GEODESICALLY TRACKING QUASI-GEODESIC PATHS FOR COXETER GROUPS

MICHAEL MIHALIK AND STEVEN TSCHANTZ

Abstract. If Λ is the Cayley graph of a Gromov hyperbolic group, then it is a fundamental fact that quasi-geodesics in Λ are tracked by geodesics. Let (W, S) be a finitely generated Coxeter system and Λ be the Cayley graph of (W, S). For general Coxeter groups, not all quasi-geodesic rays in Λ are tracked by geodesics. In this paper we classify the Λ-quasi-geodesic rays that are tracked by geodesics. As corollaries we show that if W acts geometrically on a CAT(0) space X, then CAT(0) geodesics in X are tracked by Cayley graph geodesics (taking the Cayley graph as equivariantly placed in X) and for any $A \subset S$, the special subgroup $\langle A \rangle$ is quasi-convex in X. We also show that if g is an element of infinite order for (W, S) then the subgroup $\langle g \rangle$ is tracked by a Cayley geodesic in Λ (in analogy with the corresponding result for word hyperbolic groups).

1. Introduction

Suppose G is a group with finite generating set A, and $\Lambda = \Lambda(G, A)$ is the Cayley graph of G with respect to A. If G is word hyperbolic then any quasi-geodesic in Λ is tracked by a geodesic (see [Sh]). The corresponding result for CAT(0) groups is not true. Our main goal in this paper is to classify the quasi-geodesics in the Cayley graph of a finitely generated Coxeter system that are tracked by geodesics. We define a "bracket number" for a Cayley path in terms of the wall crossings of the path and our main theorem is that a quasi-geodesic ray or line is tracked by a geodesic iff the bracket number of the ray or line is bounded. Our principal corollary to this theorem states that if (W, S) is a finitely generated Coxeter system, and W acts geometrically on a CAT(0) space X, then the CAT(0) geodesics of X are tracked by (W, S) Cayley geodesics in X. The corresponding result is not true, even for CAT(0) groups that embed as subgroups of finite index in Coxeter groups (see remark 6.5). If X is the Davis complex for (W, S) or even if W acts as a reflection group on X, the proof of the corollary is straightforward. Unfortunately, the reflection group argument has no analogue when W does not act as a reflection group on X. The principal corollary directly implies that if $A \subset S$ then the special subgroup $\langle A \rangle$ is quasi-convex in X.

2010 *Mathematics Subject Classification.* Primary 20F65, Secondary 20F55.
Key words and phrases. CAT(0) group, Coxeter group, quasi-geodesic.

If a group G acts geometrically on a CAT(0) space X and one is interested in the asymptotic properties of X, it is a considerable advantage to know that CAT(0) geodesics in X are tracked by Cayley geodesics. Clearly, the algebraic properties of G are far more apparent in Cayley geodesics than in CAT(0) geodesics. This theme is highlighted in [MRT] where local connectivity of boundaries of right angled Coxeter groups are analyzed.

The work of B. Bowditch and G. Swarup (see [S]) imply that 1-ended word hyperbolic groups have locally connected boundary. One can easily see from our tracking results that any 1-ended hyperbolic Coxeter group has locally connected boundary.

2. Coxeter Preliminaries

We use M. Davis' book [D] as a general Coxeter group reference for this section. A Coxeter system is a pair (W, S) where S is a generating set for the group W and W has presentation

$$\langle S : (s_i s_j)^{m(i,j)} \text{ for all } s_i, s_j \in S\rangle$$

where $m(i,j) \in \{1, 2, \ldots, \infty\}$, $m(i,j) = 1$ iff $i = j$ (so all generators are order 2) and $m(i,j) = m(j,i)$. If $m(i,j) = \infty$, the element $s_i s_j$ is of infinite order (and the relation $(s_i s_j)^\infty$ is left out of the presentation).

In the Cayley graph $\Lambda = \Lambda(W, S)$, for $w \in W$ and $s \in S$ there are formally two edges from w to $ws = ws^{-1}$, one labeled by the letter s and one labeled by the letter s^{-1}, but as every generator has order two, every edge is doubled. Ordinarily, it is just as easy to identify the letters s and s^{-1} and take a single edge from w to ws corresponding to the twin edges, and we will do so here. The Cayley graph is viewed as a 1-complex with unit length intervals for edges. An *edge path* in Λ is a continuous map $\beta : [0, n] \to \Lambda$ such that $n \in \mathbb{Z}^+$ and for each non-negative integer $k < n$, β maps the interval $[k, k+1]$ isometrically to an edge of Λ. Thus an edge path is determined by the sequence of its vertices, each adjacent to the previous, or alternatively, is determined by its starting point and the word of labels of its edges. Then Λ is taken with edge path metric and the distance between vertices v and vw is the length of the shortest word representing w (the length of a shortest path having label w). Similarly, for β satisfying the corresponding condition, if $\beta : [0, \infty) \to \Lambda$, then β is called a *ray* and, if $\beta : (-\infty, \infty) \to \Lambda$ then β is called a *line*. An edge path, ray, or line β is a *geodesic* if β is an isometry.

A *reflection* in W is a conjugate of an element of S. If $w \in W$ and $s \in S$, then the edge from w to ws is mapped to itself by the reflection wsw^{-1}, so that the vertices w and ws are interchanged, i.e., the edge is reflected across its midpoint. The set of those edges in Λ which are mapped to themselves in this way by some particular reflection r is called a *wall* of Λ. The walls of Λ partition the edges of Λ into disjoint sets. Notationally, we write a wall Q as $[e]$ where e is any edge of the wall Q and we define $\bar{Q}$ to be the union of the edges of Q in Λ. An edge e with label $t \in S$ belongs to a wall Q corresponding to the reflection wsw^{-1} iff a vertex of e is wq where

$qtq^{-1} = s$. An edge path β containing more than one edge from a wall cannot be geodesic, for if e and d are edges in β (with e before d, say) both in the wall given by a reflection r, and β' is the segment of β from the terminal point of e to the initial point of d, then $r\beta'$ is a path from the initial point of e to the terminal point of d and replacing (e, β', d) in β by $r\beta'$ shortens β. The converse of this observation is essentially the deletion condition for Coxeter groups; if β is not geodesic, then there are edges e and d in β belonging to the same wall. Replacing (e, β', d) with $r\beta'$ corresponds to deletion of the letters labeling edges e and d in the label of β. Moreover, in this case, we can always take a e and d in a wall where the intervening segment β' in β is a geodesic which does not cross this wall. The closure of the complement of a wall in Λ has exactly two components (which are interchanged by the reflection) called the *sides* of the wall. Two walls are *parallel* if all edges of one are on the same side of the other. If two walls are not parallel, then they *cross*. The following theorem due to B. Brink and R. Howlett (see theorem 2.8 of [BrH]) is a fundamental result concerning the wall structure of Λ.

Theorem 2.1. (Parallel Wall theorem) *Suppose (W, S) is a finitely generated Coxeter system and Λ is the Cayley graph of W with respect to S. For each positive integer n there is a constant $P(n)$ such that the following holds: given a wall Q and a point p in Λ, if the distance from p to $\bar{Q}$ is at least $P(n)$, then there exist n distinct pairwise parallel walls which separate $\bar{Q}$ from p.*

For a path β in Λ and vertex t of β let the *bracket number of t in β* be the number of walls Q such that there is an edge of Q on either side of t in β. Denote the bracket number of t in β as $B(t, \beta)$. If τ is a subpath of β the *bracket number of τ in β* is the maximum of the numbers $B(t, \beta)$ for all vertices t of τ. Denote this number $B(\tau, \beta)$. Call $B(\beta) \equiv B(\beta, \beta)$ the *bracket number of β*.

3. Wall computations

If α is an edge path in the Cayley graph Λ having consecutive vertices $a = v_0, v_1, \ldots, v_n = b$, then an *$L$-approximation to α* is an edge path β in Λ connecting a and b of the form $\beta = (\beta_1, \ldots, \beta_n)$ where each β_i is a geodesic connecting w_{i-1} to w_i for a sequence of vertices w_i each within L of the corresponding v_i. The points w_i are called *approximation points.*

Lemma 3.1. *Suppose (W, S) is a finitely generated Coxeter system with Cayley graph Λ. Then there is a function f such that, for α any edge path in Λ and β an L-approximation of α, the bracket number $B(\beta)$ is at most $f(B(\alpha), L)$ (that is, $B(\beta)$ is bounded by a constant depending only on $B(\alpha)$ and L and is otherwise independent of the particular α).*

Proof. Let the consecutive vertices of α be $a = v_0, v_1, \ldots, v_n = b$, the approximation vertices of β be $a = w_0, w_1, \ldots, w_m = b$ (so that $d(w_i, v_i) \leq L$

for all i) and β_i be the geodesic subpath of β connecting w_{i-1} to w_i, so $\beta = (\beta_1, \ldots, \beta_m)$. If x is a vertex of β_i and $B(x, \beta)$ is "large", then (as each edge belongs to exactly one wall) there is a wall Q that brackets x on β that is "far" from x and hence far from v_i. Hence it suffices to bound the distance between v_i and a wall Q that brackets x on β. The Parallel Wall theorem implies this distance is large iff there is a large set $\mathcal{Q}$ of (mutually parallel) walls that separate $\bar{Q}$ from v_i, so it suffices to bound the size of the set $\mathcal{Q}$ of walls that separate $\bar{Q}$ from v_i. Say $j < i < k$ such that e_j and e_k are edges of β_j and β_k respectively, and each of e_j, e_k belongs to the wall Q. (See figure 1.)

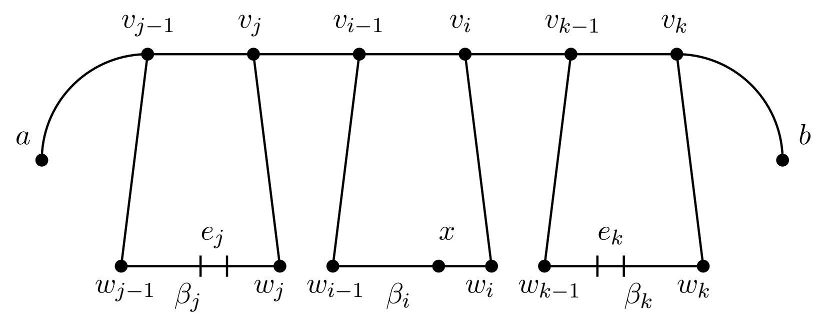

FIGURE 1.

A path δ_j, that begins at the end point of e_j follows β_j to w_j and then travels geodesically from w_j to v_j has length $\leq 3L$. If $\alpha_{j,i}$ is the subpath of α from v_j to v_i, then the path $(\delta_j, \alpha_{j,i})$ must cross each wall of $\mathcal{Q}$. Similarly define a path from e_k to v_i (which also crosses each wall of $\mathcal{Q}$). Then at most $6L$ walls of $\mathcal{Q}$ do not bracket v_i on α. This bounds the size of $\mathcal{Q}$ by $6L + B(\alpha)$. □

Lemma 3.2. *Suppose (W, S) is a Coxeter system and $\alpha = (e_1, \ldots, e_n)$ is a geodesic edge path connecting vertices a and b in $\Lambda(W, S)$ such that α does not cross the wall Q. If e_0 is an edge at a and e_{n+1} an edge at b such that e_0 and e_{n+1} belong to the wall Q then each vertex of α is within $P(1)$ of $\bar{Q}$ (where P is the function of theorem 2.1). In particular, if v is a vertex of α and v' the reflection of v across Q then $d(v, v') \leq 2P(1) + 1$.*

Proof. Otherwise, there is a wall Q' separating a vertex v of α from Q. Hence there is an edge of α between a and v that belongs to Q' and an edge of α between v and b that belongs to Q'. This is impossible as α is geodesic. □

Proposition 3.3. *Suppose (W, S) is a Coxeter system and α is an edge path of $\Lambda(W, S)$ connecting a and b. Then there is an L-approximation β to α such that each vertex of β is on some geodesic connecting a and b and such that $L \leq (2P(1) + 1)B(\alpha)$.*

Proof. Let the consecutive vertices of α be $a = v_0, \dots, v_n = b$. For $0 < i < n$ we choose an approximation point w_i for v_i as follows. Let α_i be a geodesic from a to v_i and β_i a geodesic from v_i to b. Each wall of (α_i, β_i) is crossed exactly once or twice. The number of walls crossed twice by (α_i, β_i) is

$$N_i \equiv \frac{1}{2}(d(a, v_i) + d(v_i, b) - d(a, b)) \leq B(\alpha)$$

Let e be the last edge of α_i belonging to a wall which is crossed twice by (α_i, β_i) and d the edge of β_i in the same wall as e. (See figure 2.)

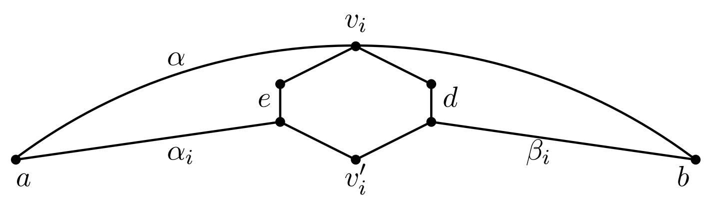

FIGURE 2.

The segment of (α_i, β_i) between e and d is geodesic. Considering the reflection of this segment across the wall containing e and d (equivalently, delete e and d from (α_i, β_i)). Then we see that v_i', the reflection of v_i, is within $2P(1)+1$ of v_i (lemma 3.2), and the distance from v_i' to a (respectively b) is less than that of v_i to a (respectively b). Hence $\frac{1}{2}(d(a, v_i') + d(v_i', b) - d(a, b)) < N_i$ and a geodesic from a to v_i' followed by a geodesic from v_i' to b crosses at most $N_i - 1$ walls twice. Continuing as above at most $N_i (\leq B(\alpha))$ such reflections are needed to move v_i to a point w_i on a geodesic between a and b, and so $d(w_i, v_i) \leq (2P(1) + 1)B(\alpha)$.

It remains to see that each vertex of a geodesic connecting w_i and w_{i+1} belongs to a geodesic connecting a and b. Consider the edge path $(\delta_i, \beta_i, \gamma_i)$ where δ_i is a geodesic connecting a to w_i, β_i is a geodesic connecting w_i to w_{i+1} and γ_i is a geodesic connecting w_{i+1} to b. The paths δ_i and γ_i only cross walls crossed by some (equivalently any) geodesic connecting a to b. If a vertex v of β_i is not on a geodesic connecting a and b then there is a wall R separating v from some (equivalently every) geodesic connecting a and b. As R separates v from a, and δ_i does not cross R, β_i must cross R between w_i and v. Similarly β_i must cross R between v and w_{i+1}. This is impossible as β_i is geodesic. □

If γ is an edge path in Λ connecting the vertices a and b, then each wall separating a and b is crossed an odd number of times by γ and each wall not separating a and b is crossed an even number of times by γ. If α is a geodesic connecting a and b then the walls separating a and b are the walls determined by the edges of α, so the walls separating a and b are in 1-1 correspondence with the edges of some (any) geodesic connecting a and b. The following observations are straightforward.

Lemma 3.4. *Suppose β is an edge path in Λ connecting the vertices a and b then the following are equivalent.*

i) Each vertex of β is on some geodesic connecting a and b
ii) Each edge of β belongs to a wall that separates a from b.
iii) Each wall crossed by β is crossed an odd number of times.
iv) For any vertices c and d of β any wall separating c and d also separates a and b.

The next result is a slightly more sophisticated version of lemma 3.2.

Lemma 3.5. *Suppose α is a geodesic edge path in Λ connecting the vertices a and b, v is a vertex of α, and a and b are each within distance A of $\bar{Q}$ for some wall Q. Then v is within distance $2A(2P(1)+1)+P(1)$ of $\bar{Q}$.*

Proof. Let a' (respectively b') be a vertex of $\bar{Q}$ within A of a (respectively b) and on the same side of Q as is a (respectively b). Let β (respectively γ) be a geodesic from a' to a (respectively b to b').

Case 1. The geodesic α does not cross Q.
In this case the path $\delta_0 \equiv (\beta, \alpha, \gamma)$ does not cross Q. Since $|\beta| \leq A$ and $|\gamma| \leq A$, a sequence of at most $2A$ deletions (the first in the path δ_0) will determine a geodesic connecting a' to b'. As noted in section 2, if at some stage δ_i is not geodesic, then we can take edges e_i and d_i of δ_i in the same wall and such that $(*)$ the subpath of δ_i between e_i and d_i is geodesic and does not cross the wall. Replacing this subpath by its reflection in the wall and deleting e_i and d_i will give a shorter path δ_{i+1} from a' to b'. Note that if δ_i does not cross Q, e_i and d_i are on the same side of Q, so the reflected subpath, being a geodesic by $(*)$, cannot cross Q, and then δ_{i+1} also does not cross Q. After $K \leq 2A$ such deletions, δ_K will be a geodesic connecting a' and b' which does not cross Q.

If v is not between e_1 and d_1 then v is a vertex of δ_1. If v is between e_1 and d_1, then v_1, the reflection of v across the wall $[e_1] = [d_1]$, is within $2P(1)+1$ of v, by lemma 3.2. (Note that the hypotheses of lemma 3.2 are satisfied since we require condition $(*.)$) In any case δ_1 contains a vertex v_1 within $2P(1)+1$ of v. If e_2 and d_2 are deleting edges of δ_1 (satisfying $(*)$), then δ_2, obtained from δ_1 by deleting e_2 and d_2, contains a vertex v_2 within $2P(1)+1$ of v_1, again by lemma 3.2, and so v_2 is within $2(2P(1)+1)$ of v. Inductively then, δ_K contains a vertex v_K within $K(2P(1)+1)$ of v. By lemma 3.2, v_K is within $P(1)$ of $\bar{Q}$ so that v is within $2A(2P(1)+1)+P(1)$ of $\bar{Q}$. This completes case 1.

Case 2. Suppose α crosses Q.
Say the edge e of α between v and b belongs to Q. Repeat the case 1 argument with δ_0 replaced by (β, α'), where α' is the subsegment of α from a to the initial point of e. Similarly if $e \in Q$ is an edge of α between a and v. Note that in both case 2 scenarios, at most A deletions are required to straighten to a geodesic, so the bound is reduced to $A(2(P(1)+1)+P(1)$. $\square$

4. Tracking Quasi-geodesics

We are interested in quasi-geodesic edge paths in Λ. Recall that an edge path β is a (λ, ϵ)-*quasi-geodesic* if for each pair of integers s and t in the domain of β, $|s-t| \leq \lambda d(\beta(s), \beta(t)) + \epsilon$. If β is a ray or line, we say β is quasi-geodesic if it satisfies the corresponding condition for some λ and ϵ. If α and β are edge paths, then β is K-*tracked* by α if each vertex of β is within K of a vertex of α. If α and β are rays or lines, we say β is tracked by α if they satisfy the corresponding condition for some K. If α is a K-approximation to β then β is K-tracked by α, but in general there is no requirement that a tracking path α be piecewise geodesic nor even that the nearest points on α to vertices of β occur in the same order. More generally, say that a set of vertices is (K-)tracked by α if every element of the set is within (some possibly unspecified bound) K of a vertex of α.

Lemma 4.1. *For $i \in \{1,2\}$ suppose β_i is a (λ_i, ϵ_i)-quasi-geodesic edge path in Λ, β_1 is K-tracked by β_2 and $\beta_1(0)$ is within K of $\beta_2(0)$. Assume both β_1 and β_2 are lines, or both are rays, or both are finite length and the terminal points of β_1 and β_2 are within K of one another. Then β_2 is $(\lambda_2(2K+1)+\epsilon_2+K)$-tracked by β_1.*

Proof. Since each vertex of β_1 is within K of a vertex of β_2, we may define an integer function a such that for each integer m (in the domain of β_1), $\beta_1(m)$ is within K of $\beta_2(a(m))$. We take $a(0) = 0$ and if β_i has n_i edges then $a(n_1) = n_2$.

The first two sets of inequalities below follow from the definitions and triangle inequalities and the third set follows from the first two.

$$\begin{aligned}\frac{|a(m+i)-a(m)|-\epsilon_2}{\lambda_2} - 2K &\leq d(\beta_2(a(m+i)), \beta_2(a(m))) - 2K \\ &\leq d(\beta_1(m+i), \beta_1(m)) \\ &\leq d(\beta_2(a(m+i)), \beta_2(a(m))) + 2K \\ &\leq |a(m+i)-a(m)| + 2K\end{aligned}$$

$$\frac{i-\epsilon_1}{\lambda_1} \leq d(\beta_1(m+i), \beta_1(m)) \leq i$$

$$\begin{aligned}\frac{i-\epsilon_1}{\lambda_1} - 2K &\leq |a(m+i)-a(m)| \\ &\leq \lambda_2(d(\beta_1(m+i), \beta_1(m)) + 2K) + \epsilon_2 \\ &\leq (i+2K)\lambda_2 + \epsilon_2\end{aligned} \tag{1}$$

The inequality $|a(i+1)-a(i)| \leq \lambda_2(2K+1)+\epsilon_2$ implies if k is between $a(i)$ and $a(i+1)$ for some i then $\beta_2(k)$ is within $\lambda_2(2K+1)+\epsilon_2+K$ of $\beta_1(i)$. In the case β_1 and β_2 are finite, the condition that terminal points are within K of one another (so that $a(n_1) = n_2$) implies that every integer in the domain of β_2 is between $a(i)$ and $a(i+1)$ for some i and this case is finished. If β_1 and β_2 are rays then $a(i)$ is non-negative and equation 1 (with $m = 0$) implies $a(i)$ is arbitrarily large for large i and again every integer in the domain of β_2 is between $a(i)$ and $a(i+1)$ for some i. If β_1 and β_2 are bi-infinite, then the $a(i)$ may be positive or negative and (again by 1) for large $|i|$, $|a(i)|$ is large, and $\lim_{i\to+\infty} a(i) = \pm\infty$ and $\lim_{i\to-\infty} a(i) = \pm\infty$. It remains to see $\lim_{i\to+\infty} a(i) \neq \lim_{i\to-\infty} a(i)$. Equality is impossible, since otherwise, for every large positive integer i, $a(-i)$ would be between $a(j)$ and $a(j+1)$ for some (depending on i) large positive integer j. But equation 1 implies $a(j)$ and $a(j+1)$ are relatively close and $a(-i)$ and $a(j)$ are far apart. □

Proposition 4.2. *Suppose β is a quasi-geodesic edge path ray in Λ and β is tracked by a geodesic, then β has bounded bracket number.*

Proof. Assume that β is a (λ, ϵ)-quasi-geodesic. Suppose α is a geodesic such that each vertex of β is within L of a vertex of α. For each integer $n \geq 0$, choose an integer $a(n)$ such that $d(\beta(n), \alpha(a(n))) \leq L$. We assume that $a(0) = 0$.

The next two sets of inequalities follow from the definitions and triangle inequalities, and the third follows from the first two.

$$a(n) - 2L \leq d(\beta(n), \beta(0)) \leq a(n) + 2L$$

$$\frac{n-\epsilon}{\lambda} \leq d(\beta(n), \beta(0)) \leq n$$

$$\frac{n-\epsilon}{\lambda} - 2L \leq a(n) \leq n + 2L$$

The proof now makes use of the following claim.

Claim 4.3. *Suppose K is an integer larger than $\lambda(4L+1)+\epsilon$. Then for any integer n, $a(n+K) > a(n)$.*

Proof. Note that if $m > \lambda(n+4L)+\epsilon$ then $a(m) > n+2L \geq a(n)$. So if $K > \lambda(4L+1)+\epsilon$, and $a(n+K) \leq a(n)$, then there is a last integer $K_1 > \lambda(4L+1)+\epsilon$ such that $a(n+K_1) \leq a(n)$. Then (see figure 3)

$$a(n+K_1+1) > a(n) \geq a(n+K_1)$$

Since $d(\beta(n+K_1), \beta(n+K_1+1)) = 1$ for all n, and $d(\beta(i), \alpha(a(i))) \leq L$ for all i, we have

$$d(\alpha(a(n+K_1)), \alpha(a(n+K_1+1))) \leq 2L+1$$

But as $\alpha(a(n))$ is between $\alpha(a(n+K_1))$ and $\alpha(a(n+K_1+1))$ on the geodesic α,

$$d(\alpha(a(n)), \alpha(a(n+K_1+1))) \leq 2L+1$$

Then $d(\beta(n), \beta(n+K_1+1)) \leq 4L+1$. But

$$d(\beta(n), \beta(n+K_1+1)) \geq \frac{1}{\lambda}(K_1+1-\epsilon) > 4L+1$$

the desired contradiction (so the claim is proved). □

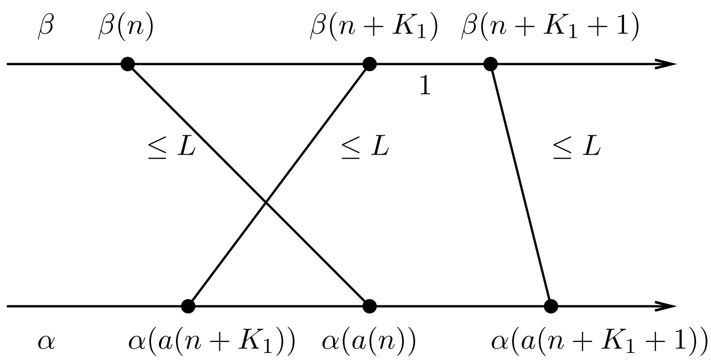

FIGURE 3.

Now suppose $v \equiv \beta(n)$ is a vertex of β with bracket number at least $2\lambda(4L+1)+2\epsilon+K$. Then there are K distinct walls, $Q_1, \ldots, Q_K$ such that for each $i \in \{1, \ldots, K\}$, there is an edge e_i of β preceding v and an edge d_i of β following v such that e_i and d_i belong to the wall Q_i, the subpath of β between e_i and d_i does not cross Q_i, e_i is not one of the $\lambda(4L+1)+\epsilon$ edges of β immediately preceding v and d_i is not one of the $\lambda(4L+1)+\epsilon$ edges of β immediately following v. I.e. $e_i = \beta([t_i, t_i+1])$ where $t_i+1 \leq n-\lambda(4L+1)-\epsilon$ and $d_i = \beta([u_i, u_i+1])$ where $u_i \geq n+\lambda(4L+1)+\epsilon$. (See figure 4.)

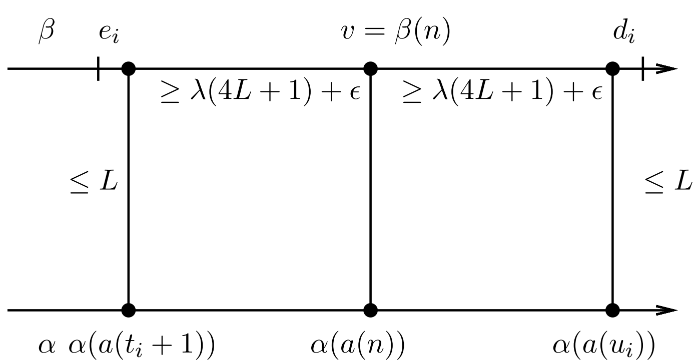

FIGURE 4.

By claim 4.3, $a(t_i+1) < a(n) < a(u_i)$. Hence, by lemma 3.5, $\alpha(a(n))$ is within $2L(2P(1)+1)+P(1)$ of the wall Q_i. For x a vertex of Λ, let $C(k)$ be the number of distinct walls that pass within k of x. Note that C is independent of vertex in Λ. Hence $K \leq C(2L(2P(1)+1)+P(1))$, bounding the bracket number of a vertex of β. ☐

5. Proof of Main Theorem

In order to prove the main theorem, we need two results, one due to B. Brink and R. Howlett [BrH], and a second, due to R. P. Dilworth [Di].

Theorem 5.1. *(Brink-Howlett) Suppose (W, S) is a finitely generated Coxeter system, and $\Lambda(W, S)$ is the Cayley graph of W with respect to S. There is a bound $F_{(W,S)}$ on the number of mutually crossing walls of Λ.*

Dilworth's theorem requires several definitions. If A is a partially ordered set (a set with reflexive, antisymmetric and transitive binary relation $\leq$ on A), then any two elements x and y are *comparable* if either $x \leq y$ or $y \leq x$. Otherwise they are in *incomparable*. A subset C of A is a *chain* when every pair of points in C is a comparable pair. A subset B of A is called an *anitchain* when every pair of points in B is an incomparable pair. The number of points in a maximal antichain is called the *width* of A.

Theorem 5.2. *(Dilworth) If A is a partially ordered set of width w, then A can be partitioned into w chains.*

Suppose x and y are vertices of $\Lambda(W, S)$ and $\mathcal{W}_{(x,y)}$ is the set of walls that separate x and y. We partially order $\mathcal{W}_{(x,y)}$ by saying $P \leq Q$ if either $P = Q$, or P and Q are parallel and P separates x from Q. Note that P and Q are parallel walls of $\mathcal{W}_{(x,y)}$, iff they are comparable. Hence P and Q are incomparable iff they cross. By proposition 5.1, the width of $\mathcal{W}_{(x,y)}$ is at most $F_{(W,S)}$. Applying Dilworth's theorem we have:

Proposition 5.3. *Suppose (W, S) is a finitely generated Coxeter system, and $\Lambda(W, S)$ is the Cayley graph of W with respect to S. For any vertices x and y of Λ the walls separating x and y can be partitioned into at most $F_{(W,S)}$ chains (where any two walls in the same chain are parallel).*

Say a path is *geodesic with respect to a set of walls* if the path crosses each wall of the set either 0 or 1 times. The following lemma is clear.

Lemma 5.4. *Suppose α is an edge path in Λ and α is geodesic with respect to the set of parallel walls $\mathcal{Q}$. If a subpath of α is replaced by a geodesic edge path, then the resulting edge path is geodesic with respect to $\mathcal{Q}$.*

Theorem 5.5. *Suppose (W, S) is a finitely generated Coxeter system with Cayley graph Λ. Then there is a function f such that, for any α a (λ, ϵ)-quasi-geodesic edge path from a to b in Λ, there is a geodesic edge path β from a to b such that α is K-tracked by β for $K = f(B(\alpha), \lambda, \epsilon)$. That is, any (λ, ϵ)-quasi-geodesic α is K-tracked by a geodesic β where the K depends only on the bracket number of α, $B(\alpha)$, and on λ and ϵ, but is otherwise independent of the particular α.*

Proof. Fix B_0, λ, and ϵ. Consider a (λ, ϵ)-quasi-geodesic α from a to b with $B(\alpha) \leq B_0$. By proposition 3.3, take α' an L-approximation to α with every vertex on some geodesic connecting a and b, where L depends only on $B(\alpha) \leq B_0$. By proposition 3.1, $B(\alpha') \leq B_1$ for a larger bound B_1 still determined only from B_0. Moreover, α' is a (λ', ϵ')-quasi-geodesic for a λ' and ϵ' determined only from L, λ and ϵ. More specifically, write α' as $(\alpha'_1, \ldots, \alpha'_q)$ where the α'_i are geodesic of length $\leq 2L+1$, with initial and terminal vertices of α'_i within L of $\alpha(i-1)$ and $\alpha(i)$ respectively. Then each interior vertex of α'_i is within $3L$ of $\alpha(i-1)$ and $\alpha(i)$. Thus an $\alpha'(i)$ and $\alpha'(j)$ are within $3L$ of vertices $\alpha(l)$ and $\alpha(m)$, respectively, such that vertices of α between these have approximation points between $\alpha'(i)$ and $\alpha'(j)$, and hence

$$\begin{aligned} |i-j| &\leq (2L+1)(|l-m|+2) \\ &\leq (2L+1)(\lambda d(\alpha(l), \alpha(m)) + \epsilon + 2) \\ &\leq (2L+1)(\lambda(d(\alpha'(i), \alpha'(j)) + 6L) + \epsilon + 2) \\ &= (2L+1)\lambda d(\alpha'(i), \alpha'(j)) + (2L+1)(\lambda(6L) + \epsilon + 2) \end{aligned}$$

from which λ' and ϵ' can be read off.

A geodesic β that K-tracks α' will $K+L$-track α. Thus it suffices to find, in terms of B_1, λ', and ϵ', a K that works for those (λ', ϵ')-quasi-geodesics α' with $B(\alpha') \leq B_1$ and having every vertex on some geodesic between its endpoints, i.e., we may as well prove the theorem for α having every vertex on some geodesic between its endpoints.

The proof is a double induction argument. For the “outside” induction, we show by induction on A with $1 \leq A \leq F \equiv F_{(W,S)}$ that, for any positive integer B_0 and constants λ and ϵ, there is a constant K_A such that, for any (λ, ϵ)-quasi-geodesic α from a to b, with $B(\alpha) \leq B_0$ and every vertex on a geodesic between a and b, if $\mathcal{W}_{(a,b)}$, the set of walls separating a and b, can be partitioned into A or fewer chains, then there is a geodesic edge path β from a to b such that α is K_A-tracked by β. By proposition 5.3, the set of walls separating any a and b can be partitioned into at most F chains, so $K = K_F$ suffices for the theorem.

Note that if $A = 1$ then all walls separating a and b are parallel. In this case, the walls separating a and b are ordered as $Q_1, \ldots, Q_m$ where for $i < j < k$, Q_j separates Q_i from Q_k. Hence, there is a unique, geodesic edge path β connecting a and b, and β crosses Q_1, then Q_2, etc. The path α only crosses the walls separating a and b (see lemma 3.4) and, in this case, is a geodesic modulo backtracking. Eliminating backtracking on α produces β. Each vertex of α is a vertex of β and the basis case is complete with $K_1 = 0$.

Assume the (outside) induction statement is true for $A < M$. For the “inside” induction, we show by induction on N with $0 \leq N \leq M$ that, for any positive integer B_0 and constants λ and ϵ, there is a constant $K_{M,N}$ such that, for any (λ, ϵ)-quasi-geodesic α from a to b, with $B(\alpha) \leq B_0$ and every vertex on a geodesic between a and b, if $\mathcal{W}_{(a,b)}$ can be partitioned into M chains $\mathcal{Q}_1, \ldots, \mathcal{Q}_M$, such that α is geodesic with respect to all but at most N of the $\mathcal{Q}_i$, then there is a geodesic edge path β from a to b such that α is $K_{M,N}$-tracked by β. The induction step for the outside induction is completed then by taking $K_M = K_{M,M}$.

Fix a B_0, λ, and ϵ. Consider a (λ, ϵ)-quasi-geodesic α from a to b, with $B(\alpha) \leq B_0$ and every vertex on a geodesic between a and b, such that M chains $\mathcal{Q}_1, \ldots, \mathcal{Q}_M$ partition the set of walls separating a and b. Again by lemma 3.4, α only crosses walls separating a and b, each an odd number of times, and each of which belongs to some $\mathcal{Q}_i$. Suppose α is geodesic with respect to all but at most N of these $\mathcal{Q}_i$. If $N = 0$, then α is geodesic, so $\beta = \alpha$ with $K_{M,0} = 0$.

Assume the (inside) induction statement is true for $N = H-1$. Assume α is such that, taking $\mathcal{Q}_i$ indexed conveniently, α is geodesic with respect to $\mathcal{Q}_i$ for $H+1 \leq i \leq M$. If α is also geodesic with respect to $\mathcal{Q}_H$, then apply the induction hypothesis. Otherwise, write α as $(e_1, \ldots, e_n)$ with consecutive vertices $a \equiv a_1, \ldots, a_n \equiv b$. Let i be the first integer such that e_i is an edge of a wall of $\mathcal{Q}_H$ where for some $j > i$, e_j and e_i are in the same wall Q. Take j the largest integer such that $e_j \in Q$. Since α crosses Q an odd number of times, the path $\alpha_{i,j} \equiv (e_i, \ldots, e_{j-1})$ (from a_i to a_j) crosses Q an even number of times. A geodesic $\beta_{i,j}$ connecting a_i to a_j does not cross Q. All walls of $\mathcal{Q}_H$ are parallel to one another, and $\beta_{i,j}$ begins and ends adjacent to (at endpoints of edges in), and on the same side of, $Q \in \mathcal{Q}_H$, Suppose Q' is a wall of $\mathcal{Q}$ other than Q, with sides S_1 (containing Q) and S_2, so β begins and ends in S_1. If $\beta_{i,j}$ crossed Q' to side S_2, then as $\beta_{i,j}$ ends in side S_1, $\beta_{i,j}$ would cross Q' a second time. But a geodesic cannot cross a wall twice and so $\beta_{i,j}$ does not cross a wall of $\mathcal{Q}_H$.. Hence a_i and a_j are not separated by a wall of $\mathcal{Q}_H$. By proposition 3.3, take $\alpha'_{i,j}$ an L-approximation to $\alpha_{i,j}$ for L determined from $B(\alpha_{i,j}) \leq B(\alpha) \leq B_0$, with $\alpha'_{i,j}$ a (λ', ϵ')-quasi-geodesic edge path connecting a_i to a_j (for some λ' and ϵ' determined from L as above), $B(\alpha'_{i,j}) \leq B_1$ (for a B_1 determined from L as above), and such that each vertex of $\alpha'_{i,j}$ is on a geodesic connecting a_i to a_j. By lemma 3.4, each wall separating a_i and a_j also separates a and b, and the walls separating a_i and a_j can be partitioned into fewer than M chains. By (outside) induction, there is a geodesic $\beta_{i,j}$ connecting a_i and a_j which K_{M-1}-tracks $\alpha'_{i,j}$ and therefore $K_{M-1} + L$-tracks $\alpha_{i,j}$ (where the K_{M-1} is determined using B_1, λ', and ϵ'). Replace $\alpha_{i,j}$ in α by $\beta_{i,j}$ to obtain a path α_1.

Now the path α_1 crosses Q exactly once at e_j. The walls of $\mathcal{Q}_H$ are ordered as $Q_1, Q_2, \ldots$ so that if $i < j$, then Q_i separates a from Q_j, and Q_j separates Q_i from b. A wall of $\mathcal{Q}_H$ preceding Q in this ordering is not crossed by α_1 after e_j. Hence if $\mathcal{Q} \subset \mathcal{Q}_H$ is the set of walls of $\mathcal{Q}_H$ preceding Q and including Q, then α_1 is geodesic with respect to $\mathcal{Q}$ and (by lemma 5.4), α_1 is geodesic with respect to each set $\mathcal{Q}_i$ for $i > H$. Now suppose e_k is the first edge of α_1 such that e_k is an edge of a wall Q of $\mathcal{Q}_H$, and for some $l > k$, $e_l \in Q$. Then e_k follows e_j on α_1, and if we assume e_l is the last edge of α_1 in Q, then as above $(e_k, \ldots, e_{l-1})$ can be replaced by a geodesic close to $(e_k, \ldots, e_{l-1})$ with the same tracking constant $K_{M-1} + L$. At each step, disjoint subpaths of the original α are replaced.

Continuing, the resulting path α_* $K_{M-1} + L$-tracks α, is geodesic with respect to $\mathcal{Q}_i$ for $H \leq i \leq M$, and crosses the same walls as α. By lemma 3.4 every vertex of α_* belongs to some geodesic connecting a to b. We show α_* is a $(\hat{\lambda}, \hat{\epsilon})$-quasi-geodesic with $B(\alpha_*) \leq B_2$ for appropriately bounded B_2, $\hat{\lambda}$ and $\hat{\epsilon}$. Let α_{**} be a $K_{M-1} + L$ approximation to α obtained by replacing each $\alpha_{i,j}$ by a path $\hat{\beta}_{i,j}$ described as follows: If the consecutive vertices of $\alpha_{i,j}$ are $w_0, \ldots, w_m$ then for $1 \leq k \leq m-1$ let $\hat{w}_k$ be a vertex of the geodesic $\beta_{i,j}$ within $K_{M-1} + L$ of w_k. For $0 \leq k \leq m-1$ connect $\hat{w}_k$ and $\hat{w}_{k+1}$ with a geodesic subpath of $\beta_{i,j}$. Concatenate these geodesics to form $\hat{\beta}_{i,j}$. Then as above α_{**} is a $(\hat{\lambda}, \hat{\epsilon})$-quasi-geodesic with $B(\alpha_*) \leq B_2$ for appropriately bounded B_2, $\hat{\lambda}$ and $\hat{\epsilon}$. Note that $\beta_{i,j}$ is obtained by eliminating backtracking in $\hat{\beta}_{i,j}$. Eliminating backtracking in a (λ, ϵ)-quasi-geodesic does not increase bracket number and the resulting path remains a (λ, ϵ)-quasi-geodesic. The path α_* is obtained from α_{**} by eliminating backtracking so that α_* is a $(\hat{\lambda}, \hat{\epsilon})$-quasi-geodesic with bracket number $\leq B_2$.

By induction hypothesis, there is a geodesic β that $K_{M,H-1}$-tracks α_*, and so β $K_{M,H}$-tracks α for $K_{M,H} = K_{M,H-1} + K_{M-1} + L$. Hence the inside induction step is established and the theorem follows. □

Note that the bound F from proposition 5.3 on the number of chains needed to partition the set of walls separating two points a and b of Λ limits the total number of times the induction steps are carried out to arrive at a geodesic.

6. Consequences of the Main Theorem

Corollary 6.1. *Suppose (W, S) is a finitely generated Coxeter system, and Λ is the Cayley graph of W with respect to S. Any infinite or bi-infinite (λ, ϵ)-quasi-geodesic edge path α with bounded bracket number B is K'-tracked by an edge path geodesic where K' is a constant only depending on λ, ϵ, B and S.*

Proof. The proof is a standard local finiteness argument in both the infinite and bi-infinite case. We give the bi-infinite case. Write α as the edge path $(\ldots, e_{-1}, e_0, e_1, \ldots)$ in Λ. Let v_i be the initial point of e_i. By theorem 5.5, there is a Λ-geodesic β_n which K-tracks $\alpha_n \equiv (e_{-n}, \ldots, e_n)$. Note that every vertex of β_n is within $2K$ of a vertex of α. For each positive integer n, some vertex x_n of β_n is within K of v_0. Hence there is an infinite number of x_n that are equal. Of this infinite subcollection of x_n, infinitely many have the same pair of edges one preceding and one following x_n on β_n, of this infinite collection of x_n there is an infinite subcollection that have the same four edges - the two preceding and the two following x_n being exactly the same. Continuing, we have a bi-infinite geodesic β and each vertex of β is within

$2K$ of a vertex of α. As α is a (λ, ϵ)-quasi-geodesic, lemma 4.1 implies each point of α is within $\lambda(4K+1)+\epsilon+2K$ of β. □

The next result follows directly from proposition 4.2 and corollary 6.1.

Corollary 6.2. *Suppose (W, S) is a finitely generated Coxeter system, and Λ is the Cayley graph of W with respect to S. Then a quasi-geodesic edge path ray in Λ is tracked by a geodesic iff it has bounded bracket number.*

A metric space (X, d) is a called a *geodesic metric space* if every pair of points are joined by a geodesic. It is *proper* if for any $x \in X$, the ball of radius r about X is compact for all positive numbers r. A group W acts *geometrically* on a space if the action is properly discontinuous, co-compact and by isometries.

Let (X, d) be a proper complete geodesic metric space. If $\triangle\, a\,b\,c$ is a geodesic triangle in X, then consider $\triangle\, \overline{a}\,\overline{b}\,\overline{c}$ in the Euclidean plane $\mathbb{E}^2$ with the same side lengths, called a *comparison triangle*. We say X satisfies the CAT(0) *inequality*, and say (X, d) is a CAT(0) space, if for any $\triangle\, a\,b\,c$ in X, and any two points p, q on $\triangle\, a\,b\,c$, the corresponding points $\overline{p}, \overline{q}$ the same distances from vertices on the sides in a comparison triangle satisfy

$$d(p,q) \leq d_{\mathbb{E}^2}(\overline{p}, \overline{q})$$

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].

Suppose (W, S) is a finitely generated Coxeter system, Λ is the Cayley graph of W with respect to S, and W acts geometrically on a CAT(0) space X. Fix a point $x \in X$, and define $\Lambda_x \subset X$ to have as vertices, the orbit Wx, and as edges, CAT(0) geodesic paths connecting wx and wsx, for $w \in W$ and $s \in S$. There is a proper W-equivariant map $P_x : \Lambda \to \Lambda_x$ so that P_x maps the identity vertex of Λ to x. This P_x is a quasi-isometry of Λ with edge path metric d_Λ into (X, d), and we ordinarily think of Λ_x as being essentially a copy of Λ in X (though for some purposes we might require at least that x is not fixed by any generator). We consider rays and lines in (X, d), and define tracking in (X, d) with respect to the metric d, all analogously to the definitions given earlier for Cayley graphs.

Intuitively, the next result says that when a Coxeter group acts geometrically on a CAT(0) space, CAT(0) geodesics are tracked by Cayley graph geodesics. This result generalizes the right angled version of the same result in [MRT].

Corollary 6.3. *Suppose (W, S) is a finitely generated Coxeter system, Λ is the Cayley graph of W with respect to S, W acts geometrically on the proper CAT(0) space X, and take $x \in X$ with $P_x : \Lambda \to \Lambda_x$ as above. Then* *any CAT(0) geodesic ray in X is tracked by (the image of) a Cayley graph geodesic in Λ_x.*

Proof. For a given CAT(0) geodesic α we find a Cayley graph geodesic β such that $P_x(\beta)$ tracks α. It suffices to find λ, ϵ, K and B such that any (finite) CAT(0) geodesic α is K-tracked by a Cayley (λ, ϵ)-quasi-geodesic with bracket number $\leq B$. Since W acts co-compactly on X, there is an integer K_1 such that every point of X is within K_1 of the orbit Wx. For each integer $0, 1, \ldots, N$ such that N is less that or equal to the length of α, choose a point $v_i x$ of Wx within K_1 of $\alpha(i)$. Let β_i be a Λ-geodesic connecting v_i to v_{i+1} and β be the Λ-edge path $(\beta_0, \beta_1, \ldots)$. Since the map $P_x : \Lambda \to \Lambda_x$ is quasi-isometric, there are numbers λ and ϵ such that any such β is a (λ, ϵ)-quasi-geodesic in Λ, and numbers D_Λ and D_X such that the length of any β_i is less than or equal to D_Λ (in Λ) and every point of such a $P_x(\beta_i)$ is within D_X of $\alpha(i)$ (in X). Certainly every point of α is within $K \equiv K_1 + 1$ of $P_x(\beta)$.

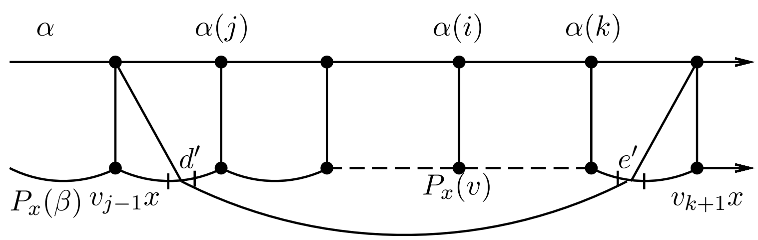

FIGURE 5.

Hence it suffices to bound the bracket number of such a β. If v is a vertex of β_i and e and d are edges of β preceding and following v respectively such that e and d belong to the same wall Q of Λ, then e is an edge of β_j and d is an edge of β_k where $j \leq i \leq k$. The mid-points e' of $P_x(e)$ and d' of $P_x(d)$ are fixed (in Λ_x and X) by the reflection $r_Q \in W$ for the wall Q. Hence the geodesic in X connecting d' and e' is fixed by r_Q. Now, d' (respectively e') is within D_X of $\alpha(j-1)$ (respectively $\alpha(k+1)$) and $P_x(v)$ is within D_X of $\alpha(i)$. By the CAT(0) inequality for quadrilaterals (in particular for the quadrilateral determined by d', e', $\alpha(j-1)$, and $\alpha(k+1)$) $\alpha(i)$ is within D_X of a point of the X-geodesic connecting d' to e' and hence $\alpha(i)$ is within D_X of a fixed point of r_Q. (See figure 5.)

Since the action of W on X is properly discontinuous, there is a bound B on the number of reflections r_Q such that r_Q does not take the ball of radius D_X centered at $\alpha(i) \in X$ (or any other point of X) off of itself. Hence there cannot be more than B walls bracketing the vertex v of β. □

Remark 6.4. *Note that the above proof is valid even when W does not act co-compactly on the CAT(0) space X, as long as the CAT(0) geodesic remains a bounded distance from Λ_x for some x.*

Remark 6.5. *Unfortunately, this result does not hold for arbitrary finitely generated CAT(0) groups. Let* T *be the Cayley graph of the free group* $F_2 = \langle x, y\rangle$*, a tree, and* R *the Cayley graph of the infinite cyclic group* $\mathbb{Z} = \langle z\rangle$*, a line, both CAT(0) spaces. Take* $G = F_2 \times \mathbb{Z}$ *acting componentwise and geometrically on the CAT(0) space* $X = T \times R$ *with metric* $d((w_1, h_1), (w_2, h_2)) = \sqrt{d_T(w_1, w_2)^2 + d_R(h_1, h_2)^2}$*. Now with* $a = xz$ *and* $b = yz$*,* $S = \{a, b, z\}$ *is a generating set of* G*. Let* Λ *be the Cayley graph with respect to this generating set and fix* $* \in X$ *to have components the identity vertices of the Cayley graphs* T *and* R*, thus determining* Λ_**. Let* α *be the CAT(0) geodesic ray from* $*$ *which passes through the points* v_n* *where* $v_0 = 1$*, and for* $n > 0$*,* $v_n = v_{n-1}x^{2^n}y^{-2^n}$*, that is, a ray having a constant second component, contained in a horizontal copy of* T*, but following exponentially increasing long sequences of* x *and* y^{-1} *edges. A Cayley graph geodesic from* $v_{n-1}*$ *to* v_n* *is labeled* $a^{2^n}b^{-2^n}$ *but this path in* Λ_* *increases to a second component value of* 2^n *at its midpoint, far from* α*. Any Cayley graph geodesic approaching* α *within some bound arbitrarily far along must also diverge arbitraily far from* α *far enough along.*

We close this remark with the observation that the Coxeter group $W = (\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2) \times (\mathbb{Z}_2 * \mathbb{Z}_2)$ *contains as a subgroup of finite index isomorphic to* $G \equiv F_2 \times \mathbb{Z}$*. So, while CAT(0) geodesics for* W *are tracked by Cayley geodesics (determined by Coxeter generators),* G *acts geometrically on a CAT(0) space where CAT(0) geodesics are not tracked by Cayley geodesics (for the obvious generators).*

The following result answers a question posed by K. Ruane.

Corollary 6.6. *Suppose* (W, S) *is a finitely generated Coxeter group with Cayley graph* Λ*, acting geometrically on the CAT(0) space* X*, and take an* $x \in X$*, and* $P_x : \Lambda \to \Lambda_x$*, as before, mapping* Λ *quasi-isometrically and* W*-equivariantly into* X*. Then for each subset* $A \subset S$*, (the image of) the subgroup* $\langle A\rangle$ *is quasi-convex in* X*.*

Proof. Let K be the tracking constant from corollary 6.3. Suppose $a_1, a_2 \in \langle A\rangle$ and α is a CAT(0) geodesic in X from $P_x(a_1)$ to $P_x(a_2)$. Let β be a Λ_x, edge path geodesic which K-tracks α. I.e. there is a $\Lambda(W, S)$ geodesic β', from a_1 to a_2 such that $P_x(\beta') = \beta$. Since $a_i \in \langle A\rangle$, the edge labels of β' are all in A. This means all vertices of β' are in $\langle A\rangle$, and so the image of α is within K of $P_x(\langle A\rangle)$. □

The next result says that elements of infinite order in a Coxeter group are tracked by geodesics in the standard Cayley graph.

Corollary 6.7. *Suppose* (W, S) *is a finitely generated Coxeter system and* $g \in W$ *is an element of infinite order. Then in the Cayley graph* $\Lambda(W, S)$ *the elements* $\{\ldots, g^{-2}, g^{-1}, 1, g, g^2, \ldots\}$ *are tracked by a Cayley graph geodesic.*

Proof. We know by G. Moussong [Mo], all finitely generated Coxeter groups are CAT(0). Let X be any CAT(0) space such that W acts geometrically on X. We write $|g| = \inf\{d(x, gx) : x \in X\}$ and define the min set of g by $\min(g) = \{x \in X : d(x, gx) = |g|\}$, those points of X moved a minimum amount by g. The min set of g contains a geodesic line l that is invariant under the action of g. Let x be any point in X and Λ_x the copy of $\Lambda(W, S)$ in X at x. Let α be an S-geodesic for g. Observe that the edge path line l_g in Λ_x determined by positive and negative iterates of α at x is a bounded distance from l. The proof of corollary 6.3 shows that l_g is a quasi-geodesic with bounded bracket number and so by corollary 6.1 is tracked by a Cayley graph geodesic. □

One of the fundamental asymptotic results for word hyperbolic groups is that 1-ended word hyperbolic groups have locally connected boundary. This result follows from a long program of results by several authors, notably B. Bowditch, and concluded by G. Swarup [S]. To give a feeling for the reach of our results, we outline an elementary proof of this fact for Coxeter groups.

Corollary 6.8. *If W is a 1-ended word hyperbolic Coxeter group then the boundary of W is locally connected.*

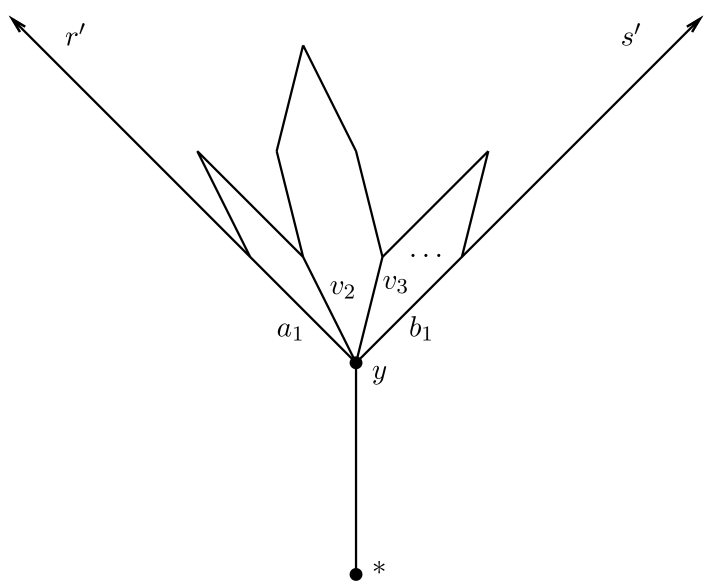

FIGURE 6.

Proof. We use an elementary form of a construction of a “filter” in [MRT] (where a partial classification of right angled Coxeter groups with locally connected boundaries is produced). Suppose W acts geometrically on the CAT(0) space X, with base point x. Let Λ_x be the copy of the Cayley graph of (W, S) at x in X with proper W-equivariant map $P_x : \Lambda(W, S) \to \Lambda_x$. Suppose r and s are “close” geodesic rays in X, with $r(0) = s(0) =$

x. Choose Λ (edge path) geodesics r' and s' at $*$ (the identity vertex of $\Lambda(W,S)$), such that $P_x(r')$ and $P_x(s')$ K-track r and s respectively. Since r and s are close in ∂X, we may assume that r' and s' have long initial segments with "close" terminal points. For simplicity we assume these initial segments agree. If y is the last vertex of this common initial segment, say the edge of r' following y has label a_1 and the edge of s' following y has lablel b_1. The presentation diagram $\Gamma(W,S)$ of (W,S) has vertex set S and an edge labeled $m(i,j)$ between distinct vertices s_i, s_j if $m(i,j) \neq \infty$. Since W is 1-ended no subset A of S with $\langle A \rangle$ a finite group separates Γ (see corollary 16 of [MT]). The set B of S-elements that label edges at y with end points closer to $*$ than y is to $*$ generates a finite subgroup of W (see lemma 4.7.2 of [D]). The set of vertices of Γ corresponding to B does not separate Γ and B does not contain a_1 or b_1. Hence there is an edge path in Γ from a_1 to b_1 avoiding B. Let the consecutive vertices of this path be $a_1 = v_1, v_2, \ldots, v_n = b_1$. If $q(i,i+1)$ is the (finite) order of $v_i v_{i+1}$ then the relation $(v_i v_{i+1})^{q(i,i+1)}$ determines a loop at $y \in \Lambda$. By way of the deletion condition, it is an easy exercise to see that for any subset C of S and geodesic α connecting vertices v_1, v_2 of Λ there is a unique closest vertex x of $v_2\langle C \rangle$ to v_1, and for any geodesic β at x in the letters of C, (α, β) is geodesic in Λ. Since the edges labeled v_i and v_{i+1} at y extend the geodesic from $*$ to y, it must be that y is the closest point of $y\langle v_i, v_{i+1} \rangle$ to $*$, and the two half loops at y making up this loop extend the Cayley geodesic from $*$ to y. Consider the subgraph F_1 of Λ determined by the edge paths r', s' and the edge loops for each $v_i v_{i+1}$ (see figure 6). Each v_i determines an edge of F_1 (with label v_i) beginning at y. At the end point of this edge there are two edges of F_1 that extend a Cayley geodesic from $*$ to y. Build a set of loops as with a_1 and b_1 for each of these pairs of edges. Then F_2 is F_1 union all new loops. Continuing we build a 1-ended subgraph $F = \cup_{i=1}^{\infty} F_i$ of Λ such that for each vertex v of F, not on the common overlap of r' and s', there is a Cayley geodesic from $*$ to v in F which passes through y. We claim that L, the limit set of $P_x(F)$ is a "small" connected set containing r and s (and so ∂X is locally connected). Certainly, r and s are in L. Since F is 1-ended and P_x is proper, L is connected. If v is a vertex of F, then there is a Cayley geodesic α_v from $*$ to v (which passes through y for all but finitely many v). If $z \in L$ then let $z_1, z_2, \ldots$ be a sequence of vertices of F such that $P_x(z_i)$ converges to z. The CAT(0) geodesic from x to $P_x(z_i)$ is K- tracked by a Cayley geodesic β_i in Λ_x. As W is word hyperbolic the Cayley geodesics $P_x(\alpha_{v_i})$ and β_i (with the same end points) must δ-fellow travel (for a fixed constant δ). In particular each β_i must pass "close" to $P_x(y)$ and so z is close to both r and s in $\partial X \equiv \partial W$. □

References

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

[BrH] B. Brink, R.B. Howlett. A finiteness property and an automatic structure for Coxeter groups. *Math. Ann.* **296(1)** (1993), 179-190.

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

[Di] R.P. Dilworth. A Decomposition Theorem for Partially Ordered Sets. *Ann. of Math.* **51** (1950), 161-166.

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

[MT] M. Mihalik, S. Tschantz. Visual decompositions of Coxeter grouip. *Groups Geom. Dyn.* **3** (2009), 173-198.

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

[Sh] H. Short (ed.) Notes on word hyperbolic groups. In *Group Theory from a Geometrical Viewpoint* (E. Ghys, A. Haefliger and A. Verjovsky ed.) World Scientific 1990, 3-64.

[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: michael.l.mihalik@vanderbilt.edu
E-mail address: steven.tschantz@vanderbilt.edu