

THE ASYMPTOTIC GEOMETRY OF RIGHT-ANGLED ARTIN GROUPS, I

MLADEN BESTVINA, BRUCE KLEINER, AND MICHAH SAGEEV

ABSTRACT. We study *atomic* right-angled Artin groups – those whose defining graph has no cycles of length ≤ 4 , and no separating vertices, separating edges, or separating vertex stars. We show that these groups are not quasi-isometrically rigid, but that an intermediate form of rigidity does hold. We deduce from this that two atomic groups are quasi-isometric iff they are isomorphic.

CONTENTS

1. Introduction	2
2. Preliminaries	9
3. Quarter-planes	13
4. Preservation of maximal product subcomplexes	20
5. Flat space	23
6. Dual disk diagrams	31
7. Rigidity of taut cycles	35
8. $\mathbb{F}$ -rigidity	39
9. The proofs of corollaries 1.7, 1.8, and 1.9	43
10. Further implications of Theorem 1.6	44
11. Quasi-isometric flexibility and the proof of Theorem 1.10	45
References	47

Date: August 17, 2007.

M.B. is supported by an NSF grant. B.K. is supported by NSF Grant DMS-0505610. M.S. is supported by ISI #580/07.

1. INTRODUCTION

1.1. Background. We recall that to every finite simplicial graph Γ , one may associate a presentation with one generator for each vertex of Γ and one commutation relation $[g, g'] = 1$ for every pair of adjacent vertices $g, g' \in \Gamma$. The resulting group is the *right-angled Artin group (RAAG) defined by Γ* , and will be denoted $G(\Gamma)$ (we will often shorten this to G when the defining graph Γ is understood). This class of groups contains the free group F_k , whose defining graph has k vertices and no edges, and is closed under taking products; in particular it contains $\mathbf{Z}^k$ and $F_k \times F_l$. Every RAAG $G(\Gamma)$ has a canonical Eilenberg-MacLane space $\bar{K}(\Gamma)$ which is a nonpositively curved cube complex (called the *Salveti complex* in [Cha]); when G is 2-dimensional $\bar{K}(\Gamma)$ is homeomorphic to the presentation complex. We let $K(\Gamma)$ denote the universal cover of $\bar{K}(\Gamma)$.

RAAG's have been studied by many authors. The solution to the isomorphism problem has the strongest form: if $G(\Gamma) \cong G(\Gamma')$ then $\Gamma \cong \Gamma'$ [Dro87], [KMLNR80]. Servatius [Ser89] conjectured a finite generating set for $\text{Aut}(G(\Gamma))$ and his conjecture was proved by Laurence [Lau95]. The group $G(\Gamma)$ is commensurable with a suitable right angled Coxeter group [DJ00]. There is an analog of Outer space in the case when Γ is connected and triangle-free [CCV]. For a nice introduction to and more information about RAAG's see Charney's survey [Cha].

Our focus is on quasi-isometric rigidity properties of right-angled Artin groups. Some special cases have been treated earlier:

- The free group $G = F_k$. Here the standard complex $K(\Gamma)$ is a tree of valence $2k$. Quasi-isometries are not rigid – there are quasi-isometries which are not at bounded distance from isometries – but nonetheless any finitely generated group quasi-isometric to a free group acts geometrically on some tree [Sta68], [Dun85], [Gro93, 1.C1] and is commensurable to a free group [KPS73]. Furthermore, any quasi-action $G' \curvearrowright K$ is quasi-isometrically conjugate to an isometric action on a tree, see [MSW03] and Section 2.

- $G = F_k \times F_l$. The model space K is a product of simplicial trees; as with free groups, quasi-isometries are not rigid. However, by [MSW03, KKL98, Ahl02, KL07], quasi-actions are quasi-isometrically conjugate to isometric actions on some product of trees. It is a standard fact that there are groups quasi-isometric to G which are not commensurable to

it (see [Wis96, BM00] for examples that are non-residually finite, or simple, respectively).

- $G = \mathbf{Z}^k$. The model space is $\mathbf{R}^k$, which is not quasi-isometrically rigid. In general, quasi-actions are not quasi-isometrically conjugate to isometric actions, although this is the case for discrete cobounded quasi-actions [Gro81, Bas72], i.e. any group quasi-isometric to $\mathbf{Z}^k$ is commensurable to it.

- $G(\Gamma)$ where Γ is a tree of diameter at least 3. Behrstock-Neumann [BN06] showed that any two such Artin groups are quasi-isometric. Using work of Kapovich-Leeb [KL97], they also showed that a finitely generated group G is quasi-isometric to such an Artin group iff it is commensurable to one.

1.2. Statement of results. Our first result is that quasi-isometries of 2-dimensional RAAG's preserve flats (recall that a flat is a subset isometric to $\mathbf{R}^2$ with the usual metric):

Theorem 1.1. *Let Γ, Γ' be finite, triangle-free graphs, and let $K = K(\Gamma)$, $K' = K(\Gamma')$. Then there is a constant $D = D(L, A)$ such that if $\phi : K \rightarrow K'$ is an (L, A) -quasi-isometry and $F \subset K$ is a flat, then there is a flat $F' \subset K'$ such that the Hausdorff distance satisfies*

$$\text{Hd}(\phi(F), F') < D.$$

The remaining results in this paper concern a special class of RAAG's:

Definition 1.2. A finite simplicial graph Γ is *atomic* if

- (1) Γ is connected and has no vertex of valence < 2 .
- (2) Γ contains no cycles of length < 5 .
- (3) Γ has no separating closed vertex stars.

A RAAG is *atomic* if its defining graph is atomic.

The pentagon is the simplest example of an atomic graph. Our main results and most of the issues in the proofs are well illustrated by the pentagon case. Conditions (1)-(3) above exclude some of the known phenomena from the examples above. For instance, Condition (2) rules out abelian subgroups of rank greater than 2 and subgroups isomorphic to $F_k \times F_l$ where $\min(k, l) > 1$. Condition (3) prevents $G(\Gamma)$ from splitting in an obvious way over a subgroup with a nontrivial center; such a splitting would lead to a large automorphism group.

Remark 1.3. The main result of [Lau95] implies that the outer automorphism group of an atomic RAAG is generated by the symmetries of Γ and inversions of generators. It also implies that an arbitrary simplicial graph Γ with no cycles of length < 5 is atomic iff $G(\Gamma)$ has finite outer automorphism group.

The following example shows that Condition (3) is necessary in order for a RAAG to be determined up to isomorphism by its quasi-isometry class.

Example 1.4. Let Γ be any connected graph with no triangles and choose a vertex $v \in \Gamma$. Let $f : G(\Gamma) \rightarrow \mathbf{Z}_2$ be the homomorphism that sends v to $1 \in \mathbf{Z}_2$ and sends all other generators to $0 \in \mathbf{Z}_2$. Then $G' = \ker(f)$ is a RAAG whose defining graph Γ' can be obtained from Γ by doubling along the closed star of $v \in \Gamma$. Concretely, when Γ is the pentagon, G and G' are commensurable but not isomorphic. Thus the atomic property for RAAG's is not a commensurability invariant (in particular, it is not quasi-isometry invariant). See Section 11 for more discussion.

Until further notice, Γ will denote a fixed atomic graph, $G = G(\Gamma)$, $\bar{K} = \bar{K}(\Gamma)$ and $K = K(\Gamma)$, and $V = V(K)$ will denote the vertex set of K . Whenever Γ' appears, it will also be an atomic graph, and the associated objects will be denoted by primes.

Before stating our main rigidity theorem we need another definition:

Definition 1.5. A *standard circle* (respectively *standard torus*) is a circle (respectively torus) in $\bar{K}$ associated with a vertex (respectively edge) in Γ . A *standard geodesic* (respectively *standard flat*) is a geodesic $\gamma \subset K$ (respectively flat $F \subset K$) which covers a standard circle (respectively torus).

Note that if $p \in K$ is a vertex, then the standard geodesics passing through p are in 1-1 correspondence with vertices in Γ , and the standard flats are in 1-1 correspondence with edges of Γ . As a consequence, the defining graph Γ is isomorphic to the associated incidence pattern.

Theorem 1.6 (Rigidity, first version). *Suppose $\phi : K \rightarrow K'$ is an (L, A) -quasi-isometry. Then there is a unique bijection $\psi : V \rightarrow V'$ of vertex sets with the following properties:*

$$(1) \ d(\psi, \phi|_V) < D = D(L, A).$$

- (2) (*Preservation of standard flats*) Any two vertices $v_1, v_2 \in V$ lying in a standard flat $F \subset K$ are mapped by ψ to a pair of vertices lying in a standard flat $F' \subset K'$.
- (3) (*Preservation of standard geodesics*) Any two vertices $v_1, v_2 \in V$ lying in a standard geodesic $\gamma \subset K$ are mapped by ψ to a pair of vertices lying on a standard geodesic $\gamma' \subset K'$.

This theorem is proved, in the language of flat spaces, as Theorem 8.10. An immediate corollary is:

Corollary 1.7. *Atomic RAAG's are quasi-isometric iff they are isomorphic.*

This follows from Theorem 1.6 because the pattern of standard geodesics and standard flats passing through a vertex $p \in V$ determines the defining graph, and by the theorem, it is preserved by quasi-isometries.

Theorem 1.6 has two further corollaries, which may also be deduced from [Lau95]:

Corollary 1.8 (Mostow-type rigidity). *Every isomorphism $G \rightarrow G'$ between atomic RAAG's is induced by a unique isometry $K \rightarrow K'$, where we identify G and G' with subsets of $\text{Isom}(K)$, $\text{Isom}(K')$.*

Corollary 1.9. *Every homotopy equivalence $\bar{K}(\Gamma) \rightarrow \bar{K}(\Gamma)$ is homotopic to a unique isometry; equivalently, the homomorphism $\text{Isom}(\bar{K}) \rightarrow \text{Out}(G)$ is an isomorphism. In particular, there is an extension*

$$1 \longrightarrow H \longrightarrow \text{Out}(G) \longrightarrow \text{Aut}(\Gamma) \longrightarrow 1$$

where $\text{Aut}(\Gamma)$ denotes the automorphism group of the graph Γ , and H consists of automorphisms that send each generator g to $g^{\pm 1}$. Thus H is isomorphic to $\mathbf{Z}_2^V$, where V is the number of vertices in Γ .

To complement the rigidity theorem, we construct examples showing that K is not quasi-isometrically rigid, and that the failure of rigidity cannot be accounted for by the automorphism group, or even the commensurator:

Theorem 1.10. *Let $\text{Comm}(G)$ denote the commensurator of G . Then both of the canonical homomorphisms*

$$\text{Aut}(G) \rightarrow \text{Comm}(G), \quad \text{Comm}(G) \rightarrow \text{QI}(G)$$

are injective, and have infinite index images.

The proof is in Section 11.

Every atomic RAAG is commensurable to groups which do not admit a geometric action on its associated $CAT(0)$ complex; see Remark 11.1 for more discussion. In Example 1.4 the obvious involution of G' cannot be realized by an isometric involution on K . Nonetheless, there is a form of rigidity for quasi-actions of atomic RAAG's, which will appear in a forthcoming paper:

Theorem 1.11. *If $H \stackrel{\rho}{\curvearrowright} K$ is a quasi-action of a group on the standard complex K for an atomic RAAG, then ρ is quasi-conjugate to an isometric action $H \curvearrowright \hat{X}$, where $\hat{X}$ is a $CAT(0)$ 2-complex.*

In fact, the 2-complex $\hat{X}$ is closely related to K .

The following remains open:

Question 1.12. *If G is an atomic RAAG and H is a finitely generated group quasi-isometric to G , are H and G commensurable? Does H admit a finite index subgroup which acts isometrically on K ?*

1.3. Discussion of the proofs.

The proof of Theorem 1.6 bears some resemblance to proof of quasi-isometric rigidity for higher rank symmetric spaces. One may view both proofs as proceeding in two steps. In the first step one shows that quasi-isometries map top dimensional flats to top dimensional flats, up to finite Hausdorff distance; in the second step one uses the asymptotic incidence of standard flats to deduce that the quasi-isometry has a special form. In both proofs, the implementation of the first step proceeds via a structure theorem for top dimensional quasiflats, but the methods are rather different. The arguments used in the second step are completely different, although in both cases they are ultimately combinatorial in nature; in the symmetric space case one relies on Tits' theorem on building automorphisms, while our proof requires the development and analysis of a new combinatorial object. Another significant difference is that quasi-isometric rigidity is false in our case, and so the rigidity statement itself is more subtle.

The proof of Theorem 1.1 goes roughly as follows. We begin by invoking a result from [BKS07] which says that, modulo a bounded subset, every quasiflat in K is at finite Hausdorff distance from a finite union of quarter-planes. Here a *quarter-plane* is a subcomplex isometric to a Euclidean quadrant, and the Hausdorff distance is controlled by the quasiflat constants. The pattern of asymptotic incidence

of quarter-planes can be encoded in the *quarter-plane complex*, which is a 1-complex analogous to the Tits boundary of a Euclidean building or higher rank symmetric space. Flats in K correspond to minimal length cycles in the quarter-plane complex. Using the fact that the image of a quasiflat under a quasi-isometry is a quasiflat, one argues that quasi-isometries induce isomorphisms between quarter-plane complexes. Hence they carry minimal length cycles to minimal length cycles, and flats to flats (up to controlled Hausdorff distance).

The outline of the proof of Theorem 1.6 goes as follows.

Step 1. ϕ maps standard flats in K to within uniformly bounded Hausdorff distance of standard flats in K' . Theorem 1.1 and standard $CAT(0)$ geometry imply that ϕ maps maximal product subcomplexes to within controlled Hausdorff distance of maximal product subcomplexes. Because Γ is atomic, every standard flat is the intersection of two maximal product subcomplexes. Since maximal product subcomplexes are preserved, their coarse intersections are also preserved, and this leads to preservation of standard flats.

Before proceeding further we introduce an auxiliary object, the *flat space* $\mathbb{F} = \mathbb{F}(\Gamma)$, which is a locally infinite $CAT(0)$ 2-complex associated with Γ . This complex coincides with the *modified Deligne complex* of Charney-Davis [CD95] (see also [Dav98, HM96]); we are giving it a different name since we expect that the appropriate analog in other rigid situations will not coincide with the modified Deligne complex. Start with the discrete set $\mathbb{F}^{(0)}$, namely the collection of standard flats in K . Join two of these points by an edge iff the corresponding flats intersect in a standard geodesic; this defines the 1-dimensional subset $\mathbb{F}^{(1)}$ of $\mathbb{F}$. It is convenient to think of $\mathbb{F}^{(0)}$ and $\mathbb{F}^{(1)}$ as the 0- and 1-skeleton of $\mathbb{F}$ even though formally this is rarely the case.

To specify the rest of $\mathbb{F}$, we recall that the standard flats passing through a vertex $p \in K$ correspond bijectively to the edges in Γ . Using this correspondence, for each $p \in K$ we cone off the corresponding subcomplex of $\mathbb{F}^{(1)}$ to obtain the 2-complex $\mathbb{F}$. With an appropriately chosen metric (see Section 5), this becomes a $CAT(0)$ complex. We use the notation $\mathbb{F}(p)$ to denote the subcomplex of $\mathbb{F}$ corresponding to the flats passing through $p \in K$. It is easy to see that Theorem 1.6 is equivalent to saying that quasi-isometries $K \rightarrow K'$ induce isometries between flat spaces:

Theorem 1.13 (Rigidity, second version). *There is a constant $D = D(L, A)$ with the following property. If $\phi : K \rightarrow K'$ is an (L, A) -quasi-isometry, then there is a unique isometry $\phi_* : \mathbb{F} \rightarrow \mathbb{F}'$ such that for*

each vertex $F \in \mathbb{F}$, the image of F under ϕ has Hausdorff distance at most D from $\phi_*(F) \in \mathbb{F}'$.

We now switch to proving Theorem 1.13. Step 1 produces a bijection $\phi_0 : \mathbb{F}^{(0)} \rightarrow \mathbb{F}'^{(0)}$ between 0-skeleta.

Step 2. The bijection ϕ_0 extends to an isomorphism $\phi_* : \mathbb{F} \rightarrow \mathbb{F}'$. We know that if two distinct standard flats $F_1, F_2 \in \mathbb{F}^{(0)}$ intersect in a standard geodesic, then $\phi_0(F_1)$ and $\phi_0(F_2)$ intersect coarsely in a geodesic, i.e. for some $D_1 = D_1(L, A)$, the intersection $N_{D_1}(\phi_0(F_1)) \cap N_{D_1}(\phi_0(F_2))$ is at controlled Hausdorff distance from a standard geodesic in K' . The property of intersecting coarsely in a geodesic defines a relation on the collection of standard flats which is quasi-isometry invariant. The remainder of the argument establishes the following:

Theorem 1.14. *Any bijection $\mathbb{F}^{(0)} \rightarrow \mathbb{F}'^{(0)}$ which preserves the relation of intersecting coarsely in a geodesic, is the restriction of an isometry $\mathbb{F} \rightarrow \mathbb{F}'$.*

This boils down to showing that if $p \in K$, then there is a point $p' \in K'$ such that all the flats passing through p are mapped by ϕ_0 to flats passing through p' . To establish this we exploit *taut cycles*, which are special class of cycles in the 1-skeleton of flat space; in the pentagon case these are just the 5-cycles. The heart of the proof is Theorem 7.1, that the map ϕ_0 carries the vertices of a taut cycle to the vertices of a taut cycle. The proof of this theorem is based on small cancellation theory. Due to the abundance of taut cycles and their manner of overlap, one deduces from this that the flats in $\mathbb{F}(p)$ are mapped by ϕ_0 to $\mathbb{F}(p')$ for some $p' \in K'$ (see Section 8).

1.4. Organization of the paper. In Section 2 we discuss some background material on RAAG's, CAT(0) spaces, and cube complexes. Sections 3 and 4 develop the geometry of quasiflats, culminating in the proof of Theorem 1.1. Sections 5, 6, and 7 build toward the proof that (the standard flats corresponding to) the vertices of a taut cycle in $\mathbb{F}$ are mapped by a quasi-isometry to (the standard flats corresponding to) the vertices of a taut cycle (cf. Theorem 7.1). This is proved by studying cycles in $\mathbb{F}^{(1)}$, and certain disk fillings of them; the argument is developed in Sections 6 and 7. Section 8 promotes taut cycle rigidity (Theorem 7.1) to full $\mathbb{F}$ -rigidity (Theorem 1.13). Sections 9, 10, and 11 work out various implications of $\mathbb{F}$ -rigidity.

2. PRELIMINARIES

2.1. The structure of the model space $K(\Gamma)$. We begin by introducing some notation and terminology connected with RAAG's.

Let Γ be a finite simplicial graph. If Γ contains no triangles, denote by $\bar{K}(\Gamma)$ the presentation complex for $G(\Gamma)$. Thus $\bar{K}(\Gamma)$ has one vertex, one oriented edge for every vertex of Γ , and one 2-cell, glued in a commutator fashion, for every edge of Γ . The closed 2-cells are tori. More generally, one may define $\bar{K}(\Gamma)$ for arbitrary Γ by adding higher dimensional cells (tori), one k -cell for every complete subgraph of Γ on k vertices. Then $\bar{K}(\Gamma)$ is a nonpositively curved complex. The universal cover is a $CAT(0)$ cube complex denoted $K(\Gamma)$. We will often use the notation $\bar{K}$ or K , suppressing the graph Γ , when there is no risk of confusion.

We label the edges of $\bar{K}$ and K by vertices of Γ , and the squares by edges of Γ . More generally, each k -dimensional cubical face of $\bar{K}$ or K is labelled by a k -tuple of vertices in Γ , or equivalently, by a face of the flag complex of Γ . If $Y \subset K$ is a subcomplex, we define the *label of Y* to be the collection Γ_Y of faces of the flag complex of Γ arising as labels of faces of Y . In this paper we will be concerned primarily with 2-dimensional complexes, when the flag complex of Γ is Γ itself.

Recall that a *full subgraph* of Γ is a subgraph $\Gamma' \subset \Gamma$ such that two vertices $v, w \in \Gamma'$ span an edge in Γ' iff they span an edge in Γ . (This is closely related to the notion of the *induced subgraph* of a set of vertices of Γ .)

If $\Gamma' \subset \Gamma$ is a full subgraph, then there is a canonical embedding $\bar{K}(\Gamma') \hookrightarrow \bar{K}(\Gamma)$, which is locally convex and locally isometric. We call the image of such an embedding a *standard subcomplex of $\bar{K}$* , or simply a *standard subcomplex*. The inverse image of a standard subcomplex under the universal covering $K \rightarrow \bar{K}$ is a disjoint union of convex subcomplexes, each of which is isometric to $K(\Gamma')$; we also refer to these as *standard subcomplexes*. A *standard product subcomplex* is a standard subcomplex associated with a full subgraph $\Gamma' \subset \Gamma$ which decomposes as a nontrivial join. A *standard flat* is a standard subcomplex $F \subset K$ which is isometric to $\mathbf{R}^2$, i.e. it is associated with a single edge in Γ . A *standard geodesic* is a standard subcomplex $\gamma \subset K$ associated with a single vertex in Γ . If $V \subset \Gamma$ is a set of vertices in Γ , then the *orthogonal complement of V* is the set of vertices $w \in \Gamma$ which are adjacent to every element of V :

$$(2.1) \quad V^\perp := \{w \in \Gamma \mid d(w, v) = 1 \text{ for every } v \in V\}.$$

Thus the join $V \circ V^\perp$ is a (bipartite) subgraph of Γ . Note that the subgroup generated by $V^\perp$ lies in the centralizer of the subgroup generated by V .

A *singular geodesic* is a geodesic $\gamma \subset K$ contained in the 1-skeleton of K . Note that standard geodesics are singular, but singular geodesics need not be standard. (As an example, consider the case when Γ is a finite set, and $\overline{K}$ is a bouquet of circles, and K is a tree. Then every geodesic is singular, and standard geodesics are those which project to a single circle.) Singular rays are defined similarly. A *quarter-plane* is a 2-dimensional subcomplex of K isometric to a Euclidean quadrant.

For the remainder of this section, we will assume that Γ is triangle-free, i.e. $\dim K \leq 2$.

Lemma 2.2. *Two standard flats $F, F' \subset K$ lie in the parallel set $\mathcal{P}(\gamma)$ of some geodesic γ iff the intersection $N_r(F) \cap N_r(F')$ is unbounded for all sufficiently large $r \in (0, \infty)$.*

Proof. Clearly, if $F \cup F' \subset \mathcal{P}(\gamma)$ then $N_r(F) \cap N_r(F')$ contains γ when $r \geq \max(d(F, \gamma), d(F', \gamma))$.

By a standard argument, since the respective stabilizers in G of F and F' act cocompactly on F and F' , the stabilizer in G of $C := N_r(F) \cap N_r(F')$ acts cocompactly on C . Therefore, if the convex set C is unbounded, it contains a complete geodesic γ , and we have $F \cup F' \subset \mathcal{P}(\gamma)$. $\square$

Note that the label Γ_E of a quarter-plane $\alpha \times \beta = E \subset K$ is a bipartite graph which is the join $\Gamma_\alpha \circ \Gamma_\beta$.

Lemma 2.3. *Suppose $C \subset K$ is a convex 2-dimensional subcomplex which splits (abstractly) as a nontrivial product of trees $C = T_1 \times T_2$; here T_i may be finite, and have no branch points. Then C is contained in a standard product subcomplex.*

Proof. C is a square complex, and isomorphic to a product. Since opposite edges of a square have the same label, the edge labelling of C descends to edge labellings of T_1 and T_2 . Note if for $i \in \{1, 2\}$, a_i is an edge label appearing in T_i , then there is a corresponding square in $C = T_1 \times T_2$, and hence the corresponding vertices of Γ are joined by an edge. Therefore C defines a complete bipartite subgraph $\Gamma' \subset \Gamma$. If we choose a vertex $v \in C$ and let P be the copy of $K_{\Gamma'}$ passing through v , then clearly $C \subset P$, since C is connected and projects to $K_{\Gamma'}$. $\square$

Lemma 2.4. *Suppose C, C' are convex subsets of a $CAT(0)$ space X , and let $\Delta := d(C, C')$ (here $d(C, C')$ denote the minimum distance). Then*

(1) *The sets*

$$(2.5) \quad Y := \{x \in C \mid d(x, C') = \Delta\}$$

$$(2.6) \quad Y' := \{x \in C' \mid d(x, C) = \Delta\},$$

are convex.

(2) *The nearest point map $r : X \rightarrow C$ maps Y' isometrically onto Y ; similarly, the nearest point map $r' : X \rightarrow C'$ maps Y isometrically onto Y' .*

(3) *Y and Y' cobound a convex subset $Z \stackrel{\text{Isom}}{\cong} Y \times [0, \Delta]$.*

(4) *If in addition X is a locally finite $CAT(0)$ complex with cocompact isometry group, and C, C' are subcomplexes, then the sets Y and Y' are nonempty, and there is a constant $A > 0$ such that if $p \in C$, $p' \in C'$, and if $d(p, Y) \geq 1$, $d(p', Y') \geq 1$ then*

$$(2.7) \quad d(p, C') \geq \Delta + A d(p, Y), \quad d(p', C) \geq \Delta + A d(p', Y').$$

Furthermore, the constant A depends only on Δ and X (but not on C and C').

Proof. Assertions (1)-(3) are standard $CAT(0)$ facts, so we only prove assertion (4).

Suppose $\{p_k\} \subset C$, $\{p'_k\} \subset C'$ are sequences such that $d(p_k, p'_k) \rightarrow \Delta$. After passing to a subsequence if necessary, we may find a sequence $\{g_k\} \subset \text{Isom}(X)$ such that the sequences of pairs

$$(g_k(C), g_k p_k), (g_k(C'), g_k p'_k)$$

converge in the pointed Hausdorff topology to pairs $(C_\infty, p_\infty), (C'_\infty, p'_\infty)$. Since there are only finitely many subcomplexes of X which are contained in a given bounded set, we will have $p_\infty \in g_k(C)$, $p'_\infty \in g_k(C')$ for sufficiently large k . Therefore $g_k^{-1}(p_\infty) \in C$, $g_k^{-1}(p'_\infty) \in C'$, and the distance Δ is realized. This shows that the sets Y_1 and Y_2 are nonempty.

We now prove the first estimate in (2.7); the second one has a similar proof. Observe that a convergence argument as above implies that there is a constant $A > 0$ such that if $x \in C$ and $d(x, Y) = 1$, then $d(x, C') > \Delta + A$. If $p \in C \setminus N_1(Y)$, then $d_{C'}$ (the distance from C') equals Δ at $r(p)$ and at least $\Delta + A$ at the point $x := \overline{pr(p)} \cap S(r(p), 1)$; here $\overline{pr(p)}$ denotes the geodesic segment with endpoints p and $r(p)$. Since $d_{C'}$ is convex, this implies (2.7). $\square$

Remark 2.8. Note that the convex sets Y and Y' in 4 of Lemma 2.4 need not be subcomplexes.

Corollary 2.9. *Suppose α, α' are asymptotic singular rays in K . Then there are singular rays $\beta \subset \alpha$, $\beta' \subset \alpha'$ which bound a flat half-strip subcomplex, and hence have the same labels.*

Proof. By passing to subrays, may assume that α and α' are subcomplexes.

Applying Lemma 2.4 to α and α' , one finds that the distance between α and α' is attained on singular subrays $\beta \subset \alpha$, $\beta' \subset \alpha'$, which bound a convex subset Z isometric to $\beta \times [0, \Delta]$.

If $\Delta = 0$ then $\beta = \beta'$ and we are done, so suppose $\Delta > 0$. We may assume also that the initial point p of the ray β is a vertex of K . We claim that Z is a subcomplex of K . To see this, let $p' \in \beta'$ be the initial point of the ray β' , and consider the segment $\overline{pp'}$. Note that there must be a square $S \subset Z$ with vertex at p , whose boundary contains the initial part of $\overline{pp'}$. Repeating this reasoning, it follows that every point $x \in \overline{pp'}$ is contained in a square lying in Z . Further repetition shows that Z is a union of squares, and hence is a subcomplex.

The corollary now follows from Lemma 2.3. $\square$

Definition 2.10. Suppose $\alpha \subset K$ is a singular ray, and $[\alpha]$ denotes its asymptote class. The *label of $[\alpha]$* is the collection $\Gamma_{[\alpha]}$ of labels determined by the asymptote class of α :

$$\Gamma_{[\alpha]} := \cap \{ \Gamma_{\alpha'} \mid \alpha' \in [\alpha] \}.$$

By Corollary 2.9, every $\alpha' \in [\alpha]$ has a subray $\alpha'' \subset \alpha'$ such that $\Gamma_{\alpha''} = \Gamma_{[\alpha]}$. If $E \subset K$ is a quarter-plane, the *label of $[E]$* is the intersection

$$\Gamma_{[E]} := \cap \{ \Gamma_{E'} \mid E' \in [E] \},$$

where the two quarter-planes are equivalent if they have a quarter-plane in common, see Lemma 3.1.

Lemma 2.11. *Suppose $P, P' \subset K$ are standard product complexes (in particular, $\dim P = \dim P' = 2$). If $P \subset N_r(P')$ for some $r \in (0, \infty)$, then $P \subset P'$.*

Proof. The distance function $d_{P'}$ is a bounded convex function on P . Since geodesic segments in P are extendible in P , it follows that $d_{P'}$ must be constant on P ; set $\Delta := d_{P'}(P)$. By a standard $CAT(0)$ fact, if $f : K \rightarrow P'$ is the nearest point retraction, then P and $r(P)$ cobound a subset isometric to $P \times [0, \Delta]$. Since $\dim P = \dim K$, we must have $\Delta = 0$, and so $P \subset P'$. $\square$

2.2. Cube complexes and hyperplanes. We recall basic terminology and facts about $\text{CAT}(0)$ cubical complexes. For more details, see [Sag95].

A cubical complex is a combinatorial cell complex whose closed cells are Euclidean n -dimensional cubes $[0, 1]^n$ of various dimensions such that the link of each vertex is a simplicial complex (no 1-gons or 2-gons). A theorem of Gromov [Gro87] then tells us that a simply connected cubical complex is $\text{CAT}(0)$ if and only if the link of each vertex is a flag complex.

Since an n -cube is a product of n unit intervals, each n -cube comes equipped with n natural projection maps to the unit interval. A *hypercube* is the preimage of $\{\frac{1}{2}\}$ under one of these projections; each n -cube contains n hypercubes. A *hyperplane* in a $\text{CAT}(0)$ cube complex X is a connected subspace intersecting each cube in a hypercube or the emptyset. Hyperplanes are said to *cross* if they intersect non-trivially; otherwise they are said to be *disjoint*.

Here are some basic facts about hyperplanes in $\text{CAT}(0)$ cube complexes which we will use throughout our arguments.

- each hyperplane is embedded (i.e. it intersects a given cube in a single hypercube).
- Each hyperplane is a track [Dun85], and hence separates the complex into precisely two components, called *half-spaces*.
- if $\{H_1, \dots, H_k\}$ is a collection of pairwise crossing hyperplanes, then $\cap_k H_k \neq \emptyset$.
- each hyperplane is itself a $\text{CAT}(0)$ cube complex.

3. QUARTER-PLANES

In this section, we define a complex using the quarter-planes in K and their asymptotic incidence. This object is analogous to the complex at infinity that one has for symmetric spaces and Euclidean buildings, and corresponds to part of the Tits boundary. The main result is Theorem 3.10, which implies that quasi-isometries between 2-dimensional RAAG complexes preserve flats.

Unless otherwise indicated, in this section Γ will be a triangle-free defining graph and K will be the associated $\text{CAT}(0)$ complex.

Lemma 3.1. *Suppose $E = \alpha \times \beta$ and $E' = \alpha' \times \beta'$ are quarter-planes in K . Then one of the following holds:*

- (1) (*Equivalent*) There is a quarter-plane $E'' \subset E \cap E'$; hence E'' has finite Hausdorff distance from both E and E' .
- (2) (*Incident*) There are constants $A, B \in (0, \infty)$ such that after relabelling the factors of the E and E' if necessary, α is asymptotic to α' and for every $p \in E, p' \in E'$,
- $$(3.2) \quad d(p, E') \geq A(d(p, \alpha) - B), \quad d(p', E) \geq A(d(p', \alpha') - B).$$
- (3) (*Divergent*) The distance function d_E grows linearly on E' , and vice-versa, i.e. there are constants $A, B \in (0, \infty)$ such that for all $x \in E, y \in E'$,
- $$(3.3) \quad d(x, E') \geq A(d(x, p) - B), \quad d(y, E) \geq A(d(y, p) - B).$$

Proof. We apply Lemma 2.4 with $C = E$ and $C' = E'$, and let $\Delta \in [0, \infty)$, $Y \subset E$, $Y' \subset E'$ and $Z \simeq Y \times [0, \Delta] \simeq Y' \times [0, \Delta]$ be as in that Lemma.

We claim that one of the following holds:

- a. $Y = Y'$ is a quarter-plane
- b. Y (respectively Y') is at finite Hausdorff distance from one of the boundary rays α, β (respectively α', β')
- c. Y, Y' are bounded.

To see this, first suppose $\Delta = 0$. Then $Y = Y' = E \cap E'$ is a convex subcomplex of both E and E' . Since K is a square complex, this implies that Y and Y' are product subcomplexes of E , and the claim follows. If $\Delta > 0$, then since Z meets Y and Y' orthogonally, it follows that Y, Y' are contained in the 1-skeleton. The claim then follows immediately.

Lemma 2.4 now completes the proof. $\square$

Henceforth we will use the terms *equivalent*, *incident*, and *divergent*, for the 3 cases in the lemma above.

Definition 3.4. We define the *quarter-plane complex* $\mathcal{Q}$ as follows. It has one vertex for each asymptote class of singular geodesic rays, and one edge joining $[\alpha]$ to $[\beta]$ for each equivalence class of quarter-planes $[\alpha \times \beta]$.

Note that this definition is motivated by the complex of chambers at infinity that one has for symmetric spaces of noncompact type, and also by the Tits boundary.

Lemma 3.5. *Suppose $E = \alpha \times \beta \subset K$ is a quarter-plane. Then precisely one of the following holds:*

- (1) $|\Gamma_{[\alpha]}^\perp| = |\Gamma_{[\beta]}^\perp| = 1$. This implies that $|\Gamma_{[\alpha]}| = |\Gamma_{[\beta]}| = 1$, which means that $\Gamma_{[E]}$ is a single edge of Γ which forms a connected component of Γ . The quarter-plane E belongs to a unique cycle in $\mathcal{Q}$, namely the 4-cycle $\mathcal{Q}_F$ associated with a unique flat $F \subset K$ labelled by $\Gamma_{[E]}$.
- (2) After relabelling the factors of E , $|\Gamma_{[\alpha]}^\perp| = 1$, $|\Gamma_{[\beta]}^\perp| > 1$. Then there is a unique equivalence class $[E']$ of quarter-planes such that $[E']$ is incident to $[E]$ at $[\alpha]$. Furthermore, any cycle $\Sigma \subset \mathcal{Q}$ containing $[E]$ must also contain $[E']$, and there is a pair of flats $F, F' \subset K$ such that the corresponding 4-cycles $\mathcal{Q}_F, \mathcal{Q}_{F'} \subset \mathcal{Q}$ intersect precisely in $[E] \cup [E']$.
- (3) $\min(|\Gamma_{[\alpha]}^\perp|, |\Gamma_{[\beta]}^\perp|) > 1$. Then there is a pair of flats $F, F' \subset K$ such that the corresponding 4-cycles $\mathcal{Q}_F, \mathcal{Q}_{F'}$ satisfy $\mathcal{Q}_F \cap \mathcal{Q}_{F'} = [E]$.

Proof. Case 1. Since $\Gamma_{[\alpha]} \subset \Gamma_{[\beta]}^\perp$ and $\Gamma_{[\beta]} \subset \Gamma_{[\alpha]}^\perp$, it follows that $|\Gamma_{[\alpha]}| = |\Gamma_{[\beta]}| = 1$. Setting $\Gamma_{[\alpha]} = \{v\}$, $\Gamma_{[\beta]} = \{w\}$, we find that w is the unique vertex adjacent to v , and vice-versa. Let $K' \subset K$ be the product subcomplex associated with the edge $\overline{vw} \subset \Gamma$, containing a sub-quarter-plane of E ; then K' is a flat which determines a cycle $\Sigma_0 \subset \mathcal{Q}$. Note that the valence at each vertex in Σ_0 is 2, and hence any cycle $\Sigma \subset \mathcal{Q}$ which intersects Σ_0 must coincide with Σ_0 .

Case 2. Let $\Gamma' \subset \Gamma$ be the join of $\Gamma_{[\alpha]}$ and $\Gamma_{[\beta]} = \Gamma_{[\alpha]}^\perp$, and let $K' \subset K$ be the copy of $K(\Gamma')$ containing a sub-quarter-plane of E . If E' is a quarter-plane adjacent to E along $[\alpha]$, then the asymptotic label of E' must be contained in Γ' , and so after passing to a sub-quarter-plane of E' if necessary, we may assume that $E' = \alpha' \times \beta' \subset K''$, where K'' is another copy of $K(\Gamma')$.

We claim that $K'' = K'$. Otherwise, K' and K'' would be disjoint, and the asymptotic singular rays α, α' would contain subrays $\bar{\alpha} \subset \alpha$, $\bar{\alpha}' \subset \alpha'$ such that $\bar{\alpha}$ and $\bar{\alpha}'$ bound a half-strip subcomplex $Z \subset K$. Since $\Gamma_{[\alpha]}^\perp = \Gamma_{[\beta]} \subset \Gamma'$, we have that Z lies in a single copy of $K(\Gamma')$, which is a contradiction.

Thus $E' \subset K'$. But then $\beta' \subset K'$ is a ray adjacent to α which is not asymptotic to β ; there is a unique such asymptote class in K' , and this implies that $[E']$ is unique.

Let $\Gamma'' \subset \Gamma$ be the join of $\Gamma_{[\beta]}$ and $\Gamma_{[\beta]}^\perp \supset \Gamma_{[\alpha]}$. Let K'' be the copy of $K(\Gamma'')$ which contains K' . Then K'' is a product $\mathbf{R} \times T$ where β

is asymptotic to the $\mathbf{R}$ -factor, and T is a tree of valence ≥ 4 because $|\Gamma_{[\beta]}| > 1$. The remaining assertions follow readily from this.

Case 3. Let $\Gamma' \subset \Gamma$ be the join of $\Gamma_{[\alpha]}$ and $\Gamma_{[\alpha]}^\perp$, and $\Gamma'' \subset \Gamma$ be the join of $\Gamma_{[\beta]}$ and $\Gamma_{[\beta]}^\perp$. After passing to a sub-quarter-plane of E if necessary, we may assume that $E \subset K' \cap K''$, where $K', K'' \subset K$ are product subcomplexes associated with Γ' and Γ'' respectively. Since $\min(|\Gamma_{[\alpha]}^\perp|, |\Gamma_{[\beta]}^\perp|) > 1$, there are pairs of flats $F'_1 \cap F'_2 \subset K'$ and $F''_1 \cap F''_2 \subset K''$ such that the intersections $F'_1 \cap F'_2, F''_1 \cap F''_2$ are half-planes whose intersection is precisely E .

□

In the remainder of this section, we will apply results from [BKS07]. We refer the reader to that paper for the definition and properties of support sets.

Lemma 3.6. *Let $Q \subset K$ be a quasiflat. There is a unique cycle $[E_1], \dots, [E_k] \subset \mathcal{Q}$ of quarter-planes in $\mathcal{Q}$, such that the union*

$$\cup_i E_i$$

has finite Hausdorff distance from Q . We denote this cycle by $\mathcal{Q}_Q$.

Proof. By [BKS07, Section 5], there is a finite collection $E_1, \dots, E_k$ of quarter-planes in K such that each E_i is contained in the support set associated with Q , and for some $r \in (0, \infty)$,

$$Q \subset N_r(\cup_i E_i).$$

Note that this collection of quarter-planes is uniquely determined up to equivalence. To see this, observe that if $E'_1, \dots, E'_l$ is another collection of quarter-planes with

$$Q \subset N_{r'}(\cup_j E'_j)$$

for some $r' \in (0, \infty)$, then for each $1 \leq i \leq k$, there must be a $1 \leq j \leq l$ such that E'_j is equivalent to E_i ; otherwise Lemma 3.1 would imply that there are points in E_i arbitrarily far from the union $\cup_j E'_j$.

We now assume that the quarter-planes $E_1, \dots, E_k$ represent distinct equivalence classes.

Pick $1 \leq j \leq k$, and consider the quarter-plane $E_j = \alpha_j \times \beta_j$. Suppose α_j is incident to i of the quarter-planes in the collection $\{E_1, \dots, E_k\}$, where $i \neq j$. Pick $R \in (0, \infty)$. By Lemma 3.1, if we choose $x \in \alpha_j$ lying sufficiently far out the ray α_i , then the ball $B(x, R)$ will intersect Q in a set which is uniformly quasi-isometric to an R -ball $B(x', R)$ lying in an i -pod, where x' is a singular point of the i -pod.

(Recall that an i -pod is a tree which is a union of i rays emanating from a single vertex.) This contradicts the fact that Q is a quasiflat. Therefore $\{E_1, \dots, E_k\}$ determines a union of cycles in $\mathcal{Q}$. But it can only contain a single cycle, because the support set of Q has only one end. $\square$

Lemma 3.7. *Every cycle $\Sigma \subset \mathcal{Q}$ arises from a quasiflat $Q \subset K$.*

Proof. Let the consecutive edges in Σ be represented by quarter-planes $E_0, \dots, E_j \subset K$, where the indices take values in the cyclic group $\mathbf{Z}_{j+1}$. Let W be the space obtained from the disjoint union $\sqcup_i E_i$ by gluing the boundary rays isometrically, in a cyclic fashion; let $\bar{E}_i$ denote the image of E_i in W under the quotient map $\pi : \sqcup E_i \rightarrow W$. With respect to the path metric, W is biLipschitz homeomorphic to $\mathbf{R}^2$, and the quotient map $E_i \rightarrow \bar{E}_i$ is a biLipschitz embedding for each i . Define $\phi : W \rightarrow K$ by setting $\phi(w) = p$, where $p \in \cup_i E_i$ is a point with $\pi(p) = w$.

Since Σ is a cycle, consecutive quarter-planes are incident, and this implies that there is a constant $C \in (0, \infty)$ such that for all $w, w' \in W$,

$$d(\phi(w), \phi(w')) \leq d(w, w') + C.$$

We claim that there are constants $L, A \in (0, \infty)$ such that for all $w, w' \in W$,

$$d(\phi(w), \phi(w')) \geq L^{-1}d(w, w') - A.$$

If this were false, there would be sequences $w_k, w'_k \in W$ such that $d(w, w') \rightarrow \infty$, and

$$\frac{d(\phi(w_k), \phi(w'_k))}{d(w_k, w'_k)} \rightarrow 0.$$

By Lemma 3.1, without loss of generality we may assume that for some $i \in \mathbf{Z}_{j+1}$, $w_k \in \text{Int}(E_i)$ and $w'_k \in \text{Int}(E_{i+1})$ for all k . Let $Y \subset E_i$, $Y' \subset E_{i+1}$ be as in Lemma 3.1, where $C = E_i$, $C' = E_{i+1}$. Then we must have

$$\limsup_{k \rightarrow \infty} \frac{\max(d(w_k, Y), d(w'_k, Y'))}{d(w_k, w'_k)} \rightarrow 0$$

because of (3.2) or (3.3). But then we may replace w_k, w'_k with sequences lying in Y and Y' respectively, and this yields a contradiction. $\square$

Definition 3.8. For each quasiflat $Q \subset K$, we let $\mathcal{Q}_Q \subset \mathcal{Q}$ denote the corresponding cycle in $\mathcal{Q}$. We denote by $\mathcal{C}$ be the collection of cycles in $\mathcal{Q}$, and by $\mathcal{N}$ the poset of subcomplexes of $\mathcal{Q}$ generated by

elements of $\mathcal{C}$ under finite intersection and union. Finally, $\mathcal{M} \subset \mathcal{N}$ will denote the collection of elements of $\mathcal{N}$ which are minimal among those of dimension 1.

We will consider the collection of subsets of X up to Hausdorff equivalence. Two subsets $A, B \subset X$ are *Hausdorff equivalent* if for some $r > 0$ we have $N_r(A) \supset B$ and $N_r(B) \supset A$. The equivalence class of A is denoted $[A]$. We define $[A] \cup [B]$ as $[A \cup B]$. It is not hard to check that this is independent of the choice of representatives. Note that $A \cap B$ is generally not Hausdorff equivalent to $N_r(A) \cap N_r(B)$, so $[A] \cap [B]$ is not defined. To remedy this, say that a collection $\{[A_i]\}$ of Hausdorff equivalence classes is *coherent* if for any finite subcollection $A_{i_1}, \dots, A_{i_k}$ there is $r_0 > 0$ such that for every $r > r_0$ the sets $N_r(A_{i_1}) \cap \dots \cap N_r(A_{i_k})$ and $N_{r_0}(A_{i_1}) \cap \dots \cap N_{r_0}(A_{i_k})$ are Hausdorff equivalent. In this situation define $[A_{i_1}] \cap \dots \cap [A_{i_k}]$ as $[N_r(A_{i_1}) \cap \dots \cap N_r(A_{i_k})]$ for large r . This concept behaves well under finite unions: if each A_{i_j} is written as a finite union of sets, and the collection of all of these is coherent, then the collection $A_{i_1}, \dots, A_{i_k}$ is coherent as well. The usual associativity and distributivity laws apply, e.g. $([A_{i_1}] \cup [A_{i_2}]) \cap [A_{i_3}] = ([A_{i_1}] \cap [A_{i_3}]) \cup ([A_{i_2}] \cap [A_{i_3}])$.

For example, the collection of (the classes of) quarter-planes in X is coherent, by Lemma 3.5.

Now consider the collection $\mathcal{QF}$ of quasiflats in X , modulo Hausdorff equivalence. By Lemma 3.6 this collection is coherent. Let $\mathcal{P}$ be the collection of subsets of X modulo Hausdorff equivalence obtained by intersecting finite collections of elements of $\mathcal{QF}$. Every element of $\mathcal{P}$ has a representative which is a finite union of quarter-planes and standard rays, so we have a natural map $\mathcal{P} \rightarrow \mathcal{N}$. This map preserves finite intersections by Lemma 3.5, and is therefore a bijection. It follows that minimal elements of $\mathcal{P}$ correspond bijectively to minimal elements of $\mathcal{N}$. After removing elements of $\mathcal{P}$ represented by collections of rays (call those *inessential*, while the other elements are *essential*) and elements of $\mathcal{N}$ that are 0-dimensional we obtain an isomorphism $\mathcal{M}_{\mathcal{P}} \rightarrow \mathcal{M}$ between the essential minimal elements of $\mathcal{P}$ and 1-dimensional elements of $\mathcal{M}$.

If $f : X \rightarrow X'$ is a quasi-isometry, then f maps quasi-flats to quasi-flats and induces a bijection $\mathcal{M}_{\mathcal{P}}(X) \rightarrow \mathcal{M}_{\mathcal{P}}(X')$. This bijection preserves inessential elements, and therefore there is an induced bijection $\mathcal{M}(X) \rightarrow \mathcal{M}(X')$.

Observe that Lemma 3.5 yields a classification of elements of $\mathcal{M}$, i.e. a minimal element of $\mathcal{N}$ consists of either 4 edges, 2 edges, or 1 edge, according to the relevant case of Lemma 3.5. Moreover, by the next lemma, if $\Sigma \in \mathcal{M}$, then we may determine how many quarter-planes it contains by determining the number of elements of $\mathcal{M}$ that are required to complete Σ to a cycle.

Lemma 3.9. *Let $E = \alpha \times \beta$ be a quarter-plane as in Lemma 3.5(2). Then there is a flat $F \subset K$ such that $F = l_A \times l_B$, l_A is asymptotic to α in the forward direction, l_B is asymptotic to β in the forward direction, and the quarter-plane $l_A^- \times l_B^-$ formed by the backward directions of l_A and l_B also satisfies that there is a unique class of quarter-planes incident to $[l_A^- \times l_B^-]$ at $[l_A^-]$.*

Proof. Since $\Gamma_{[\alpha]}^\perp = \Gamma_{[\beta]}$ consists of one label, we may pass to a sub-quarter-plane of E if necessary so that all labels along β are equal. By Lemma 2.3, E is contained in a standard product subcomplex $T_1 \times T_2$. We may replace T_2 by a line l_B that carries the label that appears along β . Now let $F = l_A \times l_B$ for a line l_A that extends α and so that all labels in T_1 appear infinitely often along l_A^- . Then $|\Gamma_{[l_B]}^\perp| > 1$ and $|\Gamma_{[l_A]}^\perp| = 1$, so the claim follows from Lemma 3.5(2). $\square$

Theorem 3.10. *Suppose Γ, Γ' are finite graphs, Γ is triangle-free, and $K = K(\Gamma)$, $K' = K(\Gamma')$ are the associated $CAT(0)$ square complexes. For every $L, A \in (0, \infty)$ there is a $D \in (0, \infty)$ such that if $f : K \rightarrow K'$ is an (L, A) -quasi-isometry and $F \subset K$ is a flat, then Γ' is triangle-free, and*

$$\text{Hd}(f(F), F') < D$$

for some flat $F' \subset K'$.

Proof. If Γ' contained a triangle, then K' would contain a 3-flat F . Then F would contain pairs of 2-flats F', F'' which are parallel but lie at arbitrarily large Hausdorff distance from one another. But applying a quasi-inverse of f to F', F'' , we would obtain pairs $Q', Q'' \subset K$ of quasi-flats with uniform constants, lying at arbitrarily large — but finite — Hausdorff distance from one another. This contradicts the fact that Q' and Q'' lie at controlled Hausdorff distance from their support sets. Thus Γ' must be triangle-free.

Let $\mathcal{Q}$ and $\mathcal{Q}'$ be the corresponding quarter-plane complexes, $\mathcal{C}, \mathcal{C}'$ be the collections of cycles in $\mathcal{Q}$ and $\mathcal{Q}'$, and $\mathcal{N}, \mathcal{N}'$ be the poset of subcomplexes of $\mathcal{Q}$ and $\mathcal{Q}'$ respectively, generated by $\mathcal{C}, \mathcal{C}'$ under finite intersection and union. As discussed above, f induces a poset

isomorphism $f_* : \mathcal{N} \rightarrow \mathcal{N}'$ which preserves dimension. Thus it induces a bijection between $\mathcal{M}$ and $\mathcal{M}'$ which also preserves the number of quarter-planes. If $\Sigma \subset \mathcal{Q}$ is a cycle, it is a union of a uniquely determined elements of $\mathcal{M}$, and these are mapped by f_* to the unique elements of $\mathcal{M}'$ which give $f_*\Sigma \subset \mathcal{Q}'$.

If $F \subset K$ is a flat, then $\mathcal{Q}_F \in \mathcal{C}$ is a 4-cycle, so $\mathcal{Q}'_{f(F)}$ is a 4-cycle in $\mathcal{Q}'$. It follows that the area of the support set S' associated with $f(F)$ grows asymptotically like πr^2 . By [BKS07, Section 3] this implies that S' is a flat, and hence $f(F)$ is at Hausdorff distance at most $D = D(L, A)$ from a flat. $\square$

4. PRESERVATION OF MAXIMAL PRODUCT SUBCOMPLEXES

Using Theorem 3.10, in this section we deduce that maximal product complexes are preserved by quasi-isometries.

Let Γ and Γ' be triangle-free graphs, and K, K' be the associated $CAT(0)$ complexes. Let $f : K \rightarrow K'$ be an (L, A) -quasi-isometry, where $\dim K \leq 2$.

Lemma 4.1. *There is a constant $D_0 = D_0(L, A) \in (0, \infty)$ with the following property. Suppose $Y \subset K$ is a subcomplex isometric to the product of a tripod with $\mathbf{R}$. Then the singular geodesic $\gamma \subset Y$ is mapped by f to within Hausdorff distance at most D_0 of a singular geodesic $\gamma' \subset K'$, and $f(Y)$ lies in the D_0 -neighborhood of the parallel set $\mathbb{P}(\gamma') \subset K'$.*

Proof. The set Y is a union

$$Y = F_1 \cup F_2 \cup F_3,$$

where the F_i 's are flats intersecting in γ . By Theorem 3.10, $f(F_i)$ lies at controlled Hausdorff distance from a unique flat $F'_i \subset K'$, and hence for $r = r(L, A) \in (0, \infty)$, the intersection

$$W := N_r(F'_1) \cap N_r(F'_2) \cap N_r(F'_3)$$

is quasi-isometric to $\mathbf{R}$. As W is convex, it contains a geodesic γ_1 . Then $\mathbb{P}(\gamma_1)$ contains $F'_1 \cup F'_2 \cup F'_3$ which implies that γ_1 is parallel to a singular geodesic γ' , where the Hausdorff distance $\text{Hd}(\gamma', \gamma_1)$ is controlled. $\square$

Theorem 4.2. *There is a constant $D = D(L, A) \in (0, \infty)$ such that if $P \subset K$ is a standard product subcomplex, then its image in K' is contained in the D -neighborhood of a standard product subcomplex of K' .*

Proof. The subcomplex P is associated with a subgraph $\Gamma' \subset \Gamma$, where Γ' is a join $\Gamma' = A \circ B$, where both A and B are nonempty.

Case 1. $|A| = |B| = 1$ and P is a single flat. By Theorem 3.10 we know that $f(P)$ lies in a neighborhood of controlled thickness around a flat $F' \subset K'$. But every flat in K' is contained in a product subcomplex by Lemma 2.3.

Case 2. $|A| = 1$, $|B| > 1$, and P is a parallel set strictly larger than a single flat. Here $P \simeq T \times \mathbf{R}$, where T is tree of valence at least 4 everywhere. Let $\gamma \subset P$ be a singular geodesic of the form $\text{pt} \times \mathbf{R}$. Then each point $p \in P$ lies in a subset $Y \subset P$ isometric to tripod $\times \mathbf{R}$, where $\gamma \subset Y$ is the singular locus of Y . Applying Lemma 4.1, we conclude that $f(Y)$ is contained in a controlled neighborhood of the parallel set of a singular geodesic $\gamma' \subset K'$, where $\text{Hd}(\gamma', f(\gamma))$ is controlled.

Case 3. $\min(|A|, |B|) > 1$, and P is a product where both factors are strictly larger than a line. Let $\mathcal{H}, \mathcal{V}$ be the collections of horizontal and vertical singular geodesics in Y . By applying Lemma 4.1, we see that these map to within controlled distance of singular geodesics in K' ; we let $\mathcal{H}', \mathcal{V}'$ be the corresponding sets of singular geodesics. Since any pair $\alpha' \in \mathcal{H}', \beta' \in \mathcal{V}'$ spans a flat plane in K' , their labels must lie in a join subcomplex of the defining graph Γ' . Let $A \subset \Gamma'$ (respectively $B \subset \Gamma'$) be the set of vertices of Γ' which arise as a label of some $\alpha' \in \mathcal{H}'$ (respectively $\beta' \in \mathcal{V}'$). Then $A \cup B$ spans a join subgraph $\Gamma_0 \subset \Gamma'$. The collection of flats $\mathcal{F}'$ spanned by pairs $\alpha' \in \mathcal{H}', \beta' \in \mathcal{V}'$, must lie in a single connected component of $p^{-1}(\bar{K}(\Gamma_0)) \subset K'$ because if $F, F' \in \mathcal{F}'$, then there is a chain of flats

$$\{F = F_1, \dots, F_k = F'\} \subset \mathcal{F}'$$

such that $F_i \cap F_{i+1}$ contains a quarter-plane for each $1 \leq i < k$. Thus $f(P)$ lies in a controlled neighborhood of $p^{-1}(\bar{K}(\Gamma_0))$. $\square$

Note that in the situation of Theorem 4.2, the image $f(P)$ need not lie at finite Hausdorff distance from a standard product subcomplex: consider the case when the defining graph of K is the star of a single vertex, and $P \subset K$ is a standard flat. However, maximal standard product subcomplexes are preserved:

Corollary 4.3. *There is a constant $D_1 = D_1(L, A) \in (0, \infty)$ such that if $P \subset K$ is a maximal standard product subcomplex, then*

$$\text{Hd}(f(P), P') < D$$

for some maximal standard product subcomplex $P' \subset K'$.

Proof. Let $g : K' \longrightarrow K$ be a quasi-inverse for f , with quasi-isometry constants controlled by (L, A) .

By Theorem 4.2, we know that $f(P) \subset N_D(P')$, where $P' \subset K'$ is a standard product complex; without loss of generality we may suppose that P' is a maximal standard product complex. Applying Theorem 4.2 to g , we conclude that $g(P') \subset N_D(P_1)$, where $P_1 \subset K$ is a standard product complex. However, this implies that P lies in a finite neighborhood of P_1 ; hence $P \subset P_1$ by Lemma 2.11, and by the maximality of P , we get $P = P_1$. It follows that P' lies in a controlled neighborhood of P , and hence $f(P)$ lies at controlled Hausdorff distance from a maximal standard product complex in K' . $\square$

Corollary 4.4. *The graph Γ has 4-cycles iff Γ' has 4-cycles.*

Proof. Γ has no 4-cycles iff the maximal join subgraphs of Γ are contained in stars of vertices iff K contains no maximal product subcomplex quasi-isometric to a product of two tri-valent trees. Thus by Corollary 4.3 the property of having 4-cycles in the defining graph is a quasi-isometry invariant property of RAAG's, among the class of RAAG's with triangle-free defining graphs. $\square$

Theorem 4.5. *Assume that Γ_1, Γ_2 are connected finite graphs with all vertices of valence > 1 and no cycles of length < 5 . Then there is a constant $D_2 = D_2(L, A) \in (0, \infty)$ such that if $S_1 \subset K_1$ is a standard flat and $f : K_1 \rightarrow K_2$ is an (L, A) -quasi-isometry, then there exists a standard flat $S_2 \subset K_2$ such that $f(S_1)$ and S_2 are at Hausdorff distance $\leq D_2$.*

In the proof we will need the following lemma. Note that maximal product subcomplexes in K_1, K_2 have the form $T \times \mathbf{R}$ for a 4-valent infinite tree T , and these are precisely the parallel sets of standard geodesics.

Lemma 4.6. *Let P, P' be two maximal product subcomplexes in K_1 (or K_2). Then precisely one of the following holds.*

- (1) $P = P'$,
- (2) $P \cap P'$ is a standard flat. Moreover, every standard flat can be represented as the intersection of two parallel sets.
- (3) There is no quarter-plane contained in Hausdorff neighborhoods of both P and P' .

Proof. If S is a standard flat then $S = P \cap P'$ where P, P' are parallel sets of two perpendicular standard geodesics in S . Now let P, P' be two arbitrary parallel sets. Let $\Delta = d(P, P')$ and let $Y \subset P$ and $Y' \subset P'$ be as in Lemma 2.4. If $\Delta > 0$ then Y, Y' are contained in the 1-skeleta of P, P' so (3) holds. If $\Delta = 0$ then $P \cap P'$ contains a vertex, say $p \in K_1$, and we may assume that P, P' are parallel sets of standard geodesics ℓ, ℓ' through p . The standard geodesics through p are in 1-1 correspondence with the vertices of Γ_1 . Say ℓ, ℓ' correspond to v, v' . If $v = v'$ then $P = P'$ and (1) holds. If v and v' are adjacent, then $P \cap P'$ is the standard flat that contains ℓ and ℓ' (and corresponds to the edge joining v to v'). If v and v' are not adjacent then $P \cap P' = \{p\}$ and (3) holds. $\square$

Proof of Theorem 4.5. Write $S_1 = P \cap P'$ where P, P' are parallel sets in K_1 . By Corollary 4.3, $f(P)$ and $f(P')$ are Hausdorff equivalent to parallel sets Q and Q' . We must have $Q \neq Q'$ since $P \neq P'$, and the coarse intersection of Q and Q' is a plane. Lemma 4.6 implies that $S_2 = Q \cap Q'$ is a standard flat. By Lemma 2.4, S_2 is the coarse intersection between Q and Q' , so it follows that $f(S_1)$ and S_2 are Hausdorff equivalent. $\square$

5. FLAT SPACE

In this section we discuss two more CAT(0) spaces associated with certain RAAG's – an alternate model space, and flat space. Flat space was introduced in [CD95], where it was called the modified Deligne complex. [Dav98] showed that it was a right-angled building whose apartments are modelled on the right-angled Coxeter group W with the same defining graph as the RAAG. In this section we give a more explicit description of the same object, and discuss some specific features that will be needed later.

5.1. Defining the exploded torus space and flat space. Let $G = G(\Gamma)$ be the right angled Artin group given by a graph Γ . To keep things simple, in the rest of the paper we will make the following assumption on Γ :

- Γ is connected, every vertex has valence > 1 , and there are no cycles of length < 5 .

As there are no 3-cycles, the group G is 2-dimensional, and the fact that there are no 4-cycles will guarantee that the “flat space” $\mathbb{F}$ discussed below is well-defined and Gromov hyperbolic.




FIGURE 1. Constructing the exploded torus space.

Let $\overline{K} = \overline{K}(\Gamma)$ be the presentation complex for G and let $K = K(\Gamma)$ be its universal cover. Then K is a square complex which is $\text{CAT}(0)$ thanks to our assumption that Γ has no 3-cycles. The complex K is the standard space associated to G .

We will now construct a space $X = X(\Gamma)$, the *exploded space*, on which G acts freely and cocompactly. It is somewhat more convenient to describe its quotient $\overline{X} = \overline{X}(\Gamma)$, a space whose fundamental group is G . Let Γ' denote the first barycentric subdivision of Γ . We define a singular fibration $p : \overline{X}^{(1)} \rightarrow \Gamma'$ so that:

- the fiber over each vertex of Γ' which corresponds to the midpoint of an edge of Γ is a 2-torus
- the fiber over any other point of Γ' is a circle.

The fibration has the following local structure. For any point $x \in \Gamma'$ which is not the midpoint of an edge of Γ , the local structure is the product structure. That is, there exists a neighborhood V of x such that $p^{-1}(V) \cong V \times S^1$. For a midpoint of an edge of Γ , the local structure is as follows. Let $A_1 \cong S^1 \times [0, 1)$ and $A_2 \cong S^1 \times [0, 1)$ be two half open annuli. Let T be a torus with two distinguished simple closed curves c_1 and c_2 meeting at a single point. Let Y be the quotient space of $T \sqcup A_1 \sqcup A_2$ obtained by identifying the boundary curve of A_i with c_i via a homeomorphism. Now for the midpoint x of an edge in Γ , there exists a neighborhood V of x such that $p^{-1}(V) \cong Y$, so that $p^{-1}(x) \cong T$ and for any other $y \in V$, $p^{-1}(y) = S^1 \times \{t\}$, a core circle of one of the annuli A_i .

Now note that since for any 2-torus, the two curves along which the annuli are attached meet at a single point, the singular fibration $p : \overline{X}^{(1)} \rightarrow \Gamma'$ has a section $f : \Gamma' \rightarrow \overline{X}^{(1)}$. Let $C(\Gamma')$ denote the cone

on Γ' . We now form the identification space $\overline{X} \equiv (\overline{X}^{(1)} \sqcup C(\Gamma'))/f$ obtained by attaching $C(\Gamma')$ to $\overline{X}^{(1)}$ along f .

By construction, we have that the fundamental group of $\overline{X}$ is precisely the right angled artin group G associated to Γ (simply choose the cone point o of $C(\Gamma)$ as the basepoint and for each vertex v of Γ let α_v be the loop which runs from o to v , along the edge $[o, v]$, around the circle associated to v , and back along $[o, v]$. Then one sees that the loops α_v generate the fundamental group and the relations are precisely the commutator relations dictated by Γ .)

We let X denote the universal cover of $\overline{X}$ and let $\rho : X \rightarrow \overline{X}$ denote the covering map. By $X^{(1)}$ denote $\rho^{-1}(\overline{X}^{(1)})$. Given a torus T in $\overline{X}$ which is the fiber over the midpoint of an edge of Γ , $\rho^{-1}(T)$ is a union of planes. Such planes are referred to as *standard flats*. For each edge e of Γ' , let $A_e \subset \overline{X}$ denote the annulus which is the closure of the union of fibers lying over the interior of e . A component in X of $\rho^{-1}(A_e)$ is called a *standard strip*.

We may obtain $\overline{K}$ from $\overline{X}$ by collapsing each annulus to a closed geodesic, and the cone on Γ' to a point; this defines a surjective piecewise linear map with contractible point inverses, and hence is a homotopy equivalence. Likewise we get a proper homotopy equivalence $X \rightarrow K$ by collapsing the interval factor of each standard strip to a point, and collapsing each copy of $C(\Gamma')$ in X to a point.

Let $\mathbb{F} = \mathbb{F}(\Gamma)$ denote the quotient of X obtained by collapsing each standard flat to a point and each standard strip $S \cong R \times I$ to an arc by projection onto the I -factor. We refer to $\mathbb{F} = \mathbb{F}(\Gamma)$ as the *flat space* associated to G . All standard flats in X are represented by points in $\mathbb{F}$, and we can think of $\mathbb{F}$ as obtained from the (discrete) set of standard flats by connecting the dots in just the right way to obtain a very useful space (as we shall see).

Denote by $\mathbb{F}^{(1)} \subset \mathbb{F}$ the image of $X^{(1)} \subset X$ under the quotient map. There is an induced action of G on $\mathbb{F}$; let $\pi : \mathbb{F} \rightarrow \overline{\mathbb{F}}$ be the quotient map. Thus $\overline{\mathbb{F}}$ can be constructed from $\overline{X}$ by collapsing the tori to points and collapsing the annuli to arcs, and therefore $\overline{\mathbb{F}}$ can be identified with the cone $C(\Gamma)$ on Γ . We will equip $\overline{\mathbb{F}}$ with the triangulation in which the base of the cone is the barycentric subdivision Γ' and $\overline{X} = C(\Gamma)$ is given the cone triangulation. Note that the image of $\mathbb{F}^{(1)}$ in $\overline{\mathbb{F}}$ is the base of the cone. We will also equip $\mathbb{F}$ with the triangulation obtained by pulling back via $\pi : \mathbb{F} \rightarrow \overline{\mathbb{F}}$.

There are three types of vertices in $\mathbb{F}$:

- vertices of $\mathbb{F}$ which are obtained from crushing flats to points are called *flat vertices* (in $\overline{\mathbb{F}}$ these correspond to the subdivision vertices in Γ'),
- vertices of $\mathbb{F}$ which are the cone points of copies of $C(\Gamma)$ are called *cone vertices*, and
- the remaining vertices are called *singular vertices*; these are the vertices which correspond to the original vertices of Γ

An arc joining two flat vertices which consists of two neighboring edges in $\mathbb{F}$ is called an *full edge* of $\mathbb{F}$.

The action of G on $\mathbb{F}$ is simplicial, the stabilizers of flat vertices are isomorphic to $\mathbf{Z} \oplus \mathbf{Z}$, the stabilizers of the full edges are infinite cyclic, and $\mathbb{F}/G \cong C = C(\Gamma)$. Thus this action endows G with the structure of a (simple) complex of groups, in which the underlying complex is C .

We endow X and $\mathbb{F}$ with a polyhedral metric as follows. Each torus fiber is given the structure of a flat square 2-torus, in which each of the designated curves along which fiber annuli are attached has length 1. The standard annuli are viewed as quotients of the unit square with a pair of opposite edges identified. Thus $\overline{X}^{(1)}$ has the structure of locally CAT(0) square complex. Therefore, we just need to metrize the attached cone $C(\Gamma')$. Note that each 2-simplex in $C(\Gamma')$ is of the form (o, m, v) , where o is the cone vertex, m is the midpoint of an edge and v is a vertex of Γ . We endow such a simplex (o, m, v) with the metric of an isosceles right triangle whose legs have length 1, so that the right angle is at the vertex v . It is now easy to check that the space X satisfies the CAT(0) condition (i.e. the links have no loops of length less than 2π). Note further that given an edge of Γ , the cone on Γ now consists of precisely two isosceles right triangles, forming a unit cube, so that in fact, $C(\Gamma')$ has the structure of a locally CAT(0) square complex.

Consequently, X is a CAT(0) square complex. Since $\mathbb{F}$ is built out of copies of $C(\Gamma')$, we can similarly endow $\mathbb{F}$ with the structure of a CAT(0) square complex (for a discussion of links see Observation 5.4 below). From now on, we will abuse notation slightly and refer to $C(\Gamma')$ simply as $C(\Gamma)$ or simply C .

Theorem 5.1 (No flats in flat space). *The flat space $\mathbb{F}$ is Gromov hyperbolic.*

Proof. Fix a small $\epsilon > 0$. Perturb the metric on $\mathbb{F}$ by replacing each Euclidean triangle (o, m, v) by a triangle with constant curvature $-\epsilon$ whose angles are $\pi/2$ at v , $\pi/4$ at m and $\pi/4 - \epsilon/2$ at o . Since Γ has

no cycles of length < 5 it follows that this is a $CAT(-\epsilon)$ metric on $\mathbb{F}$ which is quasi-isometric to the original metric. $\square$

To summarize what we have done so far, we have the following.

Proposition 5.2. *The RAAG G acts on $\mathbb{F}$ by isometries and the following holds:*

- (1) C has the structure of a $CAT(0)$ square complex, with each square having two singular vertices, one flat vertex and one cone vertex.
- (2) The quotient space $\mathbb{F}/G$ is C . We will denote the quotient map $\mathbb{F} \rightarrow C$ by π . We lift the labeling of the vertices of C (flat, cone, singular) to the vertices of $\mathbb{F}$ via π .
- (3) There is a fundamental domain $\tilde{C} \subset \mathbb{F}$ for the action of G such that $\pi : \tilde{C} \rightarrow C$ is a homeomorphism. We will identify $\tilde{C} = C$.
- (4) Stabilizers of flat vertices are $\mathbf{Z}^2$.
- (5) Stabilizers of singular vertices are $\mathbf{Z}$.
- (6) Stabilizers of cone vertices are trivial.

From the point of view of complexes of groups, we can say the following.

Observation 5.3.

- C is an ordered simplicial complex (after adding cone-to-flat diagonals in all squares), with cone vertex initial, and flat vertices terminal (i.e. the edges are oriented from cone to singular to flat).
- $\mathbb{F}/G$, as an orbispace, is C with the trivial label on the cone vertex, $\mathbf{Z}$ on the singular vertices, and $\mathbf{Z}^2$ on the flat vertices.

Observation 5.4.

- The link in $\mathbb{F}$ of every cone vertex is Γ .
- The link in $\mathbb{F}$ of every flat vertex is $\mathbf{Z} * \mathbf{Z}$ (the join of two infinite countable sets of vertices) and the stabilizer $\mathbf{Z}^2$ is acting in the obvious way: $(1, 0)$ translates one $\mathbf{Z}$ by 1 and fixes the other $\mathbf{Z}$, and $(0, 1)$ fixes the first $\mathbf{Z}$ and translates the second by 1. In particular, the complement of the set of all flat vertices in $\mathbb{F}$ is connected.
- The link in $\mathbb{F}$ of every singular vertex v is the join $L * \mathbf{Z}$ where L is the link of the image vertex in Γ . The stabilizer $\mathbf{Z}$ fixes L and translates $\mathbf{Z}$. The link of v in $\mathbb{F}^{(1)}$ is L .

Each singular vertex has two natural objects associated with it: the track and the singular star. We describe these below.

5.2. The singular star associated to a singular vertex. Fix a singular vertex $v \in \mathbb{F}$. Consider its star S_v in $\mathbb{F}^{(1)}$. By Observation 5.4, S_v can be identified with the star of v in (barycentrically subdivided) Γ . Thus S_v is the union of all edges with one vertex at v and the other vertex flat. We call S_v the *singular star* associated to v . Recall that the parallel set $P(\gamma)$ of a geodesic line γ in a CAT(0) space is the union of all geodesic lines parallel to γ . The vertex v represents a geodesic line in X and its parallel set will be denoted by P_v .

Observation 5.5. *Recall that $p : X \rightarrow \mathbb{F}$ is the quotient map. Then $P_v = p^{-1}(S_v)$. Conversely, every parallel set in X of a standard line (i.e. a line in a standard strip) arises in this way for a suitable singular vertex v . Moreover, the parallel set is the union of flats and strips that coarsely contain the given standard line. Abstractly, a parallel set in X is isomorphic to $T \times \mathbf{R}$ where T is the universal cover of the 1-complex obtained from S_v by wedging a circle to every vertex other than v .*

The image of a parallel set in X under the quotient map $X \rightarrow \mathbb{F}$ will be referred to as a *parallel set* in $\mathbb{F}$ (even though it is not the union of lines parallel to a fixed line). It is a subcomplex of $\mathbb{F}^{(1)}$, and the quotient by the stabilizer is S_v .

Observation 5.6. *The intersection between $C(\Gamma) \subset \mathbb{F}$ and a nontrivial translate of it is contained in a singular star.*

5.3. The track associated to a singular vertex. Again let $v \in \mathbb{F}$ be a singular vertex. Let $e_1, \dots, e_k$ be the edges that have one endpoint at v and the other endpoint flat (thus the singular star S_v is the union of the e_i 's). Then $C(\Gamma)$ contains k squares that have v as a vertex. These squares can be enumerated $S_1, \dots, S_k$ so that e_i is one of the sides of S_i , $i = 1, \dots, k$. Note that by construction, these squares share a common edge e that is incident to v and whose other endpoint is the cone point, so that e is distinct from the edges $e_1, \dots, e_k$. Let τ_v be the hyperplane in $C(\Gamma)$ transverse to e ; this is well-defined since $C(\Gamma)$ is a square complex, see Section 2.2. This is the *track associated to v* .

Observation 5.7. *Recall that $\pi : \mathbb{F} \rightarrow C(\Gamma)$ is the quotient map.*

- $\pi^{-1}(\tau_v)$ is a collection of hyperplanes in $\mathbb{F}$. Each hyperplane is a convex subset of $\mathbb{F}$.




FIGURE 2. Γ is the pentagon. Also pictured is the track and the singular star (pictured in black) corresponding to a singular vertex




FIGURE 3. Γ is the tripod. Also pictured is the track and the singular star corresponding to the central singular vertex

- Conversely, any hyperplane in $\mathbb{F}$ is a component of $\pi^{-1}(\tau_v)$ for a suitable singular vertex v .

Also denote by R_i the closure of the component of $S_i - \tau_v$ that contains e_i . Define the *thickened track associated to v* to be

$$R_v := R_1 \cup \cdots \cup R_k$$

Observation 5.8.

- R_v is a product neighborhood of P_v , that is, $(R_v, P_v) \cong (P_v \times I, P_v \times \{0\})$.

- Let $\tilde{R}_v$ be a component of $\pi^{-1}(R_v)$. Then $\tilde{R}_v$ contains a unique component $\tilde{P}_v$ of $\pi^{-1}(P_v)$. Moreover, $(\tilde{R}_v, \tilde{P}_v) \cong (\tilde{P}_v \times c(\mathbf{Z}), \tilde{P}_v \times \{c\})$, where $c(\mathbf{Z})$ is the cone on $\mathbf{Z}$ with cone point c .
- Under this identification, the pointwise stabilizer $\mathbf{Z}$ of $\tilde{P}_v$ acts on $\tilde{R}_v \cong \tilde{P}_v \times c(\mathbf{Z})$ by fixing $\tilde{P}_v$ and the cone point, and translating the base of the cone.

Lemma 5.9.

- Every parallel set in $\mathbb{F}$ separates $\mathbb{F}$. The set of complementary components is acted on freely and transitively by the $\mathbf{Z}$ subgroup fixing the parallel set pointwise.
- For any two distinct cone vertices in $\mathbb{F}$ there is a parallel set that separates them.

Proof. A parallel set $P_v \subset \mathbb{F}$ has a product neighborhood $P_v \times c(\mathbf{Z})$ as described in Observation 5.8. It follows from the simple connectivity of $\mathbb{F}$ that the connected components of $\mathbb{F} \setminus P_v$ are in bijective correspondence with the connected components of $(P_v \times c(\mathbf{Z})) \setminus P_v$; the latter are in bijective correspondence with the $\mathbf{Z}$ subgroup fixing P_v pointwise. For the second assertion, consider the geodesic joining the two vertices. This geodesic must cross some hyperplane, and hence some component of some $\pi^{-1}(\tau_v)$ (see Observation 5.7). The associated parallel set separates between the two cone vertices (see Observation 5.8). $\square$

5.4. Coarse distance and parallel sets. We will be interested in paths and loops in $\mathbb{F}^{(1)}$; these will correspond to sequences of standard flats in X . There is a straightforward notion of distance between standard flats. Suppose that F and F' are standard flats. The flats F and F' correspond to two vertices in $\mathbb{F}$. Note that since $\mathbb{F}^{(1)}$ is a bipartite graph, the usual path distance in $\mathbb{F}^{(1)}$ between F, F' is even; let $D(F, F')$ denote half this path distance in $\mathbb{F}^{(1)}$. We will be interested in a somewhat different notion of distance on $\mathbb{F}^{(1)}$, called *coarse distance*, which we now describe. We say that the coarse distance between F and F' is 1, $D^\infty(F, F') = 1$, if they coarsely intersect in a line, or equivalently, if belong to the same parallel set (cf. Lemma 2.2). We will say that $D^\infty(F, F') = n$, if there exists a sequence of standard flats $F = F_0, \dots, F_n = F'$ such that $D^\infty(F_i, F_{i+1}) = 1$ and n is the smallest number for which there exists such a sequence. Note that this then defines a metric on the set of flat vertices of $\mathbb{F}$. The reason this metric is natural is that it is preserved by quasi-isometries: the binary relation

on $\mathbb{F}^{(0)}$ of belonging to the same parallel set is quasi-isometry invariant, by Lemma 2.2 and Theorem 4.5.

The metric D^∞ can be seen directly in $\mathbb{F}$ as follows. Suppose that e and f are two full edges of $\mathbb{F}$ meeting at a flat vertex v . We say that e and f define a *legal turn at v* if e and f have cyclic stabilizers with trivial intersection. Otherwise, they have the same cyclic stabilizer and we say that they define an *illegal turn at v* . Now suppose that v and w are two flat vertices of $\mathbb{F}$. Suppose that α is a full edge path in $\mathbb{F}$ joining v and w . Then the *coarse length* of α , $length_\infty(\alpha)$ is computed by counting the number of legal turns along α and adding 1. Thus, the coarse distance between vertices is simply the minimal coarse length of a path between them.

If a path in $\mathbb{F}^{(1)}$ has no legal turns in it, then in fact all the flat vertices along it correspond to flats which are all contained in the same parallel set. We call such a path a *stalling path*. The union of all stalling paths containing a given singular vertex is thus a parallel set in $\mathbb{F}$; it corresponds in X to the union of all flats and strips coarsely containing a given singular line.

6. DUAL DISK DIAGRAMS

We consider $\mathbb{F}$ as a CAT(0) square complex. Let $\mathcal{H}$ denote the union of hyperplanes in $\mathbb{F}$. We call a component of $\mathbb{F} - \mathcal{H}$ a *block*. Notice that each block contains a unique vertex of $\mathbb{F}$. Suppose that α is a closed full edge path in $\mathbb{F}^{(1)}$. Then $\alpha : S^1 \rightarrow \mathbb{F}^{(1)}$ extends to a map $\Delta : D^2 \rightarrow \mathbb{F}$. Now we may make Δ transverse to the hyperplanes of $\mathbb{F}$. Thus $\mathcal{A} = \Delta^{-1}(\mathcal{H})$ is a union of embedded arcs and simple closed curves in $D = D^2$, where each arc and closed curve is a component of the preimage of a single hyperplane. As in [Sag95], one may assume that there are in fact no simple closed curves and that each pair of arcs meets in at most a single point. Since $\mathbb{F}$ is a 2-complex, we may further assume that there are no triangular regions: every collection of three arcs contains a disjoint pair. The pair $(D, \mathcal{A})$ is called a *dual disk diagram* for α . Sometimes we will abuse notation and just refer to D as the dual disk diagram. In [Sag95] it is shown that one can get from any diagram to any other through a sequence of triangular moves. Since there are no triangles, we see that the dual disk diagram is unique.


A *region* of D refers to the closure of a component of $D - (\partial D \cup \mathcal{A})$. Note that each region D is associated to a vertex of X , namely the vertex whose block the region is mapped to. By a *boundary region*

we mean a region that intersects ∂D . Regions that are not boundary regions are called *internal* regions. The union of all internal regions is called the *core* of the diagram. By a *corner* of D we mean a triangular region bounded by two arcs of $\mathcal{A}$ and a subarc of ∂D . A corner corresponds to α running around the corner of a square of $\mathbb{F}$. Since α is a path of full edges, this means that the corner of the square is a flat vertex and the turn at the vertex is a legal turn.

We now need to discuss some complications which can occur in general disk diagrams, but which will not appear in our setting, because our boundary map α may be assumed to be an embedding.

6.1. Spurs. A *spur* of D consists of a collection of nested arcs, each disjoint from all other arcs, so that one of the arcs has endpoints on neighboring boundary edges (this arc bounds a region with no curves or arcs in it) ; see figure below.

In this case, we see that the original full edge path has some backtracking. Since we will always be dealing with paths that have no backtracking, we can assume that there are no spurs.




6.2. Separating cells. Suppose that the intersection $P \cap \partial D$ of a closed cell P with the boundary of D has two distinct connected arc components γ, γ' , so that P separates D . Then the path represented by ∂D is not embedded. To see this, note that since ∂D is a path of full edges, γ and γ' must each contain a vertex, and these vertices must map to the same vertex in $\mathbb{F}$. Because we will be dealing with embedded paths, we can assume that there are no separating cells.


The upshot of the above is that if the edge loop is embedded, then the disk diagram associated to it has no spurs or separating cells. An embedded full edge loop will be called a *cycle*.

Recall that there are three types of vertices: singular, flat, and cone vertices. Since each region of D is mapped to a block, which is uniquely associated to a vertex of $\mathbb{F}$, we have three types of regions, which we also call singular, flat and cone regions.

Observation 6.1. *Every square in $\mathbb{F}$ has two diagonally opposite singular vertices, one cone vertex and one flat vertex.*



Below is an example of a dual disk diagram. The yellow cells are the boundary cells and the remaining cells form the core.



Observation 6.2. *Since corners cannot contain singular vertices, all corner regions are flat regions. Also, all cone regions are interior.*

Observation 6.3. *Observation 6.1 can be transported to the disk diagram to tell us what types the regions are; around each intersection point of two arcs in the diagram, we have two diagonally opposite singular regions, one cone region and one flat region. In particular, this tells us that (under our standing assumption that our paths are embedded) the core is always not empty.*

Observation 6.4. *If A is an arc of the disk diagram, then by Observation 6.1, we see that the regions immediately to one side of A alternate singular and cone regions, while the other side of A alternate singular and flat vertices. The latter sequence of regions gives us a stalling path in $\mathbb{F}^{(1)}$. Thus the vertices associated to this sequence of regions all lie in the same parallel set.*

6.3. Existence of 1-shells or 2-shells (by the seashore). Following McCammond and Wise, we define a boundary region of a disk to be an i -shell if it has i internal edges. In the example above, the core of the diagram has four 2-shells. Since the core of the diagram is a $C(4)$ - $T(4)$ complex, we can apply Greendlinger's lemma. We give a slightly different statement suited to our needs and we for the sake of completeness, we include a proof here (for deeper delvings into such results see [MW02]).

Lemma 6.5. *Let D be a disk tiled by polygons $P_1, \dots, P_k$ in such a way that every interior vertex is incident to at least 4 polygons and so that each polygon has at least 4 sides. Then one of the following holds.*

- (1) $k = 1$ (i.e. there is one 0-shell)

- (2) *there are at least two 1-shells*
- (3) *there is at least one 1-shell and at least two 2-shells*
- (4) *there are at least four 2-shells*

Proof. A *corner* of P_i is an unordered pair of adjacent edges. Thus an n -gon has n corners (assuming $n \geq 3$). We say that a corner of P_i is contained in ∂D if the interiors of the corresponding two edges do not intersect any other P_j 's. Denote by n_i the number of sides of P_i and by C_i the number of corners of P_i contained in ∂D .

View each P_i as a square complex with n_i squares so that curvature is concentrated at one interior vertex. The excess angle at this vertex is $(\pi/2)n_i - 2\pi$. Thus the total interior excess is at least

$$(\pi/2) \sum n_i - 2k\pi$$

(it might be larger if there are interior verices incident to > 4 polygons).

The boundary deficit is $\pi/2$ for each corner, i.e. at most

$$(\pi/2) \sum C_i$$

(it might be smaller if there are boundary vertices incident to > 2 polygons). Recall that Gauss-Bonnet says:

$$(\text{boundary deficit}) - (\text{interior excess}) = 2\pi$$

Thus

$$(\pi/2) \sum C_i - (\pi/2) \sum n_i + 2k\pi \geq 2\pi$$

and hence

$$\sum C_i \geq 4 + \sum (n_i - 4)$$

The statement now follows quickly. To each P_i assign the score of $C_i - n_i + 4$ points. The inequality says that the sum of all points is ≥ 4 . If P_i is assigned 4 points, it is a 0-shell. If it is assigned 2 points it is a 1-shell, and if it is assigned 1 point it is a 2-shell. (An interior triangle would get 1 point but we are assuming there are no triangles. All other polygons get a nonpositive number of points.) $\square$

Remark 6.6. Suppose that there are precisely two 1-shells and no 2-shells. Then the polygons can be renumbered so that P_1 and P_k are 1-shells and P_i and P_j share an edge iff $|i - j| \leq 1$. In other words, the polygons form a *ladder*.

6.4. Shells, short cuts and taut cycles. A cycle α is said to have an i -cut if there are flat vertices v, w on α such that the coarse length of any path along the cycle between v and w is greater than i , but for which there exists full edge path of length i in $\mathbb{F}^{(1)}$ joining v and w . A cycle is said to be *taut* if there exists no i -cut with $i \leq 2$.

Now suppose that D is a disk diagram for α and let D' be its core. If D' is not a single cell, then by Lemma 6.5, D' has a 1-shell or a 2-shell. As we see in the figure below, we then obtain a 1-cut or a 2-cut for α .




FIGURE 4. 1-cuts and 2-cuts arising from shells.

Thus we obtain the following lemma.

Lemma 6.7. (*Diagrams for taut cycles*) Suppose that α is a taut cycle. Then the diagram for α has a core consisting of a single cell.

7. RIGIDITY OF TAUT CYCLES

The aim of this section is to prove the the quasi-isometric preservation of taut cycles. Recall that if $\phi : X_\Gamma \rightarrow X_\Gamma$ is a quasi-isometry, Theorem 4.5 tells us that standard flats are coarsely preserved by ϕ . Since no two standard flats are coarsely equivalent, ϕ induces a bijection $\phi_\#$ on the flat vertex set of $\mathbb{F}$. Moreover, since coarse equivalence is preserved by quasi-isometries, we have that D^∞ is preserved.

Theorem 7.1 (Taut cycle rigidity). Suppose that ϕ and $\phi_\#$ are as above. Then $\phi_\#$ carries taut cycles to taut cycles. In particular, full edges that lie along taut cycles are carried to full edges.

We first prove a lemma that tells us that “quasi” cuts of length at most 3 actually give 1-cuts or 2-cuts.

Lemma 7.2 (quasi-cuts yield cuts). Let α be a cycle. Suppose that there exist non-adjacent flat vertices v and w along α , subdividing α into two paths α_1 and α_2 . Suppose further that v and w are joined by a path β so that

- (1) $\text{length}_\infty(\beta) \leq 3$
- (2) v and w are the only flat vertices of α which lie in β .

Then α is not taut.

Proof. We give the argument when $\text{length}_\infty(v, w) = 3$. The argument when $\text{length}_\infty(v, w) < 3$ is similar and indeed simpler. So we suppose that β is broken up by two flat vertices p and q into three stalling subpaths β_1, β_2 and β_3 as in Figure 5





FIGURE 5. A dual disk diagram for a stalling cut.

We let D denote a dual disk diagram for the loop $\gamma = \beta \cup \alpha_1$. Note that by assumption, this loop is indeed a cycle. We will apply Lemma 6.5 to D . Now we must consider some cases.

Case 1: β_1 and β_3 are both longer than 1. Since there are no corners along stalling sections of β , it follows that D' is not a single cell so we are in cases 2,3, or 4 of Lemma 6.5. Now a shell gives at least two consecutive corners along ∂D . It follows that if a shell produces corners along β , it is either a sequence of corners that begins at v or w and continues into α_1 or the sequence of corners is simply the pair of corners at p and q (in this case, β_2 is of length 1).

There exist 1-shells (case 2 or 3 of Lemma 6.5). Let us suppose we have a 1-shell in D' . Now a 1-shell produces a sequence of corners of length at least three, so it cannot occur at p and q . Also, both of the endpoints of the 1-cut produced by this shell cannot be along the non-stalling section of D , for otherwise α would not be taut. Thus, one of the endpoints of the 1-cut must be along the stalling sections β_1 or β_3 . But since these stalling section can only have a corner at v and w , the 1-cut provides a 2-cut between v or w and some vertex z in the non-stalling section (see Figure 6). Thus α is not taut and we are done.

FIGURE 6. A 1-cut yielding a 2-cut for α .

There exist no 1-shells (case 4 of Lemma 6.5). In this case there are at least four 2-shells. Since only two of these can produce corners at the vertices v and w , and only one can produce corners at p and q , we must have a 2-shell which produces a 2-cut in the non-stalling section, namely along α_1 . This produces a 2-cut for α and we are done.

Case 2: Only one of β_1, β_3 has length 1. Assume without loss of generality that β_3 has length 1, and β_1 is longer. We proceed as in Case 1. If there exist 1-shells, there are at least two 1-shells or a 1-shell and two 2-shells. It follows that one of them cannot be along $\beta_1 \cup \beta_2$. As before, these shells cannot occur entirely along α_1 . Thus we are in the case in which there are only two i -shells, $i \leq 2$; that is, we have two 1-shells. One of these must have its first corner at v and we get a 2-cut as in Figure 6.

If we have only 2-shells, then there are 4 of them, and hence one of them must occur entirely along the α_1 , producing a 2-cut for α_1 , contradicting the tautness of α .

Case 3: both β_1 and β_3 have length 1. Now if D does not have a corner at both v and w , then we may proceed as before. So suppose that there are corners both at v and w . We now return to our original loop $\alpha = \alpha_1 \cup \alpha_2$. Let v_1 be the flat vertex immediately adjacent to v along α_1 and let v_2 be the flat vertex immediately adjacent to v along α_2 . Let v_3 be the vertex along β immediately adjacent to v . We define w_1, w_2 and w_3 similarly as neighboring vertices of w . Now since there is a corner at v , it follows that the path $[v_3, v, v_1]$ is non-stalling. Thus, the path $[v_3, v, v_2]$ is a stalling path. Similarly, the path $[w_3, w, w_2]$ is a stalling path. (See Figure 7.)

We thus let α'_2 be the subarc of α_2 spanned by v_2 and w_2 , and let β' be $\beta \cup [v, v_2] \cup [w, w_2]$. We then consider a dual disk diagram for




FIGURE 7. One of the paths of length 2 running through v must be stalling.

$\alpha_2 \cup \beta'$ and we are back in Case 1 with α_1 replaced by α'_2 and β replaced by β' . $\square$

Proof of Theorem 7.1. Suppose that $v_1, \dots, v_n$ are the flat vertices along a taut cycle α . Let $w_i = \phi_{\#}(v_i)$. We know that $D^{\infty}(w_i, w_{i+1}) = 1$, so that we can choose a full edge geodesic β_i joining w_i to w_{i+1} , such that flat vertices along it lie in the same parallelism class. We then join the β_i 's together to get a closed loop β . We claim that this loop is a cycle. First of all, by construction, the β_i 's are embedded, and at the w_i 's we have turns, so that there is no backtracking. Secondly, if distinct β_i 's met along a vertex which was not one of the endpoints, then we could apply Lemma 7.2 to conclude that α is not taut. So now consider a dual disk diagram D for β and as usual let D' denote the core of D . We claim that D' consists of a single cell.

If D' has more than one cell, we would then have an arc a in D which separates D into two regions, each of which contains a 1-shell or 2-shell of D' . Now by Observation 6.4, there would be a stalling path joining non-adjacent vertices z_1 and z_2 of β . Now if we have $z_1 = w_i$ and $z_2 = w_j$, for some i, j , then we have that $D^{\infty}(v_i, v_j) = 1$ and we apply Lemma 7.2 to conclude that α was not taut. Now if $z_1, z_2 \notin \{w_1, \dots, w_n\}$, then let w_i be a closest such vertex to z_1 and w_j be a closest such vertex to z_2 . Letting $z'_1 = \phi_{\#}^{-1}(z_1)$ and $z'_2 = \phi_{\#}^{-1}(z_2)$. We see that the path $\gamma = [v_i, z'_1, z'_2, v_j]$ satisfies the conditions of Lemma 7.2. We can apply Lemma 7.2 in a similar manner when only one of the z_i 's is in $\{w_1, \dots, w_n\}$.

We thus have that D' consists of a single cell. This means that in fact the w_i define a closed edge loop β . Now we wish to show that this loop is taut. Suppose not. Then we have a 1-cut or a 2-cut of β .

Pulling such a cut back via $\phi_{\#}^{-1}$ gives a quasi-cut as in Lemma 7.2 thus implying that α is not taut, a contradiction. $\square$

8. $\mathbb{F}$ -RIGIDITY

We next need to argue that if F_1, F_2 are two standard flats in $\mathbb{F}$ that intersect in a line, then their images $\phi_{\#}(F_1), \phi_{\#}(F_2)$ intersect in a line as well. *A priori*, we only know that they intersect coarsely in a line, i.e. that they belong to the same parallel set. If there is a taut cycle that crosses both F_1 and F_2 then the fact that $\phi_{\#}(F_1) \cap \phi_{\#}(F_2)$ is a line follows from taut cycle rigidity. However, there are pairs of flats where there are no such taut cycles. This is of course the case if the graph contains valence 1 vertices, but there are no such vertices by our standing assumption on Γ . For a more subtle failure see the example below.

8.1. Lemmas about graphs.

Definition 8.1. Let γ be an embedded cycle in the defining graph Γ . An *i-shortcut* is an edge-path β in Γ of length i whose endpoints are in γ and whose distance in γ is $> i$.

Definition 8.2. A cycle γ is *tight* if it does not admit any 1-shortcuts nor any 2-shortcuts.

Remark 8.3. If γ is a cycle in Γ , we can lift it to a cycle $\hat{\gamma}$ in the barycentric subdivision $\Gamma' \subset \mathbb{F}$. Then γ is tight in Γ iff $\hat{\gamma}$ is taut in $\mathbb{F}$.

Example 8.4. Let Γ be the atomic graph obtained from the 1-skeleton of a dodecahedron by doubling along one of the pentagons representing a face. Then every tight cycle is contained in one of the two copies of the dodecahedron, so that a pair of edges that share a common vertex but are not contained in the same dodecahedron give rise to flats that intersect but are not both crossed by a tight cycle.

Definition 8.5. Let v be a vertex of Γ and $Lk(v)$ the link of v in Γ . The *Whitehead graph* at v is the graph $Wh(v)$ whose vertex set is $Lk(v)$ and $e, e' \in Lk(v)$ (represented by edges with initial vertex v) are connected by an edge iff there is a tight cycle in Γ that enters v along $\bar{e}$ and leaves along e' .

In the above example, the Whitehead graph at every valence 3 vertex is the complete graph (i.e. a triangle), while at the valence 4 vertices it is obtained from the complete graph by removing an edge (i.e. it is a square with one diagonal).

Lemma 8.6. *$Wh(v)$ is connected iff v is not a cut vertex.*

Proof. It is clear that $Wh(v)$ is disconnected if v is a cut vertex. For the converse, suppose $Wh(v)$ is disconnected, and let $e_1, \dots, e_p \in Lk(v)$ be the vertices of one of the components of $Wh(v)$, and let $f_1, \dots, f_q$ be all the other vertices of $Lk(v)$. Suppose v is not a cut vertex. Then there are embedded cycles in Γ that enter v along some $\bar{e}_i$ and leave along some f_j . Let C be a shortest such cycle. We claim that C is tight. To see this, suppose β is an i -shortcut for C for $i \leq 2$. Using the fact that Γ has no 4-cycle, it follows that one may replace a subpath of C with β to obtain a shorter cycle C' which contains some $\bar{e}_i$ and some f_j . Thus some e_i is connected to some f_j in $Wh(v)$, a contradiction. $\square$

Lemma 8.7. *An atomic graph contains no cut vertices, or separating closed edges.*

Proof. Let $Y \subset \Gamma$ be either a vertex or a closed edge, and suppose $v \in Y$ is a vertex. Note that if $C \subset \Gamma \setminus Y$ is a connected component, then C cannot be contained in the closed star of v , because it would then lie in a single edge of Γ , which clearly contradicts the assumption that every vertex has valence at least 2. $\square$

Lemma 8.8. *Suppose the edges of an atomic graph Γ are colored in three colors, black, white, and gray, and the following holds:*

- (1) *any two gray edges share a vertex, and*
- (2) *the edges of any tight cycle are either all black or gray, or they are all white or gray.*

Then the edges of Γ are either all black or gray, or else they are all white or gray.

Proof. Suppose not. Thus there is an edge b colored black, and there is an edge w colored white. All gray edges have a vertex v in common (by assumption (1) and because there are no triangles). Let G be the union of all (closed) gray edges. We now claim that G does not separate Γ . Indeed, take $x, y \in \Gamma - G$. If $x, y \notin St(v)$ then x, y can be joined by a path missing $St(v)$ and hence missing G . Say $x \in St(v) - G$. If x is a vertex, then it is adjacent to v and it is incident to an edge e different from $[x, v]$ (since otherwise $[x, v]$ would be a separating edge). Thus x can be joined by a path in the complement of G to the complement of $St(v)$. If x is not a vertex, then it is an interior point of some edge $[x', v]$ and it can again be joined by a path in the complement of G , first to x'

and then to the complement of $St(v)$ as before. Similar considerations apply to y and our claim is proved.

Now join b and w by a path missing G . Thus all edges along this path are black or white, and there is a vertex z along this path where color changes from black to white. Therefore all edges incident to z are black or white and at least one is black and at least one is white. By Lemma 8.6 there is a tight cycle passing through z and at z crossing one white and one black edge. But this contradicts our assumption (2). $\square$

Lemma 8.9. *Let Γ, Γ' be graphs with no vertices of valence < 2 and no cycles of length < 4 . Suppose $F : \mathcal{E}(\Gamma) \rightarrow \mathcal{E}(\Gamma')$ is a bijection of the sets of edges of Γ, Γ' such that*

- *if $e_1, e_2 \in \mathcal{E}(\Gamma)$ share a vertex then $F(e_1), F(e_2) \in \mathcal{E}(\Gamma')$ share a vertex.*

Then there is a graph isomorphism $\phi : \Gamma \rightarrow \Gamma'$ that induces F .

Proof. Let v be a vertex of Γ . Consider the set of all edges $e_1, \dots, e_k$ in Γ that contain v . By our assumptions on Γ , $k \geq 2$ and any two of the edges $F(e_1), \dots, F(e_k)$ share a vertex. Since Γ' has no cycles of length < 4 it follows that there is a unique vertex $w \in \Gamma'$ contained in all $F(e_i)$. Define $\phi(v) := w$. By construction, if v_1, v_2 span an edge then $\phi(v_1), \phi(v_2)$ span an edge. Thus $\phi : \Gamma \rightarrow \Gamma'$ is a simplicial map. Reversing the roles of Γ, Γ' provides a simplicial map $\phi' : \Gamma' \rightarrow \Gamma$ and it easy to see that $\phi\phi'$ and $\phi'\phi$ are identity. $\square$

8.2. $\mathbb{F}$ -rigidity. The following is the main theorem in the paper. It says that the flat space $\mathbb{F}$ is a quasi-isometry invariant of the RAAG G . The notation should be self-explanatory, e.g. $\mathbb{F}_1 = \mathbb{F}(\Gamma_1)$.

Theorem 8.10. *Let Γ_1, Γ_2 be two atomic graphs. Let $f : G_1 \rightarrow G_2$ be a quasi-isometry between the associated RAAGs. Then there is a label-preserving isometry $\phi : \mathbb{F}_1 \rightarrow \mathbb{F}_2$ between the associated flat spaces such that for each standard flat (i.e. a flat vertex) v we have that $f(v)$ is contained in a Hausdorff neighborhood of the flat $\phi(v)$.*

In particular, ϕ takes each cone in $\mathbb{F}_1$ isometrically to a cone in $\mathbb{F}_2$, and $\Gamma_1 \cong \Gamma_2$.

Proof. By Theorem 4.5 we have a function $\phi : \mathbb{F}_1^b \rightarrow \mathbb{F}_2^b$ defined on the set of flat vertices and satisfying the statement about the Hausdorff neighborhood. This map is a bijection. We now extend this map to an isometry.

Step 1. Let $C(\Gamma_1) \subset \mathbb{F}_1$ be the fundamental domain (see Proposition 5.2(2)). We claim that ϕ takes all flat vertices in $\Gamma_1 \subset C(\Gamma_1)$ into a translate of the fundamental domain $C(\Gamma_2) \subset \mathbb{F}_2$. By Theorem 7.1 we know that ϕ sends all flat vertices along a taut cycle in $\Gamma_1 \subset C(\Gamma_1)$ into a single cone (and this cone is unique, see Observation 5.6). Since every flat vertex in Γ_1 lies along a taut cycle it suffices to prove that there is a cone that contains the images of the flat vertices along any taut cycle in $\Gamma_1 \subset C(\Gamma_1)$.

So suppose this is not true, and that for two different taut cycles in Γ_1 the images of the flat vertices are contained in different cones $C', C'' \subset \mathbb{F}_2$. Choose a parallel set $P \subset \mathbb{F}_2$ that separates between C' and C'' (see Lemma 5.9) so that we can write $\mathbb{F}_2 = X' \cup X''$ with $C' \subset X', C'' \subset X''$ and $X' \cap X'' = P$. Now color the edges (i.e. the flat vertices) of Γ_1 into black, white or gray according to whether they are mapped into X', X'' or $X' \cap X''$. Note that any two gray edges share a vertex (the corresponding flats coarsely intersect in a line; for edges of Γ_1 this happens only when the edges share a vertex). Now by Lemma 8.8 we see that all flat vertices in Γ_1 are mapped to either X' or to X'' contradicting our choices.

Thus, after composing with an element of G_2 , we may assume that ϕ maps the flat vertices in $C(\Gamma_1)$ to flat vertices in $C(\Gamma_2)$. By considering the inverse map, ϕ restricts to a bijection $\mathcal{E}(\Gamma_1) \rightarrow \mathcal{E}(\Gamma_2)$ between the sets of edges of Γ_1 and Γ_2 .

This analysis can be applied to any translate of $C(\Gamma_1)$. As a result of Step 1, we can extend ϕ to the cone vertices so that it induces a bijection between the sets of cone vertices and any adjacent pair of a flat and a cone vertex maps to an adjacent pair.

Step 2. Now consider ϕ restricted to the “base of the cone”, that is, the map $\mathcal{E}(\Gamma_1) \rightarrow \mathcal{E}(\Gamma_2)$. Apply Lemma 8.9 to deduce that ϕ extends to an isomorphism $\phi : \Gamma_1 \rightarrow \Gamma_2$. Adjacent edges map to adjacent edges since only in this situation the corresponding flats coarsely intersect in a line, so the lemma applies.

Step 3. It remains to extend ϕ to the singular vertices. Let v be a singular vertex, say in $C(\Gamma_1)$. Then there is a unique way to define $\phi(v)$ so that ϕ preserves adjacency with the flat vertices in $C(\Gamma_1)$ (this is Step 2). We need to verify that when we regard v as a vertex in a different cone we obtain the same $\phi(v)$. There are $\mathbf{Z}$ cones that contain v (see Proposition 5.2) and all these cones also contain all flat vertices adjacent to v , so the claim follows. $\square$

9. THE PROOFS OF COROLLARIES 1.7, 1.8, AND 1.9

Proof of Corollary 1.7. Suppose Γ and Γ' are atomic graphs, and $\phi : K \rightarrow K'$ is a quasi-isometry. Let $\phi_0 : V \rightarrow V'$ be the bijection provided by Theorem 1.6. Pick $p \in K$ and let $p' := \phi_0(p)$. Then ϕ_0 establishes a bijection between the standard geodesics (respectively standard flats) passing through p and the standard geodesics (respectively standard flats) passing through p' . This induces an isomorphism $\Gamma \rightarrow \Gamma'$ of the defining graphs. $\square$

Proof of Corollary 1.8. Let $\psi : G \rightarrow G'$ be an isomorphism. Then we obtain an isometric action $G \curvearrowright K'$, which is discrete and cocompact. Hence by the fundamental lemma of geometric group theory, there is a G -equivariant quasi-isometry

$$\phi : K \longrightarrow K'.$$

Let $\phi_0 : V \longrightarrow V'$ be the bijection furnished by Theorem 1.6.

Consider a standard flat $F = \alpha \times \beta \subset K$. By Theorem 1.6, there is a standard flat $F' = \alpha' \times \beta' \subset K'$ such that ϕ_0 maps $F \cap V$ bijectively to $F' \cap V'$, and respects the product structures. Since $\phi_0|_{F \cap V}$ is equivariant with respect to the stabilizer of F , it follows that the induced mappings $(\alpha \cap V) \rightarrow (\alpha' \cap V')$ and $(\beta \cap V) \rightarrow (\beta' \cap V')$ are also $\text{Stab}(F)$ -equivariant. Since $\text{Stab}(F)$ acts transitively on $\alpha \cap V$, $\beta \cap V$, $\alpha' \cap V'$, and $\beta' \cap V'$, the map $\phi_0|_{F \cap V}$ is the restriction of an isometry $F \rightarrow F'$. Hence ϕ_0 is the restriction of an isometry $K \rightarrow K'$. $\square$

Proof of Corollary 1.9. Let $\bar{\psi} : \bar{K} \rightarrow \bar{K}'$ be a homotopy equivalence. Then $\bar{\psi}$ induces an isomorphism $G \simeq G'$ and we may lift $\bar{\psi}$ to a quasi-isometry

$$\psi : K \longrightarrow K'$$

which is G -equivariant (where we identify G and G' using the isomorphism above). By Corollary 1.8 there is a G -equivariant isometry $\phi : K \rightarrow K'$ at bounded distance from ψ , and this descends to an isometry

$$\bar{\phi} : \bar{K} \longrightarrow \bar{K}'$$

which is homotopic to $\bar{\psi}$. $\square$

10. FURTHER IMPLICATIONS OF THEOREM 1.6

At first sight, one might think that the map $\phi_0 : V \rightarrow V'$ in the conclusion of Theorem 1.6 must be the restriction of an isometry, since it preserves so much structure; however, this turns out not to be the case. In this section we single out part of the structure of the bijection ϕ_0 which efficiently distinguishes between quasi-isometries, namely we associate with each parallel set $P \subset K$ a biLipschitz homeomorphism of a copy of the integers. In the next section we will see that any biLipschitz homeomorphism can arise this way.

Let $\phi : K \rightarrow K'$ be a quasi-isometry, where K and K' are associated with atomic RAAG's, and let $\phi_0 : V \rightarrow V'$ be the map of Theorem 1.6.

Let $\mathcal{P}$ and $\mathcal{P}'$ denote the collections of maximal standard product subcomplexes in K and K' , respectively. Since K and K' are atomic, each $P \in \mathcal{P}$ or $P' \in \mathcal{P}'$ is the parallel set for a standard geodesic. Therefore P splits isometrically as a product of complexes $P = \mathbf{R} \times T$.

Definition 10.1. For each parallel set $P \in \mathcal{P}$, we define $\mathbf{R}_P$ to be the $\mathbf{R}$ -factor in the splitting

$$P = \mathbf{R} \times T,$$

and let $\mathbf{Z}_P$ to be the set of vertices of $\mathbf{R}_P$, equipped with the induced metric.

Lemma 10.2. *For every $P \in \mathcal{P}$ there is a $P' \in \mathcal{P}'$ such that ϕ_0 maps $V \cap P$ bijectively to $V' \cap P'$, preserving the product structure.*

Proof. Let $\gamma \subset P$ be a standard geodesic parallel to the $\mathbf{R}$ -factor of P . Then Theorem 1.6 implies that $V \cap \gamma$ is mapped bijectively by ϕ_0 to $V' \cap \gamma'$, for some standard geodesic $\gamma' \subset K'$. Since any two geodesics γ'_1, γ'_2 obtained this way are parallel, it follows that $\phi_0(P \cap V)$ is contained in a parallel set. Applying the same reasoning to the inverse implies that $\phi_0(V \cap P) = V' \cap P'$ for some $P' \in \mathcal{P}'$.

Since ϕ_0 preserves standard geodesics, it follows that $\phi_0|_{V \cap P}$ preserves product structure. $\square$

By abuse of notation we use ϕ to denote the induced bijection $\mathcal{P} \rightarrow \mathcal{P}'$ given by Lemma 10.2. By the lemma, for each $P \in \mathcal{P}$, we obtain a bijective quasi-isometry $\phi_{\mathbf{Z}_P} : \mathbf{Z}_P \rightarrow \mathbf{Z}_{\phi(P)}$, where the quasi-isometry constants are controlled by those of ϕ .

- Lemma 10.3.** (1) *If ϕ is an isometry, then $\phi_{\mathbf{Z}_P} : \mathbf{Z}_P \rightarrow \mathbf{Z}_{\phi(P)}$ is an isometry for every $P \in \mathcal{P}$.*
- (2) *If ϕ is induced by an element of the commensurator $\text{Comm}(G)$, then for every $P \in \mathcal{P}$, the map $\phi_{\mathbf{Z}_P} : \mathbf{Z}_P \rightarrow \mathbf{Z}_{P'}$ is equivariant with respect to cocompact isometric actions on $\mathbf{Z}_P$ and $\mathbf{Z}_{P'}$.*

Proof. Assertion (1) is immediate.

There is an isomorphism $\alpha : H \rightarrow H'$ where $H \subset G$ and $H' \subset G'$ are finite index subgroups, such that $\phi : K \rightarrow K'$ is H -equivariant, where H acts freely, cocompactly, and isometrically on K' via the isomorphism α . Then $\phi_0 : V \rightarrow V'$ is also H -equivariant. Hence if $h \in H$ stabilizes $P \in \mathcal{P}$, then it also stabilizes the parallel set $\phi(P) \in \mathcal{P}'$. In other words, the map ϕ_0 restricts to a $\text{Stab}(P) \cap H$ -equivariant mapping $V \cap P \rightarrow V' \cap \phi(P)$. Since $\text{Stab}(P) \cap H$ acts cocompactly on P , assertion (2) follows. $\square$

11. QUASI-ISOMETRIC FLEXIBILITY AND THE PROOF OF THEOREM 1.10

The homomorphism $\text{Aut}(G) \rightarrow \text{Comm}(G)$ is injective. Suppose $\alpha \in \text{Aut}(G)$. By Corollary 1.8 there is an isometry $\phi : K \rightarrow K$ which induces α , i.e. we identify G with a subgroup of $\text{Isom}(K)$. If $\alpha \in \ker(\text{Aut}(G) \rightarrow \text{Comm}(G))$, then ϕ commutes with a finite index subgroup of G , and therefore has bounded displacement,

$$\sup_{p \in K} d(\phi p, p) < \infty.$$

It follows that ϕ maps each standard flat to itself. Since the intersection of the standard flats passing through a vertex $p \in K$ is precisely p , it follows that ϕ_0 fixes every vertex, and is therefore the identity map.

The homomorphism $\text{Comm}(G) \rightarrow \text{QI}(G)$ is injective. Suppose $\alpha \in \text{Comm}(G)$. Then α can be represented by an isomorphism $f : G_1 \rightarrow G_2$, where the G_i 's are finite index subgroups of G . Therefore there is a quasi-isometry $\phi : K \rightarrow K$ which is G_1 -equivariant, where we identify G_1 with G_2 via f , and use the corresponding action on the second copy of K . If $\alpha \in \ker(\text{Comm}(G) \rightarrow \text{QI}(G))$, then ϕ is at bounded distance from the identity. Letting $\phi_0 : V \rightarrow V$ be the bijection given by Theorem 1.6, we may argue as in the preceding paragraph to conclude that $\phi_0 = \text{id}_V$. It follows that $\alpha = \text{id}$.

$\text{Aut}(G)$ has infinite index in $\text{Comm}(G)$. Pick a vertex $v \in \Gamma$, and a positive integer k . Let $\alpha : G \rightarrow \mathbf{Z}_k$ be the homomorphism which sends the generator $v \in G$ to $1 \in \mathbf{Z}_k$, and the other generators to $0 \in \mathbf{Z}_k$. Let $G_k \subset G$ be the kernel of α , and $\bar{K}_k$ be the k -fold cover of $\bar{K}$ given by $\bar{K}_k := K/G_k$. One may describe $\bar{K}_k$ as follows. Begin with $\bar{K}(\text{Star } v)$, which is a product $S^1 \times B_j$, where B_j is a bouquet of j -circles, where j is the number of vertices adjacent to v . To obtain $\bar{K}$ from $\bar{K}(\text{Star } v)$, one glues on a copy of $\bar{K}(\Gamma_1)$, where $\Gamma_1 \subset \Gamma$ is the graph obtained by deleting v and the edges incident to v . To obtain $\bar{K}_k$, one first passes to the k -fold cyclic cover of $\bar{K}(\text{Star } v)$, and then glues on k -copies of $\bar{K}(\Gamma_1)$.

Let Γ_k be the graph obtained by taking k copies of Γ , and gluing them together along the k copies of $\text{Star}(v)$. In fact $\bar{K}(\Gamma_k)$ is homotopy equivalent to $\bar{K}_k$; to see this, map $\bar{K}_k$ to $\bar{K}(\Gamma_k)$ by taking $(k-1)$ of the copies of $[0, 1] \times B_j \subset \bar{K}_k$ and collapsing them to copies of B_j (by collapsing the interval factors). In particular, G_k is isomorphic to the RAAG $G(\hat{G}_k)$. Since $\hat{\Gamma}_k$ is not atomic, this shows that the atomic condition is not commensurability invariant among RAAG's.

Note that the permutation group S_k of the set $\mathbf{Z}_k$ acts isometrically on $\bar{K}(\Gamma_k)$ by permuting the copies of $\bar{K}(\Gamma_1)$, and hence we get a homomorphism $S_k \rightarrow \text{Out}(G_k) \rightarrow \text{Comm}(G)/\text{Inn}(G)$. For each element $\alpha \in S_k$, we may lift the corresponding homotopy equivalence $\psi : \bar{K}_k \rightarrow \bar{K}_k$ to a quasi-isometry $\phi : K \rightarrow K$ which preserves a parallel set $P \subset K$ covering $\bar{K}(\text{Star } v)$. Moreover, the induced map $\phi_{\mathbf{Z}_P} : \mathbf{Z}_P \rightarrow \mathbf{Z}_P$ will be equivariant with respect to the action of $k\mathbf{Z}$ on $\mathbf{Z}_P$ by translations, and descends to the permutation of $\mathbf{Z}_k = \mathbf{Z}/k\mathbf{Z}$ corresponding to α .

Now consider the collection $\mathcal{C}$ of elements $\phi \in \text{Comm}(G)$ obtained this way, as k varies over the positive integers, and α varies over the permutation group of $\mathbf{Z}_k$. If $\text{Aut}(G)$ had finite index in $\text{Comm}(G)$, we could find a finite collection $f_1, \dots, f_i \in \mathcal{C}$ such that for each $\phi \in \mathcal{C}$ there is a $\psi \in \text{Isom}(K)$ such that $\phi \circ \psi = f_j$ for some $j \in \{1, \dots, i\}$. This means that $\psi^{-1}(P) = f_j^{-1}(P) =: \bar{P}$, and that

$$\phi_{\mathbf{Z}_P} \circ \psi_{\mathbf{Z}_{\bar{P}}} : \mathbf{Z}_{\bar{P}} \rightarrow \mathbf{Z}_{\bar{P}}$$

agrees with $(f_j)_{\mathbf{Z}_{\bar{P}}} : \mathbf{Z}_{\bar{P}} \rightarrow \mathbf{Z}_{\bar{P}}$. However, by part 1 of Lemma 10.3, $\psi_{\mathbf{Z}_{\bar{P}}}$ is an isometry. This clearly contradicts the fact that ϕ can come from any permutation α of $\mathbf{Z}_k$, for any k .

Remark 11.1. Let H denote the group of isometries of $K(\Gamma_k)$ covering the permutation action $S_k \curvearrowright \bar{K}(\Gamma_k)$. Then there is a short exact sequence

$$1 \longrightarrow G(\Gamma_k) \simeq G_k \longrightarrow H \longrightarrow S_k \longrightarrow 1.$$

Thus H is commensurable with $G(\Gamma)$. However, there is no geometric action of H on $K(\Gamma)$. This can be deduced by examining the subgroup of H which stabilizes a parallel set P , and observing that the induced action on $\mathbf{Z}_P$ is not conjugate to an isometric action.

$\text{Comm}(G)$ has infinite index in $\text{QI}(G)$. The construction is similar to the proof that $[\text{Comm}(G) : \text{Aut}(G)] = \infty$.

Pick $v \in \Gamma$, and look at the infinite cyclic cover $\bar{K}_\infty$ corresponding to the homomorphism $G \rightarrow \mathbf{Z}$ which sends v to 1 and the other generators to 0. Then $\bar{K}_\infty$ can be obtained from $\mathbf{R} \times B_j$ by gluing on infinitely many copies of $\bar{K}(\Gamma_1)$. As in the preceding paragraphs, we may produce homotopy equivalences by “permuting the copies of $\bar{K}(\Gamma_1)$ ”. By lifting these homotopy equivalences we may obtain quasi-isometries $\phi : K \rightarrow K$ which preserve a parallel set P covering $\bar{K}(\text{Star } v)$, and hence obtain a bijective quasi-isometry $\phi_{\mathbf{Z}_P} : \mathbf{Z}_P \rightarrow \mathbf{Z}_P$. It is not hard to see that any bijective quasi-isometry $\mathbf{Z}_P \rightarrow \mathbf{Z}_P$ may be obtained in this way.

If $[\text{QI}(G) : \text{Comm}(G)]$ were finite, there would only be finitely many possibilities for the $\phi_{\mathbf{Z}_P}$ ’s up to pre-composition by maps of the form ψ_P , where ψ comes from an element of the commensurator. In view of Lemma 10.3, this is clearly not the case.

REFERENCES

- [Ahl02] A. R. Ahlin. The large scale geometry of products of trees. *Geom. Dedicata*, 92:179–184, 2002. Dedicated to John Stallings on the occasion of his 65th birthday.
- [Bas72] H. Bass. The degree of polynomial growth of finitely generated nilpotent groups. *Proc. London Math. Soc. (3)*, 25:603–614, 1972.
- [BKS07] M. Bestvina, B. Kleiner, and M. Sageev. Quasiflats in $CAT(0)$ 2-complexes. Preprint, 2007.
- [BM00] M. Burger and S. Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, (92):151–194 (2001), 2000.
- [BN06] J. Behrstock and W. Neumann. Quasi-isometric classification of graph manifold groups. math.GT/0604042, 2006.
- [CCV] R. Charney, J. Crisp, and K. Vogtmann. Automorphisms of two-dimensional right-angled artin groups. math.GR/0610980.
- [CD95] R. Charney and M. W. Davis. The $K(\pi, 1)$ -problem for hyperplane complements associated to infinite reflection groups. *J. Amer. Math. Soc.*, 8(3):597–627, 1995.

- [Cha] R. Charney. An introduction to right angled Artin groups. arXiv, math.GR/0610668.
- [Dav98] M. W. Davis. Buildings are CAT(0). In *Geometry and cohomology in group theory (Durham, 1994)*, volume 252 of *London Math. Soc. Lecture Note Ser.*, pages 108–123. Cambridge Univ. Press, Cambridge, 1998.
- [DJ00] M. W. Davis and T. Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra*, 153(3):229–235, 2000.
- [Dro87] C. Droms. Isomorphisms of graph groups. *Proc. Amer. Math. Soc.*, 100(3):407–408, 1987.
- [Dun85] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.
- [Gro81] M. Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, (53):53–73, 1981.
- [Gro87] M. Gromov. Hyperbolic groups. In *“Essays in group theory”*, pages 75–263. Springer, New York, 1987.
- [Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.
- [HM96] J. Harlander and H. Meinert. Higher generation subgroup sets and the virtual cohomological dimension of graph products of finite groups. *J. London Math. Soc. (2)*, 53(1):99–117, 1996.
- [KKL98] M. Kapovich, B. Kleiner, and B. Leeb. Quasi-isometries and the de Rham decomposition. *Topology*, 37(6):1193–1211, 1998.
- [KL97] M. Kapovich and B. Leeb. Quasi-isometries preserve the geometric decomposition of Haken manifolds. *Invent. Math.*, 128(2):393–416, 1997.
- [KL07] B. Kleiner and B. Leeb. Induced quasi-actions: a remark. preprint, 2007.
- [KMLNR80] K. H. Kim, L. Makar-Limanov, J. Neggers, and F. W. Roush. Graph algebras. *J. Algebra*, 64(1):46–51, 1980.
- [KPS73] A. Karrass, A. Pietrowski, and D. Solitar. Finite and infinite cyclic extensions of free groups. *J. Austral. Math. Soc.*, 16:458–466, 1973. Collection of articles dedicated to the memory of Hanna Neumann, IV.
- [Lau95] M. R. Laurence. A generating set for the automorphism group of a graph group. *J. London Math. Soc. (2)*, 52(2):318–334, 1995.
- [MSW03] L. Mosher, M. Sageev, and K. Whyte. Quasi-actions on trees. I. Bounded valence. *Ann. of Math. (2)*, 158(1):115–164, 2003.
- [MW02] J. P. McCammond and D. T. Wise. Fans and ladders in small cancellation theory. *Proc. London Math. Soc. (3)*, 84(3):599–644, 2002.
- [Sag95] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [Ser89] H. Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.
- [Sta68] J. R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.

- [Wis96] D. Wise. Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups. Princeton thesis, 1996.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH, 155 SOUTH 1400 EAST, ROOM 233, SALT LAKE CITY, UT 84112-0090

E-mail address: `bestvina@math.utah.edu`

YALE UNIVERSITY, MATHEMATICS DEPARTMENT, PO BOX 208283, NEW HAVEN, CT 06520-8283

E-mail address: `bruce.kleiner@yale.edu`

DEPARTMENT OF MATHEMATICS, TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, HAIFA 32000, ISRAEL

E-mail address: `sageevm@techunix.technion.ac.il`