LOCAL CONNECTIVITY OF RIGHT-ANGLED COXETER GROUP BOUNDARIES

MICHAEL MIHALIK, KIM RUANE, AND STEVE TSCHANTZ

Abstract. We provide conditions on the defining graph of a right-angled Coxeter group presentation that guarantees the boundary of any $CAT(0)$ space on which the group acts geometrically will be locally connected.

0. Introduction

This is an edited version of the published version of our paper. The changes are minor, but clean up the paper quite a bit: the proof of lemma 5.8 is reworded and a figure is added showing where (in the published version) some undefined vertices of the Cayley graph are locate, the proof of theorem 3.2 is simplified by using lemma 4.2 (allowing us to eliminate four lemmas that appeared near the end of section 1).

In [BM],the authors ask whether all one-ended word hyperbolic groups have locally connected boundary. In that paper, they relate the existence of global cut points in the boundary to local connectivity of the boundary. In [Bo], Bowditch gives a correspondence between local cut points in the boundary and splittings of the group over one-ended subgroups. In [S], Swarup uses the work of Bowditch and others to prove that the boundary of such a group cannot contain a global cut point, thus proving these boundaries are locally connected.

The situation is quite different in the setting of $CAT(0)$ groups. If G is a one-ended group acting geometrically on a $CAT(0)$ space X, then ∂X can indeed be non-locally connected. For example, consider the group $G = F_2 \times \mathbb{Z}$ where F_2 denotes the non-abelian free group of rank 2, acting on $X = T \times \mathbb{R}$ where T is the tree of valence 4 (or the Cayley graph of F_2 with the standard generating set). It is easy to see that $\partial X = \Sigma(C)$ (where Σ denotes the unreduced suspension), for C a Cantor set and thus ∂X is not locally connected. Another way of viewing this group is as an amalgamated product of two copies of $\mathbb{Z} \oplus \mathbb{Z}$ over a $\mathbb{Z}$ subgroup with presentation $\langle a, b, c, d | ab = ba, cd = dc, b = c\rangle$. This group naturally acts on a $CAT(0)$ space X which is a union of planes and strips glued according to the amalgamation - in fact, X is simply the universal cover of the presentation two complex where each square is given the metric of a unit square in the plane. In [MR], we prove it is easy to construct these types of examples using amalgamated products

Date: February 26, 2016.

Key words and phrases. $CAT(0)$ group,right-angled Coxeter group, $CAT(0)$ boundary.

that are "geometric" in some sense. In particular, a more general version of the main theorem of that paper is given here as Theorem 3.3 and the above example clearly satisfies the hypotheses.

Theorem 3.3 *Suppose A, B and C are finitely generated groups and $G = A *_C B$ acts geometrically on a $CAT(0)$ space X. If the following conditions are satisfied, then ∂X is not locally connected:*

(1) $[A : C] \geq 2,\ [B : C] \geq 3$.
(2) *There exists $s \in G - C$ with $s^n \notin C$ for all $n \neq 0$ and $sCs^{-1} \subset C$.*
(3) *Cx_0 is quasi-convex in X for a basepoint x_0.*

A general question arising naturally here is whether there is a converse to this theorem. In other words, is there a relationship between the topology of the boundary and certain "geometric" splittings of the groups? It is important to note here that the boundary of a $CAT(0)$ group is not a well-defined object as is the case for a word hyperbolic group. Indeed, in [CK] there is an example of a group G acting geometrically on two $CAT(0)$ spaces with non-homeomorphic boundaries. Nonetheless, any possible boundary for this group will be non-locally connected because this group admits a "geometric" splitting as in the above theorem.

Obtaining a general converse for $CAT(0)$ groups is a very difficult question and as of yet, there is no machinery in place to prove such a result, thus it is worthwhile to test whether such a result holds for a class of $CAT(0)$ groups that have some structure to work with. In this paper, we consider this question for right-angled Coxeter groups. Right-angled Coxeter groups were shown to be $CAT(0)$ groups in [G] while general Coxeter groups were shown to be $CAT(0)$ in [Mo].

In [MR], the authors give a consequence of Theorem 3.3 above for testing whether the $CAT(0)$ Coxeter complex for a general Coxeter group has non-locally connected boundary using a presentation graph for the group. Suppose G is a Coxeter group and S is a generating set for G which gives rise to a Coxeter presentation $P = P(G, S)$ for G. The presentation can be encoded in a graph $\Gamma = \Gamma(G, S)$ called the *presentation graph*, whose vertex set is S and edge set comes from the relations of P (see Section 1 for details). Loosely, the consequence says (it is stated below as Theorem 3.2), if certain graph-theoretic conditions hold for Γ, then the boundary of the $CAT(0)$ Coxeter complex contains points of non-local connectivity. The graph-theoretic conditions are the weakest possible conditions to guarantee that the Coxeter group splits as an amalgamated product as in Theorem 3.3. where the factors in the splitting are subgroups generated by subsets of the generating set. It is Theorem 3.2 for which we are trying to obtain a converse in the right-angled case. Our first step is to show that if the graph-theoretic conditions hold, then the boundary of any $CAT(0)$ space on which the group acts will have non-locally connected boundary, not just the boundary of the $CAT(0)$ Coxeter complex. This is the contents of Theorem 3.2. An improved version of Theorem 3.3 is necessary to obtain this so we provide proofs of both theorems here in Section 3. These results can be omitted if one is only interested in the main theorem.

The statement of the main theorem is given below. To understand the statement of the main theorem, we briefly explain notation and terminology here, but give formal definitions in the paper.

If (G, S) is a Coxeter system with finite generating set, then we will denote the presentation graph by $\Gamma(G, S) = \Gamma(S) = \Gamma$ if the G and S are clear. We denote the Cayley graph by $\Lambda(G, S) = \Lambda$. We say a subset C of S is a *product separator* of Γ if $C = A \cup B$ with $A \cap B = \emptyset$, $\langle C\rangle = \langle A \cup B\rangle = \langle A\rangle \oplus \langle B\rangle$, $\langle A\rangle$ and $\langle B\rangle$ are infinite and C separates the graph Γ. A *virtual factor separator* (VFS) is a triple (C, C_1, K) where C separates Γ, $\langle C_1\rangle$ has finite index in $\langle C\rangle$, $\langle K\rangle$ is infinite and all the letters of K commute with the letters of C_1. We may reduce the general problem of interest to the case when G does not *visually* split as a direct product. This means the graph Γ cannot be decomposed into two non-empty disjoint sets where each vertex of one set is joined to every vertex of the other set and vice versa. Thus we can assume not all the letters of $S - C$ commute with all the letters in C. And finally, a *suspended separator* is a special type of VFS where $S = C \cup \{s, t\}$ and Γ (as a graph) is the suspension of C with s and t as the suspension points. In this case, $(C, C, \{s, t\})$ forms a VFS.

Main Theorem *Suppose (G, S) is a right-angled Coxeter system, G is one-ended, $\Gamma(S)$ contains no product separator and no VFS. Also assume that G does not visually split as a non-trivial direct product with infinite factors. Then G has locally connected boundary.*

Note, if G has more than one end, then it is trivial to determine if G has locally connected boundary or not so it is reasonable to assume G is one-ended. If G splits as a visual direct product, then G has non-locally connected boundary iff one of the factor groups does, so for the entire paper, we assume the group G does not visually split as a direct product unless we specifically indicate that it does. We now state Theorem 3.2 which is a consequence of Theorem 3.3 above. This result sheds light on why the other hypotheses are necessary in the Main Theorem.

Theorem 3.2 *Suppose (G, S) is a Coxeter group with presentation graph Γ.*

(1) *If Γ has a suspended separator C, then G has locally connected boundary if and only if $\langle C\rangle$ has locally connected boundary.*
(2) *If Γ has a VFS (C, C_1, K) and C is not a suspended separator, then G has non-locally connected boundary.*

To see the extent to which our Main Theorem gives a converse to Theorem 3.3, we point out the following combination of the Main Theorem and Theorem 3.2.

Corollary *Suppose (G, S) is a right-angled Coxeter system, G is one-ended and G does not visually split as a non-trivial direct product. If the presentation graph Γ contains a VFS, then G has non-locally connected boundary. If G contains no product separator and no VFS, then G has locally connected boundary.*

In Sections 4 and 5, the main technical tools for the proof are developed. From this point on, all our groups are right-angled since our methods do not work in general Coxeter groups. Section four contains two important lemmas which allow us to approximate "close" $CAT(0)$ geodesic rays by Cayley graph geodesic rays that have a common initial subsegment (lemmas 4.2 and 4.4). In general $CAT(0)$ group theory, one cannot hope to use Cayley graph geometry to study the geometry of geodesics in the $CAT(0)$ space since the natural quasi-isometry between them is badly behaved. In this paper, we prove that if the group is a right-angled Coxeter group, we can gain enough control over this quasi-isometry to use the intuition coming from the Cayley graph geometry inside the $CAT(0)$ space. This is truly the main technical difficulty in the paper.

Section 5 contains the construction of a *filter* for a pair of Cayley graph geodesics needed to prove the main theorem. In doing so, we build a collection of graphs in the plane whose edges are labeled by elements of the generating set S. There is an obvious map from any such graph into the Cayley graph Λ. A final version of these graphs will be 1-ended and the natural map from this graph into Λ will be *proper*. If the two geodesics we start with have a long common initial piece, we will show that this filter maps over to X under the natural quasi-isometry in a controlled manner.

In Section 6, we give the proof of the main theorem. Our goal is to prove local connectivity of the boundary so we start by considering two "close" geodesic rays (equivalently, two close boundary points) in the $CAT(0)$ space. From these, we obtain two approximating Cayley graph geodesic which share a long common subpiece. Using Section 5, we build a filter between these two rays which has a nice property. In particular, the graph-theoretic hypotheses allow us to conclude there is a global bound on the length of a *factor path* in the filter (see lemma 5.10). A factor path is an edge path in the Cayley graph that lies in a product subgroup (see definition 5.2) - these are exactly the paths that behave badly under quasi-isometry. Carrying this information back to X via the natural quasi-isometry allows us to construct a "small" connected set in the boundary of the space containing the two points and hence show the $CAT(0)$ space has locally connected boundary.

1. Coxeter group preliminaries

In this section, we prove several technical facts about Coxeter groups. Many of these facts can be found scattered about in the vast literature on Coxeter groups, however we state them here with our notation and with our particular viewpoint in mind. Even though Theorem 6.1 only holds for right-angled Coxeter goups, we prove the results in this section assuming G is a general Coxeter group. We will need the more general results in section 3 of this paper.

There are two types of results in this section. The results up through lemma 1.9 give geometric information about the Cayley graph of a Coxeter group. The most important lemmas for understanding the proof of the main theorem are theorem 1.1 and lemmas 1.4, 1.5, 1.7. We need these geometric results in sections 4 and 5 where we must compare the geometry of the $CAT(0)$ space X with the combinatorial properties of the Cayley graph Λ. The remaining results culminate with lemma 1.12. This final lemma says if H is a *special* subgroup of G of finite index, then G splits as a direct product in a way that you can *see* in the presentation graph of G. This lemma is used to prove lemma 5.3 which is necessary to begin building a filter for a pair of geodesic rays.

A *Coxeter group* is a group G with generating set $S = \{s_1, \dots s_n\}$ having a *Coxeter presentation* of the form $P(G,S) = \langle S \mid (s_is_j)^{m_{ij}} = 1\rangle$ for $1 \leq i,j \leq n$ with $m_{ij} = 1$ if and only if $i = j$, otherwise $2 \leq m_{ij} \leq \infty$. If $m_{ij} = \infty$, then s_i, s_j have no relation. The pair (G,S) is called a *Coxeter System*. If $m_{ij} = 2$ or ∞ for all relations in the presentation, G is called *right-angled*.

The *presentation graph* of (G,S) is the graph $\Gamma(G,S)$ with vertex set S and an edge labeled m between vertices s and t if $(st)^m = 1$ is a relation of P. For right-angled groups, we omit the m labels, thus two generators commute if and only if there is an edge between them in the presentation graph.

Assume (G,S) is a Coxeter system with S a finite set. Let $\Lambda(G,S) = \Lambda$ be the Cayley graph of G with respect to the generating set S. We assume each edge of Λ is labeled by an element of S. Each $s \in S$ is an element of order two and has a fixed point set in Λ denoted by $\mathrm{F}ix(s)$. This set consists of the midpoints of edges labeled by s with each endpoint in the subgroup of G generated by the letters of S that are adjacent to s in $\Gamma(G,S)$).

An *edge path* in Λ can be denoted as $(e_1, \dots e_n)$ where $e_i \in S$ for all i as long as the initial vertex of the path is specified. In many of our arguments the initial vertex is not important as the argument is valid for such a path at any vertex. As an example, whether or not a path is geodesic does not depend on the initial vertex. In these cases we may supress initial vertices.

Perhaps the most useful combinatorial and geometric fact about Coxeter groups is the following which is referred to as the *Deletion Condition* in this paper. It is used in this paper in any result that involves the geometry of the Cayley graph Λ.

Theorem 1.1. (The Deletion Condition). *If $w = w_1 \dots w_n$ is non-geodesic in Λ then there are $0 < i < j \leq n$ such that*

$$w = w_1 \dots w_{i-1} w_{i+1} \dots w_{j-1} w_{j+1} \dots w_n.$$

Remark. In this case, we say that w_i *deletes* with w_j. Observe that if $w_1 \dots w_{n-1}$ is geodesic, but $w_1 \dots w_{n-1} w_n$ is not geodesic, then we know w_n must delete with one of the previous letters. If the group is right-angled, then w_n deletes with w_i for some $1 \leq i < n$. In fact, $w_i = w_n$ and i is the largest integer between 1 and $n-1$ with $w_i = w_n$. Also, w_n commutes with w_j for all $i < j \leq n-1$.

Definition 1.2 (Special subgroup)**.** *Suppose (G, S) is a Coxeter system and let $V \subset S$. Then the subgroup of G generated by V is called a special subgroup of G.*

The next result shows why the special subgroups are indeed special.

Theorem 1.3. Special Subgroup Theorem. *If (G, S) is a Coxeter system and $V \subset S$ then the subgroup $\langle V \rangle$ is Coxeter with presentation $\langle V | R \rangle$ where R is the set of all relations of $P(G, S)$ involving only letters from V.*

The next two lemmas are part of the standard theory of Coxeter groups and proofs can be found in [H] for example. The first of these lemmas says that special subgroups are convex in $\Lambda(G, S)$.

Lemma 1.4. *If $U \subset S$ and $u \in \langle U \rangle$ is written as a geodesic word in U then this word is geodesic in S. Furthermore every geodesic for u can only use letters from U.*

If α is a path in Λ then define $\bar{\alpha} \in G$ to be the product of the edge labels of α. The following lemma says that if we add one more edge to a geodesic path of length n in Λ, then the new path cannot also be of length n. Also, the set of generators that make this new path have length $n-1$ only generate a finite special subgroup. This will be very important in lemmas 4.3 and 5.6.

Lemma 1.5. *If α is geodesic in Λ then the set of labels of edges e such that (α, e) is not geodesic generates a finite subgroup of G, denoted $B(\bar{\alpha})$. Furthermore, either $l(\alpha e) = l(\alpha) - 1$ or $l(\alpha e) = l(\alpha) + 1$*

The next three lemmas are geometric consequences of the convexity of the special subgroups which will be useful in the proof of the main theorem. The first of these follows directly from the Deletion Condition and lemma 1.4. It allows us to construct geodesics in Λ by combining geodesics from special subgroups with non-overlapping generating sets.

Lemma 1.6. *Suppose (G, S) is a Coxeter system, U and V are subsets of S with $U \cap V = \emptyset$. If α is a geodesic in the letters of U and β is a geodesic in the letters of V, then (α, β) is a geodesic in $\Lambda(G, S)$.*

The next lemma is essentially a more technical way of saying the special subgroups are convex. Specifically, if you start with a word of the form sut where $s, t \in S$ with t, u both in a particular special subgroup (the one generated by the letters of u). Then if the edge path sut ends in the subgroup, s must also be in that subgroup since s must equal t and this generator commutes with the letters of the word u. This form of the convexity is used in many of the results of this section and in lemmas 4.3 and 5.4.

Lemma 1.7. *Suppose $U \subset S$. If $s \in S - U$, $t \in S$ and for some $u \in \langle U \rangle$, $sut \in \langle U \rangle$, then*

(1) *$t = s$.*
(2) *$sus = u$ and*
(3) *If $u = u_1 \dots u_n$ is geodesic, (so $u_i \in U$), then $su_is = u_i$ for all i.*

Proof. As $sut \in \langle U \rangle$, $s \in \langle U \cup \{t\} \rangle$. By lemma 1.4, $t = s$ and (1) is finished.

Say $u = u_1 \dots u_n$ is geodesic, where $u_i \in U$. By lemma 1.4, this is geodesic in the entire Coxeter system. Say $sus = v_1 \dots v_m$, m minimal and $v_i \in U$.

The words $su_1 \dots u_n$ and $v_1 \dots v_m s$ are geodesics (for the same element of G) by lemma 1.6 and so $m = n$.

Now $su_1 \dots u_n s$ is not geodesic. The second s cannot delete with a u_i as $s \notin \langle U \rangle$. So s commutes with u, finishing (2).

To prove (3), we proceed by induction on n. If $n = 1$, (3) is clear from (2). Assume (3) is true for $k \leq n - 1$ and $n \geq 2$. The path $u_1 \dots u_n s$ is geodesic (as $s \notin U$), but $u_1 \dots u_n s u_n$ ($= su_1 \dots u_{n-1}$) is not geodesic. Clearly the last u_n and s do not delete as $s \neq u_n$. Say the last u_n and u_i delete. This implies (by part (2)) that s commutes with both $u_1 \dots u_{n-1}$ and $u_1, \dots u_{i-1} u_{i+1} \dots u_n$. By induction, s commutes with u_i for all i. □

It is an elementary exercise using the deletion condition to show that because of convexity of the special subgroups, we also have the notion of a unique projection onto a special subgroup (or one of it's cosets).

Lemma 1.8. *Suppose $*$ and v are vertices of Λ and $T \subset S$. Then there is a unique vertex w of $v\langle T \rangle$ closest to $*$. Furthermore, if α is a geodesic from $*$ to w and β is a geodesic in the letters of T, then (α, β) is a geodesic in Λ.*

Lemma 1.8 and the Deletion Condition now imply:

Lemma 1.9. *Suppose β is a geodesic from x to y in Λ and (β, s) is not geodesic. If t is the last letter of β then s is related to t and the edge path labeled by half of this relation, at y, is geodesic towards x.*

Definition 1.10 (2-link). *For $A \subset S$, we define $lk^2(A) = \cap_{a \in A} lk^2(a)$ where $lk^2(a)$ means all $s \in S$ that are connected to a in the graph $\Gamma(G, S)$ with an edge labeled 2. In particular, for $s \in S$, we often write $lk^2(s)$ instead of $lk^2(\{s\})$.*

We make the following observations concerning the right-angled Coxeter groups. These statements will be used in the last three sections of the paper when we restrict to the right-angled case.

Remark. Let (G,S) be right-angled and $H \subset G$. Let $C(H)$ denote the centralizer of H in G.

(1) For (G,S) right-angled with $A \subset S$, $lk^2(A) = lk(A)$ where $lk(A)$ is the usual link of A in the graph $\Gamma(G,S)$.
(2) For $s \in S$, $C(s)/\langle s\rangle = \langle lk(s)\rangle$. This follows from lemma 1.7 and the Deletion Condition. In particular, $\langle lk(s)\rangle$ is an index two subgroup of $C(s)$ in this case.

Lemma 1.11. *If V does not visually split non-trivially as a direct product, then for any geodesic α in V and $v \in V$, α can be extended to a geodesic ending with v.*

Proof: If not, choose α a geodesic such that $Y_\alpha = \{y \in S|\ \alpha$ cannot be extended to a geodesic ending with $y\}$ is as large as possible. By hypothesis, $Y_\alpha \neq \emptyset$ and by lemma 1.5, $Y_\alpha \neq S$. By the maximality of Y_α, α can be extended to an infinite geodesic (α,γ) such that the letters of $S - Y_\alpha$ occur infinitely often in γ. We show Y_α commutes with $S - Y_\alpha$ to obtain a contradiction.

Assume $\alpha = (a_1,\ldots,a_n)$ and $\gamma = (b_1,b_2,\ldots)$. Let $\beta_i = (\alpha,b_1,b_2,\ldots,b_i)$. For $y \in Y_\alpha$, (β_i,y) is not geodesic. This last y cannot delete with a b_j as $y \notin S - Y_\alpha$ and so must delete with some a_k. Choose $i < j$ such that all letters of $S - Y_\alpha$ occur in $\{b_{i+1},\ldots,b_j\}$ and the y ending (β_i,y) and (β_j,y) both delete with a_k. Then $a_k \ldots a_n b_1 \ldots b_i = a_{k+1}\ldots a_n b_1 \ldots b_i y$ and $a_k \ldots a_n b_1 \ldots b_j y = a_{k+1} \ldots a_n b_1 \ldots b_j$. So $a_{k+1}\ldots a_n b_1 \ldots b_i y b_{i+1} \ldots b_j y = a_{k+1} \ldots a_n b_1 \ldots b_j$ and $y b_{i+1} \ldots b_j y = b_{i+1} \ldots b_j$. By lemma 1.7, y commutes with $S - Y_\alpha$ as needed. $\Box$

Lemma 1.12. Finite Index Lemma *Suppose (G,S) is a Coxeter system, G is infinite, $U \subset S$, $\langle U\rangle$ and has finite index in G and $F \subset U$ is the maximal set such that $\langle F\rangle$ is finite and $\langle U\rangle$ visually splits as $\langle F\rangle \oplus \langle U-F\rangle$ then $\langle (S-U)\cup F\rangle$ is finite and G splits as $\langle (S-U)\cup F\rangle \oplus \langle U-F\rangle$.*

Proof: Clearly $\langle U-F\rangle$ has finite index in V. Hence there is a bound K such that every vertex of $\Lambda(G,S)$ is within K of $\langle U-F\rangle$. If $\langle (S-U)\cup F\rangle$ were not finite, choose a geodesic α of length $K+1$ in the letters $(S-U)\cup F$, with initial point the identity. The end point y of this geodesic is within K of a vertex $x \in \langle U-F\rangle$. Let β be a geodesic in the letters of $U-F$ from x to the identity. By the lemma 1.6, (β,α) is a geodesic from x to y, but $d(x,y) \leq K$, giving the desired contradiction. So $\langle (S-U)\cup F\rangle$ is finite.

Maximally decompose $\langle U-F\rangle$ as $\oplus_{i=1}^n \langle W_i\rangle$ with $W_i \neq \emptyset$, $\cup_{i=1}^n W_i = U-F$, and $W_i \cap W_j = \emptyset$ for $i \neq j$. By the lemma 1.11, there is a geodesic α_i with letters in W_i that begins at $1 \in \Lambda$ such that each letter of W_i occurs infinitely often in α_i. Let β be an edge path at 1 such that each element of

$(S-U)\cup F$ labels exactly one edge of β and β has length $Card((S-U)\cup F)$. By lemma 1.4, β is geodesic. Say $\alpha_i = (a_{i1}, a_{i2}, \ldots)$ and $\beta = (b_1, \ldots, b_n)$. Choose k such that all letters of W_i occur at least $n + K + 1$ times in $\alpha_{ik} = (a_{i1}, \ldots, a_{ik})$. Let x be the end point of β and y the end point of α_{ik}. Let δ be a geodesic from y to $x\langle U - F\rangle$ of length $\leq K$. Say δ ends at z. Let $\gamma = (c_1, \ldots, c_m)$ be a geodesic from x to z with $c_j \in (U - F)$ for all j. By the lemma 1.6 (α_i^{-1}, β) is geodesic and by lemma 1.8, (δ, γ^{-1}) is also geodesic. Hence $(\alpha_{ik}^{-1}, \beta, c_1)$ is not geodesic.

Say c_1 deletes wtih a_{ij}. Since $b_i \notin U - F$ for all i, we must have c_1 commuting with b_j for all j by lemma 1.7. From the Deletion Condition we see that if $(u_1 \ldots, u_n)$ and $(v_1, \ldots, v_n)$ are geodesics with the same endpoints, then $\{u_1, \ldots, u_n\} = \{v_1, \ldots, v_n\}$, thus we obtain $\{c_1, \ldots, c_m\} \subset U - F$. If we replace α_i with a geodesic in $U - F$ from c_1 to y, we see that c_2 commutes with each b_j. Continuing, each b_j commutes with each c_q. Now $c_1 \ldots c_m = ai1 \ldots a_{ik}\delta\beta^{-1}$. At most $n + K$ deletions can occur in the right hand side before a geodesic is obtained. Hence all letters of $U - F$ appear in the resulting reduced word. So all letters of $U - F$ appear in $\{c_1 \ldots, c_m\}$ forcing the letters of $U - F$ to commute with the letters of $(S - U) \cup F$ as needed. $\square$

2. $CAT(0)$ spaces and their boundaries

In this section we give definitions and basic properties of $CAT(0)$ spaces, boundaries and isometries as well as some known facts we will need in the proof of the main result.

Let (X, d) be a metric space. Then X is *proper* if closed metric balls are compact. A (*unit speed*)*geodesic* from x to y for $x, y \in X$ is a map $c : [0, D] \to X$ such that $c(0) = x$, $c(D) = y$ and $d(c(t), c(t')) = |t - t'|$ for all $t, t' \in [0, D]$. If $I \subseteq \mathbb{R}$ then a map $c : I \to X$ parametrizes its image *proportional to arclength* if there exists a constant λ such that $d(c(t), c(t')) = \lambda|t - t'|$ for all $t, t' \in I$. Lastly, (X, d) is a called a *geodesic metric space* if every pair of points are joined by a geodesic.

Definition 2.1 ($CAT(0)$)**.** *Let (X, d) be a proper complete geodesic metric space. If $\triangle\, abc$ is a geodesic triangle in X, then we consider $\triangle\, \overline{a}\overline{b}\overline{c}$ in $\mathbb{E}^2$, a triangle with the same side lengths, and call this a comparison triangle. Then we say X satisfies the $CAT(0)$ inequality if given $\triangle\, abc$ in X, then for any comparison triangle and any two points p, q on $\triangle\, abc$, the corresponding points $\overline{p}, \overline{q}$ on the comparison triangle satisfy*

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

If (X, d) is a $CAT(0)$ space, then the following basic properties hold:

(1) The distance function $d\colon X \times X \to \mathbb{R}$ is convex.
(2) X has unique geodesic segments between points.
(3) X is contractible.

For details, see [BH].

Let (X, d) be a proper C$AT(0)$ space. First, define the boundary, ∂X as a set as follows:

Definition 2.2 (Asymptotic)**.** *Two geodesic rays $c, c' \colon [0, \infty) \to X$ are said to be asymptotic if there exists a constant K such that $d(c(t), c'(t)) \leq K, \forall t > 0$ - this is an equivalence relation. The boundary of X, denoted ∂X, is then the set of equivalence classes of geodesic rays. The union $X \cup \partial X$ will be denoted $\overline{X}$. The equivalence class of a ray c is denoted by $c(\infty)$.*

There is a natural neighborhood basis for a point in ∂X. Let c be a geodesic ray emanating from x_0 and $r > 0$, $\epsilon > 0$. Also, let $S(x_0, r)$ denote the sphere of radius r centered at x_0 with $p_r : X \to S(x_0, r)$ denoting projection. 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 $\overline{X}$ such that when projected back to $S(x_0, r)$, this projection is not more than ϵ away from the intersection of that sphere with c. These sets along with the metric balls about x_0 form a basis for the *cone topology.* The set ∂X with the cone topology is often called the *visual boundary.* As one expects, the visual boundary of $\mathbb{R}^n$ is S^{n-1} as is the visual boundary of $\mathbb{H}^n$.

Definition 2.3. *Let γ be an isometry of the metric space X. The displacement function $d_\gamma \colon X \to R_+$ is defined by $d_\gamma(x) = d(\gamma \cdot x, x)$. The translation length of γ is the number $|\gamma| = inf\{d_\gamma(x) : x \in X\}$. The set of points where γ attains this infimum will be denoted Min(γ). An isometry γ is called semi-simple if Min(γ) is non-empty.*

We summarize some basic properties about this Min(γ) in the following proposition.

Proposition 2.4. *Let X be a metric space and γ an isometry of X.*

1. *Min(γ) is γ-invariant.*
2. *If α is another isometry of X, then $|\gamma| = |\alpha\gamma\alpha^{-1}|$, and Min($\alpha\gamma\alpha^{-1}$) = $\alpha \cdot$ Min(γ); in particular, if α commutes with γ, then it leaves Min(γ) invariant.*
3. *If X is $CAT(0)$, then the displacement function d_γ is convex: hence Min(γ) is a closed convex subset of X.*

Definition 2.5. *Let X be a metric space. An isometry γ of X is called*

1. *elliptic if γ has a fixed point - i.e $|\gamma| = 0$ and Min(γ) is non-empty.*
2. *hyperbolic if d_γ attains a strictly positive infimum.*
3. *parabolic if d_γ does not attain its infimum, in other words if Min(γ) is empty.*

It is clear that an isometry is semi-simple if and only if it is elliptic or hyperbolic. If two isometries are conjugate in Isom(X), then they are in the same class.

If a group Γ acts geometrically on a C$AT(0)$ space X, then the elements of Γ act as semi-simple isometries because of the cocompactness of the action.

The next theorem implies that the centralizer of an element γ in a C$AT(0)$ group Γ is again a C$AT(0)$ group because it acts geometrically on C$AT(0)$ subspace M$in(\gamma)$. A proof of this can be found in [R].

Theorem 2.6. *Suppose Γ acts geometrically on the $CAT(0)$ space X and suppose $\gamma \in \Gamma$. Then $C_\Gamma(\gamma)$ acts geometrically on the $CAT(0)$ subset $Min(\gamma)$ of X.*

Remarks. We make the following remarks which will be needed in our setting of right-angled Coxeter groups.

(1) If (G, S) is a right-angled Coxeter group acting geometrically on the C$AT(0)$ space X, then each $s \in S$ is an elliptic isometry with M$in(s) =$ F$ix(s)$ a closed, convex subset of X on which the centralizer of s acts geometrically.

(2) Since $C(s)$ acts geometrically on F$ix(s)$ and $\langle lk(s)\rangle$ is an index two subgroup of $C(s)$ (see remarks following definition 1.10) thus $\langle lk(s)\rangle$ is a C$AT(0)$ subgroup of G.

We will need the following result concerning C$AT(0)$ products. Proofs of these can be found in [BH].

Theorem 2.7. *Suppose X and Y are $CAT(0)$ spaces.Then $X \times Y$ with the product metric is also a $CAT(0)$ space. Furthermore, $\partial(X \times Y)$ is homeomorphic to the spherical join of ∂X and ∂Y. In particular, if $Y = \mathbb{R}$, then $\partial(X \times \mathbb{R})$ is homeomorphic to the suspension of ∂X.*

The following theorem is due to Milnor in [M] and will be used throughout the paper to carry information between the Cayley graph of (G, S) and the C$AT(0)$ space X.

Theorem 2.8. *If a group G with finite generating set S acts geometrically on a proper, geodesic metric space X, then the Cayley graph of G with respect to S is quasi-isometric to X under the map $g \mapsto g \cdot x_0$ where x_0 is a fixed basepoint in X.*

3. Local Connectivity

This section contains the proof of Theorem 3.2 below. The material in this section is not needed in the proof of the main theorem. 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. Clearly, if G is infinite ended, then G has non-locally connected boundary, thus we only need to consider Coxeter groups (G, S) for which the presentation graph Γ is connected.

Definition 3.1 (VFS)**.** *Suppose (G, S) is a Coxeter system. Let Γ denote the presentation graph for (G, S). A virtual factor separator (VFS) for Γ is a triple (C, C_1, K) where $C_1 \subset C$ are full subgraphs of Γ, C separates Γ, the group $\langle C_1 \rangle$ has finite index in $\langle C \rangle$, $K \subset lk^2(C_1)$, and $\langle K \rangle$ is infinite. If $S = C \cup \{s, t\}$, s and t do not span an edge of Γ and $s, t \in lk^2(C)$, then C is called a suspended separator of Γ. Note that if C is a suspended separator of Γ, then $G = \langle C \rangle \oplus \langle s, t \rangle$ and $(C, C, \{s, t\})$ is a VFS.*

Theorem 3.2. *Suppose (G, S) is a Coxeter group with presentation graph Γ.*

1. *If Γ has a suspended separator C, then G has locally connected boundary if and only if $\langle C \rangle$ has locally connected boundary.*
2. *If Γ has a VFS (C, C_1, K) and C is not a suspended separator, then G has non-locally connected boundary.*

The proof of Theorem 3.2 will be an application of the following Theorem from [MR]. The theorem cannot be applied directly because of condition three below, but lemma 4.2 of the next section implies condition three holds for special subgroups of Coxeter groups.

Theorem 3.3. *Suppose A, B and C are finitely generated groups and $G = A *_C B$ acts geometrically on a $CAT(0)$ space X. If the following conditions are satisfied, then ∂X is not locally connected:*

1. *$[A : C] \geq 2$, $[B : C] \geq 3$.*
2. *There exists $s \in G - C$ with $s^n \notin C$ for all $n \neq 0$ and $sCs^{-1} \subset C$.*
3. *Cx_0 is quasi-convex in X for a basepoint x_0.*

Remark. If C_1 is a subgroup of finite index in C, and condition (2) is replaced by (2′) below then the proof remains unchanged.

(2′) There exists $s \in G - C_1$ with $s^n \notin C$ for all $n \neq 0$ and $sC_1s^{-1} \subset C_1$.

Proof of Theorem 3.2

If C is a suspended separator, then G splits as $\langle C \rangle \oplus \langle \{s, t\} \rangle$. In this case, the element (st) is virtually central -meaning, it is central in a subgroup of finite index in G. In fact, the subgroup generated by C and the element (st) is a subgroup of finite index in G in which st is central. Theorem 3.4 of [R] implies that if X is any $CAT(0)$ space on which G acts geometrically, then the following are true. A subset Z of X is *quasi-dense* if there exists a constant $K \geq 0$ such that every $x \in X$ is within K of an element of Z.

(1) There exists a quasi-dense closed, convex subset Z of Xsuch that $\partial Z = \partial X$ (we know $\partial Z \subset \partial X$ since Z convex in X. Thus for boundary considerations, we only need to know about the space Z.
(2) Z splits isometrically as $Y \times \mathbb{R}$ where Y is a closed, convex subset of X and Y admits a geometric group action by the group $\langle C \rangle$.

Since Z splits as a product, we know $\partial Z = \partial X \equiv \Sigma(\partial Y)$ by Theorem 2.7. Thus ∂X is locally connected if and only if ∂Y is locally connected.

Notice that if Y is a $CAT(0)$ space on which $\langle C \rangle$ acts geometrically, then we can construct the $CAT(0)$ space $X = Y \times \mathbb{R}$ and a geometric action of G on X via the product action. Simply let $\langle C \rangle$ act on Y as given and let $\langle \{s,t\} \rangle$ act by the infinite dihedral action on the $\mathbb{R}$ factor. This finishes part 1 of the theorem.

Suppose (C, C_1, K) is a VFS for Γ. Then $G = \langle A \rangle *_{\langle C \rangle} \langle B \rangle$ where $A \cup B = S$ and $A \cap B = C$ since C separates Γ. The proof of this part of the theorem is a direct application of Theorem 3.3 and lemma 4.2. $\square$

4. Two important lemmas

For the remainder of the paper, we will only be dealing with right-angled Coxeter groups.

In this section, we prove two important lemmas about a right-angled Coxeter group G acting on a $CAT(0)$ space X that will allow us to transfer combinatorial information about the Cayley graph of G to the space X. The first is lemma 4.2 which allows us to approximate $CAT(0)$ geodesic rays in X with Cayley graph rays, and also implies that special subgroups of G are quasi-convex in X. The second is lemma 4.4 which says that if two $CAT(0)$ geodesic rays are sufficiently close, then the Cayley graph approximating rays can be chosen to have a long common subpeice. The method of proof used here does not extend to give the corresponding results for general Coxeter groups.

Let $\Gamma(G, S)$ denote the presentation graph for G with respect to the generating set S and $\Lambda(G, S) = \Lambda$ denote the Cayley graph of G with respect to the generating set S. We know each $s \in S$ is order 2 and has a closed, convex, fixed point set in X denoted by $\mathrm{F}ix(s)$. With a slight abuse of notation, we also denote the fixed point set of s in Λ by $\mathrm{F}ix(s)$. Recall this consists of the midpoints of edges labeled by s with each endpoint in $\langle lk(s) \rangle$.

Let x_0 be a basepoint in X. We identify a quasi-isometric copy of the Cayley graph Λ of G inside X with the orbit of x_0 under the action of G using Theorem 2.8. We will use $CAT(0)$ geodesics between adjacent vertices in the Cayley graph. With this identification, any geodesic α in Λ is assigned a peicewise $CAT(0)$ geodesic in X which we denote by α_X.

We stop to make an important observation about geodesics in Λ.

Remark. Suppose $\alpha = (e_1, e_2, \ldots, e_n)$ with each $e_i \in S$ is an edge path in Λ. The path α crosses a set of fixed point sets or *walls* as the letters of α are traversed one by one. These walls are (in order):
$\mathrm{F}ix(e_1), \mathrm{F}ix(e_1e_2e_1), \ldots, \mathrm{F}ix(e_1e_2\cdots e_{n-1}e_ne_{n-1}\cdots e_2e_1)$.

In other words, at the ith step, α crosses the wall $e_1\cdots e_{n-1}\mathrm{F}ix(e_i)$. We refer to this wall as an e_i-wall. Notice that α can have the same letter e_i occuring twice which gives (possibly) two different e_i-walls that α crosses. With this in mind, the Deletion Condition can then be interpreted geometrically as: the edge path α is geodesic in Λ if and only if it crosses each wall at most once.

Definition 4.1 (Tracking)**.** *Suppose $r : [a, b] \to X$ is a geodesic segment in X with $r(a) = x$ and $r(b) = y$ and suppose $\delta > 0$. We say the Cayley graph geodesic α, δ-tracks r (or more precisely, the image of r) if every point of α_X is within δ of a point of the image of r and the endpoints of r and α_X are within δ of each other.*

Lemma 4.2. *There exists a $\delta > 0$ such that for any geodesic ray $r : [0, \infty) \to X$ based at x_0, there exists a geodesic ray α_r in $\Lambda(G, S)$ that δ-tracks r. Furthermore, for $A \subset S$, the special subgroup $\langle A\rangle$ is quasi-convex in X.*

Proof. First assume $r : [0, d] \to X$ is a unit speed parametrization of the geodesic segment in X between x_0 and $g \cdot x_0$ for some $g \in G$. Since G acts cocompactly on X, there exists a $K > 0$ such that any point $x \in X$ is within K of an orbit point. For each $i = 1, 2, \ldots D = \lfloor d\rfloor$ (i.e. D is the largest integer less than or equal to d) choose an element $g_i \in G$ such that $d(r(i), g_i \cdot x_0) < K$. Choose Cayley graph geodesics α_0 from 1_G to g_1, α_1 from g_1 to g_2,. . .,α_D from g_D to g. Call the piecewise geodesic α. In X, the concatenated the paths corresponding to the α_i to form a piecewise Cayley graph geodesic α_X from x_0 to $g \cdot x_0$.

Claim: Each vertex of α_X is within L of a point of r where L depends only on K and the quasi-isometry constants coming from the natural quasi-isometry between G and X guaranteed by Theorem 2.8.

Indeed, suppose $v\cdot x_0$ is a vertex on α_X. Then v is either one of the g_i's or v lies on α_i for some i. In the first case, $v\cdot x_0$ is within K of r by construction, in the second case, $v \cdot x_0$ is within $\min\{d(v\cdot x_0, g_i\cdot x_0), d(v\cdot x_0, g_{i+1}\cdot x_0)\}+K$ of r. It suffices to show that the first term of this sum is bounded by an L as in the statement of the claim.

This follows from the fact that the quasi-isometry between G and X is proper - i.e. the ball of radius K in X is mapped to a ball in G whose radius is linearly distorted by the quasi-isometry. The vertex v is on a path of the form $g_{i+1}^{-1}g_i$ in G and all of these are contained in the preimage of the ball of radius $2K + 1$ in X under the quasi-isometry by assumption. This gives a bound in G on the distance v can be from the closest g_i, but now pushing back to X under the quasi-isometry gives the desired result.

Thus we now have a Cayley graph path α_X from x_0 to $g \cdot x_0$ which lies in the L neighborhood of r. We straighten this path to a geodesic which δ-tracks r. The proof is by induction on $|S| = n$.

Suppose there is an $s \in S$ which occurs in α and the corresponding s-wall W_s is crossed by α more than once. If there is no such s, then α is geodesic already. For a fixed generator s, there can be more than one s-wall which has multiple intersections with α, but these can be handled independently of one another as we point out in Remark 1 below. Start by finding the first occurence of s along α so that the corresponding s-wall W_s crosses α more than once. We consider the two cases of whether the number of times crossed is even or odd.

Even case: In this case, α can be written as $asbsc$ with the following properties:

1. The s edge following a is the first time α crosses the s-wall $W_s = \mathrm{F}ix(asa^{-1})$.
2. The s edge following b is the last time α crosses W_s.
3. The b path is all on one side of W_s.
4. The paths a and c do not cross W_s at all.

Let $M = \max\{d(x_0, s \cdot x_0) \ : \ s \in S\}$ and consider the $\mathrm{C}AT(0)$ segment $[a \cdot x_0, asbs \cdot x_0]$. The endpoints of this segment lie in $N_M(\mathrm{F}ix(asa^{-1}))$ which is a convex subset of X, therefore the entire segment lies in there. By Theorem 2.6, $\langle lk(s)\rangle$ acts geometrically on this $\mathrm{C}AT(0)$ subset of X. Now the induction hypothesis applies since $\langle lk(s)\rangle$ is a right-angled Coxeter group with no more than $n-1$ generators (s is not in this subgroup), that is isometrically embedded as a convex subset of Λ in the word metric. Thus there is a δ_{n-1} and a Cayley graph geodesic β in $\langle lk(s)\rangle$ so that β δ_{n-1}-tracks $[a \cdot x_0, asbs \cdot x_0]$ and β is also a geodesic in G. Thus we can replace the path sbs with the geodesic β.

Odd case: In this case, α can be written in the same form as above but properties 2 and 4 do not hold. In particular, 1_G and g are on opposite sides of $\mathrm{F}ix(asa^{-1})$ and so there is exactly one more place that α crosses this wall and it occurs in the c path. More specifically, the s following the b in α is the second to last crossing of this wall. We do the same procedure as above to replace b with a geodesic β which has no more occurences of s in it. The new path $a\beta c$ obtained crosses $\mathrm{F}ix(asa^{-1})$ exactly once.

Remark 1. Note that the subpath a could contain the letter s, but the corresponding s-wall is different from the s-wall used above. By assumption, any s-wall that a may cross can only be crossed once. Also, the subpath c could contain s and could cross the corresponding s-wall more than once. If so, we can do the procedure described above to this subpath independent of the procedure done for b - i.e. any s deletions that occur in c will occur within c and not with any previous s occurences by choice of b. Thus with one use of induction, we can replace pieces of α to obtain a path which has the property that any s-wall it crosses, it crosses exactly once.

Remark 2. Notice that every point of $[a \cdot x_0, asbs \cdot x_0]$ is within L of r by the $CAT(0)$ inequality.

Do this replacement procedure for each $s \in S$ to obtain a path α_r which has the property that any wall it crosses, it crosses exactly once - i.e. α_r is a Cayley graph geodesic. To see α_r δ-tracks r, we use the remarks above to conclude that $\delta_n = n\delta_{n-1} + L$ suffices. Indeed, we will have to use induction at most n times, once for each $s \in S$.

Now suppose y is any point in X. It is clear that we can find a Cayley graph geodesic α which $(\delta = \delta_n + K)$-tracks $[x_0, y]$ since we can choose an orbit point K close to y and do the above procedure.

If $r : [0, \infty) \to X$ is a geodesic ray based at x_0, then build a Cayley graph geodesic ray α_r which δ-tracks r as follows:

For each $t \in [0, \infty)$, we can find α_t which δ-tracks $[x_0, r(t)]$ by the previous step. To build α, use the local finiteness of Λ - indeed, infinitely many of the α_t share the same first edge, infinitely many of these share the same second edge, etc. This clearly builds a geodesic ray with the necessary properties.

It remains to show that special subgroups of (G, S) are quasi-convex in X. Let a be an element of $\langle A\rangle$ for A a subset of S. Then our Cayley geodesic that δ tracks $[x_0, ax_0]$ is a Cayley geodesic from 1_G to $a \in A$. This geodesic uses only letters of A by lemma 1.4. Each point of $[x_0, ax_0]$ is within δ of a vertex of our Cayley path in X and each vertex of our Cayley path is in $\langle A\rangle x_0$. $\square$

Continuing with the quasi-isometry of the Cayley graph into the $CAT(0)$ space for our right-angled Coxeter group G as before we give two improved versions of Proposition 4.2. The last proposition is the version we will need in the proof of the main theorem, however we need the following intermediate step to obtain that result.

Proposition 4.3. *Suppose g and h are at distance n apart in the Cayley graph and α is a Cayley graph geodesic from 1_G to g. Then there exist Cayley graph geodesics α' and β' from 1_G to g and h, respectively, such that each vertex of α' and β' is at most n from a vertex of α, and α' and β' have the same initial segment from 1_G to k for an element k with Cayley graph length $\ell(k) \geq \ell(g) - n$.*

Proof. We proceed by induction on n. If $n = 0$ then $g = h$ so $\alpha' = \beta' = \alpha$ works. Suppose $n > 0$ and take a Cayley graph geodesic $(s_1, s_2, \ldots, s_n)$ from g to h. Either $\ell(gs_1)$ is one more than $\ell(g)$, or else $\ell(gs_1)$ is one less than $\ell(g)$ by lemma 1.5.

In the first case, denote by α_1 the geodesic (α, s_1) from 1 to gs_1. Since h is at distance $n-1$ from gs_1, by induction hypothesis there exist α'_1 and β'_1 within $n-1$ of α_1 and sharing an initial segment of length at least $\ell(gs_1) - (n-1)$. Since (α'_1, s_1) is not geodesic, by the Deletion Condition $gs_1 = us_1v$, with u and v represented by segments of α'_1 with $s_1vs_1 = v$. In fact, by lemma 1.7, s_1 commutes with each letter of v. Replace the segment of (α'_1, s_1) corresponding to s_1vs_1 with the segment of α'_1 corresponding to

v to give a geodesic α' from 1_G to g within one of (α'_1, s_1). Suppose k_1 is represented by the common initial subpath of α'_1 and β'_1. If k_1 is given by an initial segment in u then we simply take $\beta' = \beta$. Otherwise write $k_1 = us_1w$ for w given by an initial segment of v. Then $s_1ws_1 = w$ and we take β' to be the path β'_1 with the segment corresponding to s_1w replaced by ws_1. Since uw corresponds to an initial segment of α' at most one shorter than k_1, we have α' and β' share a common initial segment of length at least $\ell(g) - n$ and since α' and β' are each at most 1 from α'_1 they are also at most one n from α.

In the second case where gs_1 is shorter than g, write $g = us_1v$ for $s_1vs_1 = v$ and s_1 commuting with each letter in v as above. Take α_1 to be the geodesic obtained by replacing the segment of (α, s_1) corresponding to s_1vs_1 by v, a geodesic to gs_1 with gs_1 a distance at most $n-1$ from h. By induction hypothesis we have α'_1 and β'_1 geodesics from 1_G to gs_1 and h, at most $n-1$ apart and sharing a common initial segment of length at least $\ell(gs_1) - (n-1) = \ell(g) - n$. Take $\alpha' = (\alpha'_1, s_1)$ and $\beta' = \beta'_1$. Then α' and β' are at most n from α and share the same common initial segment as α'_1 and β'_1 of length at least $\ell(g) - n$.□

Lemma 4.4. *There exists c and d such that for any r and s, infinite geodesic rays in X based at x_0, that are within ϵ of each other a distance M from x_0, there exist Cayley graph geodesic rays α and β which $(c\epsilon + d)$ track r and s respectively, and which share a common initial segment out to a distance $M - c\epsilon - d$ from x_0.*

Proof. Appropriate c and d are shown below. Suppose r and s are given and take α_0 and β_0 to be Cayley graph geodesics δ-tracking r and s for a δ from Proposition 4.2. Then at a $CAT(0)$ space distance of no more than $M + K$ from x_0 (where K comes from the cocompactness of the action) we have orbit points gx_0 and hx_0 on these a Cayley graph distance n apart with n bounded by a function of ϵ determined by the quasi-isometry constants and δ_0. By Proposition 4.3, we can replace the segments of α_0 and β_0 out to gx_0 and hx_0 by Cayley graph geodesic segments agreeing to a point within n of gx_0 in the Cayley graph, a distance in the $CAT(0)$ space bounded again in terms of quasi-isometry constants to be at most $(c\epsilon + d)$ from gx_0 and hence at least $(M - c\epsilon - d)$ from x_0. Since we are replacing an initial geodesic segment of α_0 by a geodesic out to gx_0 we again get an infinite geodesic ray α, actually unchanged after gx_0, and similarly we get a β from β_0. Since α and α_0 are within a distance n in the Cayley graph, by a constant determined from the quasi-isometry constants and δ, the new α tracks the original r within some $(c\epsilon + d)$ (we may as well take the larger of the c and d from here and before) and similarly for β tracking s.□

5. The Filtering Process

Recall that our ultimate goal is to start with two geodesic rays in the $CAT(0)$ space X whose endpoints are close in $\partial_\infty X$ and build a small connected set containing these two endpoints. From the two rays in X, we first obtain two geodesic rays α and β in Λ with a long common initial segment using lemma 4.4. This is the starting point for the work in this section.

Our first goal is to use the rays α and β to build a collection of graphs in the plane whose edges are labeled by elements of S. This collection will consist of overlapping *fans* (see definition 5.7) from which we will extract a planar, one-ended graph called a *filter* for α and β. One should think of this process of creating a filter as trying to fill in the Cayley graph between α and β in a systematic way so as to avoid product subgroups yet also building a sequence of geodesics between α and β.

There is an obvious map from any such filter into Λ and this map into Λ will be proper. A continuous map $f : X \to Y$ is *proper* if for each compact set $C \subset Y$, $f^{-1}(C)$ is compact in X. The most important property of this filter is that when viewed in Λ, there are no long *factor paths* (see definition 5.2 below). These are exactly the paths that behave badly under quasi-isometry so we must control the size of these paths if we hope to get any information to the $CAT(0)$ space X from this filter. Moving this information into X will ultimately enable us to achieve our goal of constructing a "small" connected set in $\partial_\infty X$ containing the two points, thus showing X has locally connected boundary.

Suppose (G, S) is a right-angled Coxeter system, G is a one ended group and $G \neq \langle A\rangle \oplus \langle B\rangle$ where A and B are non-empty disjoint subsets of S such that $A \cup B = S$. Recall that when G is right-angled, $lk^2(C)$ where C is a subset of S is the ordinary link, $lk(C)$, in Γ (see remarks following definition 1.10). We define the notion of *product separator* below. This should be compared with definition 3.1 from earlier in the paper as both will be needed here.

Definition 5.1 (Product separator). *Suppose (G, S) is a Coxeter system and Γ is the presentation graph for (G, S). A product separator for Γ is a subset C of S where C is a full subgraph, $C = A \cup B$, $\langle C\rangle = \langle A \cup B\rangle = \langle A\rangle \oplus \langle B\rangle$, both $\langle A\rangle$ and $\langle B\rangle$ are infinite and C separates Γ.*

Definition 5.2 (Factor path). *Call a path $(c_1, \ldots, c_m)$ in Λ a factor path if $\{c_1, \ldots, c_m\} \subset A \cup B$ where A and B are disjoint commuting subsets of S and $\langle A\rangle$ and $\langle B\rangle$ are infinite.*

Remark Suppose (G, S) is a Coxeter system with presentation graph Γ. G is one-ended if and only if Γ contains no complete separating subgraph, the vertices of which generate a finite subgroup. A proof of this can be found in [MT].

The following lemma allows us to easily identify when the presentation graph Γ for G has a virtual factor separator given that there are no product separators. This lemma will be used repeatedly in lemmas 5.5 and 5.6 below.

Lemma 5.3. *Suppose Γ has no product separators. Then Γ has a virtual factor separator iff there is a vertex $x \in \Gamma$ such that $lk(x)$ is a subset of a visual product with infinite factors.*

Proof. Suppose $\langle A\rangle \oplus \langle B\rangle$ is a visual product in (G, S) with infinite factors, $\langle B\rangle$ has finite index in $\langle B'\rangle$ and B' separates Γ (i.e. (B', B, A) is a virtual factor separator).

By the Finite Index Lemma (lemma 1.12), there exists $B_1 \subset B$ such that $\langle B' - B_1\rangle$ is finite and commutes with B_1. Hence $\langle B'\rangle = \langle B_1\rangle \oplus \langle B' - B_1\rangle$.

Note that $A \cup B' \neq S$, as $A \cup (B' - B_1)$ commutes with B_1 and G does not decompose as a non-trivial visual direct product. Let y be a vertex in $\Gamma - (A \cup B')$ and x be a vertex of Γ separated from y by B'. We have $x \in A$, for otherwise, $\langle A \cup B'\rangle = \langle A \cup (B' - B_1)\rangle \oplus \langle B_1\rangle$ is a product separator for Γ. If $c \in lk(x)$, and $c \notin B'$, then x and c are in the same component of $\Gamma - B'$. Hence B' separates c and y and as above $c \in A$. This means that $lk(x) \subset A \cup B' \subset \langle A \cup (B' - B)\rangle \oplus \langle B_1\rangle$.

To see the converse, suppose x is a vertex of Γ and $lk(x) \subset (A \cup B) \subset S$ where A commutes with B and both $\langle A\rangle$ and $\langle B\rangle$ are infinite. The group $\langle lk(x)\rangle$ is infinite as G has one end. If both $\langle lk(x) \cap A\rangle$ and $\langle lk(x) \cap B\rangle$ were infinite, then Γ has a product separator. If one, say $\langle lk(x) \cap A\rangle$, is finite, then $\langle lk(x) \cap B\rangle$ commutes with A and so $(lk(x), lk(x) \cap B, A)$ is a virtual factor separator. □

Remark. For the rest of this section, in addition to the basic assumptions at the beginning of the section, assume Γ has no product separator and no virtual factor separator.

Notation. The vertices of Γ are the elements of S and each edge of Λ is labeled by an element of S. If e is an edge in Λ, then we let $\bar{e} \in S \subset G$ denote the label of e. Denote the vertex sets of Γ and Λ by Γ^0 and Λ^0, respectively. Also, recall that $B(g)$ is the set of elements of S that make gs shorter than g and that this generates a finite subgroup (lemma 1.5).

To build a fan, we start with edge paths $(e_1, \ldots, e_n, a)$ and $(e_1, \ldots, e_n, b)$ beginning at a base point $*$, that are geodesic in Λ, and let $g \equiv \bar{e}_1 \cdots \bar{e}_n \in G$. Suppose $(e_i, \ldots, e_n)$ is the longest (terminal) factor path in $(e_1, \ldots, e_n)$. This may in fact be a trivial path, as in the case G is word hyperbolic.

If $\langle \bar{e}_i, \ldots, \bar{e}_n\rangle$ *is infinite* and A and B are as in the definition of factor path, then $\langle \{\bar{e}_i, \ldots, \bar{e}_n\} \cap A\rangle$ or $\langle \{\bar{e}_i, \ldots, \bar{e}_n\} \cap B\rangle$ is infinite. Without loss, assume $C \equiv \{\bar{e}_i, \ldots, \bar{e}_n\} \cap A$ and $\langle C\rangle$ is infinite. Now $\{\bar{e}_i, \ldots, \bar{e}_n\} \subset C \cup lk(C)$. Also $B \subset lk(C)$ so $\langle C\rangle$ and $\langle lk(C)\rangle$ are infinite.

Lemma 5.4. *If $\langle \bar{e}_i, \ldots, \bar{e}_n\rangle$ is infinite, then $\{\bar{e}_i, \ldots, \bar{e}_n\} \cup B(g) \subset C \cup lk(C)$. Hence $\langle C\rangle \oplus \langle lk(C)\rangle$ is a visual direct product with both factors infinite and so $C \cup lk(C)$ does not separate Γ.*

Proof. Recall, $s \in S$ is in $B(g)$ if gs is shorter than g. If $s \in B(g)$, then let j be the largest integer such that for the path $(e_1, \ldots, e_n, s)$, s deletes with e_j. If $j \geq i$, then $s \in \{\bar{e}_i, \ldots, \bar{e}_n\} \subset C \cup lk(C)$. If $j < i$, then by lemma 1.7, $s \in lk(C)$. In any case $B(g) \subset C \cup lk(C)$.□

Important: The next two lemmas will allow us to choose a path in Γ from a to b that avoids certain subsets of Γ. We use this path to build a fan between these two edge paths. We split this into two lemmas depending on whether $\langle \bar{e}_i, \ldots, \bar{e}_n \rangle$ is finite or infinite. Note that we allow the case $a = b$.

Lemma 5.5. *If $\langle \bar{e}_i, \ldots, \bar{e}_n \rangle$ is infinite, then there is an edge path τ in Γ, from a to b of length at least 2, such that other than possibly a and b, no vertex of this path is in $C \cup lk(C)$. In particular, no vertex of τ is in $\{\bar{e}_i, \ldots, \bar{e}_n\}$.*

Proof. **Case 1.** Suppose $a = b$.

By Lemma 5.3, there is a vertex $v \in lk(a)$ such that $v \notin C \cup lk(C)$. If e is the edge of Γ from a to v then our path is (e, e^{-1}).

Case 2. Suppose $a \in C \cup lk(C)$ and $b \notin C \cup lk(C)$ (or when the roles of a and b are reversed).

By Lemma 5.3, there is an edge from a to $v \notin C \cup lk(C)$. If $v \neq b$, then (as $C \cup lk(C)$ cannot separate Γ) take as the desired path, the edge from a to v followed by any edge path from v to b in $\Gamma - (C \cup lk(C))$. If $v = b$, then by Lemma 5.3 there is a vertex $w \in lk(v)$ such that $w \notin C \cup lk(C)$. If e is the edge of Γ from v to w, then take the desired path to be the edge from a to v followed by e, followed by e^{-1}.

Case 3. Suppose $a, b \in C \cup lk(C)$.

By Lemma 5.3, there is an edge e from a to $v \notin C \cup lk(C)$ and an edge d from b to $w \notin C \cup lk(C)$. Take as the desired path, the edge e followed by a path in $\Gamma - (C \cup lk(C))$ from v to w followed by d^{-1}.

Case 4. Suppose $a, b \notin C \cup lk(C)$ and $a \neq b$.

Choose as desired path, any path of length at least 2 in $\Gamma - (C \cup lk(C))$ from a to b. (Again, if there is an edge between a and b one can adjust as in Case 2.) □

Remark. It seems somewhat artificial to consider an edge followed by its inverse in our above paths, but paths of length at least 2 induce beneficial combinatorics in our constructions. These combinatorics allow for a simplified proof of lemma 5.8 at the end of the section.

Lemma 5.6. *If $\langle \bar{e}_i, \ldots, \bar{e}_n \rangle$ is finite, there is a path τ, in Γ from a to b, of length at least 2 which avoids $B(g)$.*

Proof. Recall that $B(g)$ cannot separate Γ as G is 1-ended (see remark preceding lemma 5.3). Now proceed as in lemma 5.5, case 4. □

Given any edge paths $(e_1, \ldots, e_n, a)$ and $(e_1, \ldots, e_n, b)$, we now have the path τ in Γ from a to b which avoids the appropriate subsets of Γ depending on whether $\langle \bar{e}_i, \ldots, \bar{e}_n \rangle$ is finite or infinite. Let the ordered vertices of τ be $a = t_0, t_1, \ldots, t_m = b$ ($m \geq 2$). Note that in G, $[t_i, t_{i+1}] = 1$ for all i.

We now describe the construction of a *fan* using the path τ. These fans will be the buliding blocks in the construction of a *filter* for two geodesic rays α and β in Λ.

Definition 5.7 (Fan). *The τ-fan for $(e_1, \dots e_n, a)$ and $(e_1, \dots e_n, b)$ is a planar diagram consisting of edges labeled $e_1, \dots, e_n, a, b$ plus loops labeled by the vertices of the path τ defined in Lemma 5.5 or 5.6 (see Figure 1 below). The fan-loops are four sided loops with edge labels $(t_i, t_{i+1}, t_i, t_{i+1})$ beginning at a vertex labeled $g \equiv \bar{e}_1 \cdots \bar{e}_n$. These loops correspond to the commutation relations $[t_i, t_{i+1}]$ among consecutive vertices of τ. We will simply refer to this as a fan for $(e_1, \dots, e_n, a)$ and $(e_1, \dots, e_n, b)$ if the path τ is unimportant.*

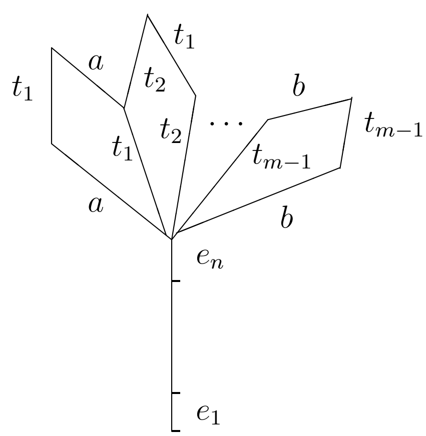

Figure 1

Recall that $t_i \notin B(g)$ for all i. Hence, by lemma 1.4, the edge paths $(e_1, \dots, e_n, t_i, t_{i+1})$ in Λ are geodesic. The edges labeled a and b at g in the τ-fan are respectively called the *left* and *right fan edges at g*. The edges labeled $t_1, \dots, t_{m-1}$ at g are called *interior fan* edges.

Next, suppose $\alpha = (e_1, \dots, e_n, a_1, a_2, \dots)$ and $\beta = (e_1, \dots e_n, b_1, b_2, \dots)$ are geodesics in Λ.

The Filter Construction: Construct a 1-ended planar graph (see Figure 2 below) with edge labels in S as follows: First construct a fan for $(e_1, \dots e_n, a_1)$ and $(e_1, \dots, e_n, b_1)$. Next overlap this fan with a fan for $(e_1, \dots e_n, a_1, a_2)$ and $(e_1, \dots, e_n, a_1, t_1)$ a fan for $(e_1, \dots, e_n, t_1, a_1)$ and $(e_1, \dots, e_n, t_1, t_2), \dots$, a fan for $(e_1, \dots, e_n, t_i, t_{i-1})$ and $(e_1, \dots, e_n, t_i, t_{i+1}), \dots$ and finally a fan for $(e_1, \dots, e_n, b_1, t_{m-1})$ and $(e_1, \dots, e_n, b_1, b_2)$.

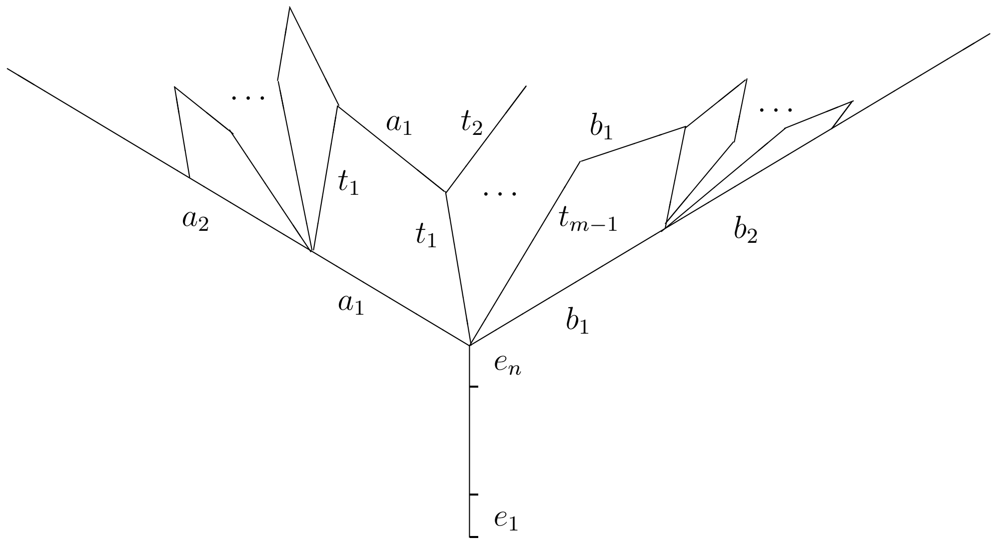

Figure 2

This gives our 1^{st} and 2^{nd}-level fans.

At this point there is the potential for two edges of our planar graph to share a vertex and have the same label. These edges are *not* identified in our planar graph. Rather, an edge path τ, constructed as in case 1 of lemma 5.5 will be used to extend our graph between these two edges.

To continue, we must specify geodesics from the base point $*$ to each vertex defined so far. In each fan-loop constructed up to this point, designate the upper left edge as a *non-tree* edge. The graph minus the non-tree edges is a tree. Take as designated geodesic from $*$ to a defined vertex, the unique geodesic of the tree.

Continuing, (at each stage designating upper left edges of fan loops as non-tree edges and constructing fans) we obtain the desired 1-ended planar graph called a *filter* for α and β. Geodesics from $*$ in this filter determine geodesics in Λ with the same edge labels. Hence there is a natural proper map from this filter to Λ. The image of this filter in Λ is a closed 1-ended subgraph.

The following are useful combinatorial facts about a filter F.

(1) Each vertex of F has exactly one or two edges directly below it. This fact would be missing if certain identifications had been allowed.
(2) If a vertex of F has exactly one edge below it, then this edge is either an interior fan edge, an α-edge or a β-edge.
(3) If a vertex of F has exactly two edges e and d below it, then one (to the right) is a left fan edge, say of the fan τ and the other (to the left) is a right fan edge of the fan λ, where $\tau \neq \lambda$ and $\bar{d} \neq \bar{e}$. Both e and d belong to a single fan loop.

(4) For each vertex of F there is a unique fan containing all edges directly above this vertex. In particular, the edges above this vertex consist of a left and right fan edge and at least one interior edge.
(5) The filter for α and β minus all non-tree edges is a tree containing α and β and all interior edges of all fans.
(6) An edge of the filter is a non-tree edge iff it is a right fan edge not on β. Right fan edges not on β are upper left edges of a fan-loop.
(7) Let T be the tree obtained from a filter F by removing all non-tree edges. There are no dead ends in T - i.e. for any vertex $v \in T$, there is an interior edge at v.

We now begin the process of showing that if we create a filter for geodesic rays α and β that share a long common initial segment, then there is a bound on the length of any factor path in this filter. Again, this will allow us to map the one-ended planar graph into X and obtain a small connected set in $\partial_\infty X$. The final result necessary is lemma 5.10 but we prove two intermediate lemmas about factor paths in filters along the way.

Lemma 5.8. *Let $\gamma \equiv (c_1, \ldots, c_n)$ be an edge path in T such that each edge is above the previous edge and at most, the initial point of γ intersects $\alpha \cup \beta$. Suppose $(c_i, \ldots, c_j)$ is a factor path of γ of length $> |S|$. Then, either $\bar{c}_{j+1}$ or $\bar{c}_{j+2}$ is not in $\{\bar{c}_i, \ldots, \bar{c}_j\}$.*

Proof. Since $j-i \geq S$, and G is right angled, $\langle \bar{c}_i, \ldots \bar{c}_j\rangle$ is infinite. Neither, c_{j+1} nor c_{j+2} is a right fan edge, as both are in T. If c_{j+1} is an interior fan edge we are done by the construction in lemma 5.5. Hence we may assume that c_{j+1} is a left fan edge and $\bar{c}_{j+1} \in \{\bar{c}_i, \ldots, \bar{c}_j\}$. Now, if c_{j+2} is an interior fan edge, once again we are finished. Hence we may assume c_{j+2} is a left fan edge.

Let u be the end point of c_{j+2} and x the initial point of c_i. As c_{j+2} is a left fan edge, there is a right fan edge a below u (see (2) and (3) of the above list of combinatorial facts).

Note that, $\bar{a} \neq \bar{c}_{j+1}$ since otherwise there is no right fan edge below the end point of (the left fan edge) c_{j+1}. Hence the situation must be one of the following two cases pictured in Figure 3.

Case 1. Suppose $\bar{a} \notin \{\bar{c}_i, \ldots, \bar{c}_j\}$.

Consider the path $(c_i, \ldots, c_{j+2})$ at z. By (4), the edge c_i at z is either an internal or left fan edge, and all other edges of this path are internal fan edges. Hence this entire path is in T. But then by the lemma 5.5 construction, $\bar{c}_{j+1} \notin \{\bar{c}_i, \ldots, \bar{c}_j\}$, contrary to our assumption.•

Case 2. Suppose $\bar{a} \in \{\bar{c}_i, \ldots, \bar{c}_j\}$.

Let k be the largest integer in $\{i, \ldots, j\}$ such that $\bar{a} = \bar{c}_k$. Note that $\bar{a} \neq \bar{c}_{j+1}$ as this was considered above . The edge path $(c_i, \ldots, c_{k-1}, c_{k+1}, \ldots, c_{j+2})$ at x is in T, as (4) implies the first edge of the subsegment $(c_{k+1}, \ldots, c_{j+2})$ (at y) is either an internal or left fan edge and each subsequent edge is an internal fan edge. Hence $(c_i, \ldots, c_{k-1}, c_{k+1}, \ldots, c_{j+2})$ at x defines a (geodesic) factor path. As the length of this geodesic is $\geq |S|+2$, $\langle \bar{c}_i, \dots, \bar{c}_{k-1}, \bar{c}_{k+1}, \dots, \bar{c}_j \rangle$ is infinite. We have $\bar{c}_{j+2} \notin \{\bar{c}_i, \dots \bar{c}_{k-1}, \bar{c}_{k+1}, \dots, \bar{c}_j\}$, by our τ-fan construction in lemma 5.5. But, we also have $\bar{c}_{j+2} \neq \bar{c}_k = \bar{a}$. $\square$

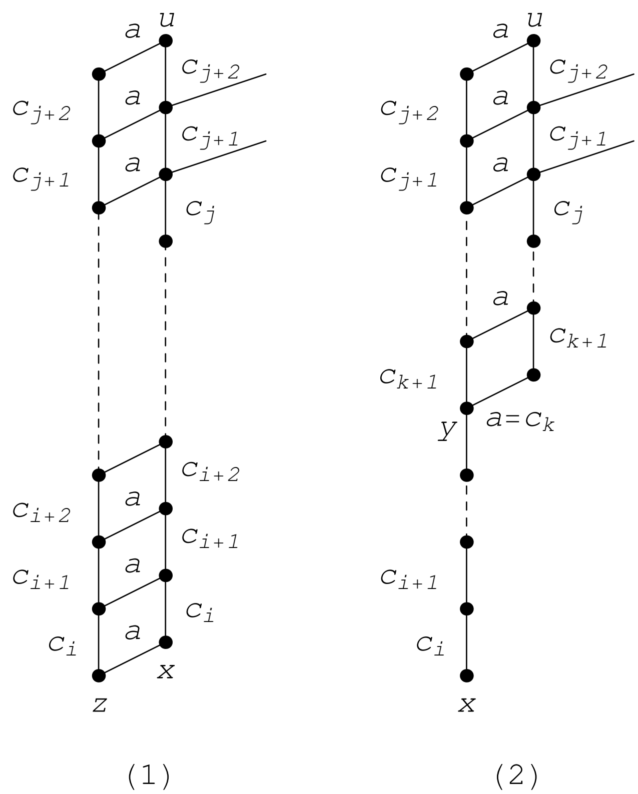

Figure 3

Lemma 5.9. *Suppose that $\gamma \equiv (c_1, \dots c_n)$ is a directed factor path in T such that γ intersects $\alpha \cup \beta$ in at most, its initial point. Then $n \leq 3|S|$.*

Proof By lemma 5.8, at least every other edge e_i following $e_{|S|}$ is such that $\bar{e}_i \notin \{\bar{e}_1, \dots \bar{e}_{i-1}\}$. $\square$

The following lemma is the main lemma of this section that will be used in the proof of main theorem. In order to obtain the necessary bounds on the lengths of factor paths, we must analyze how such a path arises in Λ. This involves studying Van-Kampen diagrams in the group. We refer the reader to [LS] for the necessary background concerning Van-Kampen diagrams for groups (abbreviated V-K diagrams).

Lemma 5.10. *For any k, there exists a δ such that, for any α and β Cayley graph geodesics between the same pair of points, if for some n, $\alpha(n)$ and $\beta(n)$ are at least δ apart, then α contains a factor subpath of length k (and also β has such a subword).*

Proof. Write $\alpha = (\alpha_0, \alpha_1)$ and $\beta = (\beta_0, \beta_1)$ for α_0 and β_0 initial subpaths of length n. Take a Cayley graph geodesic γ from $\alpha(n)$ to $\beta(n)$ and V-K diagrams for $\alpha_0\gamma = \beta_0$ and $\alpha_1 = \gamma\beta_1$ each with a minimum number of regions, such that the combination is a V-K diagram for the relation telling us that α and β represent the same element of the Coxeter group. Given any edge in this V-K diagram, we construct a chain of boxes corresponding to commuting generators in the defining presentation. Such a chain cannot cross γ twice or begin and end on α or begin or end on β since each of these is a geodesic. Such a chain cannot be a loop else one of the diagrams wouldn't have had a minimum number of regions. No two chains with the same letter cross each other. Thus each chain begins on α, ends on β, and perhaps crosses γ with as many chains crossing from α_0 to β_1 as from α_1 to β_0. Note that γ consists of an edge from each such chain. Each letter for a chain from α_0 to β_1 commutes with a letter on a chain from α_1 to β_0 (since these chains cross at some point) which gives a commutation relation between these letters. Write γ_0 for the subsequence of letters of γ that label chains extending from α_0 to β_1 and γ_1 for the subsequence of letters of γ that label chain extending from α_1 to β_0. Then in fact γ and (γ_0, γ_1) would represent the same element since each letter of γ_0 commutes with each letter of γ_1.

We take δ greater than twice the longest geodesic in a finite subgroup of the Coxeter group. Then the letters in γ_0 generate an infinite subgroup as do the letters in γ_1. Let α_0' be the subsequence of letters in α_0 belonging to chains crossing γ (a word equivalent to γ_0^{-1}). A letter of α_0 not in α_0' must commute with each earlier letter belonging to α_0' since the chain extending from such a letter extends to β_0 and crosses the chains of earlier letters in α_0'. Of the $\delta/2$ letters in α_0', one letter must occur at least $\delta/(2|S|)$ times say s_1. Between occurrences of s_1 there must be a letter not commuting with s_1 and hence also a letter of α_0', and among all such letters between successive s_1 occurrences, there must be some letter occurring at least $(\delta/(2|S|) - 1)/|S|$ times say s_2. We define s subset S' of the letters $s_1, s_2, \ldots, s_i$ occurring in α_0', and define a sequence of subwords of α_0' as follows.

First we take s_1 and s_2 in S' and take the terminal subword of α_0' after the later of the first occurrences of s_1 and s_2. Now the first occurrence of s_3 occurs either in the first half or the last half of this subword of α_0'. If the terminal segment of α_0' after the first occurrence of s_3 in the remaining word is longer the the word before this s_3, then we keep s_3 in S' and take the terminal segment after this s_3, and otherwise we omit s_3 from S' and restrict to the s_3-free initial subword of the remainder. Repeat the process for $s_4, \ldots, s_i$ - each time taking the first occurrence of the next letter and dividing the remaining word at this point keeping the larger half, putting the letter in S' if we take a terminal segment, and otherwise omitting it from S'.

In the end we are left with a subword α_0'' of α_0' of length at least $2(\delta - 2|S|)/(|S|^2 2^{|S|})$ such that the following conditions hold: each letter in the α_0'' belongs to S' and each letter in S' occurs either at the beginning of α_0'' or else earlier in α_0' than all of α_0'' as well. Let ϕ be smallest subword of α_0 containing the corresponding letters in α_0''. Then the letters in ϕ belong to S' or belong to the subset $lk(S')$ of letters commuting with all elements of S'. Since the letters in γ_1 are in $lk(S')$, $lk(S')$ generates an infinite subgroup, as does S' since s_1 and s_2 are in S'. Hence there is a subword in α_0 with letters in a product of infnite subgroups having length bounded below by a function of δ. If $\delta \geq |S|^2 2^{|S|} k + 2|S|$ then this subword will be a factor path of length at least k.

6. Proof of the Main Theorem

Theorem 6.1. *Suppose (G, S) is a right-angled Coxeter system, G is 1-ended, $\Gamma(S)$ contains no product separator and no virtual factor separator. Also assume that G does not visually split as a non-trivial direct product. Then G has locally connected boundary.*

Proof. Suppose G acts geometrically on a $CAT(0)$ space X. Choose a base point $* \in X$. Suppose r and $s : [0, \infty) \to X$ are geodesics at $*$ and $s \in N_{(M,\epsilon)}(r)$ (i.e. $d_X(r(M), s(M)) < \epsilon$).

Choose geodesics $\alpha = (e_1, \ldots, e_n, a_1, a_2, \ldots)$ and $\beta = (e_1, \ldots, e_n, b_1, b_2, \ldots)$ from Λ that δ track r and s respectively as in lemma 4.4. Recall α_X means view the path α inside X. Thus we can choose α and β so that α_X and β_X share a common initial peice that is M' long where M' comes from following the $M - c\epsilon - d$ from lemma 4.4 under the quasi-isometry between Λ and X. Note that $d(\alpha_X(M'), r(M')) < \delta$ and $d(\beta_X(M'), r(M')) < \delta$.

Construct a filter F for this pair of edge paths and consider the tree T constructed from F as in section 5. Let $f : F \to X$ be the natural map that factors through $\Lambda(G, S)$. To show local connectivity at r, it suffices to show the limit set of $f(F)$ is a "small" connected subset of ∂X containing both r and s.

The map f is proper and F is one-ended, thus the limit set of $f(F)$ is connected. It is clear that this limit set contains r and s. To see that this limit set is small, we must apply lemma 5.10 and lemma 5.9. Here is the important point.

Claim: There exists M' and ϵ' such that $f(F) \subset N_{(M',\epsilon')}(r)$ (the M' mentioned above is the same M' used here).

First we work in F, then transfer this information over to X via map f.

Let v be a vertex in F that is far away from $\alpha(n)$. In particular, we want $d(*, f(v))$ to be much bigger than M' since we are concerned with the limit set of F. Also, denote by p the point at which α and β start to separate.

We know the geodesic in F from the identity to v is the obtained by following α (and β) to p, then using the unique geodesic in the tree T (described in section 5) from p to v. Call this geodesic γ. Thus $\gamma(n) = \alpha(n) = \beta(n) = p$.

By lemma 5.9 there exists a global bound on the length of a (directed) factor path in T that only intersects $\alpha \cup \beta$ in it's initial point, call this bound k. By lemma 5.10, there is a δ_1 coming from this k such that all geodesics from the identity to v are contained in the δ_1 neighborhood of γ.

We transfer all of this information over to X via the quasi-isometry to obtain new contants. In particular, the $\mathrm{C}AT(0)$ geodesic, τ from $*$ to $f(v)$ is within δ_2 of $f(\gamma)$. Thus $d(\tau(M'), f(p)) < \delta_2$.

We know from the choice of α and β above that $d(\alpha_X(M'), r(M')) < \delta$ and $d(\beta_X(M'), r(M')) < \delta$. But $\alpha_X(M')$ is simply $f(p)$ here, thus we have $d(f(p), r(M')) < \delta$. Putting these together with the triangle inequality yields: $d(r(M'), \tau(M')) < \delta_2 + \delta$. If we use the M' mentioned above along with $\epsilon' = \delta_2 + \delta$, we have the claim. $\square$

References

[BH] M.R. Bridson, A. Haefliger. *Metric Spaces of Non-positive Curvature* (Grundl. Math. Wiss., Vol. 319, Springer, Berlin 1999).

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

[Bo] B. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.* **180** (1998), 145-186.

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

[G] M. Gromov. Hyperbolic Groups. In *Essays in Group Theory* (S. Gersten ed.), MSRI Publ. (Springer-Verlag, 1987), 75-265.

[H] J. Humphreys. *Reflection groups and Coxeter groups* (Cambridge studies in Adv. Math., Vol. 29, Cambridge Univ. Press, 1990).

[LS] R. C. Lyndon, P. E. Schupp. *Combinatorial group theory* Ergeb. Math. Grenzgeb., Springer-Verlag, 1977).

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

[MT] M. Mihalik, S. Tschantz. Visual Decompositions of Coxeter Groups. Preprint. (2005).

[M] J. Milnor. A note on curvautre and the fundamental group. *J. Diff. Geom.* **2** (1968), 1-7.

[Mo] G. Moussong. Hyperbolic Coxeter groups. PhD thesis. Ohio State University (1988).

[R] K. Ruane. Dynamics of the action of a CAT(0) group on the boundary. *Geom. Dedicata* **84** (2001), 81-99.

[S] G. Swarup. On the cut point conjecture. *Electron. Res. Announc. Amer. Math. Soc.* **2** (1996).

Department of Mathematics, Vanderbilt University, Nashville, TN 32340
E-mail address: mihalik@math.vanderbilt.edu

Department of Mathematics, Tufts University, Medford MA 02155
E-mail address: kim.ruane@tufts.edu

Department of Mathematics, Vanderbilt University, Nashville, TN 32340
E-mail address: tschantz@math.vanderbilt.edu