

The geometry of purely loxodromic subgroups of right-angled Artin groups

THOMAS KOBERDA, JOHANNA MANGAHAS, AND SAMUEL J. TAYLOR

ABSTRACT. We prove that finitely generated *purely loxodromic* subgroups of a right-angled Artin group $A(\Gamma)$ fulfill equivalent conditions that parallel characterizations of convex cocompactness in mapping class groups $\text{Mod}(S)$. In particular, such subgroups are quasiconvex in $A(\Gamma)$. In addition, we identify a milder condition for a finitely generated subgroup of $A(\Gamma)$ that guarantees it is free, undistorted, and retains finite generation when intersected with $A(\Lambda)$ for subgraphs Λ of Γ . These results have applications to both the study of convex cocompactness in $\text{Mod}(S)$ and the way in which certain groups can embed in right-angled Artin groups.

1. INTRODUCTION

1.1. **Overview.** Let Γ be a finite simplicial graph with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$, and let $A(\Gamma)$ be the corresponding right-angled Artin group. That is, we write

$$A(\Gamma) \cong \langle V(\Gamma) \mid [v_i, v_j] = 1 \text{ if and only if } \{v_i, v_j\} \in E(\Gamma) \rangle.$$

Right-angled Artin groups have occupied an important position in geometric group theory in recent years. Their intrinsic algebraic structure has been of interest since the 1980s [Dro87, Ser89, Mih68]. Right-angled Artin groups also played a key role in the study of three manifold topology, culminating in Agol's resolution of the virtual Haken conjecture [Ago13, KM12, Wis11]. Right-angled Artin groups are also a prototypical class of $\text{CAT}(0)$ groups, and have figured importantly in the study of mapping class groups of surfaces [CW04, CLM12, Kob12, KK13, KK14b, MT13].

In this article, we concentrate on this lattermost aspect of right-angled Artin group theory, where we think of right-angled Artin groups both as commonly occurring subgroups of mapping class groups of surfaces, and as algebraically similar to mapping class groups themselves.

Date: December 10, 2014.

Key words and phrases. right-angled Artin group, extension graph, convex cocompact subgroup, loxodromic isometry.

In particular, we study a class of finitely generated subgroups of right-angled Artin groups, called *purely loxodromic subgroups*. These are certain classes of finitely generated free subgroups of right-angled Artin groups which we prove are (in a very strong sense) quasiconvex in the ambient right-angled Artin group, and which have quasi-isometric orbit maps to the right-angled Artin group analogue of the curve complex, i.e. the extension graph.

From this last perspective, we show that purely loxodromic subgroups of right-angled Artin groups are analogous to convex cocompact subgroups of mapping class groups of surfaces. Convex cocompact subgroups are a class of subgroups of mapping class groups distinguished by Farb and Mosher [FM02] and which have natural and useful geometric properties. Indeed, our main theorem represents the analog to an open question for mapping class groups that has received much attention in recent years. (See Section 1.3 for details.) We shall see that, combining our results with earlier results of the second and third author, purely loxodromic subgroups of right-angled Artin groups generally provide an explicit source of convex cocompact subgroups of mapping class groups.

From our main theorem (Theorem 1.1), that strong quasiconvexity properties are implied by an element-wise condition on a finitely generated subgroup of a right-angled Artin group, several applications of general interest follow. For example, as referenced above, recent groundbreaking results imply that many geometrically significant groups embed as quasiconvex subgroups of right-angled Artin groups [Ago13, HW14]. We show that when such a group contains distorted subgroups (e.g. fiber subgroups of fibered 3-manifold groups or free-by-cyclic groups) there are elements of these subgroups that map to relatively simple elements of the right-angled Artin group. In particular, the images of these elements have non-cyclic centralizers in $A(\Gamma)$. See the discussion after Theorem 1.2 below for details.

1.2. Main results. When a right-angled Artin group $A(\Gamma)$ does not decompose as a direct product, its typical (as in [Sis11]) elements are what we call *loxodromic*. Equivalently, these are its Morse elements, its elements with contracting axes, its elements that act as rank-one isometries on the CAT(0) space associated to $A(\Gamma)$, and its elements with cyclic centralizers [Ser89, BC12, BF09]. In [KK14a], loxodromic elements are characterized as those with unbounded orbit in the action of $A(\Gamma)$ on its *extension graph* Γ^e , a hyperbolic space introduced by the first author and Kim to study embeddings between right-angled Artin groups in [KK13]. Here we study *purely loxodromic* subgroups

of $A(\Gamma)$: those in which every non-trivial element is loxodromic. Such subgroups exist in no short supply—in fact, they are typical [TT]; see Section 2.4 for details. Our main result is the following:

Theorem 1.1. *Suppose $H < A(\Gamma)$ is finitely generated and $A(\Gamma)$ does not decompose as a direct product. Then the following are equivalent.*

- (1) *Some (any) orbit map from H into Γ^e is a quasi-isometric embedding.*
- (2) *H is stable in $A(\Gamma)$.*
- (3) *H is purely loxodromic.*

A subgroup H of a finitely generated group G is *stable* if it is undistorted and quasigeodesics in G between points in H have uniformly bounded Hausdorff distance (see Section 6 for details). This property was defined by the third author and Durham in [DT14], where they also show that it implies that H is quasiconvex with respect to any word metric on G . We remark that both stability and condition (2) are strong properties about global geometry—the embedding of the subgroup in either a relevant space or the ambient group—whereas (1) can be read as a purely algebraic condition (that every element in H has cyclic centralizer). In fact, the proof goes through a fourth equivalence, *join-busting* (we give the definition in Section 1.4), which may be considered a local, combinatorial condition.

A milder condition on $H < A(\Gamma)$ is that none of its nontrivial elements conjugate into a subgroup generated by a star of Γ (see Section 2.2 for complete definitions). In that case, we say that H is *star-free*. For this more general class of subgroups of $A(\Gamma)$, we have

Theorem 1.2. *If $H < A(\Gamma)$ is finitely generated and star-free, then*

- (1) *H is a free group,*
- (2) *H is undistorted in $A(\Gamma)$, and*
- (3) *We have*

$$H_\Lambda = H \cap A(\Lambda)$$

is finitely generated, for any subgraph $\Lambda \subset \Gamma$.

The main content of Theorem 1.2 is its last two points; statement (1) is known to specialists. Statement (2) indicates that, among the multiple known constructions of distorted subgroups in right-angled Artin groups, “star-words” are a necessary common feature. In particular, it is a recent theorem of Hagen–Wise that a hyperbolic free-by-cyclic group G has a finite index subgroup G' that quasi-isometrically embeds in a right-angled Artin group $A(\Gamma)$ [HW14, HW13]. Theorem 1.2 implies that the fiber subgroup of $G' \leq A(\Gamma)$, which is always distorted, necessarily contains words that are conjugate into star subgroups.

Regarding the third statement of Theorem 1.2, it is not generally true that a finitely generated free subgroup of $A(\Gamma)$ must intersect $A(\Lambda)$ in a finitely generated group; we give a counterexample in Section 8. We call statement (3) the *Hanna Neumann property* for free subgroups of right-angled Artin groups.

1.3. Motivation from mapping class groups. In the mapping class group $\text{Mod}(S)$, typical (as in [Riv08, Mah11, Sis11]) elements are *pseudo-Anosov*. Here pseudo-Anosov mapping classes may be defined as those with unbounded orbit in the mapping class group action on its associated *curve graph* $\mathcal{C}(S)$, a famously hyperbolic space [MM99]. Pseudo-Anosov elements are also alternately characterized by being Morse, having virtually infinite cyclic centralizers, and having contracting axes in relevant spaces [Beh06, Min96].

Farb and Mosher [FM02] introduced the notion of *convex cocompact* subgroups of $\text{Mod}(S)$, which they defined as the finitely generated subgroups with quasiconvex orbit in Teichmüller space. Such subgroups are *purely pseudo-Anosov*, meaning all infinite-order elements are pseudo-Anosov. Subsequent work has led to alternate characterizations: Kent and Leininger, and, independently, Hamenstädt, proved the equivalence of convex cocompactness and (1) below, while Durham and Taylor proved the equivalence of convex cocompactness and (2).

Theorem 1.3 ([DT14, Ham05, KL08]). *A finitely generated $H < \text{Mod}(S)$ is convex cocompact if and only if either of the following equivalent conditions hold:*

- (1) *Some (any) orbit map from H into $\mathcal{C}(S)$ is a quasi-isometric embedding.*
- (2) *H is stable in $\text{Mod}(S)$.*

Postponing to Section 1.4 our comments on the apparent similarities between Theorems 1.1 and 1.3, let us first remark that Theorem 1.3 does not include the element-wise condition known to be necessary for convex cocompactness: that the group H is purely pseudo-Anosov. In fact, this is an open question that has received much recent attention.

Question 1.4 (Farb–Mosher). *Are finitely generated, purely pseudo-Anosov subgroups necessarily convex cocompact?*

So far, this question has been answered only in special cases. It is easily seen to be true for subgroups of Veech groups, which preserve a hyperbolic disk isometrically embedded in Teichmüller space [KL07]. A more significant case is resolved by [DKL12] (generalizing [KLS09])

who answer affirmatively for subgroups of certain hyperbolic 3-manifold groups embedded in $\text{Mod}(S)$.

Theorem 1.1 completes the affirmative answer for a third case, the family of mapping class subgroups studied by the second and third authors in [MT13]. That paper considers what the authors call *admissible embeddings* of $A(\Gamma)$ in $\text{Mod}(S)$, whose abundant existence is established in [CLM12]. The central result of [MT13] is that, if $A(\Gamma) < \text{Mod}(S)$ is admissible, then finitely generated $H < A(\Gamma) < \text{Mod}(S)$ is convex cocompact if and only if it is purely pseudo-Anosov and quasiconvex in $A(\Gamma)$ using the standard word metric. Since for any embedding of $A(\Gamma)$ into $\text{Mod}(S)$, pseudo-Anosov elements in the target are the images of loxodromic elements of $A(\Gamma)$, Theorem 1.1 renders the quasiconvex condition redundant:

Corollary 1.5. *If $A(\Gamma) < \text{Mod}(S)$ is an admissible embedding, then any finitely generated $H < A(\Gamma) < \text{Mod}(S)$ is convex cocompact if and only if it is purely pseudo-Anosov.*

1.4. Motivation from right-angled Artin group geometry. From another point of view, Theorem 1.1 can be read as an affirmative answer for the right-angled Artin group analogue of Question 1.4. More precisely, Question 1.4 asks whether for a finitely generated subgroup $H \leq A(\Gamma)$, the element-wise condition of being pseudo-Anosov is strong enough to guarantee the global property of being convex cocompact. Theorem 1.1 gives a positive answer to the corresponding question within a right-angled Artin group. The analogy arises from parallel conditions (1) and (2) in Theorems 1.1 and 1.3, along with the corresponding characterizations of loxodromic elements in $A(\Gamma)$ and pseudo-Anosov elements in $\text{Mod}(S)$. We refer to [KK14a] for a detailed account of the analogy between Γ^e and $\mathcal{C}(S)$.

The proof of Theorem 1.1 establishes the equivalence of (1)-(3) via a fourth condition, that H is *join-busting*. Let Γ be a connected and anti-connected graph, and let $H < A(\Gamma)$ be a finitely generated purely loxodromic subgroup. For a natural number N , we say that H is *N -join-busting* if whenever $w \in H$ is a reduced word in $A(\Gamma)$ and $\beta \leq w$ is a join subword, then $\ell_{A(\Gamma)}(\beta) \leq N$. (Join subwords are defined in Section 2.2; they are the portions of the word w that live in flats of $A(\Gamma)$.) We say H is *join-busting* if it is N -join-busting for some N . Intuitively, geodesics between elements in a join-busting subgroup can spend only bounded bursts of time in embedded $\mathbb{Z} \oplus \mathbb{Z}$ flats.

1.5. Organization of paper. Section 2 fixes notation and background about right-angled Artin groups, including in Section 2.4 a summary of

results about extension graphs and recipes for producing many purely loxodromic subgroups.

An important tool for the proof of Theorem 1.1 are what we call disk diagrams, which are dual to van Kampen diagrams and relate to the cubical structure of $A(\Gamma)$. These are defined in Section 3 and yield useful lemmas collected in Section 4.

In Section 5, we prove Theorem 5.2, that finitely generated purely loxodromic subgroups are join-busting. In Section 6, we prove that join-busting subgroups fulfill Statement (2) of Theorem 1.1 (Corollary 6.2), and we give more details on that condition, called *stability*, which is a kind of strong quasiconvexity.

We prove Theorem 1.1 in Section 7, and Theorem 1.2 in Section 8.

1.6. Acknowledgments. The first named author is partially supported by NSF grant DMS-1203964. The second named author is partially supported by NSF DMS-1204592. The third named author is partially supported by NSF grant DMS-1400498

2. BACKGROUND

2.1. Coarse geometry. For metric spaces (X, d_X) and (Y, d_Y) and constants $K \geq 1$ and $L \geq 0$, a map $f : X \rightarrow Y$ is a (K, L) -*quasi-isometric embedding* if for all $x_1, x_2 \in X$,

$$\frac{1}{K}d_X(x_1, x_2) - L \leq d_Y(f(x_1), f(x_2)) \leq Kd_X(x_1, x_2) + L.$$

A *quasi-isometric embedding* is simply a (K, L) -quasi-isometric embedding for some K, L . If the inclusion map of a subgroup H into the finitely generated group G is a quasi-isometric embedding, we say H is *undistorted* in G (this is independent of the word metric on G).

When a quasi-isometric embedding $f : X \rightarrow Y$ has the additional property that every point in Y is within a bounded distance from the image $f(X)$, we say f is a *quasi-isometry* and X and Y are *quasi-isometric*.

Where I is a subinterval of \mathbb{R} or \mathbb{Z} , we call a (K, L) -quasi-isometric embedding $f : I \rightarrow Y$ a (K, L) -*quasi-geodesic*. If $K = 1$ and $L = 0$, then $f : I \rightarrow Y$ is a *geodesic*. When we refer to a K -quasigeodesic, we mean a (K, K) -quasigeodesic as we have defined it here.

A subset $C \subset X$ is *K -quasiconvex* if for any $x, y \in C$ and any geodesic $[x, y]$ between x, y in X ,

$$[x, y] \subset N_K(C),$$

where $N_K(\cdot)$ denotes the K -neighborhood. We say C is *quasiconvex* if it is K -quasiconvex for some K . When we speak of a quasiconvex subgroup H of a group G , we have fixed a word metric on G with respect to some finite generating set (changing generating sets can change which subgroups are quasiconvex).

A quasi-geodesic γ in a metric space X is called *Morse* or *stable* if every (K, L) -quasi-geodesic with endpoints on γ is within a bounded distance of γ , depending only on K and L . For a finitely generated group G and $g \in G$, g is called a *Morse element* of G if $\langle g \rangle$ is a Morse quasigeodesic in G with respect to some (any) generating set for G . A generalization of stability to subgroups $H \leq G$ is recalled in Section 6.

2.2. RAAGs, graphs, joins, and stars. Fix a finite, simplicial graph Γ with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$, where edges are unordered pairs of distinct vertices. The *right-angled Artin group* with defining graph Γ is the group presented by

$$\langle v \in V(\Gamma) : [v, w] = 1 \text{ if and only if } \{v, w\} \in E(\Gamma) \rangle.$$

A graph Γ is a *join* if there are nonempty subgraphs $J_1, J_2 \subset \Gamma$ such that $\Gamma = J_1 * J_2$. Here, the star notation means that every vertex of J_1 is adjacent to every vertex of J_2 .

We call the generators $v \in V(\Gamma)$ the *vertex generators* of $A(\Gamma)$ and note that whenever we say that an element $g \in A(\Gamma)$ is represented by a word w , we always mean a word in the vertex generators. Call a word w representing $g \in A(\Gamma)$ *reduced* if it has minimal length among all words representing g . Throughout, we use equality “=” on words to denote equality in the group $A(\Gamma)$, i.e. $w = w'$ if and only if w and w' represent the same element in $A(\Gamma)$. If $g \in A(\Gamma)$ and w is a reduced word representing g , we say that $v \in V(\Gamma)$ is in the *support* of w , written $v \in \text{supp}(w)$, if either v or v^{-1} occurs as a letter in w . It is well-known that the support of $g \in A(\Gamma)$ is well-defined, independent of reduced representative.

We say a simplicial graph is *connected* if it is connected as a simplicial complex. The *opposite graph* of a simplicial graph Γ is the graph Γ^{opp} , whose vertices are the same as those of Γ , and whose edges are $\{\{v, w\} \mid \{v, w\} \notin E(\Gamma)\}$. A simplicial graph is *anti-connected* if its opposite graph is connected. Note that a graph is anti-connected if and only if it does not decompose as a nontrivial join.

Of particular importance are the subgroups of $A(\Gamma)$ that nontrivially decompose as direct products. Say that an induced subgraph J of Γ is a *join* (or *subjoin*) of Γ if J itself is a join subgraph. We call the induced subgroup $A(J) \leq A(\Gamma)$ (or one of its conjugates) a *join subgroup*. An

element $g \in A(\Gamma)$ is said to be a *join word* if g is conjugate into a join subgroup of $A(\Gamma)$. It is a theorem of Servatius [Ser89] that $g \in A(\Gamma)$ lives in a $\mathbb{Z} \oplus \mathbb{Z}$ subgroup of $A(\Gamma)$ if and only if g is a join subword of $A(\Gamma)$. (This is part of Theorem 2.2 below).

Let $v \in V(\Gamma)$ be a vertex. The *star* of v , written $\text{St}(v)$, is the vertex v together with the *link* of v , written $\text{Lk}(v)$, which is the set of vertices adjacent to v . A *star subgroup* is a subgroup of a right-angled Artin group generated by the vertices in a star. Note that a star of a vertex is always a join, but the converse is generally not true. A reduced word $w \in A(\Gamma)$ is a *star word* if it is conjugate into a star subgroup of $A(\Gamma)$. Note again that a star word is always a join word, but the converse is generally not true.

2.3. The Salvetti complex and its hyperplanes. In the study of the geometry of a right-angled Artin group, it is often useful to consider a canonical $K(G, 1)$ for a right-angled Artin group $A(\Gamma)$, called the *Salvetti complex* of $A(\Gamma)$, which we denote by $S(\Gamma)$. For the convenience of the reader, we recall a construction of $S(\Gamma)$.

Let $n = |V(\Gamma)|$. We consider the 1-skeleton of the unit cube $I^n \subset \mathbb{R}^n$ with the usual cell structure, and we label the 1-cells of I^n which contain the origin by the vertices of Γ . If K is a complete subgraph of Γ , we attach the corresponding sub-cube of I^n labeled by the relevant 1-cells. Doing this for every complete subgraph of Γ , we obtain a subcomplex $X(\Gamma) \subset I^n$. We write $S(\Gamma) = X(\Gamma)/\mathbb{Z}^n$, where we have taken the quotient by the usual \mathbb{Z}^n -action on \mathbb{R}^n .

It is easy to check that $S(\Gamma)$ is a compact cell complex and that $\pi_1(S(\Gamma)) \cong A(\Gamma)$. By the Milnor–Schwarz Lemma, we have that $A(\Gamma)$ is quasi-isometric to $\tilde{S}(\Gamma)$, the universal cover of $S(\Gamma)$. By construction, the labeled 1-skeleton of $\tilde{S}(\Gamma)$ is also the Cayley graph of $A(\Gamma)$ with respect to the generating set $V(\Gamma)$. Consequently, reduced words in $A(\Gamma)$ correspond to geodesics in $\tilde{S}(\Gamma)^{(1)}$, which we call *combinatorial geodesics*. We refer to distance in $\tilde{S}(\Gamma)^{(1)}$ as *combinatorial distance*. By *combinatorial quasiconvexity*, we mean quasiconvexity in $\tilde{S}(\Gamma)^{(1)}$.

It is well-known that $\tilde{S}(\Gamma)$ is a CAT(0) cube complex. This structure will be important to our proofs, so we record some relevant definitions here. In $S(\Gamma)$, each edge e_v (labeled by the generator v) is dual to a *hyperplane* H_v , determined as follows: let m be the midpoint of e_v . For every sub-cube $e_v \times I^k$ in $X(\Gamma)$, the *midcube dual to e_v* is the set $m \times I^k$; we define H_v as the image in $S(\Gamma)$ of all midcubes dual to e_v . Observe that, by construction, H_v and H_u intersect if and only if v and u commute. The lifts of hyperplanes in $S(\Gamma)$ define similarly labeled

hyperplanes dual to edges in $\tilde{S}(\Gamma)$. Each of these latter hyperplanes separates $\tilde{S}(\Gamma)$ into two convex sets. It follows that the combinatorial distance between a pair of vertices equals the number of hyperplanes in $\tilde{S}(\Gamma)$ separating those vertices.

2.4. Extension graphs and a profusion of purely loxodromic subgroups. In this section, we record some facts about the extension graph Γ^e , and we establish the existence of many purely loxodromic subgroups of right-angled Artin groups.

Let Γ be a finite simplicial graph. The *extension graph* Γ^e of Γ is the graph whose vertices are given by $\{v^g \mid v \in V(\Gamma), g \in A(\Gamma)\}$, and whose edges are given by pairs $\{\{v^g, w^h\} \mid [v^g, w^h] = 1\}$. Here, $v^g = g^{-1}vg$.

Note that the conjugation action of $A(\Gamma)$ on itself induces an action of $A(\Gamma)$ on Γ^e by isometric graph automorphisms.

Theorem 2.1 (See [KK13], [KK14a]). *Let Γ be a finite, connected, and anti-connected simplicial graph with at least two vertices.*

- (1) *The graph Γ^e is connected.*
- (2) *The graph Γ^e has infinite diameter.*
- (3) *The graph Γ^e is locally infinite.*
- (4) *The graph Γ^e is quasi-isometric to a regular tree of infinite valence. In particular, Γ^e is δ -hyperbolic.*
- (5) *The action of $A(\Gamma)$ on Γ^e is acylindrical. In particular, every nonidentity element of $A(\Gamma)$ is elliptic or loxodromic, depending on whether its orbits in Γ^e are bounded or unbounded.*
- (6) *A nontrivial cyclically reduced element of $A(\Gamma)$ is elliptic if and only if its support is contained in a subjoin of Γ . In particular, $A(\Gamma)$ contains a loxodromic element.*
- (7) *Any purely loxodromic subgroup of $A(\Gamma)$ is free.*

Hence, using Theorem 2.1, we refer to an element of $A(\Gamma)$ as *loxodromic* if it acts as a loxodromic isometry of Γ^e . Combining results of [Ser89, BC12, KK14a] we state the following characterization of loxodromic elements of $A(\Gamma)$ for easy reference to the reader.

Theorem 2.2 (Characterization of loxodromics). *Let $g \in A(\Gamma)$. The following are equivalent:*

- (1) *g is loxodromic.*
- (2) *g acts as a rank 1 isometry on $\tilde{S}(\Gamma)$.*
- (3) *The centralizer $C_{A(\Gamma)}(g)$ is cyclic.*
- (4) *g is not conjugate into a join subgroup, i.e. g is not a join subword.*

A single loxodromic element of $A(\Gamma)$ generates an infinite, purely loxodromic subgroup of $A(\Gamma)$. Items (4) and (5) of Theorem 2.1 allow us to produce many non-cyclic purely loxodromic subgroups:

Proposition 2.3. *Let $\lambda \in A(\Gamma)$ be a loxodromic element. There is an N which depends only on Γ such that the normal closure Λ_N of λ^N in $A(\Gamma)$ is purely loxodromic.*

Proof. We have that Γ^e is a δ -hyperbolic graph on which $A(\Gamma)$ acts acylindrically. By [DGO11], Theorem 2.29, we have $N \gg 0$ for which the normal closure Λ_N of λ^N is purely loxodromic and free, and this constant N depends only on the hyperbolicity and acylindricity constants of the $A(\Gamma)$ action on Γ^e . \square

We remark that, if a finite set of loxodromic elements of $A(\Gamma)$ is pairwise independent (i.e., no two elements have conjugate powers), then a standard argument in hyperbolic geometry proves that sufficiently high powers of these elements generate a purely loxodromic group. Moreover, by following the arguments in [DGO11], one can verify a generalization of Proposition 2.3: there exist powers of these elements whose normal closure (as a set) is purely loxodromic.

Proposition 2.3 shows that there are enough loxodromic elements in $A(\Gamma)$ for them to be present in every nontrivial normal subgroup:

Theorem 2.4. *Let Γ be an anti-connected graph with at least two vertices, and let $1 \neq G < A(\Gamma)$ be a normal subgroup. Then G contains a loxodromic element of $A(\Gamma)$.*

It is very easy to produce counterexamples to Theorem 2.4 if one omits the hypothesis that Γ be anti-connected.

Proof of Theorem 2.4. Let $1 \neq g \in G$ be any element, and let $\lambda \in A(\Gamma)$ be loxodromic. Let N be the constant furnished by Proposition 2.3. Consider the element $[g, \lambda^N] \in G \cap \Lambda_N$. If this commutator is the identity then g commutes with λ^N , so that these two elements have a common power. In particular, g is loxodromic. If the commutator is not the identity then it represents a nontrivial loxodromic element of G , since Λ_N is purely loxodromic. \square

We remark that Behrstock, Hagen, and Sisto [BHS] have very recently introduced a different curve complex analogue, called the contact graph, associated to RAAGs and more generally $\text{CAT}(0)$ cubical groups. Such a group's action on its contact graph may be used to define "loxodromic" elements in the group. In the RAAG case, the contact graph $\mathcal{CS}(\Gamma)$ is equivariantly quasi-isometric to the extension

graph Γ^e , except for in a few sporadic cases. Therefore both definitions of loxodromic elements coincide, and one may add to Theorem 1.1 a statement (1'), replacing Γ^e with $\mathcal{CS}(\Gamma)$.

Finally, we note forthcoming work by the third author and Tiozzo [TT], which proves that typical subgroups of appropriate $A(\Gamma)$ are purely loxodromic. More precisely, let μ be a probability measure on $A(\Gamma)$ whose support generates $A(\Gamma)$ as a semigroup, and consider k independent random walks on $A(\Gamma)$ whose increments are distributed according to μ . They prove that the subgroup generated by the n th step of each random walk is purely loxodromic with probability going to 1 as n goes to infinity.

3. DISK DIAGRAMS IN RAAGS

Our main technical tool are the disk diagrams introduced in [CW04] (as “dissections of a disk”) and further detailed in [Kim08]. We present the general theory here for the convenience of the reader.

3.1. Construction of a disk diagram. Suppose that w is a word in the vertex generators of $A(\Gamma)$ and their inverses which represents the identity element in $A(\Gamma)$. A disk diagram for w is a combinatorial tool indicating how to reduce w to the empty word. Before giving a formal definition, we show how to construct such a diagram.




FIGURE 1. Disk diagram for the identity word $abb̄aāb̄d̄cd̄abab̄āb̄d̄d̄āc̄ā$, where \bar{x} is the inverse of x .

Consider the loop in $S(\Gamma)$ determined by the word w , which we continue to denote by w . This is the map $w : S^1 \rightarrow S(\Gamma)$ that, starting

at a basepoint in S^1 , parameterizes an edge path in $S(\Gamma)$ spelling the word w . (Throughout this paper, an edge path in $S(\Gamma)$, or its universal cover, *spells* a word w if w is the word obtained by recording the labels on the oriented edges traversed by the path according to its orientation.) Since w represents the trivial word in $A(\Gamma)$, this map is null homotopic in $S(\Gamma)$. Denote by $f : D \rightarrow S(\Gamma)$ an extension of the loop w to the unit disk, and homotope $f \text{ rel } \partial D$ so that it is transverse to the hyperplanes of $S(\Gamma)$. After doing so, all preimages of hyperplanes through f consist of simple closed curves and properly embedded arcs. Using standard techniques in combinatorial topology, perform a further homotopy $\text{rel } \partial D$ to remove all simple closed curve components in the preimage and to ensure that no properly embedded arcs intersect more than once. After performing such a homotopy, the preimage of each hyperplane in $S(\Gamma)$ is a collection of properly embedded disjoint arcs.

For the hyperplane H_v in $S(\Gamma)$ dual to the edge e_v representing v , there are two co-orientations corresponding to the two orientations of the edge e_v . We label each of these two co-orientations on H_v either v or v^{-1} corresponding to the label on the oriented edge. By pulling back the co-orientation of H_v , we likewise label the co-orientation of each arc in the preimage of H_v . This is to say that an arc in the preimage of H_v is labeled with both the symbol v and a transverse arrow. Crossing the arc in the direction indicated by the arrow corresponds to reading the label v and crossing in the reverse direction reads the label v^{-1} . The oriented disk D together with this collection of labeled, co-oriented properly embedded arcs is called a disk diagram for w and is denoted by Δ . See Figure 1.

3.2. Formal definition. Our disk diagrams follow the definition in [Kim08]. For a word w representing the trivial element in $A(\Gamma)$, a *disk diagram* (or dual van Kampen diagram) Δ for w in $A(\Gamma)$ is an oriented disk D together with a collection \mathcal{A} of co-oriented, properly embedded arcs in general position, satisfying the following:

- (1) Each co-oriented arc of \mathcal{A} is labeled by an element of $V(\Gamma)$ or its inverse such that arcs with opposite orientations have inverse labels. Moreover, if two arcs of \mathcal{A} intersect then the generators corresponding to their labels are adjacent in Γ .
- (2) With its induced orientation, ∂D represents a cyclic conjugate of the word w . In other words, there is a point $*$ $\in \partial D$ such that traversing ∂D according to its orientation starting at $*$ and recording the labels of the arcs of \mathcal{A} so encountered spells the word w .

We think of the boundary of D as subdivided into edges and labeled according to the word w , such that the label on an oriented edge in this subdivision agrees with the label on the co-oriented arc of \mathcal{A} meeting this edge. In this way, each arc of \mathcal{A} corresponds to two letters of w which are represented by oriented edges on the boundary of Δ .

While not required by the definition, it is convenient to restrict our attention to *tight* disk diagrams, in which arcs of \mathcal{A} intersect at most once.

3.3. Interchangeable interpretations. In Section 3.1, we saw that any word representing the identity has a disk diagram induced by a null-homotopy $f : D \rightarrow S(\Gamma)$. Here we show how any disk diagram can be induced by such a map. This uses the duality, observed in [CW04] and detailed in [Kim08], between disk diagrams and the more familiar notion of van Kampen diagrams. For a general discussion of the latter, see [Sap14] or [LS01]. We recall the relevant details here.

A *van Kampen diagram* $X \subset S^2$ for a word w representing the identity in $A(\Gamma)$ is a simply connected planar 2-complex equipped with a combinatorial map into $S(\Gamma)$ whose boundary word is w . Such a map may be expressed by labeling the oriented 1-cells of the diagram X with vertex generators and their inverses in such a way that every 2-cell represents a (cyclic conjugate of) a relator of $A(\Gamma)$, i.e. a 2-cell of $S(\Gamma)$. The boundary word of X is the cyclic word obtained by reading the labels of ∂X , and X represents a null-homotopy of its boundary.




FIGURE 2. On left, a van Kampen diagram for the same identity word shown in Figure 1. On right, the dual disk diagram; its arcs continue to a single “vertex at infinity,” not shown (compare to Figure 1 in [Kim08]).

While X itself is not necessarily a disk, it has a small neighborhood in S^2 which deformation retracts onto X , and which is a disk D . The

composition of the retract from D to X and the map of X into $S(\Gamma)$ gives a map $f : D \rightarrow S(\Gamma)$ such that pre-images of hyperplanes are properly embedded arcs in D which induce a disk diagram for w .

It is shown in [Kim08] that any disk diagram for w is dual in S^2 to a van Kampen diagram with boundary word w , and vice versa. In particular, given a disk diagram, one obtains a van Kampen diagram by simply taking its dual complex; see Figure 2.

3.4. Surgery and subwords. Starting with a disk diagram Δ for w , one can often extract useful information about a subword of w by surgering Δ along properly embedded arcs. In details, suppose that γ is a properly embedded, oriented arc in Δ which is either an arc of \mathcal{A} or transverse to the arcs of \mathcal{A} . Traversing γ according to its orientation and recording the labels of those arcs of \mathcal{A} so crossed according to their co-orientation spells a word y in the standard generators. We say the word y is obtained from *label reading* along γ .

In particular, starting with a subword w' of w , any oriented arc γ of D which begins at the initial vertex of w' and ends at the terminal vertex of w' produces a word y via label reading such that $w = y$ in $A(\Gamma)$. To see this, note that the arc γ cuts the disk D into two disks D' and D'' , one of which (say D') determines the homotopy (and sequence of moves) to transform the word w' into y . This is to say that the disk D' along with arcs coming from Δ is a disk diagram for the word $w'\bar{y}$, and we say that this diagram is obtained via *surgery* on Δ . This simple observation will be important to our arguments.

Here let us record a few lemmas enabled by surgery. The first appears as Lemma 2.2 in [Kim08].

Lemma 3.1. *Suppose an arc of \mathcal{A} in a disk diagram Δ for the identity word w cuts off the subword w' , i.e., $w = svw'v^{-1}t$, where s, w' , and t are subwords and v, v^{-1} are letters at the ends of the arc. Then $w' \in \text{St}(v)$.*

We omit the proof, which is straightforward, and can be found in [Kim08].

If a subword in a disk diagram has the property that no two arcs emanating from it intersect, let us say that this subword is *combed* in the disk diagram. The property of being combed will be important in Sections 4 and 5.

Lemma 3.2. *Suppose w is a word representing the identity and b is a subword of w , so w is the concatenation of words a, b , and c . Let Δ be a disk diagram for w .*

Then there exists a word b' obtained by re-arranging the letters in b , such that $b' = b$ and there exists a disk diagram Δ' for $ab'c$ in which b' is combed, arcs emanating from b' have the same endpoint in the boundary subword $\bar{c}a$ as their counterpart in b , and arcs that both begin and end in $\bar{c}a$ are unchanged in Δ' .


Furthermore, there exists a word b'' obtained by deleting letters in b' , such that $b'' = b$ and there exists a disk diagram Δ'' for $ab''c$ which is precisely Δ' without the arcs corresponding to the deleted letters.

Proof. Start with any tight disk diagram Δ for w and consider the arcs emanating from b . If a pair of arcs from b intersect, they cut out a piece of the disk in which one can find an *innermost* intersection between arcs from b , meaning arcs emanating from letters x and z so that b decomposes as $sxyz$, no arc from a letter in the subword y intersects any other arc, and arcs from x and z intersect at a point p . This means that all arcs that start in y end in y . Consequently, one can divide the disk with an arc γ which is disjoint from \mathcal{A} and which divides the boundary into two components, one of which corresponds to y . Surgery using γ shows that y equals the identity. Then $sxyz = szyx$. Let $b_1 = szyx$ and let $w_1 = ab_1c$.

We now describe a tight disk diagram Δ_1 for w_1 in which the number of intersections among arcs emanating from b_1 is one less than the number of intersections among arcs emanating from b in the original disk diagram. In Δ , divide the arcs from x and z into proximal and distal segments, corresponding respectively to the part of the arc between b and the intersection point p and the part of the arc between p and its other endpoint on the boundary. Because Δ is tight, an arc intersects the proximal part of the x arc if and only if it intersects the proximal part of the z arc.

Let the arc z' consist of the proximal part of x and the distal part of z , and let x' consist of the proximal part of z and the distal part of x . It is evident that x' and z' may be perturbed in a neighborhood of p , to be disjoint. Then Δ_1 is obtained by switching boundary letters x and z , and replacing their corresponding arcs with the arcs x' and z' . See Figure 3. By repeating this process as many times as the original number of intersections between arcs from b in Δ , we obtain a disk diagram Δ_n for w_n in which the subword b_n is combed, and b_n is equal to b and obtained by rearranging its letters.

To obtain b'' and Δ'' , delete arcs with both ends in b' , by again finding innermost such arcs, which come from subwords of the form vv^{-1} . \square

FIGURE 3. On the left, part of Δ ; on the right, part of Δ_1 .

3.5. Reducing diagrams. In our applications, disk diagrams will come with some additional structure. Let h be a word in the vertex generators of $A(\Gamma)$, which is not assumed to be reduced in any sense. Let w denote a reduced word in the vertex generators which represents the same group element, i.e. $w = h$ in $A(\Gamma)$. Then, the word hw represents the identity in $A(\Gamma)$ and so it is the boundary of some disk diagram Δ . (Here \bar{w} denotes the inverse of the word w .) In this situation, we record that the boundary of Δ consists of two words h and \bar{w} and we draw the diagram as in Figure 4. We sometimes refer to a disk diagram constructed in this way as a *reducing diagram* as it represents a particular way of reducing h to the reduced word w . For such disk diagrams, ∂D is divided into two subarcs (each a union of edges) corresponding to the words h and w ; we call these subarcs the h and w subarcs, respectively.




FIGURE 4. Segmented orientation in a reducing diagram.

Suppose that Δ is a disk diagram that reduces h to the reduced word w . Since w is already a reduced word, no arc of \mathcal{A} can have both its endpoints on the w subarc of ∂D . Otherwise, one could surger the diagram to produce a word equivalent to w with fewer letters. Hence, each arc of \mathcal{A} either has both its endpoints on the subarc of ∂D corresponding to h , or it has one endpoint in each subarc of ∂D . In the former case, we call the arc (and the letters of h corresponding to its endpoints) *noncontributing* since these letters do not contribute

to the reduce word w . Otherwise, the arc is called *contributing* (as is the letter of h corresponding the endpoint contained in the h subarc of ∂D). If h as a word is partitioned into a product of subwords abc , then the *contribution of the subword b to w* is the set of letters in b which are contributing. We remark that whether a letter of h is contributing or not is a property of the fixed disk diagram that reduces h to w .

For Section 8, we loosen the definition of a reducing diagram so that the word w is only partially reduced. That is, rather than requiring w be minimal length in $A(\Gamma)$, we simply require that no arc in Δ has both endpoints in w . Then it is clear that w still has length in $A(\Gamma)$ no longer than that of h .

4. DISK DIAGRAM APPLICATIONS

In this section we consider elements of H a finitely generated, star-free subgroup of $A(\Gamma)$; such subgroups include but are more general than purely loxodromic subgroups. In Section 8 we include a short proof that star-free subgroups are necessarily free, but we do not require that knowledge here. We do note that, by finite generation, the inclusion of H into $A(\Gamma)$ is proper with respect to any word metric on both groups.

Here we employ disk diagrams, described in Section 3, to derive several useful lemmas about finitely generated, star-free subgroups. Fix a finite generating set S for H and a reduced word representative for each generator in S . An element $h \in H$ can be expressed as a geodesic word in H , that is, $h = h_1 \cdots h_n$ such that $h_i \in S$ and n is minimal. We use a disk diagram with boundary word $(h_1 \cdots h_n)h^{-1}$, where h and each h_i are written as reduced words in $A(\Gamma)$. In other words, we concatenate the reduced word representatives for the h_i to obtain a word representing $h = h_1 \cdots h_n$ and consider a reducing diagram for this word. With our choices fixed, we call such a reducing diagram for h simply a disk diagram for $h \in H$. Recall that every letter in the disk diagram is connected by an arc to its inverse somewhere else in the disk diagram, and if arcs intersect, the corresponding generators commute.

Lemma 4.1 (Cancellation diameter). *Suppose H is a finitely generated, star-free subgroup of $A(\Gamma)$. There exists $D = D(H)$ with the following property: If in a disk diagram for $h \in H$, a letter in h_i is connected to a letter in h_j , then $j - i < D$.*

Proof. Suppose in a disk diagram for $h \in H$, a letter g in h_i is connected to a letter g^{-1} in h_j . By Lemma 3.1, $h_i \cdots h_j = \sigma M \tau$, where M is in the star of g , and σ, τ are a prefix of h_i and suffix of h_j respectively.

Therefore, if the lemma is false, there a sequence of reduced-in- H words

$$h_{i(t)}^t \cdots h_{j(t)}^t = \sigma_t M_t \tau_t$$

as above, with $j(t) - i(t)$ increasing. Because Γ is finite and H is finitely generated, we may pass to a subsequence so that the M_t are in the star of the same generator, and furthermore we have constant $\sigma_t = \sigma$ and $\tau_t = \tau$, while $M_t \neq M_s$ for $s \neq t$. Then for any $s \neq t$,

$$h_{i(s)}^s \cdots h_{j(s)}^s (h_{i(t)}^t \cdots h_{j(t)}^t)^{-1} = \sigma M_s M_t^{-1} \sigma^{-1}$$

which is conjugate into a star, a contradiction. \square

Consider $h_i \cdots h_j$ a subword of $h = h_1 \cdots h_n$ reduced in H as above. For D the constant of Lemma 4.1, we have

Lemma 4.2. *The subword $h_i \cdots h_j$ may be written as $\sigma W \tau$, where the letters in σ come from $h_{i-D} \cdots h_{i-1}$ when $i > D$, and from $h_1 \cdots h_{i-1}$ otherwise; the letters in τ come from $h_{j+1} \cdots h_{j+D}$ when $j \leq n - D$ and from $h_{j+1} \cdots h_n$ otherwise; and the letters in W correspond to the contribution of $h_i \cdots h_j$ to the reduced word for h in $A(\Gamma)$.*

Proof. Consider a disk diagram for $(h_1 \cdots h_n)h^{-1}$, and let b be the subword $h_i \cdots h_j$. Lemma 3.2 provides a word $b'' = b$ and a disk diagram for $(h_1 \cdots h_{i-1} b'' h_{j+1} \cdots h_n)h^{-1}$ in which all arcs emanating from b'' terminate in $h_{j+1} \cdots h_n h^{-1} h_1 \cdots h_{i-1}$ and no two arcs from b'' intersect. Then it is clear that b'' partitions into subwords σ, W , and τ , whose arcs connect to $h_1 \cdots h_{i-1}$, h , and $h_{j+1} \cdots h_n$ respectively. Recall that the disk diagram furnished by Lemma 3.2 differs from the original only by removing or shuffling endpoints of arcs in b . Lemma 4.1 implies that the arcs that do not contribute to h end within the subwords claimed. \square

Lemma 4.3 (Bounded non-contribution). *Given H a finitely generated, star-free subgroup of $A(\Gamma)$, there exists $K = K(H)$ such that, if $h_i \cdots h_j$ is a subword of a reduced word for h in H which contributes nothing to the reduced word for h in $A(\Gamma)$, then $j - i < K$.*

Proof. Consider a vanishing subword, that is, a subword that contributes nothing to the reduced word for h in $A(\Gamma)$. By Lemma 4.1, this subword $h_i \cdots h_j$ reduces in $A(\Gamma)$ to LR , where the letters in L come from $h_i \cdots h_{i+D-1}$ and the letters in R come from $h_{j-D+1} \cdots h_j$. Therefore $\ell_{A(\Gamma)}(LR) \leq 2D \cdot \max\{\ell_{A(\Gamma)}(h_i)\}$, so there are only finitely many elements in $A(\Gamma)$ that might equal LR . Since the inclusion $H \rightarrow A(\Gamma)$ is proper, there is a $K \geq 0$ so that no word in H of length greater than K reduces to a work in $A(\Gamma)$ of length less than $2D \cdot \max\{\ell_{A(\Gamma)}(h_i)\}$. Hence, $j - i = |h_i \cdots h_j|_H \leq K$. \square

Proposition 4.4 (Undistorted). *Finitely generated star-free subgroups are quasi-isometrically embedded.*

Proof. For $h \in H < A(\Gamma)$ fulfilling the hypotheses, we have

$$\ell_H(h)/(K + 1) \leq \ell_{A(\Gamma)}(h) \leq \ell_H(h) \cdot \max\{\ell_{A(\Gamma)}(h_i)\}$$

where K is from Lemma 4.3. □

5. JOIN-BUSTING IN PURELY LOXODROMIC SUBGROUPS

Recall that H is N -join-busting if, for any reduced word w representing $h \in H$, and any join subword $\beta \leq w$, we have $\ell_{A(\Gamma)}(\beta) \leq N$. In this section we prove that purely loxodromic subgroups of $A(\Gamma)$ are N -join-busting. Before proving this result (Theorem 5.2 below), let us introduce some terminology. A word w in a join subgroup $A(J) < A(\Gamma)$ is said to be M -homogeneous with respect to J if every subword w' of w with length at least M contains every letter appearing in J . Note that w is always M -homogeneous for $M = \ell_{A(\Gamma)}(w)$.

Lemma 5.1. *Let H be a finitely generated subgroup of $A(\Gamma)$. Suppose $h_k \in H$ is a sequence of words with reduced decomposition $a_k w_k b_k$, where w_k are join subwords and $\ell_{A(\Gamma)}(w_k) \rightarrow \infty$. Then for some $M \geq 1$ there exists such a sequence in H so that the w_k are M -homogeneous with respect to the same join.*

Proof. Since there are only finitely many joins in the graph Γ , after passing to a subsequence, we may assume that the w_k are supported on a common join. If there is no M so that the sequence w_k is eventually M -homogeneous, then we can find arbitrarily large subwords of w_k supported on a proper subjoin. After again passing to a subsequence, we may assume that these subjoins are constant. In this manner, we obtain a new sequence of words $h_k^2 = a_k^2 w_k^2 b_k^2$ such that $|w_k^2| \rightarrow \infty$ and the letters of w_k^2 are contained in a common proper subjoin of the support of w_k . Now either w_k^2 are eventually M -homogeneous, for some $M \geq 1$, or we may iterate the process. Since there are only finitely many joins, each of finite size, this process terminates in a sequence of words in H which contribute M -homogeneous subwords of unbounded size and supported on the same join in Γ . □

Theorem 5.2. *Let $H < A(\Gamma)$ be finitely generated and purely loxodromic. There exists an $N = N(H)$ such that H is N -join-busting.*

Proof. Fix a finite free generating set S for H and reduced word representatives for each generator in S . Suppose the contrary of the theorem, so that we have a sequence $\{h_k\} \subset H$ of elements with reduced decompositions $h_k = a_k w_k b_k$, where the w_k are join subwords of unbounded

length. Fix disk diagrams for $(h_k)^{-1}(a_k w_k b_k)$ so that arcs emanating from w_k do not intersect each other, using Lemma 3.2. Recall that we refer to this condition when we say that the arcs for w_k are *combed* in the disk diagram. By specifying this, we also fix a particular reduced word representing w_k . Note that here our disk diagram reduces h_k , written as a concatenation of our fixed representatives from S , to the reduced word $a_k w_k b_k$.

By Lemma 5.1, we may suppose that the w_k are M -homogeneous for some uniform M , and that these contributions are supported on the same join J . Recall that this assumption implies that every subword of w_k of length at least M uses each letter coming from the subgraph J . Because, in the proof of the lemma, our final w_k are subwords of our original w_k , they retain the property that arcs emanating from them do not intersect in our fixed disk diagrams (they are combed), and so we can think of the w_k as having fixed reduced form representatives. We may also assume that $\ell_{A(\Gamma)}(w_k) \geq kM$, by passing to a subsequence.

Each h_k is represented in its disk diagram by a unique geodesic word $h_{1,k} \cdots h_{s(k),k}$, where the $h_{n,k}$ are from our fixed generating set S . For $k \geq 2$, let $h_{I(k),k}$ be the H -generator connected by an arc to the M th letter of w_k , and $h_{J(k),k}$ to the $(\ell_{A(\Gamma)}(w_k) - M + 1)$ th letter of w_k . Because the arcs from w_k are combed, $I(k) \leq J(k)$. For k sufficiently large, depending on how many times the vertices of J correspond to letters in a free generator for H , $J(k) - I(k) > 1$. In that case we may set $i(k) = I(k) + 1$ and $j(k) = J(k) - 1$ and consider the nonempty subword $h_{i(k),k} \cdots h_{j(k),k}$.

The subword $h'_k = h_{i(k),k} \cdots h_{j(k),k}$ is the largest whose contribution to w_k is M -deep, meaning that the subword contributes to w_k but none of its arcs connect to either the first or last M letters of w_k . Because the arcs of w_k are combed, all letters of w_k between the first and last sets of M do connect to h'_k , and therefore $j(k) - i(k)$ is unbounded. In particular, for sufficiently large k , we may consider the subwords

$$h''_k = h_{i(k)+D,k} \cdots h_{j(k)-D,k}$$

for D the constant from Lemma 4.1.

We may apply Lemma 4.2 and see that h''_k may be written as $\sigma_k W_k \tau_k$, where the letters for σ_k correspond to arcs in h''_k that connect to $h_{i(k)} \cdots h_{i(k)+D-1}$, and the letters for τ_k correspond to arcs in h''_k that connect to $h_{j(k)-D+1} \cdots h_{j(k)}$. The letters in W_k correspond precisely to the arcs in h''_k that connect to $a_k w_k b_k$. Any arc in h''_k that connects to either a_k or b_k must cross the arcs from either the first or the last M letters of w_k ; by our assumption of M -homogeneity, this means that

those arcs correspond to letters that commute with all of J . In particular, every letter in W_k either belongs to J (if it is contributed to w_k) or is joined with all of J (if it is contributed to either a_k or b_k).

Now observe that there are only finitely many possibilities for σ_k and τ_k , so we may pass to a subsequence where these are constant. That is, we have a sequence

$$h_t''' = \sigma W_t \tau$$

of words in H where $\ell_H(h_t''')$ is unbounded and each letter in W_t either belongs to J or is joined with J . In particular we have $h_s''' \neq h_t'''$ in this sequence, and

$$h_t'''(h_s''')^{-1} = (\sigma W_t \tau)(\tau^{-1} W_s^{-1} \sigma^{-1})$$

which is conjugate to $W_t W_s^{-1}$. But $W_t W_s^{-1}$ is a join word, since it is supported on J joined with {letters in W_t, W_s not in J }. This contradicts that H is purely loxodromic. \square

6. STABILITY OF JOIN-BUSTING SUBGROUPS

Let G be a finitely generated group and let $H \leq G$ be a finitely generated subgroup. We say that H is *stable* in G if H is undistorted in G (i.e. $H \rightarrow G$ is a quasi-isometric embedding), and for any $K \geq 1$ there is an $M = M(K) \geq 0$ such that any pair of K -quasigeodesics in G with common endpoints in H have Hausdorff distance no greater than M . This definition was introduced by the third author and Durham in [DT14]. It is clear that stable subgroups of G are quasiconvex with respect any word metric on G . In this section we prove that join-busting subgroups of $A(\Gamma)$ are stable. This will come as a corollary to Theorem 6.1 below.

Say that an element $g \in A(\Gamma)$ is N -join-busting if for any reduced word w representing g and any join subword w' of w , $\ell_{A(\Gamma)}(w') \leq N$. Of course, a subgroup in $A(\Gamma)$ is join-busting if and only if its nontrivial elements are N -join-busting for some uniform N .

Theorem 6.1. *Given $K \geq 1$ and $N \geq 0$, there exists $M = M(K, N)$ satisfying the following. Suppose $g \in A(\Gamma)$ is N -join-busting and let w be a combinatorial geodesic from v to gv . If p is a K -quasigeodesic edge path in $\tilde{S}(\Gamma)$ also joining v to gv , then the Hausdorff distance between w and p is no more than M .*

Proof. As before, we identify the edge paths p and w with the corresponding sequence of oriented edges, and hence the corresponding word representing $g \in A(\Gamma)$. As w is a geodesic path, it corresponds to a reduced word representing g , which we continue to denote by w .

We wish to bound the distance between the paths w and p . To do this, we first fix a disk diagram Δ realizing the reduction of p to the reduced form w . For concreteness, suppose that Δ is obtained from a map $f : D \rightarrow S(\Gamma)$ as in Section 3. Retaining the same notation, replace f with the lift $f : D \rightarrow \tilde{S}(\Gamma)$ such that portion of the boundary of D representing p (or w) is mapped to the paths in $\tilde{S}(\Gamma)$ spelling p (or w). We immediately see that for any two points $x, y \in \partial D$, the combinatorial distance between $f(x)$ and $f(y)$ is no greater than the number of arcs of \mathcal{A} crossed by any path in D from x to y .

Set $B = 3N + 1$. Let b and d be oriented edges in w that are separated by no fewer than $B - 2$ edges of w . We first claim that if β and δ are the arcs of \mathcal{A} with endpoints in b and d , respectively, then, no arc of \mathcal{A} can intersect both β and δ . To see this, let $w' = b \dots d$ be the length $\geq B$ subpath (subword) of w beginning with the edge b and terminating with the edge d . If there is an arc $\epsilon \in \mathcal{A}$ meeting both β and δ then the subdiagram obtained via surgery corresponding to a concatenation of a portion of β , ϵ , and δ shows that $w' \in \text{st}(\beta) \cdot \text{st}(\epsilon) \cdot \text{st}(\delta)$. Here, we have continued to use the symbols β , ϵ , and δ to denote the vertex generators labeling these arcs. Hence, w' can be written as a product of 3 join words and since $|w'| \geq B = 3N + 1$, one of these join words must have length greater than N . (Technically, we have that w' has *some* reduced spelling that is a product of 3 join words and so one of the join words in this latter spelling has length greater than N .) This contradicts the assumption that g is N -join-busting. Hence, no arc of \mathcal{A} crosses both β and δ . In particular, we see that β and δ do not intersect.

This implies that we can decompose the geodesic w into subpaths as follows:

$$w = e_1 \cdot w_1 \cdot e_2 \cdot w_2 \dots e_{n-1} \cdot w_n,$$

where $\ell(w_i) = B - 1$ for $1 \leq i \leq n - 1$, $\ell(w_n) \leq B - 1$, $(n - 1)B < \ell(w) \leq nB$, and each e_i is a single edge with the property that if we denote by α_i the arc of \mathcal{A} with endpoint in e_i , then no arc of \mathcal{A} crosses any two of the α_i . In particular, α_i and α_j are disjoint for $i \neq j$.

Since the $\alpha_i \in \mathcal{A}$ have their other endpoint in edges of the path p , we have an induced decomposition:

$$p = p_0 \cdot e_1 \cdot p_1 \cdot e_2 \cdot p_2 \dots e_{n-1} \cdot p_n.$$

Set $l_i = \ell(\alpha_i)$ to be the number of arcs of \mathcal{A} that α_i intersects (i.e. the combinatorial length of the path α_i), and set $l_0, l_n = 0$. We have the following basic fact:

$$(1) \quad \ell(p_k) \geq l_k + l_{k+1} - \ell(w_k),$$

for all $0 \leq k \leq n$. (Here we have set w_0 to be the empty word.) To see this, note that any arc of \mathcal{A} which crosses α_k must either cross the path p_k or w_k , since the four paths in the above inequality form of subdisk of Δ and no arcs cross both α_k and α_{k+1} . As the same is true for arcs crossing α_{k+1} , this proves the inequality.

Now set $M = 2((K + 1)(B + 1) + K)$ and let e_{i+1}, \dots, e_{j-1} be any maximal consecutive sequence for which $l_{i+1}, \dots, l_{j-1} > M$. For future reference, note that any e_k with $l_k > M$ lives in such a sequence. If no such e_i exist then it is easy to see that the paths p and w have Hausdorff distance bounded in terms of B, M , and K , and we are done.

For this maximal sequence, consider the subpath of w

$$w' = e_i \cdot w_i \cdot e_{i+1} \cdot w_{i+1} \dots e_{j-1} \cdot w_{j-1} \cdot e_j$$

and note that $l_i, l_j \leq M$ by maximality. Further, $(j - i - 1)B \leq \ell(w') \leq (j - i)(B + 1)$. The subpath w' induces the following subpath of p

$$p' = e_i \cdot p_i \cdot e_{i+1} \cdot p_{i+1} \dots e_{j-1} \cdot p_{j-1} \cdot e_j,$$

and using the fact that p is a K -quasigeodesic and $l_i, l_j \leq M$, we see that

$$\begin{aligned} \ell(p') &\leq K(2M + \ell(w')) + K \\ &\leq K\ell(w') + 2MK + K. \end{aligned}$$

However, combining Inequality (1) together with the assumption that $M < l_k$ for $i + 1 \leq k \leq j - 1$ gives

$$\begin{aligned} (2) \quad 2M(j - i - 1) &\leq 2 \sum_{k=i+1}^{j-1} l_k \\ (3) \quad &\leq \ell(p') + \ell(w') \\ (4) \quad &\leq (K + 1)\ell(w') + 2MK + K. \end{aligned}$$

Assuming $j - i - 1 \geq 1$ and dividing, we have

$$\begin{aligned} M &\leq \frac{K + 1}{2} \cdot \frac{\ell(w')}{j - i - 1} + K \cdot \frac{M}{j - i - 1} + \frac{K}{j - i - 1} \\ &\leq \frac{K + 1}{2} \cdot (B + 1) \frac{j - i}{j - i - 1} + K \cdot \frac{M}{j - i - 1} + \frac{K}{j - i - 1} \\ &\leq (K + 1)(B + 1) + K \cdot \frac{M}{j - i - 1} + K. \end{aligned}$$

Since $M = 2((K + 1)(B + 1) + K)$, the above inequality implies that

$$M \leq \frac{1}{2}M + K \cdot \frac{M}{j - i - 1},$$

but this is only possible if $j - i \leq 2K + 1$. This imposes the uniform bound, $\ell(w') \leq (\frac{K}{2} + 1)(B + 1)$. Returning to Inequality (2) with this new information, we see that

$$\begin{aligned} \max_{i \leq k \leq j} l_k &\leq \sum_{k=i+1}^{j-1} l_k \\ &\leq \frac{1}{2}(K + 1)\ell(w') + MK + \frac{1}{2}K. \\ &\leq \frac{1}{2}(K + 1)(2K + 1)(B + 1) + MK + \frac{1}{2}K. \end{aligned}$$

If we denote this last quantity by L , we see that L depends only on K and N . Since any $1 \leq k \leq n$ with $l_k > M$ occurs in such a sequence investigated above, we have $l_k \leq L$. Hence, for any $1 \leq k \leq n$,

$$\ell(p_k) \leq K(2L + (B - 1)) + K.$$

Hence, any vertex on the path p lies at distance at most

$$\frac{1}{2}(K(2L + (B - 1)) + K) + L$$

from the geodesic w . Further, any vertex on w lies at distance at most $\frac{1}{2}(B - 1) + L$ from the path p . This completes the proof of the theorem by setting

$$S = \frac{1}{2}(K(2L + (B - 1)) + K) + L,$$

where $B = 3N + 1$ and

$$L = \frac{1}{2}(K + 1)(2K + 1)(B + 1) + MK + \frac{1}{2}K$$

for $M = 2((K + 1)(B + 1) + K)$. \square

Corollary 6.2 (Join-busting implies stable). *Suppose $H \leq A(\Gamma)$ is a finitely generated, join-busting subgroup. Then H is stable in $A(\Gamma)$. In particular, H is combinatorially quasiconvex in $A(\Gamma)$.*

Proof. Suppose that H is N -join-busting. By Proposition 4.4, H is undistorted in $A(\Gamma)$. Hence, to prove stability of H in $A(\Gamma)$ it suffices to show that for any elements $h_1, h_2 \in H$ there is a geodesic γ in $A(\Gamma)$ joining h_1 to h_2 such that if α is a K -quasigeodesic in $A(\Gamma)$ also joining h_1 to h_2 , then the Hausdorff distance between α and γ is bounded by M , depending only on K and N .

To see this, note that $g = h_1^{-1}h_2 \in H$ is N -join-busting by definition. By Theorem 6.1, if we denote by w a geodesic in $A(\Gamma)$ from 1 to g , then the Hausdorff distance between w and $h_1^{-1} \cdot \alpha$ is bounded by M ,

where M depends only on N and K . (Here, we are using the identification between the Cayley graph of $A(\Gamma)$ with its vertex generators and $\tilde{S}(\Gamma)^{(1)}$.) Setting $\gamma = h_1 \cdot w$ completes the proof of the corollary. \square

7. PROOF OF THEOREM 1.1

In this section we complete the proof of Theorem 1.1.

Proof. Suppose $H < A(\Gamma)$ is finitely generated. First let us prove that (3) implies (1), that is, if H is purely loxodromic then the orbit map of H into Γ^e is a quasi-isometric embedding. By Theorem 5.2, we know that, for any element of H written as a geodesic word in $A(\Gamma)$, the length of any join subword is bounded. As star subwords are join subwords, this implies that star-length is bi-Lipschitz to $A(\Gamma)$ -length for elements in H . On the other hand, orbit distance in Γ^e is quasi-isometric to star-length [KK14a]. Moreover, Proposition 4.4 says that $|\cdot|_H$ is bi-Lipschitz to $|\cdot|_{A(\Gamma)}$ for elements in H . Therefore we may conclude that, for H a purely loxodromic, finitely generated subgroup of $A(\Gamma)$, the orbit map of $(H, |\cdot|_H)$ into Γ^e is a quasi-isometric embedding. The converse is immediate: if $H \rightarrow \Gamma^e$ is a quasi-isometric embedding, no element of H can be elliptic, because these have bounded orbit, and right-angled Artin groups are torsion-free.

We have seen in Corollary 6.2 that join-busting implies stability; this and Theorem 5.2 proves that (3) implies (2). The converse follows from the observation that, if a subgroup is stable, all its \mathbb{Z} -subgroups are stable, and therefore generated by loxodromic elements. To see the latter claim, consider the cyclic subgroup generated by an elliptic element w , whose centralizer contains some nontrivial c with $\langle c, w \rangle \cong \mathbb{Z} \oplus \mathbb{Z}$. If we represent w and c by reduced words, then the edge paths corresponding to $c^n w^n c^{-n}$ and w^n are a pair quasigeodesics whose Hausdorff distance grows with n . Therefore $\langle w \rangle$ is not stable. \square

8. THE HANNA NEUMANN PROPERTY FOR STAR-FREE SUBGROUPS

In this section we complete the proof of Theorem 1.2. Let $v \in V(\Gamma)$, and let

$$G_v = \langle \{w \mid w \in V(\text{St}(v))\} \rangle,$$

i.e. the subgroup of $A(\Gamma)$ generated by the vertices in the star of v . We say that such a group G_v is a *star subgroup* of $A(\Gamma)$. Let $H < A(\Gamma)$ be a subgroup. We say that H is *star-free* if for every $1 \neq g \in H$, we have that g is not conjugate into a star subgroup of $A(\Gamma)$. In other words, there is no choice $h \in A(\Gamma)$ and $v \in V(\Gamma)$ such that $g^h \in G_v$. In the

language of [KK14a], H is star-free if for every $1 \neq g \in H$, a cyclically reduced conjugate of g always has star length at least two.

Note that a purely loxodromic subgroup of $A(\Gamma)$ is star-free, but the converse is false. For example, one can take

$$F_2 \times F_2 \cong \langle a, c \rangle \times \langle b, d \rangle.$$

If Γ is a square labeled cyclically by the vertices $\{a, b, c, d\}$, we have that

$$A(\Gamma) \cong \langle a, c \rangle \times \langle b, d \rangle.$$

Note that $A(\Gamma)$ has no loxodromic elements in this case, but any cyclically reduced word with full support is star-free. The subgroup $\langle abcd \rangle$ is cyclic, and every element is star-free. The following proposition gives the exact relationship between star-free and purely loxodromic subgroups of $A(\Gamma)$:

Proposition 8.1. *Let $H \leq A(\Gamma)$ be a finitely generated subgroup. Then H is purely loxodromic if and only if H is star-free and combinatorially quasiconvex.*

Proof. If H is purely loxodromic, then it is star-free by definition and quasiconvex (with respect to any generating set) by Theorem 1.1.

Now suppose that H is star-free and combinatorially quasiconvex but not purely loxodromic. As conjugating H preserves these conditions, we may suppose that there is a join $J \leq \Gamma$ such that

$$H_J = H \cap A(J) \neq \langle 1 \rangle.$$

Since both H and $A(J)$ are combinatorially quasiconvex in $A(\Gamma)$, H_J is combinatorially quasiconvex in $A(J)$. Since J is a join, $J = X * Y$ where X and Y are nonempty subgraphs of Γ each of which is contained in some star of Γ . This implies that $H_J \leq A(J) = A(X) \times A(Y)$ and

$$H_J \cap A(X) = H_J \cap A(Y) = \langle 1 \rangle,$$

because H is star-free.

Let $K \geq 0$ be the quasiconvexity constant for $H_J \leq A(J) = A(X) \times A(Y)$. If $h_i \in H_J$ are distinct elements whose word lengths go to infinity, then $h_i = x_i y_i$ and the lengths of either $x_i \in A(X)$ or $y_i \in A(Y)$ also go to infinity. Suppose $|x_i|_{A(J)} \rightarrow \infty$. Then since the concatenation $[1, x_i] \cup [x_i, x_i y_i]$ is a geodesic in $A(J) = A(X) \times A(Y)$ and H_J is K -quasiconvex, there must be a sequence $p_i \in A(X)$ and $q_i \in A(Y)$ with $p_i q_i \in H_J$ and $|q_i|_{A(Y)} \leq K$. As there are only finitely many elements of $A(Y)$ of length no greater than K , there are $i \neq j$ with $q_i = q_j$. Then

$$1 \neq p_i q_i (p_j q_j)^{-1} = p_i p_j^{-1} \in H_J \cap A(X),$$

a contradiction. \square

Star-free subgroups are always free, even without the assumption of finite generation:

Proposition 8.2. *Let $H < A(\Gamma)$ be a star-free subgroup. Then H is free.*

In particular, we obtain part (1) of Theorem 1.2. The proof of this proposition is identical to the proof of Theorem 53 in [KK14a] [cf. Theorem 2.1, item (7)]. We recall a proof for the convenience of the reader.

Proof of Proposition 8.2. Suppose $H \neq 1$. Choose a vertex $v \in V(\Gamma)$, and let Γ_v be the subgraph of Γ spanned by $V(\Gamma) \setminus \{v\}$. Then we have

$$A(\Gamma) \cong \langle \text{St}(v) \rangle *_{\langle \text{Lk}(v) \rangle} A(\Gamma_v),$$

so we get an action of $A(\Gamma)$ on the tree associated to this splitting. Since H is star-free, the action of H has trivial edge stabilizer, so that H splits as a free product of subgroups of $A(\Gamma_v)$ with possibly an additional free factor. The proposition follows by induction on $|V(\Gamma)|$. \square

Henceforth in this section, we will fix a finite free generating set S for H as a subgroup of $A(\Gamma)$. We also fix reduced word representatives for each element of S .

Note that Proposition 4.4 gives part (2) of Theorem 1.2. So, we are left with establishing part (3). In the remainder of this section, we will prove:

Theorem 8.3. *Let Γ be a finite simplicial graph, $\Lambda < \Gamma$ an induced subgraph, and let $H < A(\Gamma)$ be a finitely generated star-free subgroup. Then the group $H_\Lambda = H \cap A(\Lambda)$ is also finitely generated.*

Note that Theorem 8.3 does not exclude the possibility that H_Λ is the identity, and this will indeed be the case if Λ is contained in the star of a vertex of Γ . Theorem 8.3 is false without the assumption that H is star-free. For instance, consider the group $A(P_3) = F_2 \times \mathbb{Z}$, where the factors are generated by $\{a, b\}$ and z respectively. Consider the subgroup $H = \langle az, b \rangle$, which is a free subgroup of $A(P_3)$, and let us look at its intersection with $\langle a, b \rangle$. We have that a word w in az and b lies in $\langle a, b \rangle$ if and only if the a exponent of w is zero, i.e. w is in the kernel of the homomorphism $H \rightarrow \mathbb{Z}$ defined by $az \rightarrow 1$ and $b \rightarrow 0$. It follows that $H \cap \langle a, b \rangle$ is infinitely generated.

In the proof of Theorem 8.3, we use the fact that we can control cancellation distance within star-free subgroups in order to produce an explicit finite generating set for H_Λ .

8.1. Cancellation patterns. Before proceeding, we must define cancellation patterns of vertices within words in the star-free subgroup H . We will use the disk diagram terminology outlined in Section 3. Let $h \in H$ be an element viewed as a word in the free generating set S for H , and fix a reducing diagram Δ for h which realizes h as the reduced word w in $A(\Gamma)$. Using Lemma 3.2, we will assume the the disk diagram Δ and the reduced word w have been chosen so that w is combed in Δ . Thus, each letter of h has a co-oriented arc in Δ attached to it, which either has both endpoints in the subarc of $\partial\Delta$ labelled by h (in which case the two vertices cancel), or which has one endpoint in the subarc of $\partial\Delta$ labelled by w (in which case the vertex does not cancel and in which case we call the arc *contributing*). Arcs in the reducing diagram which are not contributing will be called *non-contributing* arcs. Every arc in Δ is labelled by a vertex of Γ or its inverse, depending on the co-orientation of the arc.




FIGURE 5. Reducing diagram indicating subword h_0 . Light blue and dark blue arcs have both ends in h .

Let h_0 be an H -subword of h , which is to say that if

$$h = s_1 \cdots s_n,$$

where each $s_i \in S$, we have that

$$h_0 = s_i \cdots s_j$$

for some $1 \leq i \leq j \leq n$. We will often use the notation $h_0 \leq h$ for a subword of h .

We now fix the H -subword $h_0 \leq h$ and a reducing diagram Δ for h . Suppose $v \in V(\Gamma)$. A v -cancellation partition for h_0 is a partition of the occurrences of v and v^{-1} in h_0 , viewed as an unreduced word in the vertex generators of $A(\Gamma)$, into eight sets

$$\{L_v^\pm, R_v^\pm, I_v^\pm, N_v^\pm\},$$

where a set with a positive exponent only contains copies of the element v and a set with a negative exponent only contains copies of the element v^{-1} . Intuitively, $h_0 \leq h \in H$ is a subword in which occurrences of v cancel in the reducing diagram Δ . The sets L_v^+ and R_v^+ correspond to the copies of v which cancel to the left of h_0 and to the right if h_0 , while the sets L_v^- and R_v^- play the same respective roles for copies of v^{-1} . The sets I_v^\pm correspond to copies of v and v^{-1} which cancel internally within h_0 , and N_v^\pm corresponds to copies of v and v^{-1} which do not cancel at all when h is reduced within $A(\Gamma)$. We will say that two words have the *same cancellation partition* if the corresponding sets

$$\{L_v^\pm, R_v^\pm, I_v^\pm, N_v^\pm\}$$

have the same cardinality. If the word h_0 is not clear from context, we will write

$$\{L_v^\pm(h_0), R_v^\pm(h_0), I_v^\pm(h_0), N_v^\pm(h_0)\}$$

to eliminate any ambiguities.

In a reducing diagram Δ for h , the copies of v and v^{-1} in L_v^+ and L_v^- have non-contributing arcs emanating from them that move off to the left of h_0 . Similarly, the copies of v and v^{-1} in R_v^+ and R_v^- have arcs emanating from them that move off to the right of h_0 . The copies of v and v^{-1} in $I_v^+ \cup I_v^-$ are connected by arcs with both endpoints in h_0 , and the copies of v and v^{-1} in $N_v^+ \cup N_v^-$ have vertical arcs connecting them to w .

Observe that group inversion relates a cancellation partition for a subword $h_0 \leq h$ with a cancellation partition for the subword $h_0^{-1} \leq h^{-1}$, with respect to the reducing diagram for h^{-1} induced by the inversely labeled mirror image of the reducing diagram for h . Specifically, we have

$$I_v^\pm(h_0) = I_v^\mp(h_0^{-1}), N_v^\pm(h_0) = N_v^\mp(h_0^{-1}),$$


and similarly

$$R_v^\pm(h_0) = L_v^\mp(h_0^{-1}), L_v^\pm(h_0) = R_v^\mp(h_0^{-1}).$$

We are now ready to define the cancellation pattern for a subword $h_0 \leq h$. We delete all the arcs in the diagram Δ which do not have an endpoint in h_0 , and for the remaining arcs, we retain the original co-orientation. Furthermore, we label each arc by the vertex of its terminal endpoint in h_0 . We then draw two unoriented, simple arcs ι and τ (for “initial” and “terminal”) in Δ which satisfy the following conditions:

- (1) The arcs ι and τ connect the subarc of $\partial\Delta$ labelled by h to the subarc of $\partial\Delta$ labelled by w .

- (2) Neither ι nor τ intersect any arcs of h that contribute to the reduced word w . (This is possible since the word w is combed in Δ .)
- (3) If a non-contributing arc has exactly one endpoint in h_0 , then the cancellation arc intersects exactly one of ι and τ exactly once.
- (4) There are no triple intersections of arcs.

FIGURE 6. Choice of ι and τ .

The *cancellation pattern* $\mathcal{P}(h_0)$ of $h_0 \leq h$ with respect to the chosen reducing diagram Δ is the closed region of the reducing diagram bounded by the boundary arcs, ι , and τ , together with the co-oriented, labelled non-contributing arcs. Two words h_1 and h_2 have the *same cancellation pattern* if $\mathcal{P}(h_1)$ is homeomorphic to $\mathcal{P}(h_2)$ via a homeomorphism which is a label and co-orientation-preserving homeomorphism of the union of the non-contributing arcs, and which preserves the boundary arcs of the cancellation pattern. Informally, two words have the same cancellation pattern if the cancellations of letters in the two words happen “in the same way.” It is clear from the pigeonhole principle that a fixed word $h_0 \leq h$ in H can have only finitely many different cancellation patterns as h varies among all possible words in H which contain h_0 as a subword, up to the equivalence relation of having the same cancellation pattern.

Fixing a reducing diagram Δ for h and a subword $h_0 \leq h$, a cancellation pattern of h_0 is related to a cancellation pattern for $h_0^{-1} \leq h^{-1}$ in a manner analogous to that of cancellation partitions. Again, we can build a reducing diagram for h^{-1} by taking the mirror image of the reducing diagram for h , inverting all the labels on the boundary arcs, and switching all the co-orientations on the non-contributing arcs. Thus, the natural choice for the cancellation pattern $\mathcal{P}(h_0^{-1})$ in h^{-1} is just the mirror image of $\mathcal{P}(h_0)$, with the labels ι and τ switched, the labels on

boundary arcs inverted, and the co-orientations of the non-contributing arcs switched.




FIGURE 7. Cancellation pattern $\mathcal{P}(h_0)$.

We remark briefly that when we define a cancellation pattern for a subword $h_0 \leq h$, we keep track of all the non-contributing arcs, not just the ones corresponding to a particular vertex of Γ . This point is essential in the proof of Theorem 8.3.

8.2. The Hanna Neumann Property.

Proof of Theorem 8.3. By induction on $|V(\Gamma)| - |V(\Lambda)|$, it suffices to prove the result for Λ equal to Γ with one vertex, say v , deleted. Given $h \in H_\Lambda = H \cap A(\Lambda)$, we will decompose h into a product of two words $h = h_1 \cdot h_2$, where $h_1, h_2 \in H_\Lambda$, where h_1 has H -length bounded by $K = K(H)$, and where h_2 has H -length shorter than the maximum of K and the H -length of h . Here, the H -length of an element of H is the word length with respect to the free generating set S . This will prove the result, since it will show that h is a product of elements of H_Λ of bounded H -length.

Observe that the statement $h \in H_\Lambda$ is just the statement that in some (and hence in every) reducing diagram for h , every occurrence of v and v^{-1} cancels. In other words, we have

$$N_v(h_0)^+ \cup N_v(h_0)^- = \emptyset$$

in the cancellation partition for every subword h_0 of h .

Writing

$$h = s_1 \cdots s_n,$$

Lemma 4.1 says that if $v \in \text{supp } s_i$ cancels with $v^{-1} \in \text{supp } s_j$, then $|j - i| \leq D$, where D is a constant which depends only on H . For every subword $h_0 \leq h$, we will fix a reducing diagram $\Delta(h_0)$.

We consider the set $\{\alpha_i\}$ of initial H -segments of h , which is to say that if

$$h = s_1 \cdots s_n,$$

then

$$\alpha_i = s_1 \cdots s_i.$$

Given an initial segment α_i , we call $\omega_i = \alpha_i^{-1}h$ the terminal H -segment of h . For notational convenience, we will assume that $D \leq i \leq n - D$. Since D is fixed and depends only on H , this assumption is valid, since there are only finitely many elements of H of H -length bounded by D .

For each i , we write b_i and e_i for the terminal H -segment of α_i of H -length D , and the initial H -segment of ω_i of H -length D , respectively. That is to say, we have

$$b_i = s_{i-D+1} \cdots s_i,$$

and

$$e_i = s_{i+1} \cdots s_{i+D}.$$

Observe that if $n \gg 0$ then there are two indices, say i and j with $i < j$, for which the segments b_i and b_j are equal as words in H , and for which the cancellation patterns are the same.




FIGURE 8. Part of reducing diagram with $b_i = b_j$ and e_j indicated; only arcs corresponding to v are shown, although other arcs appear in their cancellation patterns. However, by Lemma 4.1, no arcs traverse these three cancellation patterns, since they have H -length D .

Since we may assume j and i to be bounded in a way which depends only on H (since we produce i and j via a pigeonhole principle argument), we have that $\alpha_j \cdot \alpha_i^{-1}$ has universally bounded H -length, say

$K = K(H)$. Observe that both $\alpha_j \cdot \alpha_i^{-1}$ and $\alpha_i \cdot \omega_j$ have H -length shorter than the maximum of that of h and K . Moreover, we claim that $\alpha_j \cdot \alpha_i^{-1}$ and $\alpha_i \cdot \omega_j$ both lie in H_Λ . This last claim will suffice to establish the result.

We consider $\alpha_j \cdot \alpha_i^{-1}$ first. Observe that we may write $\alpha_j = \beta_j \cdot b_i$ and $\alpha_i = \beta_i \cdot b_i$ as reduced words in H , since $b_i = b_j$. By Lemma 4.1, in any reducing diagram for h and for each vertex $v \in V(\Gamma)$, any occurrence of v or v^{-1} in α_i or in α_j which cancels to the right must lie in b_i . Furthermore, since these are initial segments of $h \in H_\Lambda$, no occurrence of v nor v^{-1} cancels to the left. Thus, the only v -cancellation arcs for α_i and α_j which do not have both endpoints in α_i and α_j are precisely the ones coming from the sets R_v^\pm associated to b_i .

We may now construct a reducing diagram for

$$\beta_j b_i \cdot b_i^{-1} \beta_i^{-1} = \alpha_j \cdot \alpha_i^{-1}$$

by placing the cancellation patterns $\mathcal{P}(b_i)$ and $\mathcal{P}(b_i^{-1})$ next to each other, pairing the τ -arc of $\mathcal{P}(b_i)$ with the ι -arc of $\mathcal{P}(b_i^{-1})$. We then splice the cancellation arcs together by matching arcs with compatible labels and co-orientations. By the choice of D , no non-contributing arc with an endpoint in β_j traverses $\mathcal{P}(b_i)$, and no non-contributing arc with an endpoint in β_i^{-1} traverses $\mathcal{P}(b_i^{-1})$. The result is that every occurrence of v and v^{-1} in $\alpha_j \cdot \alpha_i^{-1}$ cancels in $A(\Lambda)$, which is to say that $\alpha_j \cdot \alpha_i^{-1} \in H_\Lambda$.

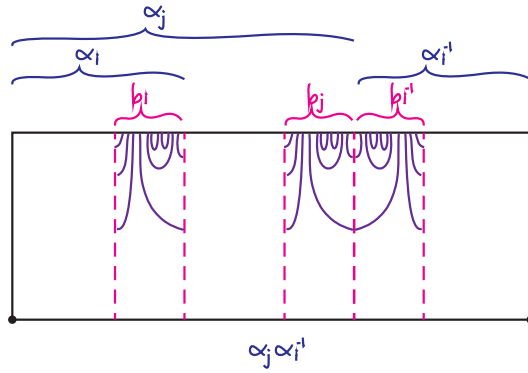


FIGURE 9. Part of reducing diagram for $\alpha_j \cdot \alpha_i^{-1}$.

The case of $\alpha_i \cdot \omega_j$ is similar but simpler. Consider the subword $b_i \cdot e_j$ of $\alpha_i \cdot \omega_j$. Since $b_i = b_j$ as words in H and since the cancellation patterns of b_i and b_j are the same as subwords of h , we have that every $v \in R_v^+(b_i)$ can be cancelled with a $v^{-1} \in L_v^-(e_j)$, and similarly every

$v^{-1} \in R_v^-(b_i)$ can be cancelled with a $v \in L_v^+(e_j)$. Observe that every copy of v and v^{-1} which does not lie in $R_v^\pm(b_i)$ lies in $I_v^\pm(\alpha_i)$. Similarly, every copy of v and v^{-1} which does not lie in $L_v^\pm(e_j)$ lies in $I_v^\pm(\omega_j)$. Thus, every occurrence of v and v^{-1} cancels in $\alpha_i \cdot \omega_j$, so that this element also lies in H_Λ . \square




FIGURE 10. Part of reducing diagram for $\alpha_i \cdot \omega_j$.

REFERENCES

- [Ago13] Ian Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning.
- [BC12] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356.
- [Beh06] Jason Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geom. Topol. **10** (2006), 1523–1578.
- [BF09] Mladen Bestvina and Koji Fujiwara, *A characterization of higher rank symmetric spaces via bounded cohomology*, Geom. Funct. Anal. **19** (2009), no. 1, 11–40.
- [BHS] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces I: curve complexes for cubical groups*, preprint.
- [CLM12] Matt T. Clay, Christopher J. Leininger, and Johanna Mangahas, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. **6** (2012), no. 2, 249–278.
- [CW04] John Crisp and Bert Wiest, *Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups*, Algebr. Geom. Topol. **4** (2004), 439–472.
- [DGO11] François Dahmani, Vincent Guirardel, and Denis Osin, *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*, preprint (2011), arxiv:1111.7048.
- [DKL12] Spencer Dowdall, Richard P. Kent, IV, and Christopher J. Leininger, *Pseudo-Anosov subgroups of fibered 3-manifold groups*, preprint (2012), arxiv:1208.2495.

- [Dro87] Carl Droms, *Graph groups, coherence, and three-manifolds*, J. Algebra **106** (1987), no. 2, 484–489.
- [DT14] Matthew Gentry Durham and Samuel J. Taylor, *Convex cocompactness and stability in mapping class groups*, preprint (2014), arxiv:1404.4803.
- [FM02] Benson Farb and Lee Mosher, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. **6** (2002), 91–152 (electronic).
- [Ham05] Ursula Hamenstadt, *Word hyperbolic extensions of surface groups*, preprint (2005), arXiv:math/0505244.
- [HW13] Mark F. Hagen and Daniel T. Wise, *Cubulating hyperbolic free-by-cyclic groups: the irreducible case*, preprint (2013), arxiv:1311.2084.
- [HW14] ———, *Cubulating hyperbolic free-by-cyclic groups: the general case*, preprint (2014), arxiv:1406.3292.
- [Kim08] Sang-hyun Kim, *Co-contractions of graphs and right-angled Artin groups*, Algebr. Geom. Topol. **8** (2008), no. 2, 849–868.
- [KK13] Sang-hyun Kim and Thomas Koberda, *Embedability between right-angled Artin groups*, Geom. Topol. **17** (2013), no. 1, 493–530.
- [KK14a] ———, *The geometry of the curve graph of a right-angled Artin group*, Internat. J. Algebra Comput. **24** (2014), no. 2, 121–169.
- [KK14b] ———, *An obstruction to embedding right-angled Artin groups in mapping class groups*, Int. Math. Res. Not. IMRN (2014), no. 14, 3912–3918.
- [KL07] Richard P. Kent, IV and Christopher J. Leininger, *Subgroups of mapping class groups from the geometrical viewpoint*, In the tradition of Ahlfors-Bers. IV, Contemp. Math., vol. 432, Amer. Math. Soc., Providence, RI, 2007, pp. 119–141.
- [KL08] Richard P. Kent, IV and Christopher J. Leininger, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. **18** (2008), no. 4, 1270–1325.
- [KLS09] Richard P. Kent, IV, Christopher J. Leininger, and Saul Schleimer, *Trees and mapping class groups*, J. Reine Angew. Math. **637** (2009), 1–21.
- [KM12] Jeremy Kahn and Vladimir Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. (2) **175** (2012), no. 3, 1127–1190.
- [Kob12] Thomas Koberda, *Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups*, Geom. Funct. Anal. **22** (2012), no. 6, 1541–1590.
- [LS01] Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Classics in Mathematics, Springer-Verlag, Berlin, 2001.
- [Mah11] Joseph Maher, *Random walks on the mapping class group*, Duke Math. J. **156** (2011), no. 3, 429–468.
- [Mih68] K. A. Mihailova, *The occurrence problem for direct products of groups*, Mat. Sb. (N.S.) **75** (1968), no. 117, 199–210.
- [Min96] Y.N. Minsky, *Quasi-projections in Teichmüller space*, Journal für die Reine und Angewandte Mathematik (1996), 121–136.
- [MM99] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
- [MT13] Johanna Mangahas and Samuel J. Taylor, *Convex cocompactness in mapping class groups via quasiconvexity in right-angled artin groups*, preprint (2013), arxiv:1306.5278.

- [Riv08] Igor Rivin, *Walks on groups, counting reducible matrices, polynomials, and surface and free group automorphisms*, Duke Math. J. **142** (2008), no. 2, 353–379.
- [Sap14] Mark Sapir, *A Higman embedding preserving asphericity*, J. Amer. Math. Soc. **27** (2014), no. 1, 1–42.
- [Ser89] Herman Servatius, *Automorphisms of graph groups*, J. Algebra **126** (1989), no. 1, 34–60.
- [Sis11] Alessandro Sisto, *Contracting elements and random walks*, preprint (2011), arxiv:1112.2666.
- [TT] Samuel J. Taylor and Giulio Tiozzo, in preparation.
- [Wis11] Daniel T. Wise, *The structure of groups with a quasiconvex hierarchy*.

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, 20 HILLHOUSE AVE,
NEW HAVEN, CT 06520, USA

E-mail address: thomas.koberda@gmail.com

DEPARTMENT OF MATHEMATICS, UNIVERSITY AT BUFFALO, 244 MATHEMATICS
BUILDING, BUFFALO, NY 14260, USA

E-mail address: mangahas@buffalo.edu

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, 20 HILLHOUSE AVE,
NEW HAVEN, CT 06520, USA

E-mail address: s.taylor@yale.edu