical lemmas to identify geodesic rays in X with certain rays in a right-angled Coxeter group which acts on X.

Theorem 3.9. *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.*

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 the previous theorem. 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 3.10. 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 [10].

Lemma 3.11. *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 3.12. *There exist $c, c' > 0$ such that, for any infinite geodesic rays r and s and X based at x_0 that remain ϵ-close to each other on their initial segments of length M, there are Cayley graph geodesic rays α and β which $(c\epsilon + c')$-track r and s respectively, and which share a common initial segment of length $M - c\epsilon - c'$.*

4. Local connectivity and a basic filter construction

Definition 4.1. 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 4.2. 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.

Remark 4.3. The right-angled Coxeter group W is one-ended if and only if $\Gamma(W,S)$ contains no complete separating subgraph (i.e., a subgraph whose vertices generate a finite group in W). For a proof of this, see [11].

Remark 4.4. If e is an edge in the Cayley graph $\Lambda(W,S)$, we let $\overline{e} \in S$ denote the label of e. Recall 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 2.13).

Remark 4.5. 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.

As discussed in the introduction, the meat of Theorem 6.14 lies in showing 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 3.12 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. We ultimately 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.

We begin by demonstrating the construction of a rather basic filter. 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 the paths $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ 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 $(c + c')$-track two CAT(0) geodesics r and s in X (as in Lemma 3.12), and let x_m denote the endpoint of $(e_1, \ldots, e_m)$. Set $a_1 = \overline{e}_{m+1}$ and $b_1 = \overline{d}_{m+1}$. 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 the (labeled) planar diagram of Figure 3 that maps naturally into Λ (respecting labels) and then to X:

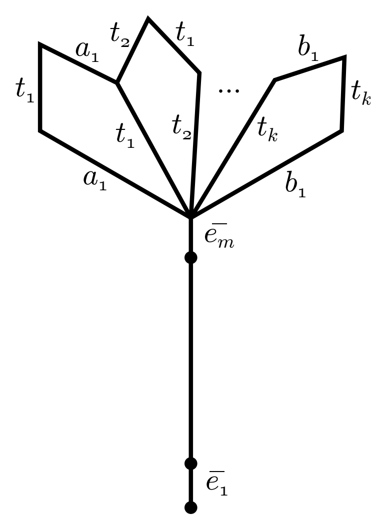

FIGURE 3

As in [10], we call Figure 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 2.15). 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 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. We refer to the image of a filter, in Λ or in X, again as a filter.

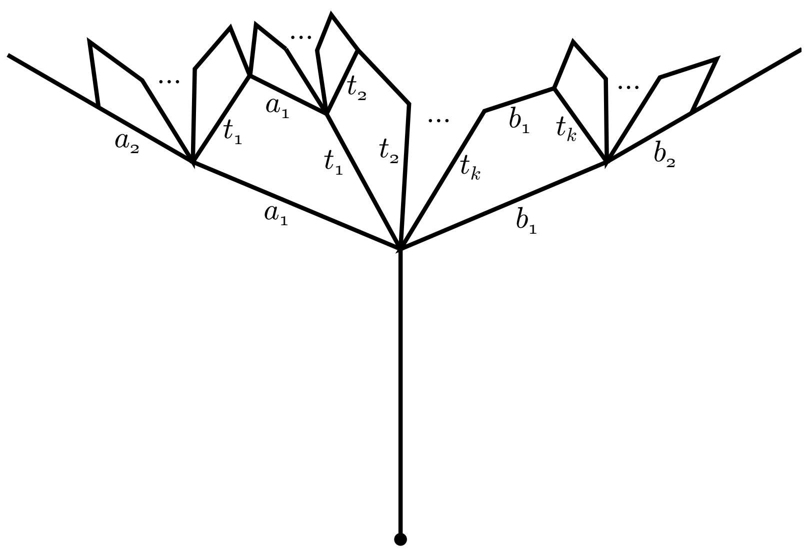

FIGURE 4

Lemma 4.6. *Suppose W acts geometrically on a CAT(0) space X. Let F be a filter for the $\Lambda(W, S)$-geodesics $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ and $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$, and let $C(F)$ be the image of F in X (via the natural map $F \to \Lambda \to X$). Then the limit set of $C(F)$ is a connected subset of ∂X.*

Proof. Let x be a base point in X, and let $B_n(x)$ denote the open ball of radius n about x. Let $\overline{X}$ be the compact metric space $X \cup \partial X$. Let $\overline{C(F)}$ denote the closure of $C(F)$ in $\overline{X}$. Since $C(F)$ is connected, $\overline{C(F)}$ is connected. Since $C(F)$ is one-ended, $\overline{C(F)} - C(F)$ (the limit set of $C(F)$) is contained in a component of $\overline{C(F)} - B_n(x)$, denoted C_n, for each $n > 0$. Then $\overline{C(F)} - C(F) = \cap_{n=1}^{\infty} C_n$ is the intersection of compact connected subsets of a metric space and is therefore connected. □

Thus the limit set of our filter in ∂X is a connected set containing our original rays r and s. The problem, then, is that this limit set may not be small.

5. Constructing directions

In order 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 2.21, 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. As before, let (W, S) be a right-angled Coxeter system where W is one-ended and acts geometrically on a CAT(0) space X, suppose $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ 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 $(c+c')$-track two CAT(0) geodesics r and s in X (as in Lemma 3.12), and x_i is the endpoint of $(e_1, \ldots, e_i)$. The base point of our filter will be x_m. Finally, set $N = |S|$ and note that $N \geq 4$ since W is one-ended.

Remark 5.1. For the rest of this paper, 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 5.2. Construct a geodesic from the vertex x to $*$ in Λ as follows: let α_1 be a longest geodesic with edge labels in the finite group $\langle B(x \to *)\rangle$, and let y_1 be the endpoint of α_1 based at x. 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 to $*$. We call this a *back combing* geodesic from x to $*$.

Remark 5.3. We have the following properties of a back combing geodesic $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ from x 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 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 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 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 to $*$. If e_j is an edge of α_j sharing a wall with τ_1 and γ_1, (for some $1 \leq j \leq r$), then for each $1 \leq i \leq j$, there is an edge e_i of α_i sharing a wall with τ_1 and γ_1 such that $\bar{e}_i$ and $\bar{e}_{i+1}$ do not commute for $1 \leq i < j$.

Proof. (of 5.3 (3)) The walls of β_1 are a subset of the walls of α_1 by definition. Assume the walls of β_j are a subset of the walls of α_j for all $j < i$. Suppose the wall of the edge e of β_i is not a wall of α_i. By (1) we assume e is the first edge of β_i. If e belongs to a wall of α_k where $k < i$, then we assume the wall of e is the wall of the first edge of α_k. Consider a van Kampen diagram for the loop formed by $(\alpha_1, \ldots, \alpha_k)$ and $(\beta_1, \ldots, \beta_s, \gamma_2)$. As each wall of β_{i-1} is a wall of α_{i-1}, the wall of e crosses (in the diagram) each wall of β_{i-1} (contrary to (2)).

Instead, the wall of e is a wall of α_k with $k > i$. But then, each wall of α_{k-1} crosses the wall of e, again contrary to (2). □

Vertices in our filter for r and s will be end points of geodesics in Γ that pass through x_m. We will always assume that x_m and $*$ will be far apart (since r and

s will be close in ∂X), and so we are only interested in vertices of Γ that are far from $*$.

At this point we fix the vertex x in Γ with $d(x,*) > 7N^2$, and $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ a back combing from x to $*$. Let $\alpha_{7N+1} = (u_1, u_2, \ldots, u_d)$ (note $d < N$), and for $1 \le i \le d$, let U_i be a shortest Λ-geodesic based at x such that the last edge of U_i is in the same wall as u_i (so by Lemma 2.24, U_i extends to a geodesic from x to $*$). There may be several such geodesics, but they all have the same set of walls and so are rearrangements of one another.

Lemma 5.4. *If (γ_1, γ_2) is a Λ-geodesic from x to $*$ with $|\gamma_1| \ge 7N^2$, then γ_1 can be rearranged to begin with exactly U_i, for some $1 \le i \le d$.*

Proof. Let $\gamma_1 = (t_1, t_2, \ldots, t_s)$, where $s \ge 7N^2$. 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 $1 \le i \le d$ (such a j exists from Remark 5.3 (3) and because the lengths of $\alpha_1, \ldots, \alpha_{7N}$ are each less than N). By Lemma 2.24, U_i can be extended to a geodesic ending at the endpoint of γ_1. □

We now have a finite number $d < N$ of "directions", given by our U_i, in which a bigon can be wide at x. Our next goal is to reduce this collection to at most two directions while retaining the conclusion of Lemma 5.4. In Proposition 5.5, we refine our list of U_i through a 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. All the while, Lemma 5.4 remains valid for the new U_i.

Lemma 5.6 is proved within the proof of our proposition. It is a fundamental combinatorial consequence of our no $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ hypothesis which allows us to reduce to at most two directions.

Proposition 5.5. *Let x and $*$ be vertices of $\Lambda(W,S)$, with $d(x,*) = k > 7N^2$. Then either there are Λ-geodesics $U_1^x \neq U_2^x$, based at x, with all the following properties:*

1. *$|U_1^x| > 6N$ and $|U_2^x| > 6N$;*
2. *U_1^x and U_2^x begin geodesics from x to $*$;*
3. *If γ_1 begins a geodesic from x to $*$ and $|\gamma_1| \ge 7N^2$, then $|\gamma_1|$ can be rearranged to begin with at least one of U_1^x or U_2^x;*
4. *U_1^x and U_2^x can be rearranged as $(U_0^x, \hat{U}_1^x)$ and $(U_0^x, \hat{U}_2^x)$ respectively, such that $\langle lett(U_0^x)\rangle$ is finite (so $|U_0^x| < N$, $|\hat{U}_1^x| > 5N$, and $|\hat{U}_2^x| > 5N$), and each wall of $\hat{U}_1^x$ crosses each wall of $\hat{U}_2^x$.*
5. *If η is a geodesic from $*$ to x, (η, γ) is a geodesic extension of η, and γ' is a rearrangement of (η, γ) whose $(k+1)^{st}$ vertex is of distance at least $14N^2$ from x, then the down edge path at x for the diamond (Lemma 2.21) for (η, γ) and γ' can be rearranged to begin with exactly one of U_1^x or U_2^x.*
6. *If η is a geodesic from $*$ to x and $j \in \{1, 2\}$, then there is a geodesic extension (η, γ_j) of η and a rearrangement γ'_j of (η, γ_j) whose $(k+1)^{st}$ vertex is of distance at least $16N^2$ from x, such that the down edge path at x for the diamond (Lemma 2.21) for (η, γ_j) and γ'_j can be rearranged to begin with U_j^x.*

or, there is a Λ-geodesic U_1^x, based at x, with the following properties:

(7) *U_1^x contains two edges with unrelated labels;*
(8) *U_1^x begins a geodesic from x to $*$;*
(9) *If η is a geodesic from $*$ to x, (η, γ) is a geodesic extension of η, and γ' is a rearrangement of (η, γ) whose $(k+1)^{st}$ vertex is of distance at least $16N^2$ from x, then the down edge path at x for the diamond (Lemma 2.21) for (η, γ) and γ' can be rearranged to begin with exactly U_1^x (note this will be trivially satisfied if no such γ' exists).*

Proof. Remark 2.22 (6), implies that if parts (5), (6) and (9) of the lemma hold for some path η from $*$ to x, then they hold for any rearrangement of η. Hence we fix η throughout the proof. We begin with a back combing $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ from x to $*$ and a collection of directions U_i as in Lemma 5.4. So, $\alpha_{7N+1} = (u_1, u_2, \ldots u_d)$ and R begins at $7N$.

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. In this case, let τ_a be a shortest Λ-geodesic based at x whose last edge is in the same wall as a. Applying Lemma 2.24 separately to U_i and U_j implies both U_i and U_j can be rearranged to begin with τ_a. Our process is as follows:

(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) A shortest geodesic based at x and ending with an edge in the wall of a extends (by Lemma 2.24) to a rearrangement of U_i and extends to a rearrangement of U_j (and so extends to a geodesic to $*$). Redefine U_i to be this shortest geodesic and redefine u_i (an edge of α_{R+1} in the same wall as the last edge of the original U_i) to be a (an edge of α_R in the same wall as the last edge of the new U_i).
(3) Eliminate U_j from the list of U_ℓ and note that any geodesic from x to $*$ beginning with the old U_i or U_j can be rearranged to begin with the new U_i (so the new U_i effectively replaces both in the conclusion of Lemma 5.4).
(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 and ending with an edge in the wall of b_ℓ, and redefine u_ℓ to be b_ℓ. Again an original U_l can be rearranged to begin with a rearrangement of the new U_l so that Lemma 5.4 remains valid under this replacement.
(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, then they are rearrangements of one another. 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 $R-1$ and observe that Lemma 5.4 remains valid for the new U_i.

When this process stops, no two U_i can R-overlap, and each u_i is an edge of α_{R+1} sharing a wall with the last edge of U_i. Since U_i is a shortest geodesic with last edge in the wall of u_i, every geodesic from x to the end point of U_i ends with the last edge of U_i. By the minimality of U_i and Remark 5.3 (3), if c is an edge of U_i in a wall of α_R, then $\bar{u}_i$ and $\bar{c}$ do not commute. Note that when this process stops, $6N < R \leq 7N$. Hence once the following lemma is proved parts (1), (2) and (3) of the proposition are clear.

Lemma 5.6. *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, the commuting elements $\overline{a}_1$, $\overline{a}_2$ and $\overline{a}_3$ are distinct. But $\overline{a}_i$ does not commute with $\overline{u}_i$ for $i = 1, 2, 3$, and the pairs $(\overline{a}_i, \overline{u}_i)$ 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 to be U_1. Note that U_1 and U_2 have length at least $6N$. If $U_1 \neq U_2$, there are two possible further reductions that can be made, each of which leave us with a single direction.

- $\boxed{1}$ If there is no geodesic extension of η that can be rearranged to form a bigon of width $16N^2$ with the down edge path of the diamond at x (Lemma 2.21) containing every wall of U_2, then the U_2 direction is never a consideration in our filter construction "above" x and we remove it from consideration here by redefining U_2 to be U_1, and similarly for U_1 (leaving a single direction). If no geodesic extension of η can lead to a wide bigon in either direction, then we will see that an arbitrary filter (built as in the example in the previous section) has a "small" connected limit set in ∂X (of the type in the conclusion of the main theorem).
- $\boxed{2}$ If U_1 and U_2 share two walls with unrelated labels, consider $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ our back combing from x to $*$. By (5) of Remark 5.3 there are edges a_2 in α_2 and a_1 in α_1 so that both U_1 and U_2 have edges in the same wall as a_1 and a_2. Redefine $U_1 = U_2$ to be a shortest geodesic at x containing an edge in the same wall as a_2. Note that this shortest geodesic contains an edge in the wall of a_1 and both of the original U_1 and U_2 are geodesic extensions of this shortest geodesic (so Lemma 5.4 remains valid for $\{U_1\}$).

Part (6) of the proposition follows from $\boxed{1}$ (since otherwise we would have $U_1^x = U_2^x$). Part (4) of the proposition follows from the next remark.

Remark 5.7. If $U_1 \neq U_2$ after reductions $\boxed{1}$ and $\boxed{2}$, then by Lemma 2.21 and Remark 2.22, U_1 and U_2 can be rearranged as $(U_0, \hat{U}_1)$ and $(U_0, \hat{U}_2)$ respectively, such that $\langle lett(U_0)\rangle$ is finite (so $|U_0| < N$, $|\hat{U}_1| > 5N$, and $|\hat{U}_2| > 5N$), and each wall of $\hat{U}_1$ crosses each wall of $\hat{U}_2$. Intuitively, our two directions are almost orthogonal.

Part (5) of the proposition follows from the next remark, concluding the part of the proposition where we assume $U_1^x \neq U_2^x$. For the next remark, note that x is the $(k+1)^{st}$ vertex of η (since $*$ is the first).

Remark 5.8. If $U_1 \neq U_2$, (η, γ) is a Λ-geodesic and γ' is some rearrangement of (η, γ) whose $(k+1)^{st}$ vertex is of distance at least $14N^2$ from x, then the down edge path τ at x of the diamond (Lemma 2.21) for these two geodesics can be rearranged to begin with either U_1 or U_2, by Lemma 5.4. Both cannot initiate rearrangements of τ, since otherwise there is a $(\mathbb{Z}_2 * \mathbb{Z}_2)^2$ in $\langle lett(\tau)\rangle$, and the diamond at x containing τ determines a $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ in (W, S).

If the reduction process reduces our collection to a single U_1 and $\boxed{2}$ is not part of the reduction, then $|U_1| > 6N$ and certainly (7) of the proposition follows. In this case, (8) and (9) of the proposition follow for exactly the same reasons (2) and (5) do, respectively.

If $\boxed{2}$ is part of the reduction process, then the edges of the remaining U_1 in the same walls as a_1 and a_2 of $\boxed{2}$ satisfy the conclusion of (7). Since the final U_1 is an initial segment of rearrangements of the two directions (call them U_1' and U_2') before applying $\boxed{2}$ and $\{U_1', U_2'\}$ satisfy (2) and (6), we have (8) and (9) satisfied after $\boxed{2}$ reduces the collection $\{U_1', U_2'\}$ to the final $\{U_1\}$. This finishes the proof of the proposition. □

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 $(c + c')$-tracking two CAT(0) geodesics in X, and x_i is the endpoint of $(e_1, \ldots, e_i)$ for all i.

Remark 5.9. As with the filter in Section 4 any geodesic ray at $*$ in our filter (viewed in Λ) will have one of the following forms:

1. $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$;
2. $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$;
3. $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots, e_i, \ell_1, \ell_2, \ldots)$ for some Λ-edges ℓ_j, $\ell_1 \neq e_{i+1}$ (with possibly $i = m$);
4. $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots, d_i, \ell_1, \ell_2, \ldots)$ for some Λ-edges ℓ_j, $\ell_1 \neq d_{i+1}$.

As mentioned in the beginning of this section, our concern is that the CAT(0) geodesics from vertices of these Λ-geodesics to $*$ may be far from x_m, and we wish to build a filter so that this does not occur. However, the Λ-geodesics $(e_1, e_2, \ldots, e_m, e_{m+1}, e_{m+2}, \ldots)$ and $(e_1, e_2, \ldots, e_m, d_{m+1}, d_{m+2}, \ldots)$ already track CAT(0) geodesics, and so we will not need to concern ourselves with these rays. Any property satisfied by rays of type (3) will, by symmetry, hold for rays of type (4). Thus, to simplify notation, for the remainder of our constructions we will only consider rays of type (3). We will also use the convention that any geodesic ray notated as in (3) has $\ell_1 \neq e_{i+1}$.

The next proposition is the main result of the section. It follows mostly from Proposition 5.5 and Remark 5.11, and will be used not only to establish the fact that there are at most two directions towards $*$ from any vertex of a filter for our original rays r and s, but that directions for adjacent vertices are tightly connected to one another. Recall x_i is the end point of $(e_1, e_2, \ldots, e_i)$ and the paths $U_1^{x_i}$ and $U_2^{x_i}$ are defined in Proposition 5.5 for $i \geq m > 7N^2$.

Proposition 5.10. *For $i \geq m$, suppose $\lambda = (\ell_1, \ell_2, \ldots, \ell_n)$ is a geodesic extension of $(e_1, e_2, \ldots, e_i)$ (with $\ell_1 \neq e_{i+1}$). For $1 \leq j \leq n$, let $\lambda_j = (\ell_1, \ldots, \ell_j)$, and let v_j be the endpoint of λ_j. If $U_1^{x_i} \neq U_2^{x_i}$, we have that either for each $1 \leq j \leq n$, there are Λ-geodesics $U_1^{x_i}(\lambda_j) \neq U_2^{x_i}(\lambda_j)$, based at v_j, with the following properties:*

1. *$|U_1^{x_i}(\lambda_j)| > 6N$ and $|U_2^{x_i}(\lambda_j)| > 6N$;*
2. *$U_1^{x_i}(\lambda_j)$ and $U_2^{x_i}(\lambda_j)$ begin geodesics from v_j to $*$;*
3. *If γ_1 begins a geodesic from v_j to $*$ and $|\gamma_1| \geq 7N^2$, then $|\gamma_1|$ can be rearranged to begin with at least one of $U_1^{x_i}(\lambda_j)$ or $U_2^{x_i}(\lambda_j)$;*
4. *$U_1^{x_i}(\lambda_j)$ and $U_2^{x_i}(\lambda_j)$ can be rearranged as $(U_0^{x_i}(\lambda_j), U_1^{x_i'}(\lambda_j))$ and $(U_0^{x_i}(\lambda_j), U_2^{x_i'}(\lambda_j))$ respectively, such that $\langle lett(U_0^{x_i}(\lambda_j))\rangle$ is finite (so $|U_0^{x_i}(\lambda_j)| < N$, $|U_1^{x_i'}(\lambda_j)| > 5N$, and $|U_2^{x_i'}(\lambda_j)| > 5N$), and each wall of $U_1^{x_i'}(\lambda_j)$ crosses each wall of $U_2^{x_i'}(\lambda_j)$;*
5. *If β_j is a geodesic from $*$ to v_j, (β_j, γ) is a geodesic extension of β_j, and γ' is a rearrangement of (β_j, γ) whose $(i+j+1)^{st}$ vertex is of distance at least*

$14N^2$ from v_j, then the down edge path at v_j for the diamond (Lemma 2.21) for (β_j,γ) and γ' can be rearranged to begin with exactly one of $U_1^{x_i}(\lambda_j)$ or $U_2^{x_i}(\lambda_j)$;

(6) $U_1^{x_i}(\lambda_j)$ has at least $6N-3$ walls in common with $U_1^{x_i}(\lambda_{j-1})$, and $U_2^{x_i}(\lambda_j)$ has at least $6N-3$ walls in common with $U_2^{x_i}(\lambda_{j-1})$;

or, for some $j \leq n$, there are Λ-geodesics $U_1^{x_i}(\lambda_j) = U_2^{x_i}(\lambda_j)$, based at v_j, such that:

(7) $U_1^{x_i}(\lambda_j)$ contains two edges with unrelated labels;

(8) $U_1^{x_i}(\lambda_j)$ begins a geodesic from v_j to $*$;

(9) If β_j is a geodesic from $*$ to v_j, (β_j,γ) is a geodesic extension of β_j, and γ' is a rearrangement of (β_j,γ) whose $(i+j+1)^{st}$ vertex is of distance at least $16N^2$ from v_j, then the down edge path at v_j for the diamond (Lemma 2.21) for (β_j,γ) and γ' can be rearranged to begin with exactly $U_1^{x_i}(\lambda_j)$ (possibly because no such γ' exists);

(10) $U_1^{x_i}(\lambda_k) = U_2^{x_i}(\lambda_k)$ for $k \geq j$.

(11) For each $k \geq j$, $U_1^{x_i}(\lambda_k)$ is a shortest geodesic based at v_k containing an edge in every wall of $U_1^{x_i}(\lambda_j)$.

Proof. 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 and the direction(s) at the endpoint of λ. We 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 $v \equiv v_1$ denote the endpoint of ℓ_1.

(1) If $U_1^{x_i} = U_2^{x_i}$ and $\bar{\ell}_1$ commutes with $lett(U_1^{x_i})$, 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 $\bar{\ell}_1$ commutes with $lett(U_1^{x_i})$, then $\bar{\ell}_1 \notin lett(U_1^{x_i})$, since $(\ell_1^{-1}, U_1^{x_i})$ is geodesic. Also note that in this case, $U_1^{x_i}(\ell_1)$ is a shortest geodesic based at v containing an edge in each wall of $U_1^{x_i}$.

(2) If $U_1^{x_i} = U_2^{x_i}$ and $\bar{\ell}_1$ does not commute with $lett(U_1^{x_i})$, then set $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1) = ({\ell_1}^{-1}, U_1^{x_i})$. Note that in this case, $U_1^{x_i}(\ell_1)$ is a shortest geodesic based at v containing an edge in each wall of $U_1^{x_i}$.

(3) If $U_1^{x_i} \neq U_2^{x_i}$, consider directions U_1^v, U_2^v given by Proposition 5.5. If there is only one direction $U_1^v = U_2^v$, set $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1) = U_1^v$. If there are two directions U_1^v and U_2^v, but there is no geodesic extension of $(e_1, e_2, \ldots, e_i, \ell_1)$ that can lead to a $16N^2$ wide bigon in the U_2^v direction at v, then set $U_1^{x_i}(\ell_1) = U_2^{x_i}(\ell_1) = U_1^v$ (and equivalently for U_1^v). If there is no geodesic extension that can lead to a wide bigon in either direction, then we will see that building arbitrary fans, as in the example in the previous 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_{ℓ_1} of $(e_1, e_2, \ldots, e_i, \ell_1)$ so that a rearrangement of $(e_1, e_2, \ldots, e_i, \ell_1, \gamma_{\ell_1})$ gives a $16N^2$ wide bigon at v whose down edge path of the diamond at v (Lemma 2.21) begins with U_1^v. This bigon must be more than $14N^2$ wide at x_i, and so by (5) of Propositon 5.5, 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 is $U_1^{x_i}$ set $U_1^{x_i}(\ell_1) = U_1^v$ and $U_2^{x_i}(\ell_1) = U_2^v$. Otherwise, set $U_1^{x_i}(\ell_1) = U_2^v$ and $U_2^{x_i}(\ell_1) = U_1^v$. It will be made clear by Lemma 5.12 that this choice does not depend on the choice of γ_{ℓ_1}.

We now define $U_1^{x_i}((\ell_1,\ell_2))$ and $U_2^{x_i}((\ell_1,\ell_2))$ by replacing $U_1^{x_i}$, $U_2^{x_i}$, and x_i by $U_1^{x_i}(\ell_1)$, $U_2^{x_i}(\ell_1)$, and v (respectively) 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 $(e_1,e_2,\dots,e_i)$ 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 5.11. For $i\geq m$, let λ be a geodesic extension of $(e_1,e_2,\dots,e_i)$ with $U_1^{x_i}(\lambda)\neq U_2^{x_i}(\lambda)$, and let v be the endpoint of λ. Items (1)-(6) of Proposition 5.5 hold with $U_1^{x_i}(\lambda),U_2^{x_i}(\lambda)$ replacing $U_1^{x_i},U_2^{x_i}$ and v replacing x_i, since $\{U_1^{x_i}(\lambda),U_2^{x_i}(\lambda)\}=\{U_1^v,U_2^v\}$. If $U_1^{x_i}(\lambda)=U_2^{x_i}(\lambda)$, (7) and (8) of Proposition 5.5 still hold, while (9) may not (since $U_1^{x_i}(\lambda)$ may be long).

Lemma 5.12. *Suppose $i\geq m$, λ geodesically extends $(e_1,e_2,\dots,e_i)$, e is an edge with $(e_1,e_2,\dots,e_i,\lambda,e)$ 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 ($\equiv U_1^{x_m}$) and $U_1(\ell_1)$ ($\equiv U_1^{x_m}(\ell_1)$), as in the first step of our $U_i(\lambda)$ construction. Let $\beta=(e_1,e_2,\dots,e_m)$ and γ_{ℓ_1} 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 $(\beta,\ell_1,\gamma_{\ell_1})$ 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 2.21. Note $|\tau|\geq 8N^2$. By Remark 5.11 (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 $(\alpha_1,\dots,\alpha_{6N},\dots)$ 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 does not commute with $\overline{b}_1$. If neither wall $w(b_1)$ nor $w(b_2)$ contains an edge of $U_1(\ell_1)$, 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 $w(b_2)$ contains an edge of $U_1(\ell_1)$, 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 5.12 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 $U_2^{x_i}$. Then, by (4) of Proposition 5.5, $\langle lett(U_1^{x_i}(\ell_1))\rangle$ must 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. If this were the case, then we would not have had $U_1^{x_i}(\ell_1)\neq U_2^{x_i}(\ell_1)$, and the proof of the proposition is finished. □

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 5.13. *Suppose $i\geq m$, and $(\lambda_1,\lambda_2,\lambda_3)$ is a geodesic extending $(e_1,e_2,\dots,e_i)$ (not passing through x_{i+1}) with endpoint v. Let τ be another Λ-geodesic from $*$ to 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-distance at least R from τ. If the down edge path of the diamond at z_J for τ and $(e_1, e_2, \ldots, e_i, \lambda_1, \lambda_2, \lambda_3)$ 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. Let $\beta = (e_1, e_2, \ldots, e_m)$. 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 2.21. (See Figure 5.)

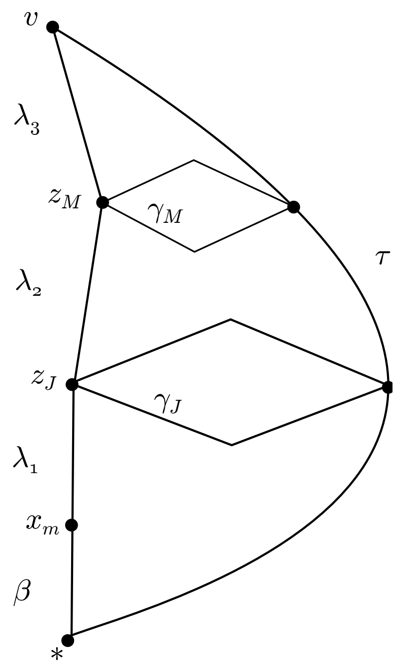

FIGURE 5

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 (5) of Proposition 5.10. By (6) of Proposition 5.10 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 2.21, 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. □

6. Filter construction

From this point on we let $\beta = (e_1, e_2, \dots, e_m)$. The proof of the next lemma basically follows that of Lemma 5.5 of [10].

Lemma 6.1. *Let λ be a geodesic based at x_i extending $(\beta, e_{m+1}, \dots, e_i)$ 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 $(\beta, e_{m+1}, \dots, e_i, \lambda)$ 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 $(\beta, e_{m+1}, \dots, e_i, \lambda)$ can be rearranged to end with γ_2, for any $b \in B(v \to *)$, either $b \in lett(\gamma_2)$ or $b \in \mathrm{lk}(lett(\gamma_2))$. Hence

$$\langle \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)\rangle = \langle \mathrm{lk}(lett(\gamma_2))\rangle \times \langle B(v \to *) - \mathrm{lk}(lett(\gamma_2))\rangle.$$

To see that $\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ does not separate $\Gamma(W, S)$, observe that otherwise, either W 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 $\langle \mathrm{lk}(lett(\gamma_2))\rangle$ is infinite. Note that if $s \in \Gamma$, then $\langle \mathrm{lk}(s)\rangle$ is infinite, since W is one-ended, and so

$$\mathrm{lk}(s) \not\subset \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *),$$

since otherwise $(\mathrm{lk}(s), \mathrm{lk}(s) - B(v \to *))$ is a virtual factor separator. We have two cases:

If $s = t$, then there is a vertex $a \in \Gamma$ adjacent to s with $a \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$. If e is the edge between s and a, we use the path e followed by e^{-1}.

If $s \neq t$, then if $s, t \notin \mathrm{lk}(lett(\gamma_2))$, such a path exists since $\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ does not separate 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 $s \in \mathrm{lk}(lett(\gamma_2))$, then let $a, b \notin \mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ be vertices of Γ adjacent to s and t respectively, and connect a to b outside $\mathrm{lk}(lett(\gamma_2)) \cup B(v \to *)$ as before.

□

Remark 6.2. Edge paths in Γ of the form (e, e^{-1}) and (e, e^{-1}, e) may seem unorthodox, but as in [10], they are combinatorially useful in the filter construction.

Remark 6.3. 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 in Section 4, and notice that Lemma 6.1 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 will now specify our choice of γ at each vertex v in the filter.

Recall once more 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 $(\beta, d_{m+1}, d_{m+2}, \dots)$ $(c + c')$-track two CAT(0) geodesics in X as in Lemma 3.12, and x_i denotes the endpoint of $(\beta, e_{m+1}, \dots, e_i)$, for $i \geq m$.

Definition 6.4. 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 3.11).

Definition 6.5. Suppose λ is a geodesic extending $(\beta, e_{m+1}, \ldots, e_i)$ 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 $(\beta, e_{m+1}, \ldots, e_i, \lambda)$ 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 6.6. 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 (5) of Proposition 5.5. As this is the situation we will usually consider, we may drop the R-value if $R \geq 14N^2$ and simply say that z is wide in one of these directions at x_i.

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, $d_\Lambda(v,w) \geq d_X(C(v), C(w))$. Let $\delta_0 = (\max\{1, \delta_1, c + c'\})$, where δ_1 is the tracking constant from Lemma 3.11 and c, c' are the tracking constants from Lemma 3.12.

Let λ be a geodesic extending $(\beta, e_{m+1}, \ldots, e_i)$ 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 2.24, $A^i(\lambda)$ 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 6.1.

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

Lemma 6.7. *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 $U_1^{x_i}(\lambda_1)$. Let e be an edge that geodesically extends $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2)$ with $\bar{e} \notin lk(lett(\tau))$. Suppose γ is a geodesic extension of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e)$ and γ' is a rearrangement of $(\beta, e_{m+1}, \ldots, e_i, \lambda_1, \lambda_2, e, \gamma)$. Then, if the down edge path at the endpoint of λ_1 contains edges in all the walls of $U_1^{x_i}(\lambda_1)$, then no edge in the wall $w(e)$ 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' .*

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 $\bar{e}$ does not commute with $\bar{c}_1$. 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

$\bar{c}_1$ does not commute with $\bar{c}_2$. The walls $w(c_2)$ and $w(e)$ are on opposite sides of $w(c_1)$ (see Remark 2.18), 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 6.8. Note that Lemma 6.7 does not require that $U_1^{x_i}(\lambda_1) \neq U_2^{x_i}(\lambda_1)$ or $U_1^{x_i}((\lambda_1, \lambda_2)) \neq U_2^{x_i}((\lambda_1, \lambda_2))$. If $U_1^{x_i}(\lambda_1) = U_2^{x_i}(\lambda_1)$, then by (11) of Proposition 5.10, τ (as defined in Lemma 6.7) has the same walls as $U_1^{x_i}((\lambda_1, \lambda_2)) = U_2^{x_i}((\lambda_1, \lambda_2))$.

We now return to the filter construction. Set $a_1 = \bar{e}_{m+1}$ and $b_1 = \bar{d}_{m+1}$. 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 6.1) from a_1 to b_1 in $\Gamma(W, S)$, where each $t_i \notin \text{lk}(lett(A^m)) \cup B(x_m \to *)$. We construct a fan in Λ as before:

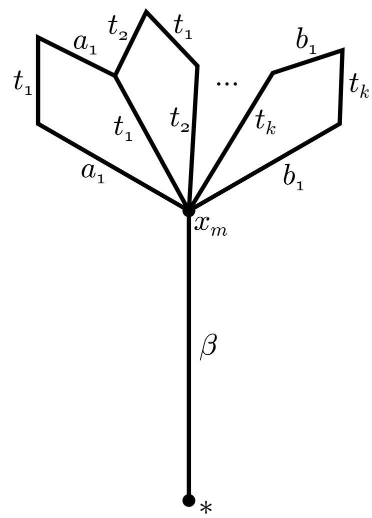

FIGURE 6

Definition 6.9. 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 $1 \leq i \leq k$. Continue constructing the filter by constructing fans avoiding $\text{lk}(lett(A^m((w_i)))) \cup B((\beta, w_i))$ at the endpoint of each w_i, avoiding $\text{lk}(lett(A^{m+1})) \cup B(x_{m+1} \to *)$ at x_{m+1}, and avoiding the appropriate subset of Γ at y_{m+1} (recall by Remark 5.9, we will not consider filter geodesics having d_i edges, as they will be treated analogously to filter geodesics along the e_i edges). Each of these fans is called a *second-level* fan, and each vertex of distance 2 from x_m is called a *second-level* vertex (and will be the base vertex of a third-level fan).

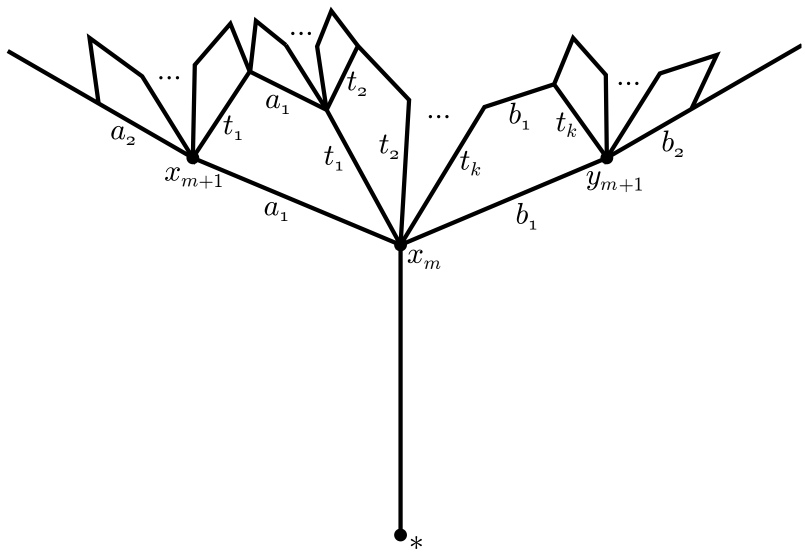

FIGURE 7

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 7. We do not identify these edges; instead, we will construct an edge path between them as described in Lemma 6.1 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.

Definition 6.10. We choose 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 alternate right/left 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 $(\beta, d_{m+1}, d_{m+2}, \dots)$. Removing the non-tree edges of the filter leaves a tree.

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 $(t_1, \dots, t_n)$ 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 the filter F, we say a and e are in the same *filter wall* if there is such a path $(t_1, \dots, t_n)$ in F.

Remark 6.11. The following are useful facts about a filter F for two such geodesics ((1)-(5) from [10]):

(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 a single fan loop contains both.
(4) F minus all non-tree edges is a tree containing $(\beta, e_{m+1}, e_{m+2}, \dots)$ and $(\beta, d_{m+1}, d_{m+2}, \dots)$ 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 $(\beta, e_{m+1}, e_{m+2}, \dots)$ or $(\beta, d_{m+1}, d_{m+2}, \dots)$ 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}.

Lemma 6.12. *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 $(14N^2)$ 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. 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, and that $\delta_0 \geq \delta_1 Q$, where Q is a quasi-isometry constant such that $d_\Lambda(u,x) \leq (d_X(C(u), C(x)))Q$ for $u, x \in \Lambda$ of distance at least N. Also recall that our CAT(0) metric is scaled so that $d_\Lambda(v,w) \geq d_X(C(v), C(w))$

Claim 1: The Cayley geodesic ρ_v is at least $75N^2$ wide at w.

We show that otherwise, Lemma 6.12 holds. Let w' be the vertex of ρ_v satisfying $d_\Lambda(*, w) = d_\Lambda(*, w')$, and let x_i' be the vertex of ρ_w satisfying $d_\Lambda(*, x_i) = d_\Lambda(*, x_i')$. Let w'' be a point on the CAT(0) geodesic from $C(v)$ to $C(*)$ within δ_1 of $C(w')$. Then $d_X(C(w), w'') \leq 75N^2\delta_0 + \delta_1$. Let x'' be a point on the CAT(0) geodesic from $C(w)$ to $C(*)$ within δ_1 of $C(x_i')$. By considering a Euclidean comparison triangle for $\triangle(C(w), w'', C(*))$ (Definition 3.2), we see there is z on the CAT(0) geodesic from w'' to $C(*)$ (and hence on the CAT(0) geodesic from $C(v)$ to $C(*)$), such that $d_X(x'', z) < 75N^2\delta_0 + \delta_1$. So $d_X(C(x_i'), z) < 75N^2\delta_0 + 2\delta_1$. As $d_X(C(x_i), C(x_i')) \leq 21N^2\delta_0$, the X-distance from $C(x_i)$ to the CAT(0) geodesic connecting $C(v)$ and $C(*)$ is $\leq 96N^2\delta_0 + 2\delta_1$ as claimed in Lemma 6.12.

Claim 2: The vertex v is wide in the $U_1^{x_i}(\lambda_w)$ direction at w.

Consider Figure 8, with diamonds for the geodesics λ, ρ_v, and ρ_w as in Lemma 2.21:

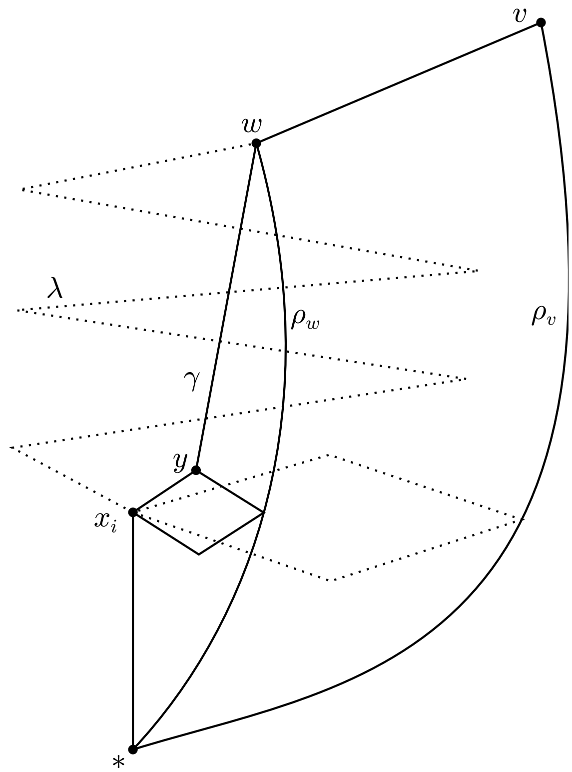

FIGURE 8

Let y be the endpoint of the up edge path of the diamond at x_i 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 the walls of γ_1 appear in the same order as they do on ρ_w (since each wall of γ_1 is also a wall of ρ_w). Let γ be any geodesic from x_i to y followed by γ_1. By Lemma 2.21, 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). Using an argument identical to that in Claim 1, we obtain that γ is of Λ-distance at least $54N^2$ 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 5.13 (with λ_1 trivial), v is $75N^2$ 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, again by Lemma 5.13, 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 $U_2^{x_i}(\lambda_{z_1})$ (by (5) of Proposition 5.10), else by Lemma 2.25, γ and ρ_v are within $6N$ of one another. 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 5.13, 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 3: The vertex v cannot be wide in the $U_1^{x_i}(\lambda_w)$ direction at w.

Let z be any vertex on λ between w and v, and let λ_z denote the initial segment of λ ending at z. Because w was chosen as the last vertex on λ between $20N^2\delta_0$ and $21N^2\delta_0$ wide at x_i, and every vertex after w is at least $21N^2\delta_0$ wide at x_i, $A_i^{x_i}(\lambda_z)$ is a shortest geodesic based at z containing an edge in each wall of $U_1^{x_i}(\lambda_w)$. Because we constructed the fan based at z avoiding $lk(lett(A_i^{x_i}(\lambda_z)))$, none of the interior fan edges have labels in $lk(lett(A_i^{x_i}(\lambda_z)))$, and so by Lemma 6.7, none of these interior fan edges can have walls appearing on the up edge path of a $U_1^{x_i}(\lambda_w)$ diamond at w. Therefore, 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. 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 based at 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 6.11. Thus v is not wide in the $U_1^{x_i}(\lambda_w)$ direction at x_i, which gives the desired contradiction. □

Lemma 6.13. *If λ is a geodesic in the tree T with endpoint v that extends $(\beta, e_{m+1}, \dots, e_i)$ 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 $120N^2\delta_0$ of $C(x_i)$.*

Proof. If $U_1^{x_i} \neq U_2^{x_i}$, let λ_w be the longest initial segment of λ such that $U_1^{x_i}(\lambda_w) \neq U_2^{x_i}(\lambda_w)$, and let w be the endpoint of λ_w. By Lemma 6.12, the CAT(0) geodesic between $C(w)$ and $C(*)$ comes within X-distance $101N^2\delta_0$ of $C(x_i)$. Let y be the vertex of λ after w (with λ_y the initial segment of λ ending at y), so $U_1^{x_i}(\lambda_y) = U_2^{x_i}(\lambda_y)$ and these paths satisfy (7)-(11) of Proposition 5.10. The CAT(0) geodesic between $C(y)$ and $C(*)$ contains a point y' within X-distance $102N^2\delta_0$ of $C(x_i)$.

If $U_1^{x_i} = U_2^{x_i}$, let $y = x_i$, $y' = C(x_i)$, λ_y be trivial, and set $U_j^{x_i}(\lambda_y) = U_j^{x_i}$. Note that $U_1^{x_i} = U_2^{x_i}$ satisfy (7)-(11) of Proposition 5.10, and the CAT(0) geodesic between $C(y)$ and $C(*)$ contains the point $y' = C(y) = C(x_i)$.

In either case, we have that y is a vertex of λ where $U_1^{x_i}(\lambda_y) = U_2^{x_i}(\lambda_y)$ and y' is a point on the CAT(0) geodesic from $C(y)$ to $C(*)$ that is within X-distance $102N^2\delta_0$ of $C(x_i)$. Note that if the CAT(0) geodesic between $C(v)$ and $C(*)$ is more than $18N^2\delta_0$ from $C(y)$, then v (equivalently λ) is at least $16N^2$ wide in the $U_1^{x_i}(\lambda_y)$ direction at y. We have the following cases (from (9) of Proposition 5.10):
Case 1: No geodesic extension of $(\beta, e_{m+1}, \dots, e_i, \lambda_y)$ leads to a bigon $16N^2$ wide at y.

In this case, λ is not $16N^2$ wide in any direction at y so there is $k \geq i$ such that $d_\Lambda(y, \rho_v(k)) \leq 16N^2$. Then there is a point v' on the CAT(0) geodesic from $C(v)$ to

$C(*)$ within X-distance δ_0 of $C(\rho_v(k))$ and so within $16N^2+\delta_0$ of $C(y)$. Consider a Euclidean comparison triangle for $\triangle(v', C(y), C(*))$. By Definition 3.2, some point on the CAT(0) geodesic between v' and $C(*)$ (and so on the geodesic between $C(v)$ and $C(*)$) is within X-distance $16N^2+\delta_0$ of y' and so within X-distance $120N^2\delta_0$ 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 y, then it is wide in the $U_1^{x_i}(\lambda_y)$ direction at y.

From Lemma 6.7 and Remark 6.8, we know that any interior fan edge on λ after y cannot have its wall on the up edge path of a $U_1^{x_i}(\lambda_y)$ diamond at y. If the first edges of λ after y are a right fan edge followed by a left fan edge, the left fan edge shares a wall with an interior fan edge based at y, and so the left fan edge also cannot have an edge in its wall on the up edge path of a $U_1^{x_i}(\lambda_y)$ diamond at y. 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 y). Thus by (6) of Remark 6.11, λ cannot be $16N^2$ wide in the $U_1^{x_i}(\lambda_y)$ direction at y, so some point on the CAT(0) geodesic between $C(v)$ and $C(*)$ is within X-distance $120N^2\delta_0$ of $C(x_i)$. □

Theorem 6.14. *Suppose (W,S) is a one-ended right-angled Coxeter system containing 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 if and only if $\Gamma(W,S)$ does not contain a virtual factor separator.*

Proof. If (W,S) has a virtual factor separator, then by [10], 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 c' are the tracking constants from Lemma 3.12, δ_1 is the tracking constant from Lemma 3.11, and $\delta_0 = (max\{1, \delta_1, c+c'\})$. Recall $C : \Lambda(W,S) \to X$ is W-equivariant. Assume for simplicity $C(*) = *$. Choose M large enough so that for all $m \geq M - c - c'$, if s is an X-geodesic ray based at $*$ within $122N^2\delta_0$ 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 3.12, r and s can be δ_0-tracked by Cayley geodesics α_r and α_s sharing an initial segment of length at least $M - c - c'$. Let $m = M - c - c'$ 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, there is a point v' on the X-geodesic from $C(v)$ to $*$ within $120N^2\delta_0$ of $C(x_i)$ (or $C(y_i)$), where $i \geq m$. There is also a point x' on the X-geodesic from $C(x_i)$ to $*$, within $2\delta_0$ of $C(x_m)$, since α_r and α_s are δ_0-tracking paths for r and s, respectively. Considering a Euclidean comparison triangle for $\triangle(C(x_i), v', *)$ gives a point v'' on the X-geodesic from v' to $*$ (and hence on the geodesic from $C(v)$ to $*$) which passes within $120N^2\delta_0$ of x', and therefore v'' is within $122N^2\delta_0$ 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. □

7. Two Interesting Examples

Let (W, S) be the (one-ended) right-angled Coxeter system with presentation graph Γ give by Figure 9:

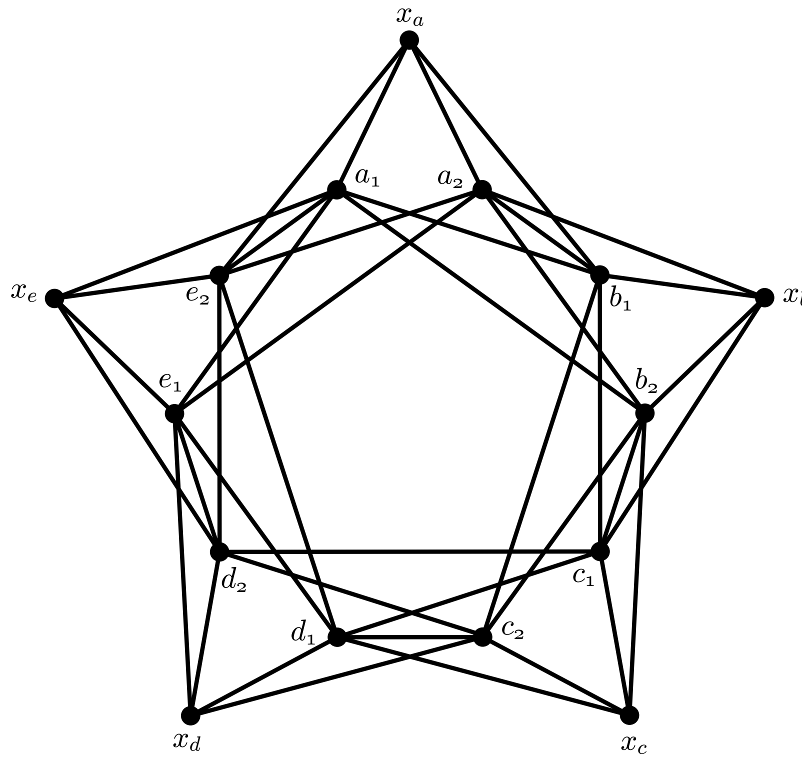

Figure 9

For what follows, let $A = \{a_1, a_2\}$, $B = \{b_1, b_2\}$, $C = \{c_1, c_2\}$, $D = \{d_1, d_2\}$ 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 [9] gives specific conditions for when the boundary of a right-angled Coxeter group is non-locally connected:

Corollary 7.1. *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 $(\text{lk}(v) \cap \text{lk}(w), \text{lk}(v) \cap \text{lk}(w))$ is a virtual factor separator for $\Gamma(W, S)$.

Let $G_1 = \langle S - x_a \rangle$. Note that $\text{lk}(e_1) \cap \text{lk}(e_2) = A \cup D \cup \{x_e\}$ separates e_2 from the rest of $\Gamma(G_1, S - \{x_a\})$, so G_1 has non-locally connected boundary, with

$(e_1e_2)^\infty$ 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 $(e_1e_2)^\infty$ a point of non-local connectivity. Note that we now have $W = G_1 *_Q G_2$, 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.

NOTE - NOT in published version: If any vertex is removed from this example then the resulting graph has virtual factor separators. If a_1 is removed then then $E \cup \{d_2\}$ separates x_e from the rest of the graph and E commutes with D. If x_a is removed then $A \cup C \cup \{b_2\}$ separates $\{b_1, x_b\}$ from the rest of the graph and $A \cup C$ commutes with B. By symmetry the same holds for all vertices.

For our second example consider the right-angled Coxeter group (G, S) with presentation graph of Figure 10.

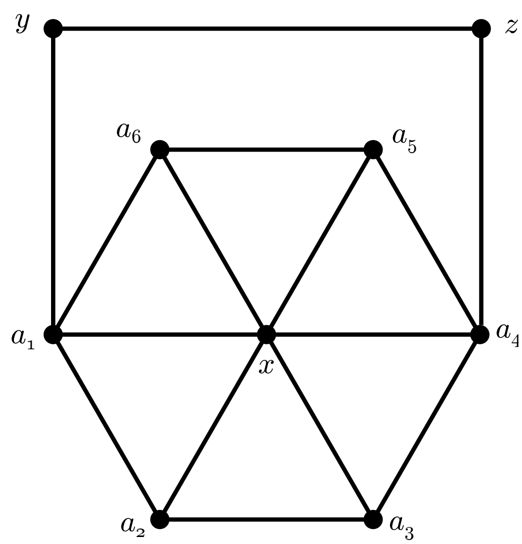

FIGURE 10

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.

8. A Final comment

If the hypothesis that no $(\mathbb{Z}_2 * \mathbb{Z}_2)^3$ is removed in an attempt to classify all right-angled Coxeter groups with locally connected boundary, much of what we develop in this paper carries through. Finitely many directions (as opposed to two) can be defined to measure how large the limit set of a filter becomes. As with our development, if the filter starts to become large in a certain direction at a vertex, it is possible to avoid that direction with subsequent vertices. But when there are only two directions, as in this paper, we are able to show that when we go from being slightly wide in one direction to slightly wide in the other, then the filter did not get too wide in either direction. It seems that when there are more than two directions, CAT(0) geometry of right-angled Coxeter groups is not well enough understood yet to accomplish this.

References

[1] M. Bestvina and G. Mess. *The boundary of negatively curved groups.* J. Amer. Math. Soc. 4 (1991), 469-481.
[2] N. Bourbaki. *Groupes et algèbres de Lie.* Chapitres 4, 5, et 6, Hermann, Paris, 1968.
[3] M.R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature.* Grundlehren Math. Wissensch. 319. Springer-Verlag, 1999.
[4] G. Conner, M. Mihalik and S. Tschantz. *Homotopy of Ends and Boundaries of CAT(0) Groups.* Geometriae Dedicata Volume 120, Number 1, June 2006, 1-17.
[5] C. Croke and B. Kleiner. *Spaces with non-positive curvature and their ideal boundaries.* Topology 39 (2000), 549-556.
[6] M.W. Davis. *The geometry and topology of Coxeter groups.* London Mathematical Society Monographs Series Vol. 32, Princeton University Press, Princeton, NJ, 2008.
[7] T. Januszkiewicz, J. Swiatkowski. *Hyperbolic Coxeter groups of large dimension.* Comment. Math. Helv. 78 (2003), 555-583.
[8] R. Lyndon and P. Schupp. *Combinatorial Group Theory.* Springer-Verlag, New York, 2001.
[9] M. Mihalik and K. Ruane. *CAT(0) groups with non-locally connected boundary.* J. London Math. Soc. (2) 60 (1999), 757-770.
[10] M. Mihalik, K. Ruane, and S. Tschantz. *Local connectivity of right-angled Coxeter group boundaries.* J. Group Theory 10 (2007), 531-560.
[11] M. Mihalik and S. Tschantz. *Visual Decompositions of Coxeter Groups.* Groups, Geometry, and Dynamics 3 (2009), 173-198.
[12] J. Milnor. *A note on curvature and the fundamental group.* J. Differential Geometry 2 (1968), 1-7.
[13] G. Moussong. *Hyperbolic Coxeter groups.* Ph.D. Thesis. Ohio State University (1988).
[14] G. Swarup. *On the cut point conjecture.* Electron. Res. Announc. Amer. Math. Soc. 2 (1996), 98-100 (electronic).

Department of Mathematics, Vanderbilt University, Nashville, TN 37240
E-mail address: w.camp@vanderbilt.edu, michael.l.mihalik@vanderbilt.edu