RIGHT-ANGLED ARTIN SUBGROUPS OF RIGHT-ANGLED COXETER AND ARTIN GROUPS

PALLAVI DANI AND IVAN LEVCOVITZ

Abstract. We determine when certain natural classes of subgroups of right-angled Coxeter groups (RACGs) and right-angled Artin group (RAAGs) are themselves RAAGs. We first consider subgroups of RACGs generated by elements that are products of two non-commuting RACG generators; these were introduced in LaForge's thesis [LaF17]. Within the class of 2-dimensional RACGs, we give a complete characterization of when such a subgroup is a finite-index RAAG system. As an application, we show that any 2-dimensional, one-ended RACG with planar defining graph is quasi-isometric to a RAAG if and only if it contains an index 4 subgroup isomorphic to a RAAG. We also give applications to other families of RACGs whose defining graphs are not planar. Next, we show that every subgroup of a RAAG generated by conjugates of RAAG generators is itself a RAAG. This result is analogous to a classical result of Deodhar and Dyer for Coxeter groups. Our method of proof, unlike in the classical case, utilizes disk diagrams and is geometric in nature.

1. Introduction

Let Γ be a finite simplicial graph with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$. The *right-angled Artin group* (RAAG for short) associated to Γ is the group A_Γ given by the following presentation:

$$A_\Gamma = \langle V(\Gamma) \mid st = ts \text{ for all } (s,t) \in E(\Gamma)\rangle$$

This article is concerned with the following question. Given a finite set S of elements in a group, when is the group generated by S isomorphic to a RAAG in the "obvious" way (i.e. with S as the "standard" RAAG generating set)? To make this precise, we define the notion of *RAAG system*.

Definition 1.1 (RAAG system)**.** Let G be any group with generating set S. Let Δ be the graph whose vertex set is in bijection with S, and which has an edge between distinct $s, t \in S \equiv V(\Delta)$ if and only if s and t commute. We call Δ the *commuting graph* associated to S. There is a canonical homomorphism $\phi : A_\Delta \to G$ extending the bijection $V(\Delta) \to S$. We say that (G, S) is a *RAAG system* if ϕ is an isomorphism. In particular, $(A_\Gamma, V(\Gamma))$ is a RAAG system for any RAAG A_Γ.

The *right-angled Coxeter group* (RACG for short) associated to the finite simplicial graph Γ is the group W_Γ given by the presentation:

$$W_\Gamma = \langle V(\Gamma) \mid s^2 = 1 \text{ for all } s \in V(\Gamma), st = ts \text{ for all } (s,t) \in E(\Gamma)\rangle$$

This work of the first author was supported by a grant from the Simons Foundation (#426932, Pallavi Dani) and by NSF Grant No. DMS-1812061. This work of the second author was supported by the Israel Science Foundation and in part by a Technion fellowship.

In this article we study subgroups G generated by particular natural subsets S of right-angled Coxeter and Artin groups, and we give characterizations for when (G, S) is a RAAG system or a finite-index RAAG system.

A theorem of Davis–Januskiewicz states that every RAAG is commensurable to some RACG [DJ00]. This leads to the following question addressing the converse:

Question 1.2. Which RACGs are commensurable to RAAGs?

A RACG that is commensurable to a RAAG is, in particular, quasi-isometric to a RAAG. By considering different quasi-isometry invariants, one sees that the converse to the Davis–Januskiewicz theorem above is far from being true. For instance, there are many RACGs that are one-ended hyperbolic (such as virtual hyperbolic surface groups), while no RAAG is both one-ended and hyperbolic. Furthermore, RAAGs have linear, quadratic or infinite divergence [BC12], whereas the divergence of a RACG can be a polynomial of any degree [DT15]. Restricting to RACGs of at most quadratic divergence is still not enough guarantee being quasi-isometric to a RAAG. For instance, the Morse boundary of a RAAG with quadratic divergence is always totally disconnected (see [CS15, CH17]), while the Morse boundary of a RACG of quadratic divergence can have nontrivial connected components [Beh19]. The above examples show that there are numerous families of RACGS which are not quasi-isometric, and hence not commensurable to any RAAG. Within the subclass of one-ended RACGs with planar, triangle-free defining graphs, Nguyen–Tran give a complete characterization of those quasi-isometric to RAAGs [NT17, Theorem 1.2]. However, Question 1.2 is still open even in this setting.

One approach to proving that a RACG is commensurable to a RAAG is to look for finite index subgroups that are isomorphic to RAAGs. We focus on a class of subgroups of RACGs, introduced by LaForge in his PhD thesis [LaF17], that are logical candidates for being RAAGs. Given a RACG defined by a graph Γ and two non-adjacent vertices $s, t \in V(\Gamma)$, it readily follows that st is an infinite order element of W_Γ. There is then a correspondence between edges of the complement graph Γ^c with such infinite order elements of Γ. Given a subgraph Λ of Γ^c, which, without loss of generality, we assume contains no isolated vertices, let G be the subgroup generated by $E(\Lambda)$ (thought of as infinite order elements of W_Γ). A natural question is:

Question 1.3. When is $(G, E(\Lambda))$ a finite-index RAAG system?

If $(G, E(\Lambda))$ is indeed a RAAG system, then G is called a *visual RAAG subgroup* of W_Γ. In his thesis, LaForge obtained some necessary conditions for such subgroups G to be visual RAAGs.

We say that W_Γ is *2-dimensional* if Γ is triangle-free. Our first main theorem gives an exact characterization of the finite-index visual RAAG subgroups of 2-dimensional RACGs in terms of graph theoretic conditions:

Theorem A. *Let W_Γ be a 2-dimensional right-angled Coxeter group. Let Λ be a subgraph of Γ^c with no isolated vertices, and let G be the subgroup generated by $E(\Lambda)$. Then the following are equivalent.*

1. *$(G, E(\Lambda))$ is a RAAG system and G is finite index in W_Γ.*
2. *$(G, E(\Lambda))$ is a RAAG system and G has index either two or four in W_Γ (and exactly four if W_Γ is not virtually free).*

(3) *Λ has at most two components and satisfies conditions $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$.*

The conditions $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$ in the above theorem are algorithmically checkable graph theoretic conditions on Γ and Λ. See Section 3 for precise definitions of these conditions.

In Section 5 we provide several applications to concrete families of RACGs. In particular we prove:

Theorem B. *Let W_Γ be a 2-dimensional, one-ended RACG with planar defining graph. Then W_Γ is quasi-isometric to a RAAG if and only if it contains an index 4 subgroup isomorphic to a RAAG.*

A complete description of which RACGs considered in Theorem B are quasi-isometric to RAAGs is given by Nguyen-Tran [NT17]. Theorem B shows these are actually commensurable to RAAGs.

We also give two families of RACGs defined by non-planar graphs which contain finite-index RAAG subgroups (see Corollaries 5.1 and 5.2). These cannot be obtained by applying the Davis–Januskiewicz constructions to the defining graphs of the RAAGs they are commensurable to. For the family in Corollary 5.1, we use work of Bestvina–Kleiner–Sageev on RAAGs [BKS08], to conclude the RACGs are quasi-isometrically distinct. We believe that the methods from this article may be used to study commensurability in RACGs.

The proof of Theorem A consists of two main parts. One part involves obtaining an understanding of when G is of finite index, leading to conditions $\mathcal{F}_1$ and $\mathcal{F}_2$. To obtain these, we use *completions of subgroups*, introduced in [DL19]. The other aspect consists of obtaining criteria to recognize when $(G, E(\Lambda))$ is a RAAG system. To do so, we prove the following theorem by careful analysis of disk diagrams:

Theorem C. *Let W_Γ be a right-angled Coxeter group. Let Λ be a subgraph of Γ^c with no isolated vertices and at most two components. Then the subgroup $(G, E(\Lambda)) < W_\Gamma$ is a RAAG system if and only if $\mathcal{R}_1$–$\mathcal{R}_4$ are satisfied.*

Conditions $\mathcal{R}_1$, $\mathcal{R}_2$, and a condition more or less equivalent to $\mathcal{R}_3$ were known to be necessary conditions for $(G, E(\Lambda))$ to be a RAAG system by work of LaForge [LaF17]. We show in Example 3.12 that they are not sufficient. We introduce a fourth graph-theoretic condition $\mathcal{R}_4$ to obtain a complete characterization of all visual RAAG subgroups defined by subgraphs of Γ^c with at most two components. The bulk of the proof of Theorem C consists of showing that the conditions $\mathcal{R}_1$–$\mathcal{R}_4$ are sufficient.

Note that, unlike in Theorem A, there is no assumption on the dimension of the RACGs in Theorem C. On the other hand, there is an additional assumption in Theorem C, namely that the subgraph Λ of Γ can have at most two components.

When Λ contains more than two components, the situation becomes much more complex. We show that additional graph-theoretic conditions are necessary to generalize the Theorem C to this setting (see Lemma 3.31 and Lemma 3.33). Remarkably, a consequence of these conditions is that if Γ is triangle-free and $(G, E(\Lambda))$ is a finite-index RAAG system, then Λ can have at most two components. This fact is crucial to the proof of Theorem A, which does not have any assumption on the number of components of Λ. Additionally, we are aware that even more conditions are necessary than those in this article, but we do not have a complete conjectural list of conditions that would be sufficient to characterize visual RAAGs.

We next turn our attention to RAAG subgroups of RAAGs. A classical theorem on Coxeter groups, proven independently by Deodhar [Deo89] and Dyer [Dye90], states that reflection subgroups of Coxeter groups (i.e., those generated by conjugates of generators) are themselves Coxeter groups. We define a *generalized RAAG reflection* to be an element of a RAAG A_Δ that is conjugate to a generator in $V(\Delta)$. We prove a result for subgroups of RAAGs generated by generalized reflections analogous to the one for Coxeter groups:

Theorem D. *Let $\mathcal{T}$ be a finite set of generalized RAAG reflections in the RAAG A_Γ. Then the subgroup $G < A_\Gamma$ generated by $\mathcal{T}$ is a RAAG.*

Kim-Koberda show that given a set of generalized RAAG reflections $\mathcal{T}$ with corresponding commuting graph Δ, there exists a subgroup of the ambient RAAG (generated by sufficiently high powers of the elements of $\mathcal{T}$) which is isomorphic to the RAAG A_Δ [KK13]. However, it remained open whether the subgroup G generated by $\mathcal{T}$ is itself a RAAG, and Theorem D precisely answers this question. Note that $(G, \mathcal{T})$ may not itself be a RAAG *system* and in general, G is not isomorphic to the RAAG A_Δ. However, we explain how to give an algorithm to obtain a set of generalized RAAG reflections $\mathcal{T}'$ such that $(G, \mathcal{T}')$ is a RAAG system. Consequently, this exactly determines the isomorphism class of the RAAG G.

We prove Theorem D using a characterization of RAAG systems in terms of the deletion condition, given by Basarab [Bas02]. We use disk diagrams to show that subgroups generated by generalized RAAG reflections satisfy the criteria in Basarab's characterization. Our geometric approach is very different from that of Deodhar (which utilizes root systems) and that of Dyer (which is algebraic and uses cocycles) towards proving the corresponding result for Coxeter groups.

Acknowledgements. The authors would like to thank Jingyin Huang for suggesting the question that led to Theorem D, Kevin Schreve for a comment that led to Corollary 4.9, and Hung Tran for encouraging us to look at the examples considered in Theorem B. Additionally, we would like to thank Jason Behrstock, Garret LaForge, and Kim Ruane for helpful conversations.

2. Background

2.1. **Basic terminology and notation.** Let G be a group with generating set S. We say that $w = s_1 \dots s_n$, with $s_i \in \left(S \cup S^{-1}\right)$ for $1 \leq i \leq n$, is a *word over* S or a *word in* G. If the words w and w' represent the same element of G, then we say that w' is an *expression* for w and write $w' \doteq w$. We say the word $w = s_1 \dots s_n$ is *reduced* (or *reduced over* S for emphasis) if given $w' = t_1 \dots t_m \doteq w$, it follows that $n \leq m$.

2.2. **Right-angled Coxeter and Artin groups.** Coxeter groups can be characterized as those groups which are generated by involutions and which satisfy the deletion condition, see Definition 2.1 below (for a proof of this fact, see [Dav15, Theorem 3.3.4]). By work of Basarab [Bas02], RAAGs can be characterized in a similar manner (see Theorem 2.2 below). This characterization will be utilized in Section 6.

Definition 2.1 (Deletion Condition)**.** Let G be a group generated by S. We say that (G, S) satisfies the *deletion condition* if, given any word w over S, either w is reduced or $w = s_1 \dots s_k$ and there exist $1 \leq i < j \leq k$ such that $s_1 \dots \hat{s}_i \dots \hat{s}_j \dots s_k$ is an expression for w.

The result below directly follows from a result of Basarab.

Theorem 2.2 ([Bas02])**.** *Let G be a group generated by S such that $S \cap S^{-1} = \emptyset$ and $1 \notin S$. Then (G, S) is a RAAG system if and only if each of the following holds:*

1. *Every s in S has infinite order.*
2. *(G, S) satisfies the deletion condition.*

Proof. If (G, S) is a RAAG system, then G is torsion-free (see [Cha07]), so (1) holds. Furthermore, (G, S) satisfies (2) by [Bas02, Corollary 1.4.2] (see also [Bah05, pg. 31, ex. 17] for a simpler proof in this setting). The converse also follows from a direct application of [Bas02, Corollary 1.4.2]. □

We now define certain moves which can be performed on a word that produce another expression for it. These moves provide a solution to the word problem for RAAGs and RACGs (see Theorem 2.4 below).

Definition 2.3 (Tits moves)**.** Let G be a group generated by S. Let $w = s_1 \dots s_n$ be a word over S. If s_i and s_{i+1} commute for some $1 \leq i < n$, then the word $s_1 \dots s_{i-1} s_{i+1} s_i s_{i+2} \dots s_n$ is an expression for w obtained by a *swap operation* performed to w, which *swaps* s_i and s_{i+1}. If $s_i = s_{i+1}^{-1}$ for some $1 \leq i < n$, then $s_1 \dots s_{i-1} s_{i+2} \dots s_n$ is an expression for w is obtained by a *deletion operation* performed to w. A *Tits move* is either a swap operation or a deletion operation. We say a word is *Tits reduced* if no sequence of Tits moves can be performed to the word to obtain an expression with fewer generators.

Theorem 2.4 below shows that RAAGs and RACGs admit a nice solution to the word problem. This solution to the word problem for RACGs is a well known result of Tits [Tit69], a version of which holds more generally for all Coxeter groups. The result below in the setting of RAAGs follows from a theorem of Basarab [Bas02, Theorem 1.4.1] which generalizes Tits' result.

Theorem 2.4 ([Tit69, Bas02])**.** *Let A_Γ be either a RAAG or a RACG. Then the following hold:*

1. *If w_1 and w_2 are reduced words over $V(\Gamma)$ representing the same element of G, then w_2 can be obtained from w_1 by Tits moves.*
2. *Given any word w over $V(\Gamma)$, a reduced expression for w can be obtained by applying Tits moves to w.*

We will often not refer directly to the above theorem, and we will instead simply say that a given RAAG or RACG *admits a Tits solution to the word problem.*

The next two lemmas are well known and will often be implicitly assumed.

Lemma 2.5. *Let A_Γ either be a RAAG or RACG. Then $s, t \in V(\Gamma)$ commute as elements of A_Γ if and only if (s, t) is an edge of Γ.*

Proof. One direction of the claim follows from the definitions of a RAAG and a RACG. If A_Γ is a RACG, then the other direction follows from [BB05, Prop 4.1.2].

On the other hand, suppose that A_Γ is a RAAG, and let $s, t \in V(\Gamma)$ be non-adjacent vertices. Suppose, for a contradiction that $w = sts^{-1}t^{-1} \simeq 1$. Let D be a disk diagram with boundary w (see Section 2.3 for a reference for disk diagrams). This disk diagram contains exactly two intersecting hyperplanes: one labeled by s and one labeled by t. However, this is a contradiction as a pair of hyperplanes whose labels are non-adjacent vertices of Γ cannot intersect. □

Lemma 2.6. *Let W_Γ be a RACG, and let $s, t, q, r \in V(\Gamma)$ be such that s and t do not commute, and r and q do not commute. We have that $(st)(qr) \doteq (qr)(st)$ if and only if one of the following holds:*

(1) *There is a square in Γ formed by s, q, t, r.*
(2) *$t = q$ and $s = r$.*
(3) *$t = r$ and $s = q$.*

Proof. Clearly each of (1), (2) and (3) implies that $(st)(qr) \doteq (qr)(st)$.

To prove the converse, suppose that $(st)(qr) \doteq (qr)(st)$. Suppose first that $t = q$, and consequently $stqr \doteq sr$. As s and t do not commute and q and r do not commute, this is only possible if $r = t$. Thus, (2) holds.

If $s = q$, as $qrts \doteq tsqr$, we can apply the same argument to conclude that $t = r$, showing (3) holds. By similar arguments, if $s = r$ then $t = q$, and if $t = r$ then $s = q$. Thus, we may assume that s, t, q and r are all distinct vertices of Γ. In this case we can again conclude by Tits' solution to the word problem, that as $stqr \doteq qrst$ then s, q, t and r form a square in Γ. □

2.3. **Disk Diagrams.** We give a brief background on disk diagrams as they are used in our setting, and we refer the reader to [Sag95] and [Wis11] for the general theory of disk diagrams over cube complexes. We then give some preliminary lemmas that will be needed in later sections.

Let A_Δ be a RAAG, and let $w = s_1 \dots s_n$, with $s_i \in V(\Delta)$, be a word equal in W_Γ to the identity, i.e. $w \doteq 1$. There exists a van Kampen diagram D with boundary label w, and we call this planar 2-complex a *disk diagram in A_Δ with boundary label w*. We now describe some additional properties of D in our setting. The edges of D are oriented and labeled by generators in $V(\Delta)$. A *path in D* is a path p in the 1-skeleton of D, traversing edges $e_1, \dots, e_m$, and the label of p is the word $a_1 \dots a_m$ where, for each $1 \le i \le m$, a_i is the label of e_i if e_i is traversed along its orientation, and a_i^{-1} is the label of e_i if e_i is traversed opposite to its orientation. Every cell in D is a square that has a boundary path with label $aba^{-1}b^{-1}$ for some commuting generators a and b in $V(\Delta) \cup V(\Delta)^{-1}$.

There is a base vertex $p \in \partial D$ and an orientation on D, such that the smallest closed path δ which traverses the boundary of D in the clockwise orientation starting at p and traversing every edge outside the interior of D has label w. We call δ the *boundary path* of D. Note that if D contains an edge e not contained in a square, then necessarily δ traverses e exactly twice.

If W_Γ is a RACG and w is a word over $V(\Gamma)$ equal in W_Γ to the identity, then we define a disk diagram D in W_Γ with boundary w similarly. However, as each generator in $V(\Gamma)$ is an involution, we do not need to orient the edges of D.

Let D be a disk diagram and $q = [0,1] \times [0,1]$ be a square in D. The subset $\{\frac{1}{2}\} \times [0,1] \subset q$ (similarly, $[0,1] \times \{\frac{1}{2}\} \subset q$) is a *midcube*. The midpoint of an edge in D is also defined to be a *midcube*. A *hyperplane* in D is a minimal non-empty collection H of midcubes in D with the property that given any midcube $m \in H$ and a midcube m' in D such that $m \cap m'$ is contained in an edge of D, it follows that $m' \in H$. We say that H is dual to an edge e if the midpoint of e is in H.

Since opposite edges in every square in D have the same label, it follows that every edge intersecting a fixed hyperplane H has the same label. We call this the *label of the hyperplane*. Since adjacent sides of a square have distinct labels which commute, it follows that no hyperplane self-intersects, and if two hyperplanes intersect, then their labels correspond to distinct, commuting generators.

Definition 2.7 (Maps preserving boundary combinatorics)**.** Let D and D' be disk diagrams, and let δ and δ' respectively be their boundary paths. Let $E = \{e_1, \dots, e_m\}$ (resp. $E' = \{e'_1, \dots, e'_n\}$) be the edges traversed by δ (resp. δ'). More precisely, e_i (resp. e'_i) is the ith edge traversed by δ (resp. δ') for each i. Observe that every hyperplane of D is dual to two edges $e_j, e_k \subset E$ for some $j \neq k$. (It could be that $e_j = e_k$ thought of as edges of D.) A similar statement holds for D'.

Let $F \subset E$ and $F' \subset E'$, and let $\psi : F \to F'$ be a bijection. We say that ψ *preserves boundary combinatorics* if for every pair of edges $e, f \in F$ which are dual to the same hyperplane of D, their images $\psi(e)$ and $\psi(f)$ are dual to the same hyperplane of D'.

Note that if Ψ preserves boundary combinatorics, then Ψ^{-1} does as well.

A pair of hyperplanes H and H' in a disk diagram D form a *bigon* if they intersect in at least two distinct points. The following lemma, first proven in [Sag95, Theorem 4.3], guarantees that we can always choose a disk diagram without bigons. The boundary combinatorics statement below is guaranteed by the proof of this fact in [Wis11, Lemma 2.3, Corollary 2.4].

Lemma 2.8 ([Sag95] [Wis11])**.** *Given a disk diagram D with boundary label w, there exists a disk diagram D' also with boundary label w such that D' does not contain any bigons. Moreover, the natural bijection between the edges traversed by the boundary paths of D and D' induced by the label w preserves boundary combinatorics.*

Remark 2.9. In light of Lemma 2.8, for the rest of this paper we will always assume that any disk diagrams we consider do not have bigons.

Remark 2.10. Let α be a path with label $s_1 \dots s_n$ in some disk diagram. The "edge of α with label s_i" is understood to be the ith edge α traverses (even though there may be several edges of α with the same label as this edge. A similar statement holds when we refer to subpaths of α.

Given a disk diagram with boundary label w, we will often want to produce a new disk diagram with boundary label w', where w' is obtained from w by a Tits move, and such that boundary combinatorics are preserved on appropriate subsets of the boundary paths. The following lemma exactly describes how we can perform these operations.

Lemma 2.11. *Let D be a disk diagram over the group W, where W is either a RACG or a RAAG. Suppose the boundary path of D traverses the edges $e_1, \dots e_n$ and has label $w = s_1 \dots s_n$.*

(1) *If s_r and s_{r+1} (taken modulo n) are distinct and commute for some $1 \leq r \leq n$, then there is a disk diagram D' whose boundary path traverses the edges $e'_1, \dots, e'_n$ and has label $s_1 \dots s_{i+1} s_i \dots s_n$. Furthermore, the map ψ preserves boundary combinatorics, where ψ is defined by $\psi(e_r) = e'_{r+1}$, $\phi(e_{r+1}) = e'_r$, and $\psi(e_j) = e'_j$ for $j \neq r, r+1$.*

(2) *If* $s_r = s_{r+1}^{-1}$ *(taken modulo* n*) for some* $1 \leq r \leq n$*, then there is a disk diagram* D' *with boundary label* $s_1 \dots s_{r-1}s_{r+2} \dots s_n$*. Moreover, the natural map from edges traversed by the boundary path of* D' *to edges traversed by the boundary path of* D *preserves boundary combinatorics.*

(3) *Given any generator (or inverse of a generator)* s *and any* r*, with* $1 \leq r \leq n$*, it follows that there exists a disk diagram* D' *with boundary label* $s_1 \dots s_r(ss^{-1})s_{r+1} \dots s_n$*. Moreover, the natural map from edges traversed by the boundary path of* D *to the edges traversed by the boundary path of* D' *preserves boundary combinatorics.*

Proof. We first prove (1). Let q be a square whose edges are labeled consecutively by s_r, s_{r+1}, s_r^{-1}, s_{r+1}^{-1}. We form the disk diagram D' by identifying consecutive edges of q labeled by s_r and s_{r+1} to the edges of ∂D labeled by s_r and s_{r+1} (these edges must be distinct as $s_r \neq s_{r+1}$). The claim is readily checked.

We next prove (2). Let e and f be the edges of ∂D labeled respectively by s_r and s_{r+1}. Suppose first that e and f are distinct. In this case, form the disk diagram D' by identifying e and f, i.e. "fold" these edges together. On the other hand, if $e = f$, then as D has boundary label w, it must follow that e is a spur, i.e. an edge attached to D that is not contained in any square and which contains a vertex of valence 1. In this case we can remove the edge e from D to obtain D'. In either case, the claim is readily checked.

To show (3), form D' by inserting a spur edge with label s to the vertex traversed by the boundary path of D between s_r and s_{r+1}. □

3. Visual RAAG subgroups of right-angled Coxeter groups

In this and the next section we study visual RAAG subgroups of RACGs, as described in the introduction. We begin by describing some notation that will be used throughout these sections.

Let Γ be a graph, and let W_Γ be the corresponding right-angled Coxeter group. Let Γ^c denote the complement of Γ, that is, the graph with the same vertex set as Γ, which has an edge between two vertices if and only if the corresponding vertices are not adjacent in Γ. Let Λ be a subgraph of Γ^c with no isolated vertices, i.e., one in which every vertex of Λ is contained in some edge.

We form a new graph $\Theta = \Theta(\Gamma, \Lambda)$ which we think of as a graph containing the edges of both Γ and Λ. More formally, $V(\Theta) = V(\Gamma)$ and $E(\Theta) = E(\Gamma) \cup E(\Lambda)$. Note that as $E(\Lambda) \subset \Gamma^c$, it follows that Θ is simplicial. We refer to edges of Θ that correspond to edges of Γ (resp. Λ) as Γ-edges (resp. Λ-edges).

A Λ-edge between vertices a and b corresponds to an inverse pair of infinite order elements of W_Γ, namely ab and ba. By a slight abuse of terminology, we will use the term Λ-edge to refer to one of these elements and vice versa. We identify $E(\Lambda)$ with a subset of W_Γ by arbitrarily choosing one of the two infinite order elements corresponding to each Λ-edge, and we define G^Θ to be the subgroup of W_Γ generated by $E(\Lambda)$. As we are dealing with subgroups generated by $E(\Lambda)$, there is no loss in generality in assuming that Λ has no isolated vertices. The goal of this section is to study when $(G, E(\Lambda))$ is a RAAG system.

Let Δ be the commuting graph corresponding to $E(\Lambda)$ (as defined in the introduction), and let A_Δ be the corresponding RAAG. Recall that, by definition, $(G^\Theta, E(\Lambda))$ is a RAAG system if and only if the natural homomorphism $\phi : A_\Delta \to$

G^Θ extending the bijection between $V(\Delta)$ and $E(\Lambda)$ is an isomorphism. As ϕ is always surjective, we would like to understand when ϕ is injective.

For the remainder of this section, we fix $\Gamma, \Lambda, \Theta, A_\Delta$, and ϕ as above. Furthermore, we will use the following terminology. The path γ in Θ *visiting vertices* $x_1, x_2, \ldots, x_n$ is defined to be the path which starts at x_1, passes through the remaining vertices in the order listed, and ends at x_n. We say that γ is simple if $x_i \neq x_j$ for $i \neq j$, and that γ is a loop if $x_1 = x_n$. Finally, γ is a cycle if it is a loop with $n \geq 3$, such that $x_i \neq x_j$ unless $\{i, j\} = \{1, n\}$. We call a path (resp. cycle) in Θ consisting only of Γ-edges a Γ-path (resp. Γ-cycle). We define Λ-paths and Λ-cycles similarly.

We begin by describing some graph theoretic conditions on Θ which are consequences of either G^Θ being a RAAG or of $(G^\Theta, E(\Lambda))$ being a RAAG system.

Conditions $\mathcal{R}_1$ and $\mathcal{R}_2$, defined below, when combined, are equivalent to LaForge's star-cycle condition. LaForge proves that $\mathcal{R}_1$ and $\mathcal{R}_2$ are necessary conditions for $(G^\Theta, E(\Lambda))$ to be a RAAG system [LaF17, Lemma 8.2.1]. We include proofs here for completeness.

Definition 3.1 (Condition $\mathcal{R}_1$)**.** We say that Θ satisfies *condition* $\mathcal{R}_1$ if it does not contain a Λ-cycle.

Lemma 3.2 ([LaF17])**.** *If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Θ satisfies $\mathcal{R}_1$.*

Proof. Suppose Θ does not satisfy $\mathcal{R}_1$. Then it contains a Λ-cycle, say with vertices $a_1, \ldots, a_k$, where $k \geq 3$, such that for each i (mod k), a_i is connected to a_{i+1} by a Λ-edge. Let g_i be the generator of A_Δ (or its inverse) corresponding to the (oriented) Λ-edge $a_i a_{i+1}$, and observe that for all $i \neq j$ we have that $g_i \neq g_j^{-1}$. This, together with the fact that RAAGs satisfy the deletion condition (see Theorem 2.2), implies that $g = g_1 g_2 \ldots g_k$ is a non-trivial element of A_Δ. Moreover, $\phi(g) = (a_1a_2)(a_2a_3)\ldots(a_ka_1) = 1$, so g is in the kernel of ϕ, and therefore ϕ is not injective. □

Definition 3.3 (Condition $\mathcal{R}_2$)**.** We say that Θ satisfies *condition* $\mathcal{R}_2$ if each component of $\Lambda \subset \Theta$ (with the natural inclusion) is an induced subgraph of Θ.

Lemma 3.4 ([LaF17])**.** *If G^Θ is a RAAG, then Θ satisfies $\mathcal{R}_2$.*

Proof. Suppose Θ does not satisfy $\mathcal{R}_2$, and let u and v be a pair of vertices in a component Λ, such that u and v are adjacent in Θ. It follows that u and v are connected by a Γ-edge, and therefore, they commute. Since u and v are in the same component of Λ, there is a simple Λ-path from u to v whose vertices (in order) are $u = a_1, \ldots, a_k = v$. Note that $k \geq 3$, since Θ is a simplicial graph. For $1 \leq i \leq k-1$, let g_i be the generator of A_Δ (or its inverse) corresponding to the Λ-edge $a_i a_{i+1}$, and let $g = g_1 g_2 \ldots g_{k-1}$. The element g is a non-trivial element of A_Δ, as RAAGs satisfy the deletion condition by Theorem 2.2.

We now have that $g^2 = \big((a_1a_2)(a_2a_3)\ldots(a_{k-1}a_k)\big)^2 = (a_1a_k)^2 = (uv)^2 = 1$, since u and v commute. This implies that G^Θ has torsion. Thus, G^Θ cannot be a RAAG as RAAGs are torsion-free (see [Cha07]). □

Before stating the next result, we need to introduce some terminology.

Definition 3.5. (2-component paths and cycles). We say the Γ-path γ in Θ is a *2-component path* if γ visits vertices (in order) $c_1, d_1, c_2, d_2 \ldots, c_n, d_n$ for some

$n \geq 1$, where d_n could be trivial if $n > 1$, the c_i's all lie in a single component Λ_c of Λ, and the d_i's all lie in a single component $\Lambda_d \neq \Lambda_c$ of Λ. If it is important to emphasize the components visited by γ, we will call it a $\Lambda_c\Lambda_d$-path.

A *2-component loop* is a 2-component path visiting $c_1, d_1, \dots, c_n, d_n, c_{n+1}$ such that $c_1 = c_{n+1}$. A *2-component cycle* is a 2-component loop which is a Γ-cycle. A 2-component cycle of length four will be called a *2-component square*.

Definition 3.6. (Λ-convex hull) We define the Λ-*convex hull* of a set $X \subset V(\Theta)$ to be the convex hull of X in Λ.

Definition 3.7 (Condition $\mathcal{R}_3$)**.** We say that Θ satisfies *condition* $\mathcal{R}_3$ if the following holds for every 2-component square in Θ. Consider a 2-component square in Θ visiting vertices c_1, d_1, c_2, d_2, where $c_1, c_2 \in \Lambda_c$, $d_1, d_2 \in \Lambda_d$, and Λ_c, Λ_d are distinct components of Λ. Then the graph Γ contains the join of $V(T_c)$ and $V(T_d)$, where T_c and T_d are the Λ-convex hulls of $\{c_1, c_2\}$ and $\{d_1, d_2\}$ respectively. (See Figure 3.1.)

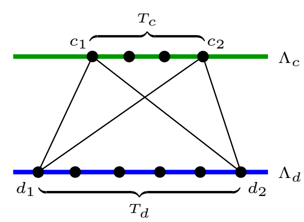

FIGURE 3.1. In the figure, the colored parts consist of Λ-edges, and the black parts consist of Γ-edges. The condition $\mathcal{R}_3$ says that if Θ contains a black square as shown, then every vertex of T_c is joined by a Γ-edge to every vertex of T_d.

We will often need to utilize an expression for a word in W_Γ which is the product of Λ-edges. This construction is the content of the following definition.

Definition 3.8 (Λ-edge words)**.** Suppose Θ satisfies condition $\mathcal{R}_1$, and let $w = (a_1a_1')(a_2a_2')\dots(a_na_n')$ be a word in W_Γ such that a_i and a_i' are in the same Λ-component of Θ for each $1 \leq i \leq n$. As Θ satisfies $\mathcal{R}_1$, there is a unique simple Λ-path from a_i to a_i'. Let $a_i = a_1^i, \dots, a_{m_i}^i = a_i'$ be the vertices visited by this path. Form the word:

$$w' = \Big((a_1^1a_2^1)(a_2^1a_3^1)\dots(a_{(m_1-1)}^1a_{m_1}^1)\Big)\dots\Big((a_1^na_2^n)(a_2^na_3^n)\dots(a_{(m_n-1)}^na_{m_n}^n)\Big)$$

We call w' the Λ-*edge word* associated to w. Note that w' is well-defined, $w \simeq w'$, and w' is a product of Λ-edges.

Remark 3.9. If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Θ satisfies $\mathcal{R}_1$ by Lemma 3.2. In particular, given a word w in W_Γ as in the above definition, the Λ-edge word associated to w is well-defined.

Lemma 3.10. *If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Θ satisfies $\mathcal{R}_3$.*

Remark 3.11. (Comparison of Lemma 3.10 with Laforge's chain-chord condition.) In [LaF17, Lemma 8.2.3], LaForge introduced a necessary condition called the chain-chord condition which, if interpreted in the language of joins and 2-component cycles, is close to our condition $\mathcal{R}_3$. We note that there are errors in the statement and proof of [LaF17, Lemma 8.2.3].

Proof of Lemma 3.10. Suppose there is a 2-component square γ in Θ visiting vertices c_1, d_1, c_2, d_2 as in condition $\mathcal{R}_3$. Let Λ_c and Λ_d be the components of Λ respectively containing $\{c_1, c_2\}$ and $\{d_1, d_2\}$. Let T_c and T_d be the Λ convex hulls respectively of $\{c_1, c_2\}$ and $\{d_1, d_2\}$. By Lemma 3.2, there is a unique simple Λ-path from c_1 to c_2 (resp. d_1 to d_2) and this path is equal to T_c (resp. T_d).

Let w denote the commutator $[c_1c_2, d_1d_2]$. The existence of γ tells us that c_1 and c_2 both commute with d_1 and d_2, so w represents the identity in W_Γ.

Let w_1, w_2 and w' be the Λ-edge words associated to respectively c_1c_2, d_1d_2 and w. As ϕ is injective w' represents the trivial element of A_Δ, and there is a disk diagram D over A_Δ with boundary label w'. We warn that the edges of D are labelled by Λ-edges, i.e., generators of A_Δ. We will analyze hyperplanes of this diagram.

Let $p_{w_1}, p_{w_2}, p_{w_1^{-1}}$ and $p_{w_2^{-1}}$ be the paths in ∂D with labels w_1, w_2, w_1^{-1} and w_2^{-1} respectively. For $i \in \{1, 2\}$, the word w_i (thought of as a word over $V(\Delta) = E(\Lambda)$) does not contain any repeated letters in $V(\Delta)$. Consequently, a hyperplane is dual to at most one edge of p_{w_1} (resp. $p_{w_2}, p_{w_1^{-1}}$ and $p_{w_2^{-1}}$). Furthermore, w_1 and w_2 are words over respectively $E(T_c)$ and $E(T_d)$. As Λ_c and Λ_d are distinct components of Λ, a hyperplane dual to an edge of p_{w_1} must be dual to an edge of $p_{w_1^{-1}}$ and vice versa. A similar statement holds for hyperplanes dual to p_{w_2} and $p_{w_2^{-1}}$.

It follows that every hyperplane dual to p_{w_1} intersects every hyperplane dual to p_{w_2}. Consequently, every Λ-edge in the word w_1 commutes with every Λ-edge in the word w_2. Since ϕ is a homomorphism, the Coxeter group elements corresponding to these Λ-edges must commute as well. By Lemma 2.6 each vertex of T_c commutes with each vertex of T_d. The result follows. □

The next example shows that the conditions obtained so far are not sufficient for $(G^\Theta, E(\Lambda))$ to be a RAAG system.

Example 3.12. Let Γ be a hexagon, and let Λ be the graph with two components shown on the left side in Figure 3.2. It is clear that $\mathcal{R}_1, \mathcal{R}_2$, and $\mathcal{R}_3$ are satisfied. However, Lemma 3.15 below implies that $(G^\Theta, E(\Lambda))$ is not a RAAG system.

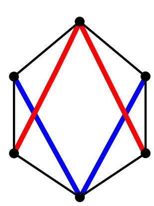

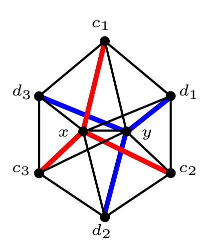

FIGURE 3.2. The graph on the left concerns Example 3.12 and the graph of the right concerns Example 3.13.

Example 3.12 shows that at least one additional condition is needed in order to obtain a characterization of visual RAAGs. It is tempting to conjecture that, given any $\Lambda_c\Lambda_d$-cycle with corresponding Λ-convex hulls T_c and T_d, the graph Γ contains the join of $V(T_c)$ and $V(T_d)$ (as is the case when the cycle has length four, by Lemma 3.10 above). However, the following example shows this is not necessarily true for longer cycles.

Example 3.13. In Figure 3.2, let Θ be the graph shown where Γ-edges are black and Λ edges are colored. Observe that Λ has two components, colored red and blue. Consider the 2-component cycle visiting vertices $c_1, d_1, c_2, d_2, c_3, d_3, c_1$. Then T_c is the entire red tree and T_d is the entire blue tree. However, Γ does not contain the join of $V(T_c)$ and $V(T_d)$. (For example, there is no edge in Γ connecting c_1 and d_2.) On the other hand, $(G^\Theta, E(\Lambda))$ is a RAAG system in this case. (See Corollary 5.1 for a proof.)

Despite the fact that $\mathcal{R}_3$ does not generalize to a necessary condition on longer cycles in the obvious way, the following weaker statement does turn out to be necessary to guarantee that $(G^\Theta, E(\Lambda))$ is a RAAG system, and is missing from [LaF17].

Definition 3.14 (Condition $\mathcal{R}_4$)**.** We say that Θ satisfies *condition* $\mathcal{R}_4$ if the following holds. Let γ be any $\Lambda_c\Lambda_d$-cycle in Θ visiting vertices $c_1, d_1, c_2, d_2 \ldots, c_n, d_n, c_1$ for some $n \geq 2$. Let T_c and T_d be the Λ-convex hulls of $\{c_1, \ldots, c_n\}$ and $\{d_1, \ldots, d_n\}$ respectively. Then every edge of γ is contained in a 2-component square of Θ with two vertices in T_c and two vertices in T_d. (See Figure 3.3.)

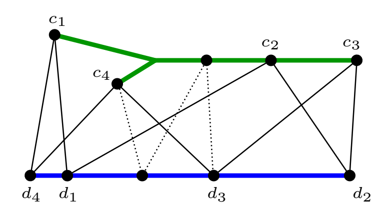

FIGURE 3.3. This figure illustrates condition $\mathcal{R}_4$. The green subgraph is T_c and the blue subgraph is T_d. The condition says that any edge in the 2-component cycle (shown in solid black edges) is part of a square of Γ with two vertices in T_c and two in T_d. This is illustrated here for the edge from d_3 to c_4. The dotted lines are Γ-edges which are not necessarily in the 2-component cycle.

The next lemma shows $\mathcal{R}_4$ is necessary for $(G^\Theta, E(\Lambda))$ to be a RAAG system.

Lemma 3.15. *If* $(G^\Theta, E(\Lambda))$ *is a RAAG system, then* Θ *satisfies* $\mathcal{R}_4$*.*

Proof. Let γ be a $\Lambda_c\Lambda_d$-cycle visiting vertices $c_1, d_1, \ldots, c_n, d_n, c_1$, and let T_c and T_d be as in Definition 3.14. Let w be the word below.

$$w = (d_n d_1)(c_1 c_2)(d_1 d_2)(c_2 c_3)(d_2 d_3) \cdots (d_{n-1} d_n)(c_n c_1) \tag{3.1}$$

Then $w \simeq 1$ in W_Γ. To see this, note that for each i, we know that c_i commutes with d_{i-1} and d_i (where i is taken mod n). Using this we can cancel the c_i for $i > 1$ in pairs to get

$$w \simeq d_n d_1 c_1 d_1 d_2 d_2 d_3 \cdots d_{n-1} d_n c_1 \simeq d_n c_1 d_n c_1 \simeq 1$$

Let w' be the Λ-edge word associated to w. Let D be a disk diagram over A_Δ with boundary label w'. As in the proof of Lemma 3.10, edges of D are labeled by Λ-edges, which are thought of as generators of A_Δ.

Color the part of the boundary of D and the hyperplanes coming out of it green if they correspond to Λ-edges from Λ_c and blue if they correspond to Λ-edges from Λ_d. Now we see from the structure of w', that ∂D alternates between green and blue stretches, and a stretch of a given color corresponds to a simple path in the corresponding component of Λ. It follows that a hyperplane of a given color must start and end in different stretches of that color.

Let $L = |E(T_c)|$ denote the number of Λ-edges in T_c. We will prove that condition $\mathcal{R}_4$ holds for γ by induction on (n, L). The conclusion of the lemma is obvious for γ corresponding to $(2, L)$ for any L, since the cycle itself is a square. This includes the base case, when $n = 2$ (i.e. γ is a square) and T_c is an edge. Now let $n > 2$, and assume the claim is true for all (n', L') such that either $n' < n$ or $n' = n$ and $L' < L$.

By Lemma 3.2, T_c and T_d are trees. Now suppose c_j is a leaf of T_c, and let xc_j be the Λ-edge incident to c_j in T_c. Since the $c_i \neq c_j$ for all $i \neq j$ (by the definition of a 2-component cycle), we know that xc_j occurs exactly once in w' (as part of the subword of w' representing $c_{j-1}c_j$) and c_jx occurs exactly once in w' (as part of the subword representing c_jc_{j+1}). It follows there is a unique hyperplane H labeled xc_j which is dual to both the path whose label is an expression for c_jc_{j+1} and the path whose label is an expression for $c_{j-1}c_j$. Moreover, the subword w'' of w' between these two subwords is the product of Λ-edges which is an expression for $d_{j-1}d_j$. It follows that every hyperplane dual to the path in ∂D labeled w'' must intersect the hyperplane H. By Lemma 2.6, both x and c_j commute (in W_Γ) with each letter of $V(\Gamma)$ used in the word w''. In particular, d_{j-1} and d_j each commute with x.

Now there are two possibilities. Suppose first that $x = c_t$ for some $t \neq j$. Since $t \neq j$ and $n > 2$ (which implies that γ has more than 4 edges), it follows that either c_td_{j-1} or c_td_j is a diagonal of γ. We can use this diagonal to cut γ into two 2-component cycles γ_1 and γ_2 as follows. Assume c_td_j is a diagonal δ of γ (the other case is analogous), and let β_1 and β_2 be the two components of γ obtained by removing the vertices labeled c_t and d_j. Set $\gamma_1 = \beta_1 \cup \delta$ and $\gamma_2 = \beta_2 \cup \delta$. Note γ_1 and γ_2 each have strictly fewer vertices than γ. For $i = 1, 2$ let T_c^i and T_d^i be the components of the Λ-convex hull of γ_i contained respectively in Λ_c and Λ_d. By the induction hypothesis, we see that every edge in γ_i is part of a square in Γ with two vertices in $T_c^i \subset T_c$ and two in $T_d^i \subset T_d$. Since each edge of γ is either an edge of γ_1 or of γ_2, we the claim follows for this case.

On the other hand, suppose that $x \neq c_i$ for any $1 \leq i \leq n$. Consider the new 2-component cycle γ' obtained from γ by replacing the edges $d_{j-1}c_j$ and c_jd_j with $d_{j-1}x$ and xd_j. As $x \neq c_i$ for any $1 \leq i \leq n$, this does not violate the requirement that 2-component cycles do not repeat vertices. Let T_c' and T_d' be the components of the Λ-convex hull of γ' contained respectively in Λ_c and Λ_d. Since c_j is a leaf of T_c, it follows that $|E(T_c')| < |E(T_c)|$, and we also have that $|V(\gamma')| = |V(\gamma)| = n$. We now apply the induction hypothesis to conclude that each edge of γ' is part of a square of Γ with two vertices in $T_d' = T_d$ and two vertices in $T_c' \subset T_c$. This means that this property holds automatically for all edges of γ, except possibly $d_{j-1}c_j$ and c_jd_j. However, these edges are part of the square in T_c with vertices x, c_j, d_{j-1} and d_j. Thus, the claim follows for this case as well. □

We summarize Lemma 3.2, Lemma 3.4, Lemma 3.10 and Lemma 3.15 into the following proposition:

Proposition 3.16. *If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Θ satisfies $\mathcal{R}_1 - \mathcal{R}_4$.*

If Λ has at most two components, then it turns out that there are no additional obstructions to $(G^\Theta, E(\Lambda))$ being a RAAG system. More precisely:

Theorem 3.17. *Suppose Λ has at most two components. Then $(G^{\Theta}, E(\Lambda))$ is a RAAG system if and only if $\mathcal{R}_1$–$\mathcal{R}_4$ are satisfied.*

Proof outline. Proposition 3.16 constitutes one direction of the theorem. The following strategy will be used to prove that $\mathcal{R}_1$–$\mathcal{R}_4$ imply that $(G^{\Theta}, E(\Lambda))$ is a RAAG system. We wish to show that the image of every non-trivial element of A_Δ under ϕ is non-trivial in W_Γ.

Let $g \in A_\Delta$ be a non-trivial element. Let $v = v_1 v_2 \cdots v_n$ be a reduced word over the set of the generators of A_Δ, which represents g. By the definition of A_Δ, we have that $\phi(v_i)$ is a Λ-edge of Θ, for $1 \leq i \leq n$. Then $u = \phi(v_1)\phi(v_2)\cdots\phi(v_n)$ is a concatenation of Λ-edges which represents $\phi(g)$. Towards a contradiction, we assume that u represents the identity element of W_Γ. Then there is a disk diagram D whose boundary label (read clockwise starting from some basepoint) is u.

To complete the proof of Theorem 3.17, we will put u in a certain normal form which will be defined in terms of the configuration of hyperplanes in D. The proof will then consist of analyzing hyperplanes in this disk diagram, to show that if $\mathcal{R}_1$–$\mathcal{R}_4$ are satisfied, then the normal form is violated.

Before we embark on the proof, we need to develop some preliminaries on disk diagrams, and on transferring information from the disk diagram D to the graph Θ. In what follows, we assume that D is a disk diagram whose boundary is a word in Λ-edges. Note that unlike in the proofs of Lemmas 3.10 and 3.15, we are now working in W_Γ rather than A_Δ, so the hyperplanes are labeled by generators of W_Γ rather than by Λ-edges.

We associate a color to each component of Λ. Each hyperplane then inherits the color corresponding to the component of Λ in which its label lies. Thus, the two edges of ∂D constituting a Λ-edge are dual to hyperplanes of the same color.

Observation 3.18. If Θ satisfies $\mathcal{R}_2$, then no two hyperplanes of the same color intersect. This is because if two hyperplanes intersect, then their labels are distinct and commute, and so are connected by a Γ-edge. Thus they cannot be in the same component of Λ, since each component of Λ is an induced subgraph of Θ, by $\mathcal{R}_2$.

The hyperplanes of D can be partitioned into "closed chains of hyperplanes," as described in Definition 3.19 below. Although the proof of Theorem 3.17 only uses disk diagrams whose boundary labels are words in Λ-edges, the definition below applies to slightly more general disk diagrams, as this will be needed in Section 6.

Definition 3.19. (Chains of hyperplanes) Let D be a disk diagram whose boundary ∂D contains a connected subpath η, such that the label of η is a word in Λ-edges. Let $H_0, \dots, H_n$ be a sequence of distinct hyperplanes in D. Let e_i and f_i be the edges on ∂D that are dual to H_i. We say that $\{H_0, \dots, H_n\}$ is a *chain* in D, if for all $0 \leq i < n$, the edges f_i and e_{i+1} are contained in η and are dual to the same Λ-edge of η. Note that e_0 and f_n can be dual to edges not contained in η.

Additionally, if e_0 and f_n are contained in the same Λ-edge of η, then we say that $\{H_0, \dots, H_n\}$ is a *closed chain*. (Figure 3.4 shows several closed chains.)

Since the two hyperplanes dual to a Λ-edge have the same color, each chain also inherits a well-defined color.

Observation 3.20. If the label of ∂D is a word in Λ-edges, then it readily follows that every hyperplane is contained in a unique closed chain. Thus, there is a partition of the hyperplanes of a given color into closed chains.

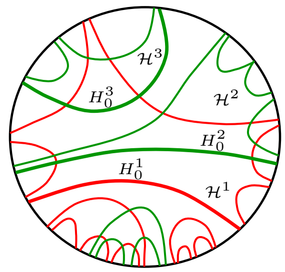

FIGURE 3.4. The figure illustrates the procedure for finding $\mathcal{H}$ and H_0 in Lemma 3.22. Each closed chain in the sequence is labeled in its interior. The hyperplanes H_0^i are shown in bold.

We say that a chain $\mathcal{K}$ intersects a hyperplane H, if some $K \in \mathcal{K}$ intersects H. We say that chains $\mathcal{H}$ and $\mathcal{K}$ intersect if $\mathcal{K}$ intersects some $H \in \mathcal{H}$. We will need the following observation:

Observation 3.21. If a hyperplane H intersects a closed chain $\mathcal{K}$, then it intersects $\mathcal{K}$ in exactly two distinct hyperplanes. To see this, note that given a hyperplane $K \in \mathcal{K}$, the hyperplanes in $\mathcal{K} \setminus \{K\}$ all lie in a single component of $D \setminus K$. It follows that if H intersects $\mathcal{K}$ more than twice, it must intersect some hyperplane of $\mathcal{K}$ twice. This contradicts the fact that D has no bigons (see Remark 2.9).

In order to define a normal form for the word u from the proof outline, we will need to choose a closed chain in D with some special properties:

Lemma 3.22. *Let u and D be as in the proof outline. There exists a closed chain $\mathcal{H}$ of D, containing a distinguished hyperplane H_0, such that given any closed chain $\mathcal{K} \neq \mathcal{H}$, one of the following holds:*

(1) *$\mathcal{K}$ and $\mathcal{H} \setminus \{H_0\}$ lie in different components of $D \setminus H_0$.*
(2) *$\mathcal{K}$ intersects H_0.*

Proof. We inductively construct a sequence of closed chains $\mathcal{H}^1, \mathcal{H}^2, \dots$ with distinguished hyperplanes $H_0^1, H_0^2, \dots$, such that for all $i > 1$, we have:

(i) $\mathcal{H}^i$ and $\mathcal{H}^{i-1} \setminus \{H_0^{i-1}\}$ lie in the same component of $D \setminus H_0^{i-1}$, and
(ii) H_0^{i-1} and $\mathcal{H}^i \setminus \{H_0^i\}$ lie in different components of $D \setminus H_0^i$.

Let $\mathcal{H}^1$ and H_0^1 be arbitrary. Now for any j, if $\mathcal{H}^j$ and H_0^j do not satisfy the conclusion of the lemma, then there must exist another closed chain $\mathcal{H}^{j+1}$ which lies entirely in C_j, where C_j is the component of $D \setminus H_0^j$ containing $\mathcal{H} \setminus \{H_0^j\}$. (Figure 3.4 illustrates this for $j = 1, 2$.) There is a unique hyperplane in $\mathcal{H}^{j+1}$ satisfying condition (ii) above with $i = j+1$, and we set this equal to H_0^{j+1}. Thus, we can produce a longer sequence of closed chains with properties (i) and (ii).

By construction, there is a nesting of components $C_1 \supset C_2 \supset C_3 \dots$, and it follows that $H_0^1, H_0^2, \dots$ are distinct hyperplanes in D. As D has finitely many hyperplanes, this process can only be repeated finitely many times. Thus, $\mathcal{H}^j$ satisfies the claim for some j. ∎

The following two observations enable us to transfer information from the disk diagram D to the graph Θ.

Observation 3.23. (Chains in D give Λ-paths in Θ.) Let $\mathcal{K} = \{K_0, \dots, K_l\}$ be a chain in D, and for $0 \leq i \leq l$, let k_i be the label of K_i. Then by the definition of a chain, K_i and K_{i+1} are dual to the same Λ-edge in ∂D for each i, so there is an edge in Λ between k_i and k_{i+1}. It follows that $\mathcal{K}$ naturally defines a Λ-path in Θ visiting vertices $k_0, k_1, \dots, k_l$. Moreover, if $\mathcal{K}$ is a closed chain, then the corresponding Λ-path is a loop.

Observation 3.24. (Pairs of intersecting closed chains give 2-component loops in Θ.) Consider two closed chains which intersect, say a red chain $\mathcal{K}$ and a green chain $\mathcal{L}$. Let $K_1 \in \mathcal{K}$ and $L_1 \in \mathcal{L}$ be intersecting hyperplanes. By Observation 3.21, the hyperplane L_1 intersects $\mathcal{K}$ in a second hyperplane $K_2 \neq K_1$. Similarly, K_2 intersects $\mathcal{L}$ in a second hyperplane L_2. Proceeding in this way, we obtain a polygon with at least four sides, with sides alternating between red and green hyperplanes.

Since an intersecting pair of hyperplanes corresponds to an edge of Γ, a 2-colored polygon of the type we just constructed defines a 2-component loop in Θ. We warn that the 2-component loop obtained from a 2-colored polygon in D may not be a 2-component cycle. (Note that a 2-component cycle is a 2-component loop in which all of the vertices are distinct, and there are at least two vertices in each component.)

We are now ready to prove the theorem.

Proof of Theorem 3.17. As discussed, we need to show that if $\mathcal{R}_1$–$\mathcal{R}_4$ are satisfied, then ϕ is injective. Let g, v, u, and D be as in the proof outline above. By Lemma 2.8 we may assume that D has no bigons.

If an element has non-trivial image under ϕ, then so does every element of its conjugacy class. Thus, we may assume that g is of minimal length in its conjugacy class, where the length of an element is defined to be the minimal length of a word representing it.

We partition the hyperplanes of D into disjoint closed chains. (See Observation 3.20.) By Lemma 3.22, we can choose a chain $\mathcal{H}$, with distinguished hyperplane H_0, such that given any other chain $\mathcal{K}$, either H_0 separates $\mathcal{K}$ from $\mathcal{H} \setminus \{H_0\}$, or $\mathcal{K}$ intersects H_0. Let $a_0, a_1, \dots, a_s$ be the labels of the hyperplanes of $\mathcal{H}$, starting from H_0, and proceeding in order in the clockwise direction around ∂D. Then the Λ-edges $a_0a_1, a_1a_2, \dots, a_{s-1}a_s, a_sa_0$ appear in ∂D in that order, possibly interspersed with some other Λ-edges.

Let p denote the vertex on ∂D which is the endpoint of the Λ-edge from $\mathcal{H}$ labeled a_sa_0, read clockwise. (See Figure 3.5.) Let w be the word labeling ∂D clockwise, starting from p. Then w is a cyclic conjugate of u. Let x be the corresponding cyclic conjugate of v. Since v was chosen to be reduced, and since g (the element of A_Δ represented by v) is of minimal length in its conjugacy class by assumption, it follows that x is reduced.

We now show that we can modify D in such a way that the resultant boundary label is a word representing $w = \phi(x)$ which is in a certain normal form:

Claim 3.25. There exists a disk diagram $\tilde{D}$ such that the following hold.

1. There is a closed chain $\widetilde{\mathcal{H}}$ in $\tilde{D}$ which has a distinguished hyperplane $\tilde{H}_0$ satisfying the criterion in Lemma 3.22. The labels of the hyperplanes of $\widetilde{\mathcal{H}}$ starting from $\tilde{H}_0$ and proceeding clockwise, are $a_0, \dots, a_s$ (i.e. they are the same labels as the labels of the hyperplanes in $\mathcal{H}$).

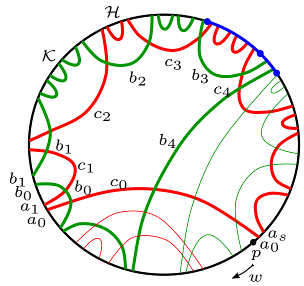

FIGURE 3.5. This example illustrates the proof of Theorem 3.17. The chain $\mathcal{H}$ satisfying Claim 3.25 is shown in thick red lines. In particular, no Λ-edge from $\mathcal{H}$ (except possibly the last one) commutes with the Λ-edge appearing after it in ∂D. The chain $\mathcal{K}$, which contributes the first Λ-edge not in $\mathcal{H}$ (after a_0a_1), is shown in thick green lines. The polygon formed by the intersection of $\mathcal{H}$ and $\mathcal{K}$ induces a 2-component loop which visits (in this example) $c_0, b_0, c_1, b_1, c_2, b_2, c_3, b_3, c_4, b_4, c_0$. The blue subpaths of ∂D are the subpaths defined in Claim 3.27, with $i = 4$.

(2) Let $\tilde{p}$ be the endpoint of the Λ-edge a_sa_0 from $\widetilde{\mathcal{H}}$, and let $\tilde{w}$ be the word labeling ∂D in the clockwise direction starting from $\tilde{p}$. Then $\tilde{w} = \phi(\tilde{x})$, where $\tilde{x}$ is a reduced word in A_Δ obtained from x by Tits swap moves.
(3) The Λ-edges from $\widetilde{\mathcal{H}}$ appear as far right as possible in $\tilde{w}$. More formally, the word $\tilde{w}$ has no subword of the form $a_ia_{i+1}bb'$ such that a_ia_{i+1} is one of the Λ-edges coming from $\widetilde{\mathcal{H}}$ with $0 \le i \le s$ (with indices mod s), $a_ia_{i+1} \ne bb'$, and a_ia_{i+1} commutes with bb'.

Proof. We construct $\tilde{D}$ iteratively, starting with D. If $\mathcal{H}, p, w$, and x are as defined above, then the first two conditions in the claim are satisfied. If (3) is not satisfied, then w has a subword $a_ia_{i+1}bb'$ as in (3). Since a_sa_0 is the last Λ-edge of w, we conclude that $a_ia_{i+1} \ne a_sa_0$.

Note that $a_ia_{i+1} \ne bb'$ (by condition (3)) and $a_ia_{i+1} \ne (bb')^{-1}$ (since x is reduced and $w = \phi(x)$). Then it follows from Lemma 2.6, that each of a_i and a_{i+1} commutes with each of b and b'. By applying Lemma 2.11(1) four times, we obtain a new disk diagram D' such that the label of $\partial D'$ is obtained from the label of ∂D by swapping the Λ-edges a_ia_{i+1} and bb'. Moreover, the natural map ψ from the edges of ∂D to the edges of $\partial D'$ (defined in Lemma 2.11(1)) preserves boundary combinatorics. By applying Lemma 2.8 if necessary, we may assume that D' has no bigons, so hyperplanes in D' intersect at most once.

Since boundary combinatorics are preserved, ψ induces a bijection between the hyperplanes dual to ∂D and those dual to $\partial D'$. Since the transition from D to D' involves swapping a pair of Λ-edges, the label of $\partial D'$ is still a product of Λ-edges, and so the hyperplanes of D' can be partitioned into closed chains of hyperplanes. Moreover, ψ induces a bijection between the closed chains of hyperplanes in D and D'.

If $\mathcal{H}'$ and H_0' denote the images of $\mathcal{H}$ and H_0 respectively under ψ, it is clear that the labels of the hyperplanes of $\mathcal{H}'$, starting from H_0' and proceeding clockwise, are

$a_0, \dots a_s$. We now prove that $\mathcal{H}'$ together with H_0' still satisfies the criterion in Lemma 3.22 which is required in (1).

Let $\mathcal{K}'$ be a closed chain in D', and let $\mathcal{K}$ be its preimage in D. Our choice of $\mathcal{H}$ implies that either H_0 separates $\mathcal{K}$ from $\mathcal{H} \setminus \{H_0\}$, or $\mathcal{K}$ intersects H_0. In the former case, H_0' still separates $\mathcal{K}'$ from $\mathcal{H}' \setminus \{H_0'\}$. This is because the swap performed does not involve any hyperplanes from chains which do not intersect H_0, since (as noted above) $a_i a_{i+1} \neq a_s a_0$.

On the other hand, suppose that $\mathcal{K}$ intersects H_0. By Observation 3.21, there are exactly two hyperplanes K_1 and K_2 in $\mathcal{K}$ which intersect H_0. If K_j, for $j = 1, 2$, is not dual to the Λ-edge labeled by bb', then the image of K_j intersects H_0'. Moreover, if $i \neq 0$, then it follows that the images of K_1 and K_2 in D' intersect the hyperplane H_0'. Thus, we only need to consider the case where the Λ-edge $a_0 a_1$ is swapped, and (up to relabeling) K_1 is dual to b and K_2 is dual to b'. In this case, K_1 and K_2 are dual to the same Λ-edge. It follows that no hyperplane in $\mathcal{K} \setminus \{K_1, K_2\}$ is contained in the same component of $D \setminus H_0$ as $\mathcal{H} \setminus \{H_0\}$. Thus, in D', no hyperplane of $\mathcal{K}'$ is contained in the same component of $D' \setminus H_0'$ as $\mathcal{H}' \setminus \{H_0'\}$. We have shown that $\mathcal{H}'$, with distinguished hyperplane H_0', satisfies the conclusion of Lemma 3.22.

Let p' be the vertex on $\partial D'$ which is the endpoint of the Λ-edge from $\mathcal{H}'$ labeled $a_s a_0$. Since the swap performed did not involve $a_s a_0$, the label w' of $\partial D'$, read clockwise from p', is obtained from w by swapping a single pair of Λ-edges, and its preimage in x' in A_Δ is obtained from x by swapping one pair of generators. This shows (2).

We have established that D', together with $\mathcal{H}'$, satisfies (1) and (2) of Claim 3.25. If (3) still fails, we may repeat the process above. Since each individual iteration involves moving one Λ-edge from the image of $\mathcal{H}$ to the right, this process eventually stops. After finitely many iterations, we arrive at a disk diagram $\tilde{D}$ such that all three conditions hold. □

For the rest of the proof we assume, without loss of generality, that $D, \mathcal{H}, p, w$, and x satisfy the conclusion of Claim 3.25.

We now analyze closed chains which intersect $\mathcal{H}$. First consider the case that there are no such chains. This includes the case when Λ has a single component. Since $\mathcal{H}$ is a closed chain, it defines a loop in Λ. (See Observation 3.23.) On the other hand, since no chains intersect $\mathcal{H}$, the union of the edges of ∂D dual to the hyperplanes of $\mathcal{H}$ is a continuous subpath (with label $(a_0a_1)(a_1a_2)\cdots(a_sa_0)$). Applying the following claim to this subpath, we conclude that the Λ-path defined by $\mathcal{H}$ is simple, a contradiction. (The claim will be used again later in this proof.)

Claim 3.26. Let ν be a subpath of ∂D labeled by a product of Λ-edges. Suppose there exists a closed chain $\mathcal{X}$, such that each edge of ν is dual to a hyperplane in $\mathcal{X}$. It follows that the label of ν is $(x_1x_2)\cdots(x_{n-1}x_n)$, where $x_1, x_2, \dots, x_n$ are the labels of the hyperplanes of $\mathcal{X}$ dual to ν, in order. Furthermore, the Λ-path through vertices $x_1, \dots x_n$ is simple.

Proof. The claim about the label of ν is immediate. If the path through vertices $x_1, \dots, x_n$ is not simple, then there is a Λ-loop through vertices $x_i, x_{i+1}, \dots, x_{i+j} = x_i$ for some i, j. By $\mathcal{R}_1$, the image of this loop in Λ is a tree. Let x_r be a leaf of this tree, with $i < r < j$. It follows that $x_{r-1} = x_{r+1}$. Consequently, the label of ν (and therefore of the word w) has a subword $(x_{r-1}x_r)(x_rx_{r-1})$. This is a contradiction, as it implies that the preimage x of w in A_Δ is not reduced. □

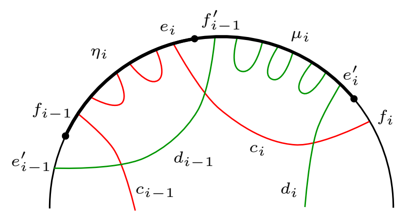

FIGURE 3.6. The paths η_i and μ_i from Lemma 3.27 are shown in bold, delineated by dots. We remark that if $i = k$, then there could be additional hyperplanes not in $\mathcal{K}$ or $\mathcal{H}$ between the endpoint of μ_k and the start of the edge f_i.

Thus, we may assume that there is at least one chain intersecting $\mathcal{H}$. In particular, Λ has two components: say a red component Λ_a which contains the labels of $\mathcal{H}$, and a green component Λ_b. By Claim 3.25, each chain intersecting $\mathcal{H}$ intersects H_0. Let $\mathcal{K}$ be the "first" such chain, in the sense that the first Λ-edge from a chain other than $\mathcal{H}$ appearing in w to the right of a_0a_1 is from $\mathcal{K}$. (See Figure 3.5.) By $\mathcal{R}_2$ and Observation 3.18, we conclude that $\mathcal{K}$ is green. Let $b_0, \dots b_{s'}$ be the labels of the hyperplanes of $\mathcal{K}$, where b_0b_1 is the label of the first Λ-edge from $\mathcal{K}$ appearing in w to the right of a_0a_1.

Now consider the 2-colored polygon in D whose sides alternate between hyperplanes in $\mathcal{H}$ and $\mathcal{K}$, as described in Observation 3.24. Let $c_0, d_0, \dots c_k, d_k$ be the labels of these sides, where $c_0 = a_0$, $d_0 = b_0$, and $c_0, \dots c_k$ (resp. $d_0, \dots, d_k$) is a subsequence of $a_0, \dots a_s$ (resp. of $b_0, \dots, b_{s'}$). (See Figure 3.5.)

The following technical claim about the hyperplanes dual to certain subpaths of ∂D associated to this 2-colored polygon will be needed in what follows:

Claim 3.27. For $0 \leq i \leq k$, let e_i and f_i (resp. e'_i and f'_i) be the edges dual to the hyperplane of $\mathcal{H}$ labeled c_i (resp. the hyperplane of $\mathcal{K}$ labeled d_i), where e_i (resp. e'_i) appears before f_i (resp. f'_i) reading clockwise from p.

For $i > 0$, let η_i be the subpath of ∂D from (and including) f_{i-1} to (and including) e_i, and let μ_i be the subpath of ∂D from the endpoint of η_i to (and including) e'_i. (Figure 3.6.) Then every edge of η_i (resp. μ_i) is dual to a hyperplane in $\mathcal{H}$ (resp. $\mathcal{K}$).

Proof. Suppose there is some hyperplane L dual to an edge e of η_i such that the closed chain $\mathcal{L}$ containing L is not equal to $\mathcal{H}$. From the definition of η_i, we conclude that e is on the same side of H_0 as $\mathcal{H} \setminus \{H_0\}$ in D, and by our choice of $\mathcal{H}$ (and Lemma 3.22), it follows that $\mathcal{L}$ intersects H_0. Therefore, L is green by Observation 3.18.

Let K denote the hyperplane of $\mathcal{K}$ labeled d_{i-1}. Then K separates e from $\mathcal{K} \setminus \{K\}$, so $L \notin \mathcal{K}$, i.e. $\mathcal{L} \neq \mathcal{K}$. If $d_{i-1} = b_0$, i.e. if K does intersect H_0, then our choice of $\mathcal{K}$ implies that K is the first hyperplane not in $\mathcal{H}$ dual to ∂D after the Λ-edge a_0a_1, so such an L cannot exist. On the other hand, if K does not intersect H_0, then K separates e from H_0. So, in order to intersect H_0, the chain $\mathcal{L}$ must also intersect $\mathcal{K}$, which is a contradiction, since $\mathcal{L}$ and $\mathcal{K}$ are both green.

Now suppose $L \in \mathcal{L} \neq \mathcal{K}$ is dual to an edge e of μ_i. Since μ_i is only defined for $i > 0$, it is on the same side of H_0 as $\mathcal{H} \setminus \{H_0\}$, and consequently, the same holds for e. Therefore, we conclude as before that L is green.

Additionally, we conclude as before that the hyperplane $K \in \mathcal{K}$ labeled d_{i-1} does not intersect H_0. Now consider the subchain of $\mathcal{K}'$ of $\mathcal{K}$ consisting of the hyperplanes dual to all but the last edge e_i' of μ_i. Since K does not intersect H_0, it follows that e is separated from H_0 by some hyperplane in $\mathcal{K}'$. Thus, in order to intersect H_0, $\mathcal{L}$ must intersect $\mathcal{K}$, which is again a contradiction. □

The 2-colored polygon obtained above gives a 2-component loop in Θ, as described in Observation 3.24. A priori this loop may not be a 2-component cycle, i.e., it is possible that $c_i = c_j$ or $d_i = d_j$ for some i, j. However, we now show that it contains a cycle. We will then be able to apply $\mathcal{R}_4$ to this cycle to make progress towards obtaining a contradiction to the normal from in Claim 3.25.

Claim 3.28. Consider the 2-component loop in Θ visiting $c_0, d_0, \dots c_k, d_k, c_{k+1} = c_0$ defined above. There exist $0 \leq l \leq k-1$ and $m \geq 2$, such that one of the two following subsequences of vertices (with indices taken mod 2) defines a 2-component cycle in Θ:

(1) $c_l, d_l, c_{l+1}, \dots, d_{l+m-1}, c_{l+m} = c_l$
(2) $d_l, c_{l+1}, d_{l+1} \dots, c_{l+m}, d_{l+m} = d_l$

Proof. Observe that since $c_{k+1} = c_0$, the following set is non-empty:

$$\{j \mid c_i = c_{i+j} \text{ or } d_i = d_{i+j} \text{ for some } 0 \leq i < k-1 \text{ and } 1 \leq j \leq k+1\}$$

Let m denote its minimum value. We first show that $m \geq 2$, or equivalently that, for each $0 \leq i \leq k-1$, both $c_i \neq c_{i+1}$ and $d_i \neq d_{i+1}$ are true. Suppose $c_i = c_{i+1}$ for some i. Consider the path η_{i+1} from Claim 3.27. It is labeled by Λ-edges, and every edge in it is dual to a hyperplane from $\mathcal{H}$. Then by Claim 3.26, it follows that η_{i+1} defines a simple Λ-path from the vertex c_i to the vertex c_{i+1}. However, this contradicts the assumption that $c_i = c_{i+1}$. This proves that for all $0 \leq i \leq k-1$, we have $c_i \neq c_{i+1}$. The proof that $d_i \neq d_{i+1}$ is similar.

Now if l is such that $c_l = c_{l+m}$ (the case when $d_l = d_{l+m}$ is similar), then it readily follows from the minimality of m that the vertices $c_l, d_l, c_{l+1}, \dots c_{l+m-1}, d_{l+m-1}$ are distinct, and therefore define the desired cycle. □

Continuing the proof of the theorem, we can now assume Θ has a 2-component cycle γ as in (1) from Claim 3.28. (The case in which Θ has a 2-component cycle as in (2) is similar.) Let T_c and T_d be the Λ-convex hulls of $\{c_l, \dots, c_{l+m-1}\}$ and $\{d_l, \dots d_{l+m-1}\}$ respectively. Then T_c and T_d are trees by $\mathcal{R}_1$. Let c_j be a leaf of T_c with $c_j \neq c_0$. Then c_j labels a hyperplane $H_t \in \mathcal{H}$ for some $t \neq 0$, so that $a_t = c_j$. Similarly, c_{j-1} labels a hyperplane H_{t-r} of $\mathcal{H}$, while d_{j-1} and d_j label hyperplanes $K_{t'}$ and $K_{t+r'}$ respectively of $\mathcal{K}$, where $d_{j-1} = b_{t'}$ and $d_j = b_{t'+r'}$

Consider the paths η_j and μ_j defined in Claim 3.27. The last Λ-edge of η_j is $a_{t-1}a_t$. By Claim 3.27, the first edge of μ_j is dual to a hyperplane in $\mathcal{K}$. It follows that this must be $K_{t'}$, with label $b_{t'}$, for otherwise $K_{t'}$ would separate this edge from $\mathcal{K} \setminus K_{t'}$. It follows that the first Λ-edge of μ_i is $b_{t'}b_{t'+1}$, and that the word w has a subword $a_{t-1}a_t b_{t'} b_{t'+1}$.

To complete the proof, we will show that the presence of this subword violates the normal form established in Claim 3.25(3). Note that since the labels of $\mathcal{H}$ and $\mathcal{K}$ are from different components of Λ, it is immediate that $a_{t-1}a_t \neq b_{t'}b_{t'+1}$. We now show that $a_{t-1}a_t$ and $b_{t'}b_{t'+1}$ commute.

The 2-component cycle γ in Θ contains an edge with endpoints a_t and $b_{t'}$. Applying $\mathcal{R}_4$ to this edge, we conclude that there is a 2-component square visiting

$a_t, b_{t'}, a$, and b, where $a \in T_c$ and $b \in T_d$. Next, applying $\mathcal{R}_3$ to this 2-component square, we see that $b_{t'}$ commutes with the vertices of the Λ-convex hull of $\{a_t, a\}$. Claim 3.27 and Claim 3.26 together imply that the path η_j induces a simple Λ-path visiting vertices $a_{t-r}, a_{t-r+1}, \dots, a_t$. Consequently, the vertices along this path, and in particular a_{t-1}, are in T_c. Moreover, a_{t-1} is the unique vertex of T_c adjacent to a_t, since $a_t = c_j$ is a leaf of T_c. It follows that a_{t-1} is contained in the Λ-convex hull (which is the same as the T_c-convex hull) of $\{a_t, a\}$. Thus, a_{t-1} and $b_{t'}$ commute. The same reasoning, applied to the edge of γ with endpoints a_t and $b_{t+r'}$, implies that a_{t-1} and $b_{t+r'}$ commute.

Using the Γ-edges whose existence is implied by these two additional commutation relations, we obtain a 2-component square visiting $a_t, b_{t'}, a_{t-1}, b_{t'+r'}$. Applying $\mathcal{R}_3$ to this square, we conclude that a_t and a_{t-1} commute with each vertex in the Λ-convex hull of $\{b_{t'}, b_{t'+r'}\}$. By Claim 3.26, we see that the path μ_j from Claim 3.27 defines a simple Λ-path visiting $b_{t'}, b_{t'+1}, \dots b_{t'+r'}$. It follows that $b_{t'+1}$ is in the convex hull of $\{b_{t'}, b_{t'+r'}\}$, and consequently, a_t and a_{t-1} commute with $b_{t'+1}$.

Putting together the commutation relations established in the previous paragraphs, we conclude that $a_{t-1}a_t$ commutes with $b_{t'}b_{t'+1}$. This contradicts the fact that we have chosen D so that it satisfies (3) of Claim 3.25. □

3.1. **Three or more Λ-components.** In the case that Λ contains at most two components, Theorem 3.17 shows that $\mathcal{R}_1 - \mathcal{R}_4$ are necessary and sufficient conditions that guarantee $(G^\Theta, E(\Lambda))$ is a RAAG system. In this subsection, we do not place any restriction on the number of components of Λ. We give an additional necessary Condition $\mathcal{R}_5$ for $(G^\Theta, E(\Lambda))$ to be a RAAG system, and Example 3.30 shows this condition is independent of conditions $\mathcal{R}_1 - \mathcal{R}_4$. The authors are aware that *even more* conditions are required in order to generalize Theorem 3.17 to this setting. These extra conditions are not included here, as they are complicated and the authors do not believe to yet possess the complete list of the necessary and sufficient conditions for this generalization.

We further show in this subsection that if Θ contains certain subgraphs and $(G^\Theta, E(\Lambda))$ is a RAAG system, then Γ must necessarily contain a triangle. These results are needed in the next section.

Definition 3.29 (Condition $\mathcal{R}_5$)**.** We say that Θ satisfies *condition* $\mathcal{R}_5$ if the following holds. Let Λ_a, Λ_c and Λ_d be distinct components of Λ. Suppose we have vertices $a, a' \in \Lambda_a$, $c, c' \in \Lambda_c$ and $d, d' \in \Lambda_d$, such that Θ contains a 2-component square visiting c, d, c' and d'. Furthermore, suppose that c and c' are each adjacent to a in Γ and that d and d' are each adjacent to a' in Γ. (See Figure 3.7.) Let T_a, T_c and T_d be the Λ-convex hulls of $\{a, a'\}$, $\{c, c'\}$ and $\{d, d'\}$ respectively. Then given any Λ-edge xx' of T_a, the graph Γ contains either the join of $\{x, x'\}$ with $V(T_c)$ or the join of $\{x, x'\}$ with $V(T_d)$.

The following is a concrete example showing that when Λ has more than two components, the conditions $\mathcal{R}_1 - \mathcal{R}_4$ are not sufficient to guarantee that $(G^\Theta, E(\Lambda))$ is a RAAG system.

Example 3.30. Let Γ be the graph whose vertex set is $\{a, a', c, c', d, d'\}$ and whose edge set is the set of black edges in Figure 3.7. Let $\Lambda \subset \Gamma^c$ consist of exactly three Λ-edges: aa', cc', dd'. Then $\Theta = \Theta(\Gamma, \Lambda)$ satisfies conditions $\mathcal{R}_1 - \mathcal{R}_4$ and does not satisfy condition $\mathcal{R}_5$. By Lemma 3.31 below, $(G^\Theta, E(\Lambda))$ is not a RAAG system.

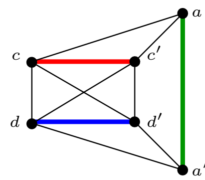

FIGURE 3.7. This figure illustrates condition $\mathcal{R}_5$. The red, blue and green segments are respectively T_c, T_d and T_a. Condition $\mathcal{R}_5$ states that any Λ-edge contained in the green segment must either commute with every Λ-edge in the red segment or must commute with every Λ-edge in the blue segment.

We now show that condition $\mathcal{R}_5$ is necessary.

Lemma 3.31. *If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Θ satisfies condition $\mathcal{R}_5$.*

Proof. By Theorem 3.17, we may assume that Θ satisfies conditions $\mathcal{R}_1 - \mathcal{R}_4$. Let $a, a' \in \Lambda_a$, $c, c' \in \Lambda_c$ and $d, d' \in \Lambda_d$ be as in Definition 3.29. Define the words $z_a = a'a$, $z_c = cc'$, $z_d = dd'$ and $z = [z_a z_c z_a^{-1}, z_d]$. By the commuting relations imposed in Definition 3.29, it follows that $z \doteq 1$ in W_Γ. Let w_a, w_c, w_d and w be the Λ-edge words corresponding respectively to z_a, z_c, z_d and z. Let D be a disk diagram over A_Δ with boundary label w.

Let γ_c, ζ_c, γ_d and ζ_d be the paths in ∂D labeled respectively by w_c, w_c^{-1}, w_d and w_d^{-1}. Note that no hyperplane is dual to two distinct edges of γ_c (resp. ζ_c, γ_d and ζ_d). This follows as z_c is a word in unique Λ-edges. Thus, every hyperplane dual to γ_c (resp. γ_d) is also dual to ζ_c (resp. γ_d).

Let α be a path in ∂D between γ_c and γ_d (which is labeled by w_a). Again, no hyperplane is dual to two distinct edges of α. Let xx' be a Λ-edge of T_a, and let H be the unique hyperplane dual to α with label xx'. Note that either H intersects every hyperplane dual to γ_c or H intersects every hyperplane dual to γ_d. Furthermore, every Λ-edge of T_c (resp. T_d) is the label of a hyperplane dual to γ_c (resp. γ_d). The claim now follows from Lemma 2.6, and the fact that intersecting hyperplanes correspond to commuting generators of A_Δ. □

The following corollary shows that if Θ contains a configuration like that in the hypothesis of condition $\mathcal{R}_5$, then Γ must contain a triangle. This corollary is a warm-up to the more complicated Lemma 3.33.

Corollary 3.32. *If $(G^\Theta, E(\Lambda))$ is a RAAG system and Θ contains a set of vertices $\{a, a', b, b', c, c'\}$ satisfying the hypothesis of $\mathcal{R}_5$, then Γ contains a triangle.*

Proof. Let $P = \{a, a', c, c', d, d'\}$ be a subset of vertices of Θ satisfying the hypothesis of $\mathcal{R}_5$. We call such a P a *configuration* in Θ. Keeping the same notation as in Definition 3.29, we call the number of vertices of T_a the *complexity* of P, and we prove the claim by induction on complexity. Note that $a = a'$ is possible in the hypothesis of $\mathcal{R}_5$, so the lowest possible complexity is $N = 1$. The corollary follows in this case, as Γ then contains a triangle spanned by the vertices $a = a'$, c and d.

Now let $N > 1$ and suppose the claim is true for all configurations P of smaller complexity. As $N > 1$, there is a vertex y such that ay is a Λ-edge of T_a. By Lemma 3.31, either y is adjacent in Γ to both c and c', or y is adjacent in Γ to both d and d'. In either case, we see that Θ contains a configuration of smaller complexity. □

The next lemma shows that if Θ contains certain subgraphs which generalize the configurations in the hypothesis of $\mathcal{R}_5$, then Γ must contain a triangle.

Lemma 3.33. *Let Λ_a, Λ_c and Λ_d be distinct components of Λ. Suppose Θ has a $\Lambda_a\Lambda_c$-path visiting $c_1, a_1, c_2, \dots, a_{n-1}, c_n$, and a $\Lambda_a\Lambda_d$-path visiting $d_1, a'_1, d_2, \dots, a'_{m-1}, d_m$, where $c_i \in \Lambda_c$, $d_i \in \Lambda_d$, and $a_i, a'_i \in \Lambda_a$ for all appropriate i. Further suppose that Θ contains a 2-component square visiting c_1, d_1, c_n and d_n. (See Figure 3.8). If $(G^\Theta, E(\Lambda))$ is a RAAG system, then Γ has a triangle.*

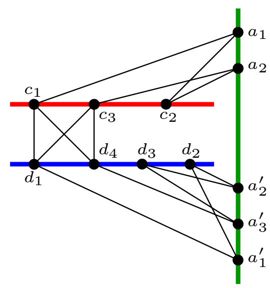

FIGURE 3.8. This figure illustrates the configuration described in Lemma 3.33 in the case $n = 3$ and $m = 4$. The black edges are edges of Γ. The red, green and blue parts consist of Λ-edges, and are all contained in Λ. The different colors indicate that they are in three distinct components of Λ.

Proof. By Theorem 3.17 and Lemma 3.31, we may assume that Θ satisfies conditions $\mathcal{R}_1 - \mathcal{R}_5$. Let $A = \{a_1, \dots, a_{n-1}, a'_1, \dots, a'_{m-1}\}$, $C = \{c_1, \dots, c_n\}$ and $D = \{d_1, \dots, d_m\}$ be vertices of Θ as in the statement of the lemma. We call such a triple (A, B, C) a configuration of Θ. Let T_a, T_c and T_d be the Λ-convex hulls of A, C and D respectively. We define the complexity of (A, C, D) to be the integer $N = |C| + |D| + |T_a|_E + |T_c|_E + |T_d|_E$, where $|X|_E$ denotes the number of edges in a graph X. The proof will be by induction on complexity of configurations.

By hypothesis, we have that $n, m \geq 2$ and $|T_c|_E, |T_d|_E \geq 1$. If $n = m = 2$ then Γ contains a triangle by Corollary 3.32. In particular, the base case follows.

We now fix a configuration $A = \{a_1, \dots, a_{n-1}, a'_1, \dots, a'_{m-1}\}$, $C = \{c_1, \dots, c_n\}$ and $D = \{d_1, \dots, d_m\}$ as above of complexity N, and we assume that the result holds for configurations of smaller complexity. By the previous paragraph, we may also assume (up to relabeling) that $n > 2$. We prove the lemma by showing that either Γ contains a triangle or Θ contains a configuration of smaller complexity.

Define α_{ac} and α_{ad} to respectively be the hypothesized $\Lambda_a\Lambda_c$-path and $\Lambda_a\Lambda_d$-path. We may assume that α_{ac} and α_{ad} are simple paths, for if not, we would be able to excise a loop to obtain a configuration of smaller complexity.

We claim that for all $1 < i < n$, we may assume that c_i does not lies on the simple path in Λ_c from c_1 to c_n. For suppose there exists such a vertex c_i. By $\mathcal{R}_3$, it follows that c_i commutes with both d_1 and d_m. There then exists a $\Lambda_a\Lambda_c$ path visiting $c_1, a_1, \dots, a_{i-1}, c_i$, and it follows that Θ contains a configuration of smaller complexity (obtained by replacing α_{ac} with this new path). Thus we may make this assumption without loss of generality. Furthermore, as $n \geq 3$, there exists an integer j such that c_j is a leaf vertex of T_c and such that $1 < j < m$. We fix such a vertex c_j.

Define the word z_c to be:

$$z_c = (a_1'a_1)(c_1c_2)(a_1a_2)(c_2c_3)(a_2a_3)\cdots(a_{n-2}a_{n-1})(c_{n-1}c_n)(a_{n-1}a_1')$$

and define the word z_d, depending on the value of m, to be:

$$\begin{cases} z_d = d_1d_2 & \text{if } m = 2 \\ z_d = (d_1d_2)(a_1'a_2')(d_2d_3)(a_2'a_3')\cdots(a_{m-2}'a_{m-1}')(d_{m-1}d_m)(a_{m-1}'a_1') & \text{if } m > 2 \end{cases}$$

In W_Γ we have that $z_c \simeq a_1'c_1c_na_1'$ and $z_d \simeq a_1'd_1d_ma_1'$. Let $z = [z_c, z_d]$. Note that as c_1 and c_n commute with d_1 and d_m in W_Γ we have:

$$z \simeq [a_1'c_1c_na_1', a_1'd_1d_ma_1'] \simeq a_1'[c_1c_n, d_1d_m]a_1' \simeq 1$$

Let w_c, w_d and w be the Λ-edge words associated to z_c, z_d and z respectively. Let D be a disk diagram over A_Δ with boundary label w. Let γ_c, ζ_c, γ_d and ζ_d be the subpaths of ∂D labeled respectively by w_c, w_c^{-1}, w_d and w_d^{-1}.

Let yc_j be the Λ-edge of T_c incident to c_j. Since α_{ac} does not repeat vertices and since c_j is a leaf of T_c, it follows that w_c contains exactly two occurrences of the letter y contained in the subword labeled by $(yc_j)(a_{j-1}x_1)(x_1x_2)\cdots(x_la_j)(c_jy)$, where the x_i's are vertices in Λ_a. In particular, there are exactly four edges of ∂D (two on γ_c and two on ζ_c) labeled by either yc_j or c_jy. Correspondingly, there are exactly two hyperplanes, H and H' in D labeled yc_j.

We claim that we may assume that H is dual to both γ_c and ζ_c, and the same is true for H'. For suppose otherwise, and suppose that H is dual to two edges of γ_c. (The case of H' is similar.) It follows that any hyperplane dual to the subpath of γ_c labeled by $(a_{j-1}x_1)(x_1x_2)\ldots(x_la_j)$ (which lies between the endpoints of H) must intersect H. Thus, in particular, $(a_{j-1}x_1)$ and (x_la_j) commute with yc_j, and applying Lemma 2.6, we conclude that y commutes with both a_{j-1} and a_j. We now show that we can replace α_{ac} with a new $\Lambda_a\Lambda_c$-path from c_1 to c_n such that $|T_a|_E$ is reduced, and thus Θ contains a smaller complexity configuration. If y is not equal to any c_k for any $1 \leq k \leq m$, then we obtain this path by simply replacing c_j with y in α_{ac}. On the other hand, if $y = c_k$ for some k, then we replace α_{ac} with the $\Lambda_a\Lambda_c$ path visiting $c_1, a_1, \ldots, c_k, a_j, c_{j+1}, a_{j+1} \ldots a_{n-1}, c_n$ if $k < j$ and perform a similar replacement if $k > j$. In either case, we have produced a configuration of smaller complexity. Thus, we now assume that each of H and H' is dual to both γ_c and ζ_c.

Let Q and Q' be the hyperplanes in D dual to the edges of γ_c labeled by $a_{j-1}x_1$ and x_la_j respectively. If both Q and Q' intersect $H \cup H'$, then we can conclude, as above, that y commutes with both a_{j-1} and a_j. We can then find a smaller complexity configuration as in the previous paragraph. Thus, we can assume that either Q or Q' is dual to both γ_c and ζ_c. We assume that Q has this property (the case of Q' is similar).

We now examine hyperplanes dual to γ_d and ζ_d. If $m = 2$, then the unique hyperplane whose label contains d_1 is dual to both γ_d and ζ_d, and this hyperplane intersects both H and Q. Thus, d_1 commutes with both c_j and a_{j-1}. Since c_j commutes with a_{j-1}, it follows that Γ contains a triangle. On the other hand, if $m > 2$ by the same reasoning as before, we can assume there is a leaf vertex $d_{j'}$

of T_d and a hyperplane with label $y'd_{j'}$ that intersects both γ_d and ζ_d. This then implies that $d_{j'}$ commutes with both c_j and a_{j-1} and consequently, Γ contains a triangle. The lemma now follows. □

4. Finite index visual RAAGs

As in the previous section, given a simplicial graph Γ and a subgraph Λ of Γ^c with no isolated vertices, we set $\Theta = \Theta(\Gamma, \Lambda)$, and let G^Θ be the subgroup generated by $E(\Lambda)$. Our goal is to characterize graphs $\Lambda \subset \Gamma^c$ such that $(G^\Theta, E(\Lambda))$ is a RAAG system and G^Θ has finite index in W_Γ.

Suppose the graph Γ contains a vertex s which is Γ-adjacent to every other vertex of Γ. We say that s is a *cone vertex*. In this case, it easily follows that $W_{\Gamma\setminus s}$ has index 2 in W_Γ and that s cannot be contained in any Λ-edge.

We now recall a construction from [DL19] which will help us compute the index of G^Θ. The construction is general, but for simplicity, and as it is all that we use, we choose to only describe it in the context where Γ is triangle-free. We refer the reader to [DL19] for full details.

Let Γ be a triangle-free graph. We say a cell complex is Γ-labeled if every edge of the complex is labeled by a vertex of Γ. Let X be a Γ-labeled complex. Suppose two edges of X have the same label and a common endpoint. A *fold operation* produces a new complex from X by naturally identifying these two edges.

Suppose now that f_1 and f_2 are edges of X which share a common vertex u and whose labels $s_1, s_2 \in V(\Gamma)$ have an edge between them in Γ. Let c be a 2-cube with edges c_1, c_2, c_3 and c_4 such that $c_i \cap c_{i+1}$ is a vertex of c for each $i \mod 4$. We label c_1 and c_3 by s_1, and c_2 and c_4 by s_2. A *square attachment operation* produces a new complex from X by attaching c to X by identifying c_1 to f_1 and c_2 to f_2. Note that, unlike in [DL19], we do not need to define cube attachments for higher dimensional cubes, as we are in the case that Γ is triangle-free.

Finally, given a collection of 2-cubes in X with common boundary, we can produce a new complex from X by naturally identifying every 2-cube in this collection to a single 2-cube. In this case, we say a *cube identification operation* was performed to X.

We define a Γ-labeled complex Ω_0 associated to G^Θ as follows. First, we enumerate the Λ-edges as $s_1t_1, \dots, s_nt_n$, where s_i and t_i are the two endpoints of the ith Λ-edge. We set Ω_0 to be a bouquet of n circles, each of which is subdivided into two edges, such that the ith circle has label s_it_i.

Next, we describe a series of complexes built iteratively from Ω_0. These are

$$\Omega_0 \to \Omega_1 \to \Omega_2 \to \dots$$

For each $i > 0$, the complex Ω_i is obtained by either a fold, square attachment or square identification operation performed to Ω_{i-1}. Furthermore, we assume that the order of operations is as follows: first all possible fold and square identifications are performed, then all possible square attachment operations are applied to the resulting complex, and these processes are alternated (see [DL19] for details).

Let Ω be the direct limit of such a sequence. We call Ω a *completion* of G^Θ. In [DL19] we show that properties of Ω reflect those of the subgroup G^Θ.

The index of G^Θ can be determined by properties of Ω. We say that a vertex u of a Γ-labeled complex has *full valence* if for any vertex $s \in \Gamma$, there is an edge incident to u with label s. Below we present a version of [DL19, Theorem 6.9] together with [DL19, Lemma 6.8] under the hypotheses which we will need:

Theorem 4.1. *Let Γ be a triangle-free graph with no cone vertex. A subgroup $G < W_\Gamma$ has finite index in W_Γ if and only if Ω is finite and every vertex of Ω has full valence. Furthermore, if G is indeed of finite index, then its index is exactly the number of vertices of Ω.*

We introduce two new properties below which will help us characterize when G^Θ has finite index in W_Γ.

Definition 4.2 (Conditions $\mathcal{F}_1$ and $\mathcal{F}_2$)**.** We say that $\Theta = \Theta(\Gamma, \Lambda)$ satisfies *condition $\mathcal{F}_1$* if given any $s \in V(\Theta)$ which is not a cone vertex of Γ, it follows that s is the endpoint of some Λ-edge. We say that Θ satisfies *condition $\mathcal{F}_2$* if given any distinct components Λ_s, Λ_t of Λ, and vertices s of Λ_s and t of Λ_t, there is a $\Lambda_s\Lambda_t$-path in Θ from s to t.

Remark 4.3. Suppose Γ is connected, Λ contains exactly two components and that $\Theta = \Theta(\Gamma, \Lambda)$ satisfies $\mathcal{R}_2$ and $\mathcal{F}_1$. Then Θ satisfies $\mathcal{R}_2$. For given any two vertices contained in different components of Λ, as Γ is connected, there is a Γ-path between them. Furthermore, this has to be a 2-component path as Θ satisfies $\mathcal{R}_2$, and the two Λ-components this path visits have to be the ones containing the chosen vertices (as there are only two Λ components). This remark will prove to be useful when verifying whether certain graphs satisfy $\mathcal{F}_2$.

Remark 4.4. Suppose $\Theta = \Theta(\Gamma, \Lambda)$ satisfies $\mathcal{F}_2$, and let Λ_1 and Λ_2 be distinct Λ-components. Then there exists an $\Lambda_1\Lambda_2$-path between any two distinct vertices of Λ_1. To see this, let s and s' be distinct vertices of Λ_1, and let t be vertex of Λ_2. By $\mathcal{F}_2$ there is a $\Lambda_1\Lambda_2$-path from s to t, and similarly there is an $\Lambda_1\Lambda_2$-path from t to s'. Combining these two paths gives an $\Lambda_1\Lambda_2$-path form s to s'.

Lemma 4.5. *Let Γ be a triangle-free graph with no cone vertex, and let Λ be a subgraph of Γ^c with no isolated vertices, such that $(G^\Theta, E(\Lambda))$ is a RAAG system. If Λ has at most $k \leq 2$ components and Θ satisfies $\mathcal{F}_1$ and $\mathcal{F}_2$, then G^Θ is of index 2^k in W_Γ.*

We remark that this proof readily generalizes to the case of arbitrary k. However, we only need the case $k \leq 2$.

Proof. Let Ω_0 be the Γ-labeled complex defined above, and let Ω' be the complex obtained from Ω_0 by all possible fold operations.

Suppose first that Λ has one component. As Λ is connected, it is easily seen that Ω' consists of two vertices with an edge labeled by s between them for $s \in V(\Lambda)$. As Λ satisfies $\mathcal{R}_2$ by Proposition 3.16, no two vertices of Λ have an edge between them in Γ. Thus, no square attachments can be performed to Ω', and it follows that $\Omega = \Omega'$. Hence, Ω is finite and has exactly two vertices.

Note that by the description of $\Omega = \Omega'$ above, every vertex of Ω is adjacent to every edge of Ω. Also note that by condition $\mathcal{F}_1$, for every vertex $s \in \Gamma$ there is some edge in Ω labeled by s. From these two facts we deduce that every vertex of Ω has full valence. Thus, G^Θ has index 2 in W_Γ by Theorem 4.1.

Now suppose that Λ has two components Λ_1 and Λ_2. In this case, Ω' is readily seen to be a complex consisting of three vertices, u, v_1, v_2, with an edge from u to v_i labeled s corresponding to each vertex s of Λ_i, for $i = 1, 2$. By condition $\mathcal{F}_1$, the vertex u has full valence. Furthermore, by $\mathcal{R}_2$, for each $i \in \{1, 2\}$, no two edges of Ω' that are each adjacent to both v_i and u have labels which are adjacent in Γ.

Let Ω'' be the complex obtained from Ω' by performing all possible square attachment operations to Ω', and let Ω''' be the complex obtained from Ω'' by all possible fold and square identification operations. In particular, $\Omega'' = \Omega_l$ and $\Omega''' = \Omega_k$ for some $0 \leq l \leq k$. Let s, s' be distinct vertices of Λ_1, and let t be any vertex of Λ_2. By condition $\mathcal{F}_2$, there is a $\Lambda_1\Lambda_2$-path whose vertices are $s, t_1, s_1, t_2, s_2, \dots, t_m, s_m, t$ where $s_i \in \Lambda_1$ and $t_i \in \Lambda_2$ for all $1 \leq i \leq m$. Similarly, there is a $\Lambda_1\Lambda_2$-path whose vertices are $s', t'_1, s'_1, t'_2, s'_2, \dots, t'_n, s'_n, t$ where $s'_i \in \Lambda_1$ and $t'_i \in \Lambda_2$ for all $1 \leq i \leq n$. Thus, Ω'' must contain length two paths, which do not intersect u, from v_1 to v_2 with each of the following labels

$$t_1 s, t_1 s_1, t_2 s_1, t_2 s_2, \dots, t_m s_{m-1}, t_m s_m, t s_m,$$

and similarly length two paths, which do not intersect u, from v_1 to v_2 with each of the following labels

$$t'_1 s', t'_1 s'_1, t'_2 s'_1, t'_2 s'_2, \dots, t'_n s'_{m-1}, t'_m s'_m, t' s_m$$

It follows that the middle vertices of all these paths get folded to a single vertex v_3 in Ω'''. This analysis can be done for any $s, s' \in \Lambda_1$. Similar paths can also be produced for any $t, t' \in \Lambda_2$. It then follows that Ω''' consists of exactly 4 vertices: u, v_1, v_2 and v_3. Furthermore, there is an edge with label s between v_1 and v_3 for each $s \in \Lambda_1$, and there is an edge with label t between v_2 and v_3 for each $t \in \Lambda_2$. Thus, every vertex of Ω''' can be seen to have full valence. Additionally, by condition $\mathcal{R}_2$, no additional square attachment operations can be performed to Ω'''. Hence, $\Omega = \Omega'''$. It follows that G^Θ has index exactly four in W_Γ. $\square$

The next lemma shows that $\mathcal{F}_1$ and $\mathcal{F}_2$ are necessary conditions for G^Θ to be a finite-index subgroup.

Lemma 4.6. *Let Γ be a triangle-free graph with no cone vertex, and let Λ be a subgraph of Γ^c with no isolated vertices, such that $(G^\Theta, E(\Lambda))$ is a RAAG system. If $G = G^\Theta$ is of finite index in W_Γ, then Λ satisfies $\mathcal{F}_1$ and $\mathcal{F}_2$.*

Proof. We first check that condition $\mathcal{F}_1$ holds. Let Ω be a completion of $G := G^\Theta$ as described in the beginning of this section. Theorem 4.1 implies in particular that given any vertex $s \in \Gamma$ there is an edge of Ω with label s. This implies the vertex s is contained in some Λ-edge. Thus, $\mathcal{F}_1$ must hold.

We now check condition $\mathcal{F}_2$. Let $s \in \Lambda_s$ and $t \in \Lambda_t$ be as in the definition of condition $\mathcal{F}_2$ (Definition 4.2). If s commutes with t, then there is an edge in Γ between s and t, and we are done. So we may assume that s and t do not commute.

As G is of finite index, it follows that there exist $g_1, \dots, g_n \in W_\Gamma$ such that $W_\Gamma = Gg_1 \sqcup Gg_2 ... \sqcup Gg_n$. Let $w_1, \dots, w_n$ be reduced words representing $g_1, \dots, g_n$, and let $K = \max\{|w_1|, ..., |w_n|\}$. Define the word $h = s_1 t_1 s_2 t_2 ... s_{K+4} t_{K+4}$ where $s_i = s$ and $t_i = t$ for all $1 \leq i \leq K + 4$. It readily follows from Tits' solution to the word problem (see Theorem 2.4) that h is reduced. Furthermore, we can write $h \simeq ww'$, where w and w' are words in W_Γ such that $w' = w_i$ for some $1 \leq i \leq n$ and w is a product of Λ-edges representing an element of G. We can form a disk diagram in W_Γ with boundary label $hw'^{-1}w^{-1}$. Let α_h, α_w and $\alpha_{w'}$ respectively be the corresponding paths along the boundary of D with labels respectively h, w and w^{-1}.

Note that as h is reduced, no hyperplane intersects α_h twice. Also note that any pair of hyperplanes emanating from α_h cannot intersect as s and t do not commute. As $|h| > |w'| + 4$, it follows that the hyperplanes $H_{s_1}, H_{t_1}, H_{s_2}$ and H_{t_2}, dual respectively to the first four edges of α_h (namely those labeled by s_1, t_1, s_2 and t_2), must each intersect α_w. It must now be the case that there exists a chain of hyperplanes (see Definition 3.19) $H_{s_1} = H_0, H_1, \dots H_m = H_{s_2}$ and another chain of hyperplanes $H_{t_1} = H'_0, H'_1 \dots, H'_n = H_{t_1}$. These two chains intersect, and by reasoning similar to that in Observation 3.24, it follows that there is a $\Lambda_s\Lambda_t$-path from s to t. ☐

Lemma 4.7. *Let Γ be a triangle-free graph. Let Λ be a subgraph of Γ^c with no isolated vertices, such that $(G^\Theta, E(\Lambda))$ is a RAAG system and G^Θ has finite index in W_Γ. If Γ contains a cone vertex, then Λ contains exactly one component. If W_Γ is not virtually free, then Λ contains exactly two components. Otherwise, Λ contains at most two components.*

Proof. Suppose first that Γ contains a cone vertex $s \in \Gamma$. We may assume that Γ does not consist of a single edge, as Λ would be empty in that case. As Γ is triangle-free in addition, there can be at most one cone vertex. Since Γ is triangle-free, it follows that $\Gamma' = \Gamma \setminus s$ is a graph with no edges and is therefore virtually free. Furthermore, every Λ-edge is contained in Γ', and G^Θ is a finite-index subgroup of $W_{\Gamma'}$. By Lemma 4.6, we conclude that $\Theta' = \Theta(\Gamma', \Lambda)$ satisfies condition $\mathcal{F}_2$. In particular, there is an Γ'-edge between any two Λ components. As Γ' does not have any edges, Λ has exactly one component and the claim follows in this case.

We may now assume that Γ does not contain a cone vertex. Furthermore, by Lemma 4.6 we may assume that $\Theta = \Theta(\Gamma, \Lambda)$ satisfies $\mathcal{F}_1$ and $\mathcal{F}_2$, and that Θ satisfies $\mathcal{R}_1 - \mathcal{R}_4$ by Proposition 3.16.

Suppose now that no two distinct Λ-edges commute. It follows that G^Θ is isomorphic to a free group, and since G^Θ is of finite index, W_Γ is virtually free. Suppose, for a contradiction, that Λ has three distinct components Λ_1, Λ_2 and Λ_3. Let s and t be distinct vertices of Λ_1. By Remark 4.4 there is an $\Lambda_1\Lambda_2$-path α_1 from s to t which we can assume does not repeat vertices. Similarly, there is an $\Lambda_1\Lambda_3$-path α_2 from s to t which does not repeat vertices. Observe that $s, t \in \alpha_1 \cap \alpha_2$. Starting at s and traveling along α_1, let x be the first vertex after s such that $x \in \alpha_1 \cap \alpha_2$. Then the subpath α'_1 of α between s and x contains exactly two vertices of $\alpha_1 \cap \alpha_2$. Let α'_2 be the subpath of α_2 between s and x. Note that $|\alpha'_1|, |\alpha'_2| \geq 2$, as every other vertex of α_1 is in Λ_2 and $\alpha_2 \cap \Lambda_2 = \emptyset$. It follows that $c = \alpha'_1 \cup \alpha'_2$ is a cycle in Γ. Let c' be a sub-cycle of c which is an induced subgraph of Γ. If c' has three vertices, then this contradicts Γ being triangle-free. On the other hand, if c' has more than 3 vertices, then this contradicts W_Γ being virtually free. Thus, Λ can have at most two components and the claim follows in this case.

Suppose now there exist Λ-edges a_1a_2 and b_1b_2 which commute, with $a_1a_2 \neq (b_1b_2)^{\pm 1}$. These Λ-edges must be in different components of Λ by condition $\mathcal{R}_2$ and Lemma 2.6. In this case, W_Γ is not virtually free as it contains a subgroup isomorphic to $\mathbb{Z}^2$. Suppose, for a contradiction, that Λ contains at least three distinct Λ-edge components Λ_1, Λ_2 and Λ_3. Without loss of generality, we may assume that $a_1b_1 \in \Lambda_1$ and that $a_2b_2 \in \Lambda_2$. We will obtain a contradiction by showing that Γ must contain a triangle.

By Lemma 2.6, a_1, a_2, b_1, b_2 form a square in Γ. By Remark 4.4, there is an $\Lambda_1\Lambda_3$-path from a_1 to a_2. Similarly, there is a $\Lambda_2\Lambda_3$-path from b_1 to b_2. Thus, Γ

contains the configuration described in the statement of Lemma 3.33. That lemma then implies that Γ contains a triangle, a contradiction. $\square$

Theorem 4.8. *Let W_Γ be a 2-dimensional right-angled Coxeter group. Let Λ be a subgraph of Γ^c with no isolated vertices, and let G^Θ be the subgroup of W_Γ generated by the Λ-edges. Then the following are equivalent.*

(1) *$(G^\Theta, E(\Lambda))$ is a RAAG system and G^Θ has finite index in W_Γ.*
(2) *$(G^\Theta, E(\Lambda))$ is a RAAG system and G^Θ has index either two or four in W_Γ (and exactly four if W_Γ is not virtually free).*
(3) *Λ has at most two components and satisfies conditions $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$.*

Proof. Clearly (2) implies (1).

To see the remaining implications, suppose first that Γ contains a cone vertex s. Then $\Gamma' = \Gamma \setminus s$ is a graph with no edges, and $W_{\Gamma'}$ is an index two subgroup of Γ. Suppose that (1) holds. By Lemma 4.7, Λ has exactly one component. By Theorem 3.17 and Lemma 4.6, $\Theta' = \Theta(\Gamma', \Lambda)$ satisfies conditions $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$. Consequently, Θ satisfies these conditions as well. Thus (3) holds. By Lemma 4.5, we know that $(G^{\Theta'}, E(\Lambda))$ is a RAAG system of index 2 in $W_{\Gamma'}$, and thus $(G^\Theta, E(\Lambda))$ is a RAAG system of index four in W_Γ. Therefore (2) holds. Finally, if (3) holds then (1) holds by Theorem 3.17 and Lemma 4.5.

Now suppose that Γ does not have a cone vertex. If (1) holds, then by Lemma 4.7, Λ has exactly two components if W_Γ is not virtually free and at most two components otherwise. Thus (2) holds by Lemma 4.5. By Theorem 3.17 and Lemma 4.6, (3) holds. Finally if (3) holds, then (1) follows by Theorem 3.17 and Lemma 4.5. $\square$

Corollary 4.9. *Let W_Γ be a 2-dimensional right-angled Coxeter group. Let Λ be a subgraph of Γ^c with no isolated vertices such that the subgroup $(G, E(\Lambda))$ is a finite index RAAG system. Then either:*

(1) *The graph Γ does not contain any edges and $E(\Lambda)$ is a spanning tree in Γ^c. In particular, W_Γ is virtually free.*
(2) *The group W_Γ is not virtually free. Furthermore, the vertices of Γ can be 2-colored by red and blue (i.e., each edge of Γ connects a red vertex and a blue vertex) and G is isomorphic to the kernel of the homomorphism $\Psi : W_\Gamma \to \mathbb{Z}_2 \times \mathbb{Z}_2 = \langle r, b \mid r^2 = b^2 = 1\rangle$ which maps red and blue generators of $V(\Gamma)$ to r and b respectively.*

Proof. By Theorem 4.8, Λ has at most two components. Suppose first that Λ contains exactly one component. Again by Theorem 4.8, the graph Θ satisfies $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{F}_1$. From these conditions, it follows that Γ cannot contain any edges and that $E(\Lambda)$ is a spanning tree in Γ^c. As Γ does not contain any edges, W_Γ is virtually free.

Suppose now that Λ has exactly two components. We color the vertices of one component red and the vertices of the other component blue. By $\mathcal{R}_2$, each edge of Γ connects a red vertex and a blue vertex, i.e., we have a 2-coloring of Γ. Furthermore, by the definition of Ψ, every Λ-edge (thought of as an element of G) is in the kernel of Ψ. As G is generated by such elements, it follows that $G < \ker(\Psi)$. By Theorem 4.8, G has index 4 in W_Γ. As $\ker(\Psi)$ has index 4 as well, it follows that G is isomorphic to $\ker(\Psi)$. $\square$

5. Applications

In this section we give concrete families of RACGs containing finite-index RAAG subgroups. These cannot be obtained by applying the Davis–Januskiewicz constructions to the defining graphs of the RAAGs they are commensurable to.

5.1. Non-planar RACGs commensurable to RAAGs.

In this subsection, we construct two families of RACGs with non-planar defining graphs containing finite-index RAAG subgroups. These will serve as a warm-up for Theorem 5.5.

We begin by constructing a family of quasi-isometrically distinct RACGs defined by the sequence of graphs Γ_n (shown in Figure 5.1) which are commensurable to RAAGs whose defining graphs are cycles.

Corollary 5.1 (to Theorem 4.8)**.** *For $n \geq 3$, let Γ_n be the graph obtained by starting with a $2n$-gon whose vertices (in cycle order) are $c_1, d_1, c_2, d_2, \ldots, c_n, d_n$ and adding two vertices x and y, such that y is adjacent to c_i for each i, x is adjacent to d_i for each i, and x is adjacent to y (see Figure 5.1). Then*

1. *The right-angled Coxeter group W_{Γ_n} has a subgroup of index four that is isomorphic to (and hence is commensurable to) the right-angled Artin group $A_{C_{2n}}$, where C_{2n} is a cycle of length $2n$.*
2. *W_{Γ_n} is not quasi-isometric to W_{Γ_m} for $m \neq n$.*

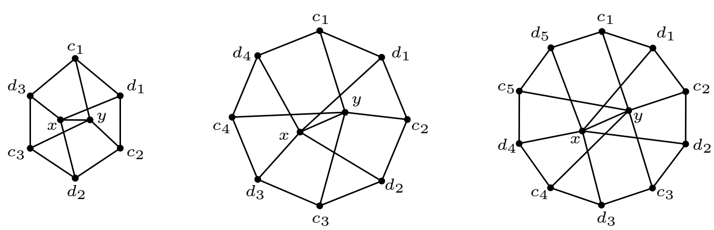

Figure 5.1. The figure illustrates the graphs Γ_n defined in Corollary 5.1 for $n = 3, 4, 5$.

Proof. Fix $n \geq 3$, and let Γ denote Γ_n. We define a graph $\Lambda \subset \Gamma^c$ as follows. Let Λ_x be the star graph consisting of the union of the edges of Γ^c from x to c_i for each i. Let Λ_y be the star graph consisting of the edges of Γ^c from y to d_i for each i. Let $\Lambda = \Lambda_x \cup \Lambda_y$. (See Figure 3.2 for an illustration of Λ in the case $n = 3$.)

We show below that $\Theta = \Theta(\Gamma, \Lambda)$ satisfies $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$. Then it will follow from Theorem 4.8, that $(G^\Theta, E(\Lambda))$ is a RAAG system, and that G^Θ has index four in W_Γ. Moreover, it is easily checked that the commuting graph Δ associated to Λ (as defined in Section 2.2) is isomorphic to C_{2n}. Consequently, G^Θ is isomorphic to $A_{C_{2n}}$. Thus, this will show (1).

It is easy to verify $\mathcal{F}_1$, $\mathcal{R}_1$, and $\mathcal{R}_2$. Then by Remark 4.3, it follows that $\mathcal{F}_2$ holds as well. We now check $\mathcal{R}_3$. First note that there are exactly three squares in Γ containing the edge c_1d_1, and each of these satisfies the property in $\mathcal{R}_3$. Now the fact that every square contains an edge of the $2n$-gon, together with the symmetry of the diagram, implies that $\mathcal{R}_3$ holds.

To check $\mathcal{R}_4$, let γ be a $\Lambda_x\Lambda_y$-cycle and let e be an edge of γ. By symmetry, we can assume that e is either c_1d_1, c_1y or xy. Suppose first that $e = c_1d_1$. Then γ contains either d_nc_1 or yc_1. In both cases, the Λ-convex hull of the vertices of γ contains y. Similarly, γ contains either d_1x or d_1c_2, and in both cases the Λ-convex hull of γ contains x. As c_1d_1 is contained in the square c_1d_1xy, and x, y are in the appropriate convex hulls, it follows that $\mathcal{R}_4$ holds for the Γ-cycle γ and edge e.

Suppose now that $e = c_1y$. It follows that γ contains either yx or yc_i for some $i > 1$. In each case, x is in the Λ-convex hull of γ. Furthermore, γ contains either c_1d_1 or c_1d_n. In the former case, the square c_1d_1xy contains e and has vertices in the Λ-convex hull of the vertices in γ. In the latter case the same argument applies to the square c_1d_nxy.

Finally, suppose that $e = xy$. By symmetry, we may assume that γ contains yc_1. Furthermore, yc_1 must be followed by either c_1d_2 or c_1d_n in γ. Then, as in the previous paragraph, either the square c_1d_1xy or the square c_1d_nxy contains e and has vertices in the Λ-convex hull of γ. Thus $\mathcal{R}_4$ is satisfied in all cases.

We have thus established that (1) holds, by showing that Θ satisfies $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$. Consequently, for each n, we know that W_{Γ_n} is commensurable, and in particular quasi-isometric, to $A_{C_{2n}}$. Claim (2) then follows from [BKS08]. □

Next, we give a family of RACGs whose defining graphs are not planar and are commensurable to RAAGs which are not atomic (as defined in [BKS08]).

Corollary 5.2. *Given $n \geq 3$ and $k \geq 1$, let Δ_{nk} be the graph obtained by taking k copies of Γ_n (defined in Figure 5.2), and identifying them all along the subgraph induced by $V(\Gamma_n) \setminus \{a_0\}$. Thus Δ_{nk} has vertices $a_1, a_2, \dots, a_n, b_1, \dots, b_n$ and also $a_{01}, \dots, a_{0k}$. (The left side of Figure 5.3 shows Δ_{42}.) Then $W_{\Delta_{nk}}$ contains an index four subgroup that is isomorphic to a RAAG.*

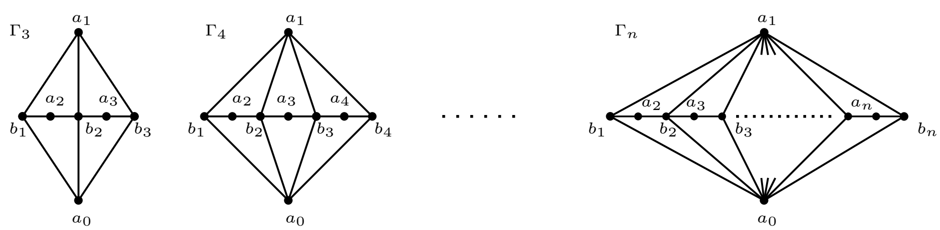

FIGURE 5.2. The figure defines the family of graphs Γ_n, for $n \geq 3$, used in Corollary 5.2.

Proof. Fix $n \geq 3, k \geq 1$ and let $\Delta = \Delta_{nk}$. We define Λ, a subgraph of Δ^c consisting of two components. The first component Λ_a is the union of the edges of Δ^c of the form a_1a_i, where $2 \leq i \leq n$ and a_1a_{0j} for $1 \leq j \leq k$. The second component Λ_b is the path in Δ^c visiting $b_1, b_2, \dots, b_n$. (See the right side of Figure 5.3 for an illustration of the case $n = 4, k = 2$).

Let $\Theta = \Theta(\Delta, \Lambda)$ be as in the previous sections. We verify the properties $\mathcal{R}_1$–$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$. It will then follow from Theorem 4.8 that the subgroup generated by $E(\Lambda)$ is an index four visual RAAG subgroup.

The conditions $\mathcal{R}_1, \mathcal{R}_2$ and $\mathcal{F}_1$ are immediate. Condition $\mathcal{F}_2$ holds by Remark 4.3. We now check $\mathcal{R}_3$. Each square in Δ is of one of the following forms:

(1) $b_ia_{i+1}b_{i+1}a_1$ for $1 \leq i \leq n-1$
(2) $b_ia_{i+1}b_{i+1}a_{0j}$ for $1 \leq i \leq n-1, 1 \leq j \leq k$
(3) $b_ia_1b_{i'}a_{0j}$ for $1 \leq i < i' \leq n, 1 \leq j \leq k$
(4) $b_ia_{0j}b_{i'}a_{0j'}$ for $1 \leq i < i' \leq n, 1 \leq j \leq k$

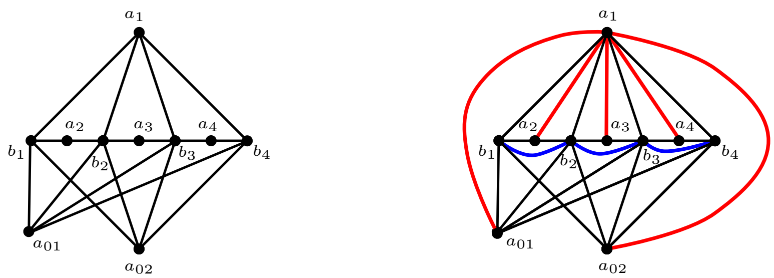

FIGURE 5.3. The figure shows Δ_{42} on the left, and the two components of Λ for the graph Δ_{42} on the right. The component Λ_a is shown in red and the component Λ_b is shown in blue.

Condition $\mathcal{R}_3$ follows immediately for the first type of square, as the appropriate Λ-convex hulls do not contain any additional vertices not included in the vertex set of the square. For the second type of square, the convex hull in Λ_b does not contain any additional vertices, but the convex hull in Λ_a contains the additional vertex a_1, as this vertex lies on the Λ_a-path between a_{i+1} and a_{0j}. Since a_1 is adjacent to b_i and b_{i+1}, the condition $\mathcal{R}_3$ is verified for this type of 2-component square. For the third type, the Λ-convex hull of $\{a_1, a_{0j}\}$ does not contain any additional vertices of Λ, and the Λ-convex hull of $\{b_i, b_{i'}\}$ contains the additional vertices $b_{i+1}, \dots, b_{i'-1}$. Since a_1 and a_{0j} are adjacent to each of these, $\mathcal{R}_3$ is verified for this type of 2-component square as well. Finally, for the last type, the Λ-convex hull of $\{a_{0j}, a_{0j'}\}$ contains the additional vertex a_1, and the Λ-convex hull of $\{b_i, b_{i'}\}$ contains the additional vertices $b_{i+1}, \dots, b_{i'-1}$. Once again, it is easily verified that $a_0, a_{1j}, a_{1j'}$ are each adjacent to each of $b_i, \dots, b_{i'}$. Thus $\mathcal{R}_3$ is verified.

Finally, we check $\mathcal{R}_4$. Let γ be a $\Lambda_a\Lambda_b$-cycle and let e be an edge of γ. First suppose e is of the form a_ib_i for $2 \leq i \leq n$. In this case, γ necessarily passes through b_{i-1} and some a, where either $a = a_1$ or $a = a_{0j}$, for some $1 \leq j \leq k$. Thus the $\Lambda_a\Lambda_b$-square $b_{i-1}a_ib_ia$ satisfies the criterion in $\mathcal{R}_4$, since it contains e, and the two vertices a and b_{i-1} are contained in the Λ-convex hull of the vertices of γ. The case where e is of the form a_ib_{i-1} for $2 \leq i \leq n$ is similar.

Now suppose e is of the form a_1b_i for some $1 \leq i \leq n$. Then γ necessarily passes through an edge of one of the following forms: b_ia_{0j} for some $1 \leq j \leq k$, b_ia_i or b_ia_{i+1}. In the first of these cases, γ necessarily also passes through a vertex $b_{i'}$ for some $i' \neq i$, and $a_1b_ia_{0j}b_{i'}$ is the desired square. If the edge is of the form b_ia_i (resp. b_ia_{i+1}) then γ must also pass through b_{i-1} (resp. b_{i+1}), and the desired square is $a_1b_ia_ib_{i-1}$ (resp. $a_1b_ia_{i+1}b_{i+1}$). The case where e is of the form $a_{0j}b_i$ for some $1 \leq i \leq n$ and $1 \leq j \leq k$ is similar. This completes the verification of $\mathcal{R}_4$, and the corollary follows. □

Remark 5.3. We note that the RAAGs obtained in the above corollary do not have a tree for defining graph when $k \geq 2$ and $n \geq 3$. This is easy to check by computing the associated commuting graph.

5.2. **2-dimensional RACGs with planar defining graph.** In [NT17], Nguyen–Tran characterize exactly which one-ended, 2-dimensional RACGs defined by non-join CFS graphs are quasi-isometric to RAAGs. In this subsection, we use their work in conjunction with Theorem 4.8 to prove Theorem B from the introduction. Note that CFS is a graph-theoretic condition introduced in [DT15] to characterize

RACGs with at most quadratic divergence. We omit the definition, as it is not needed here.

Remark 5.4. Any one-ended, 2-dimensional RACG that is quasi-isometric to a RAAG must have CFS defining graph. This follows as one-ended RAAGs have either linear or quadratic divergence [BC12], and the defining graph of a 2-dimensional RACG with linear or quadratic divergence is CFS [DT15].

Recall that a graph Σ is a *suspension* if Σ decomposes as a join $\Sigma = \{a_1, a_2\} \star B$ where a_1 and a_2 are non-adjacent vertices. We also say that Σ is the *suspension* of the graph B. We use the notation $\Sigma_k(a, b)$ to denote the suspension graph $\{a_1, a_2\} \star \{b_1, \dots, b_k\}$, and we say that a_1 and a_2 are the *suspension vertices*.

Let Γ be a graph which is connected, triangle-free, CFS and planar. Suppose that a planar embedding from Γ into the sphere $\mathbb{S}^2$ is fixed. In [NT17], Nguyen–Tran construct a tree T (this is the *visual decomposition tree* of Section 3 of that paper) associated to Γ with the following properties. The vertices of T are in bijection with maximal suspension subgraphs of Γ. As Γ is triangle-free, every maximal suspension of Γ is of the form $\Sigma_k(a, b)$, where both $\{a_1, a_2\}$ and $\{b_1, \dots, b_k\}$ are each sets of disjoint vertices of Γ, and $k \geq 3$ if T contains at least two vertices. Moreover, every vertex of Γ is contained in some suspension corresponding to a vertex of T. Two vertices of T corresponding to suspensions $\Sigma = \Sigma_k(a, b)$ and $\Sigma' = \Sigma_l(c, d)$ are connected by an edge if $\Sigma \cap \Sigma'$ is a 4-cycle C which separates $\mathbb{S}^2$ into two non-trivial components B_1 and B_2, such that $\Sigma_1 \setminus C \subset B_1$ and $\Sigma_2 \setminus C \subset B_2$. Moreover, it must follow (by the maximality of the suspensions) that $C = \{a_1, c_1, a_2, c_2\}$, i.e. C contains exactly the suspension vertices of Σ and Σ'.

If Γ (with the above assumptions) is a join, then it readily follows that Γ is quasi-isometric to a RAAG whose defining graph is a tree of diameter at most 2. Nguyen–Tran show that if Γ is not a join, then W_Γ is quasi-isometric to a RAAG if and only if every vertex $v \in T$ has valence strictly less than k [NT17, Theorem 1.2], where $\Sigma_k(a, b)$ is the maximal suspension in Γ corresponding to v. Moreover, they show that such RAAGs have defining graph a tree of diameter at least 3. Below, we prove such RACGs are in fact commensurable to RAAGs.

Theorem 5.5. *Let W_Γ be a 2-dimensional, one-ended RACG with planar defining graph Γ. Then W_Γ is quasi-isometric to a RAAG if and only if it contains an index 4 subgroup isomorphic to a RAAG.*

Proof. One direction of the theorem is obvious. Thus, we prove that if W_Γ satisfies these hypotheses and is quasi-isometric to a RAAG, then W_Γ contains an index 4 subgroup isomorphic to a RAAG. We do this by constructing a subgraph $\Lambda \subset \Gamma^c$ with two components and satisfying the hypotheses of Theorem 4.8.

Fix a planar embedding of Γ into the sphere $\mathbb{S}^2$. Note that by Remark 5.4 and the hypotheses of the theorem, it follows that Γ is triangle-free, CFS and planar. Thus, there exists a visual decomposition tree T associated to Γ as described above. Furthermore, as W_Γ is quasi-isometric to a RAAG, it follows from [NT17, Theorem 1.2] that the valence of a vertex of T corresponding to the maximal suspension $\Sigma_k(a, b)$ is less than k.

Henceforth, to simplify notation, the word suspension will always refer to a maximal suspension, and will consequently correspond to a vertex of T. Given a suspension $\Sigma_k(a, b) = \{a_1, a_2\} \star B$ we say that a labeling $\{b_1, \dots, b_k\}$ of the vertices of B is *cyclic* if the following holds. If C is a 4-cycle spanning the vertices

$\{a_1, b_i, a_2, b_{i+1}\}$ for some $1 \leq i \leq k$ or spanning the vertices $\{a_1, b_1, a_2, b_k\}$, then every vertex of $\Sigma \setminus C$ is contained in a common component of $\mathbb{S}^2 \setminus C$. Observe that if E is a cycle corresponding to an edge of T incident to the vertex of T given by $\Sigma_k(a, b)$, then the planarity of Γ implies that E is one of the cycles C mentioned in the previous sentence.

Let N be the number of vertices of T. Let $T_1 \subset \cdots \subset T_N = T$ be a nested sequence of subtrees of T such that T_1 consists of a single vertex of T and T_i has exactly i vertices. Such choices are clearly possible. For each $1 \leq i \leq n$, let Γ_i be the subgraph of Γ spanned by every suspension that corresponds to a vertex of T_i. Note that $\Gamma_i \subset \Gamma_{i+1}$ for all $1 \leq i < N$ and that $\Gamma_N = \Gamma$. We define a nested sequence of graphs $\Lambda_1 \subset \cdots \subset \Lambda_N$ such that for each $1 \leq i \leq N$, $\Lambda_i \subset \Gamma_i^c$ and the following hold:

(1) Let C be a 4-cycle corresponding to an edge of T that is incident to T_i. Then each pair of non-adjacent vertices in C is contained in a common edge of Λ_i.
(2) The graph Λ_i contains exactly two components, and $\Theta_i = \Theta(\Gamma_i, \Lambda_i)$ satisfies conditions $\mathcal{R}_1 - \mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$.

The theorem clearly follows from this claim by using the graph $\Lambda = \Lambda_N \subset \Gamma^c$.

We first define Λ_1 corresponding to the vertex $T_1 = \{v\}$. Let $\Sigma = \Sigma_k(a, b) = \{a_1, a_2\} \star \{b_1, \ldots, b_k\}$ be the suspension corresponding to v, and assume that $\{b_1, \ldots, b_k\}$ is cyclic. As the valence of v in T is less than k, by possibly relabeling, we can assume that the 4-cycle $\{a_1, b_1, a_2, b_k\}$ does not correspond to an edge of T. We define one component of Λ_1 to be the edge (a_1, a_2), and the other component of Λ_1 to consist of the edges $(b_1, b_2), (b_2, b_3), \ldots, (b_{k-1}, b_k)$. By the observation above and our choice of labeling, any 4-cycle C corresponding to an edge of T incident to v is of the form $\{a_1, b_i, a_2, b_{i+1}\}$ for some $1 \leq i \leq k-1$. Thus condition (1) follows. Condition (2) is readily verified.

Suppose now that we have defined the graph Λ_{n-1} corresponding to the tree T_{n-1} satisfying conditions (1) and (2). We now define Λ_n.

Let u be the unique vertex in $T_n \setminus T_{n-1}$, and let u' be the unique vertex of T_{n-1} that is adjacent to u. Let $\Sigma = \Sigma_k(a, b) = \{a_1, a_2\} \star \{b_1, \ldots, b_k\}$ and $\Sigma' = \Sigma_l(c, d) = \{c_1, c_2\} \star \{d_1, \ldots, d_l\}$ be the suspension graphs corresponding to u and u' respectively. Furthermore, suppose these labelings are cyclic. It follows that $E = \{a_1, c_1, a_2, c_2\}$ is the 4-cycle corresponding to the edge in T between u and u'. By possibly relabeling, we can assume that $c_1 = b_1$, $c_2 = b_k$, $a_1 = d_1$ and $a_2 = d_l$. As Λ_{n-1} satisfies (1) above, (a_1, a_2) and (c_1, c_2) are edges of Λ_{n-1}.

As the valence of u is less than k, there exist some $1 \leq j < k$ such that the 4-cycle $\{b_j, a_1, b_{j+1}, a_2\}$ does not correspond to an edge of T. We define $\Lambda_n \subset \Gamma_n^c$ to contain every edge of $\Lambda_{n-1} \subset \Gamma_{n-1}^c \subset \Gamma_n^c$ and additionally the edges:

$$(b_1, b_2), (b_2, b_3), \ldots, (b_{j-1}, b_j), (b_{j+1}, b_{j+2}), \ldots, (b_{k-1}, b_k)$$

This corresponds to adding one or two line segments each to a distinct vertex of Λ_{n-1}. As Λ_{n-1} contains two components (by (2)) and does not contain any cycles (by $\mathcal{R}_1$), it follows that Λ_n contains two components and satisfies $\mathcal{R}_1$ as well. Furthermore, (1) and condition $\mathcal{F}_1$ (for Θ_n) follow from directly from our choices. Condition $\mathcal{F}_2$ then follows from Remark 4.3.

We now check $\mathcal{R}_2$. Let $x, y \in \Lambda_n$ be vertices contained in the same component of Λ_n. If x and y are both contained in Λ_{n-1}, then the claim follows as Λ_{n-1} satisfies

$\mathcal{R}_2$ and no new edges are added between vertices of Γ_{n-1} in forming Γ_n. If x and y are both contained in Σ, then by construction, they must lie in the same factor of the join Σ and there is no edge between them. The only case left to check is that x and y lie in different components of $\mathbb{S}^2 \setminus E$. However, in this case there is no edge between x and y as E separates x from y in the planar embedding.

We now check that $\mathcal{R}_3$ holds. Let C be a 2-component square in Λ_n. As E separates every vertex of $\Sigma \setminus E$ from every vertex in $(\Gamma_{n-1} \setminus E) \subset \Gamma_n$, it follows that either C lies in $\Gamma_{n-1} \subset \Gamma_n$ or C lies in Σ. In the first case the claim follows as Θ_{n-1} satisfies $\mathcal{R}_3$ (and noting that the convex hull of C in Λ_n lies in Λ_{n-1}). In the latter case, the claim is easily verified.

We now check $\mathcal{R}_4$. Let P be a 2-component cycle in Γ_n. If P lies entirely in Γ_{n-1} then every edge of P satisfies the condition in $\mathcal{R}_4$ as Θ_{n-1} satisfies $\mathcal{R}_4$. If P lies entirely in Σ, then $\mathcal{R}_4$ is easily verified. Thus, we may assume that P decomposes into two subpaths P_1 and P_2 such that $P_1 \subset \Gamma_{n-1}$ and $P_2 \subset \Sigma \setminus E$. As P does not repeat vertices, it follows that P_2 consists of just two edges (a_1, b_q) and (a_2, b_q) for some $2 \leq q \leq k$. As the valence of u' is less than l, there exists some $1 \leq q' < l$ and corresponding 4-cycle $\{c_1, d_{q'}, c_2, d_{q'+1}\}$ such that every vertex of Γ_n is contained in a common component of $\mathbb{S}^2 \setminus C$. From this, we see that a_1 and a_2 are in different components of $\Gamma_{n-1} \setminus \{c_1, c_2\}$. Thus, P_1 must either contain c_1 or c_2. Suppose that P_1 contains c_1 (the other case is similar). The path P_1 does not contain both the edge (a_1, c_1) and the edge (a_2, c_1), for if it did, then P would either the equal 4-cycle $\{a_1, c_1, a_2, b_q\}$ or contain it as a sub-cycle. In the former case $P \subset \Sigma$, a case we have already ruled out, and in the latter case, P necessarily repeats a vertex (which is not allowed). We now define a cycle P' depending on which edges P_1 contains. We set $P' = (P_1 \setminus (a_1, c_1)) \cup (a_2, c_1)$ if $(a_1, c_1) \subset P_1$, $P' = (P_1 \setminus (a_2, c_1)) \cup (a_1, c_1)$ if $(a_2, c_1) \subset P_1$, and $P' = P_1 \cup (a_1, c_1) \cup (a_2, c_1)$ if P_1 does not contain either of (a_1, c_1) and (a_2, c_1). In each case, it follows that P' is a cycle in Γ_{n-1} containing every edge of P_1, except possibly (a_1, c_1) and (a_2, c_1). Additionally, every vertex of P' is a vertex of P, so the Λ-convex hull of P' is contained in the Λ-convex hull of P. From this and as Γ_{n-1} satisfies $\mathcal{R}_4$, it follows that every edge of P that is contained in P_1 satisfies $\mathcal{R}_4$ as well. Finally, every edge of $P \setminus P_1$ can been seen to satisfy $\mathcal{R}_4$ by using the 4-cycle $\{a_1, b_q, a_2, c_1\}$.

We have thus checked that (2) holds. The theorem now follows. □

6. Generalized reflection subgroups of RAAGs

Let A_Γ be a RAAG. A *generalized RAAG reflection* is a conjugate of an element of $V(\Gamma)$, i.e. wsw^{-1} for some $s \in V(\Gamma) \cup V(\Gamma)^{-1}$ and w a word in A_Γ. Let $\mathcal{T}$ be a set of reduced generalized RAAG reflections. We say that $\mathcal{T}$ is *trimmed* if $\mathcal{T} \cap \mathcal{T}^{-1} = \emptyset$, and if given any two distinct generalized RAAG reflections wsw^{-1} and $w's'w'^{-1}$ in $\mathcal{T}$, no expression for w' has prefix ws^{-1} or prefix ws. The following lemma follows from a straightforward adaptation of the proof of [DL19, Lemma 10.1] to the setting of RAAGs.

Lemma 6.1. *Let $\mathcal{T}$ be a set of generalized RAAG reflections in the RAAG A_Γ, and let G be the subgroup generated by $\mathcal{T}$. Then G is generated by a trimmed set of generalized RAAG reflections which can be algorithmically obtained from $\mathcal{T}$.*

The main result of this section is the following:

Theorem 6.2. *Let $\mathcal{T}$ be a finite set of generalized RAAG reflections in A_Γ. Then the subgroup $G < A_\Gamma$ generated by $\mathcal{T}$ is a RAAG. Moreover, if $\mathcal{T}$ is trimmed, then $(G, \mathcal{T})$ is a RAAG system.*

We will use the characterization of RAAGs in Theorem 2.2 to show that G is a RAAG. We first prove a series of lemmas about disk diagrams of a special type, namely, ones whose boundary labels are words over a trimmed set of generalized RAAG reflections.

The setup for these lemmas is as follows and will be fixed for the rest of this section. We fix a trimmed set $\mathcal{T}$ of reduced generalized RAAG reflections in A_Γ. Let $z = r_1 \dots r_n$ be an expression for the identity element where $r_i = w_i s_i w_i^{-1} \in \mathcal{T}$ for each $1 \le i \le n$. Let D be a disk diagram whose boundary ∂D is labeled by z. For $1 \le i \le n$, let p_{r_i} be the subpath of ∂D which is labeled by r_i. Furthermore let p_{w_i} and $p_{w_i^{-1}}$ denote the subpaths of ∂D labeled w_i and w_i^{-1} respectively, and let e_i denote the edge labeled s_i. Let H_i be the hyperplane dual to e_i, and let $\mathcal{H} = \{H_i\}_{i=1}^n$ be the collection of all such hyperplanes. Note that as r_i is a reduced word, no hyperplane is dual to two edges of p_{r_i} for any i.

In all of the following lemmas, arithmetic is taken modulo n.

Lemma 6.3. *For each $1 \le i \le n$, the hyperplane H_i does not intersect a hyperplane dual to p_{w_i} or a hyperplane dual to $p_{w_i^{-1}}$*

Proof. Suppose H_i intersects a hyperplane K that is dual to an edge f of p_{w_i}. Without loss of generality, we may assume that f is the edge closest to e_i out of all possible choices for K. As no hyperplane is dual to two edges of p_{r_i}, it follows that every hyperplane dual to an edge of p_{w_i} which lies between e_i and f must intersect K. Thus, w_i has suffix the word $t_1 \dots t_m$, where t_1 is the label of K and t_1 commutes with s_i, as well as with t_j for $2 \le j \le m$. This readily implies that r_i is not reduced, for in $r_i = w_i s_i w_i^{-1}$, an occurrence of the RAAG generator t_1 in w_i can be canceled with an occurrence of t_1^{-1} is w_i^{-1}. However, this is a contradiction as r_i is reduced by assumption. The argument for hyperplanes dual to $p_{w_i^{-1}}$ is analogous. □

Lemma 6.4. *For each $1 \le i \le n$, the hyperplane H_i is not dual to $p_{w_{i+1}}$, $p_{w_{i+1}^{-1}}$, $p_{w_{i-1}^{-1}}$ or $p_{w_{i-1}}$.*

Proof. For a contradiction, suppose H_i is dual to an edge f of $p_{w_{i+1}}$. By Lemma 6.3, every hyperplane dual to an edge of $p_{w_i^{-1}}$ must also be dual to $p_{w_{i+1}}$. Write $s_i w_i^{-1} = t_1 \dots t_m$ and $w_{i+1} = k_1 \dots k_l$ where $t_j \in V(\Gamma)$ for $1 \le j \le m$ and $k_j \in V(\Gamma)$ for $1 \le j \le l$. The structure of the hyperplanes in D implies that w_{i+1} has an expression which begins with $t_m^{-1} \dots t_1^{-1} = w_i s_i^{-1}$. This is a contradiction as $\mathcal{T}$ is trimmed. A similar argument shows that H_i is not dual to $p_{w_{i-1}^{-1}}$.

Suppose now that H_i is dual to $p_{w_{i-1}}$. By Lemma 6.3, it follows that H_{i-1} is dual to p_{w_i}. However, this is not possible by the same argument as above. Similarly, H_i cannot be dual to $p_{w_{i+1}^{-1}}$. □

The proof of the following lemma is similar to that of the previous one.

Lemma 6.5. *If $H_i = H_{i+1}$ for some $1 \le i \le n$ then $r_i \doteq r_{i+1}^{-1}$.* □

Lemma 6.6. *If H_i intersects H_{i+1}, then r_i and r_{i+1} commute. Furthermore, there is a disk diagram D' with boundary label $r_1 \dots r_{i-1} r_{i+1} r_i r_{i+2} \dots r_n$, such that* *the natural bijection, from e_i, e_{i+1} and the edges traversed by the subpath of the boundary path of D labeled by $r_{i+2} \dots r_n r_1 \dots r_{i-1}$ to the edges traversed by the corresponding subpaths of the boundary path of D' with the same labels, preserves boundary combinatorics.*

Proof. Suppose H_i intersects H_{i+1}. By Lemma 6.3, every hyperplane dual to $p_{w_i^{-1}}$ is either dual to $p_{w_{i+1}}$ or intersects H_{i+1}. Similarly, every hyperplane dual to $p_{w_{i+1}}$ is either dual to $p_{w_i^{-1}}$ or intersects H_i. It then readily follows that w_i has a reduced expression ba_1 and w_{i+1} has a reduced expression ba_2, where a_1, a_2 and b are words, such that the generators in the word $a_1 s_i$ are all distinct from and commute with the generators in the word $a_2 s_{i+1}$. Consequently, r_i commutes with r_{i+1}.

We now construct the disk diagram D'. By Tits' solution to the word problem, the expression ba_1 (resp. ba_2) can be obtained from w_i (resp. w_{i+1}) by sequentially permuting adjacent letters. Thus, by repeatedly applying Lemma 2.11(1), we obtain a disk diagram with boundary label:

$$r_1 \dots r_{i-1} \; (ba_1 s_i a_1^{-1} b^{-1}) \; (ba_2 s_{i+1} a_2^{-1} b^{-1}) \; r_{i+2} \dots r_n$$

By repeatedly applying Lemma 2.11(2), we can "cancel" $b^{-1}b$ and obtain a disk diagram with boundary label:

$$r_1 \dots r_{i-1} \; (ba_1 s_i a_1^{-1}) \; (a_2 s_{i+1} a_2^{-1} b^{-1}) \; r_{i+2} \dots r_n$$

Then, by repeatedly applying Lemma 2.11(1), we obtain a disk diagram with label:

$$r_1 \dots r_{i-1} \; (ba_2 s_{i+1} a_2^{-1}) \; (a_1 s_i a_1^{-1} b^{-1}) \; r_{i+2} \dots r_n$$

By Lemma 2.11(3), we obtain a disk diagram with boundary label:

$$r_1 \dots r_{i-1} \; (ba_2 s_{i+1} a_2^{-1} b^{-1}) \; (ba_1 s_i a_1^{-1} b^{-1}) \; r_{i+2} \dots r_n$$

Finally, by repeatedly applying Lemma 2.11(1), we obtain a disk diagram D' with boundary label:

$$r_1 \dots r_{i-1} r_{i+1} r_i r_{i+2} \dots r_n$$

Note that in each of these steps, the desired boundary combinatorics are preserved. □

Lemma 6.7. *For every $1 \leq i \leq n$, there exists some $j \neq i$ such that $H_i = H_j$.*

Proof. Suppose we have a disk diagram with boundary label $z = r_1 \dots r_n$ such that, for some $1 \leq i \leq n$, the hyperplane H_i is dual to an edge f of ∂D where $f \neq e_j$ for all $1 \leq j \leq n$. We call any disk diagram which has such an H_i a *pathological diagram* with *pathology caused by* H_i. Given such a diagram, we define p to be a path along ∂D between e_i and f, which does not include e_i and f. We also let $\mathcal{H}'$ denote the set of H_j such that e_j is contained in p.

Given a pathological disk diagram D we may choose a hyperplane H_i causing the pathology together with a path p so that the set $\mathcal{H}'$ is minimal among all possible choices of H_i and p. After such a choice, we call $|\mathcal{H}'|$ the complexity of D. We will prove that pathological diagrams are not possible by induction on the complexity c of such a diagram. The base case, when $c = 0$, already follows from Lemma 6.4.

Now suppose we are given a pathological disk diagram D with pathology caused by H_i such that its complexity is $c = |\mathcal{H}'| > 0$, and suppose by induction there do not exist pathological disk diagrams of complexity smaller than c.

The edge $f \neq e_i$ of ∂D that is dual to H_i lies in a path $p_{r_{i'}}$ in ∂D labeled by $w_{i'} s_{i'} w_{i'}^{-1}$ for some $1 \leq i' \leq n$ where $i \neq i'$. Let Q denote the hyperplane $H_{i'}$.

Note that Q may or may not be in $\mathcal{H}'$. We prove our claim by considering two cases:

Case 1: Every hyperplane in $\mathcal{H}'$ intersects H_i.

We first observe that $\mathcal{H}'$ is non-empty (since the complexity of D is positive) and does not consist of of Q alone (by Lemma 6.4). Therefore, we may choose $K \in \mathcal{H}' \setminus Q$ such that no hyperplane in $\mathcal{H}' \setminus Q$ intersects H_i between $K \cap H_i$ and $H_i \cap e_i$. Let $1 \leq l \leq n$ be such that K is dual to $e_l \subset p_{r_l} \subset p$. Then for each j with $i < j < l$, the hyperplane H_j intersects $K = H_l$. Thus, by repeatedly applying Lemma 6.6, we can produce a new disk diagram with boundary label $r_1 \dots r_l r_i \dots r_{l-1} r_{l+1} \dots r_n$. Furthermore, this new disk diagram is still pathological and has complexity smaller than D. However, this is not possible by our induction hypothesis.

Case 2: Some hyperplane $K \in \mathcal{H}'$ does not intersect H_i.

We can choose such a hyperplane K to be innermost, i.e. choose $K \in \mathcal{H}'$ such that K does not intersect H_i and such that any hyperplane of $\mathcal{H}'$ dual to the subpath of p between the edges dual to K intersects K. Since H_i and p were chosen to attain the complexity of D, it follows that K does not cause a pathology, and is dual to distinct edges e_l and $e_{l'}$ in p, where $1 \leq l, l' \leq n$. By relabeling the r_j's if necessary, we may assume that $l < l'$, and that the subpath of ∂D from e_l to $e_{l'}$ is contained in p. By repeatedly applying Lemma 6.6, we can produce a new pathological disk diagram D' with label $r_1 \dots r_{l-1} r_{l+1} \dots r_{l'-1} r_l r_{l'} \dots r_n$ and where some hyperplane, which we still denote by K, is dual to both the edge labeled by e_l and the one labeled by $e_{l'}$. By Lemma 6.5, $r_l \doteq r_{l'}^{-1}$. Furthermore, by repeatedly applying Lemma 2.11(1) if necessary, we may assume that $r_l = r_{l'}^{-1}$ in the label of $\partial D'$.

We now produce a new disk diagram D'' by identifying the consecutive paths in $\partial D'$ labeled by r_l and $r_{l'}$, i.e. we fold these two paths together. If $K \neq Q$, then we have produced a new pathological disk diagram with complexity $c-2$, contradicting the induction hypothesis. On the other hand, if $K = Q$, note that the image of H_i in D'' must intersect the path labeled by $r_i \dots r_{l-1} r_{l+1} \dots r_{l'-1}$ in $\partial D''$. Moreover we claim that it cannot be dual to an edge labeled by e_j for $i < j \leq l' - 1$. Suppose it is dual to an edge labeled e_j. It follows that the hyperplane H_j in D is dual to an edge f' in p, such that $f' \neq e_k$ for any k, and such that the images of f and f' are identified in D''. This is a contradiction, as it implies that H_j causes a pathology of lower complexity than H_i. Thus, the image of H_i in D'' causes a pathology of complexity at most $c - 2$, which is again a contradiction. □

We are now ready to prove Theorem 6.2.

Proof of Theorem 6.2. As G can be generated by a trimmed set of generalized RAAG reflections (by Lemma 6.1), we assume without loss of generality that $\mathcal{T}$ is trimmed. We will show that $(G, \mathcal{T})$ is a RAAG system by applying Theorem 2.2. Note that $\mathcal{T} \cap \mathcal{T}^{-1} = \emptyset$ as $\mathcal{T}$ is trimmed. We check each condition of that theorem, by proving the corresponding two claims:

i) Every $r \in \mathcal{T}$ has infinite order.

By definition, r is equal to a reduced word wsw^{-1} with $s \in V(\Gamma) \cup V(\Gamma)^{-1}$ and w a word in W_Γ. It follows that $ws^n w^{-1}$ is an expression for r^n. Moreover, as r is reduced, it readily follows from Theorem 2.4 that $ws^n w^{-1}$ is reduced as well.

Hence, r has infinite order.

ii) Given any word $w = a_1 \dots a_m$, with $a_i \in \mathcal{T}$, either w is reduced over $\mathcal{T}$ or there is an expression for w of the form $a_1 \dots \hat{a}_i \dots \hat{a}_j \dots a_m$.

Suppose $w = a_1 \dots a_m$ is not reduced over $\mathcal{T}$. Let $w' = b_1 \dots b_k$, with $b_i \in \mathcal{T}$ and $k < m$, be an expression for w which is reduced over $\mathcal{T}$. Form a disk diagram D with boundary label ww'^{-1}.

We relabel the generalized reflections in the word ww'^{-1} by setting $r_i = a_i$ for $1 \leq i \leq m$, and $r_{m+i} = b_{k-i+1}^{-1}$ (the ith generalized RAAG reflection in w'^{-1}) for $1 \leq i \leq k$. By Lemma 6.7, every $H \in \mathcal{H}$ is only dual to edges of ∂D labeled by s_i for some i, where $r_i = w_i s_i w_i^{-1}$. As $m > k$, there exists some hyperplane $H \in \mathcal{H}$ that is dual to two edges of the subpath p of ∂D labeled by w. Furthermore, we may choose an innermost such $H \in \mathcal{H}$, in the sense that every hyperplane in $\mathcal{H} \setminus H$ intersects p at most once.

Let e_l and $e_{l'}$ be the edges dual to H where $l < l' \leq m$. By repeatedly applying Lemma 6.6, we produce a disk diagram whose boundary label is

$$r_1 \dots \hat{r}_l \dots r_{l'-1} r_l r_{l'} \dots r_n,$$

such that a hyperplane of $\mathcal{H}$ is still dual to the images of the edges e_l and $e_{l'}$ under the natural map between the boundaries of the disk diagrams. By Lemma 6.5, $r_l = r_{l'}^{-1}$. Thus, $r_1 \dots \hat{r}_l \dots \hat{r}_{l'} \dots r_n$ is an expression for ww'^{-1}. Consequently, $r_1 \dots \hat{r}_l \dots \hat{r}_{l'} \dots r_m = a_1 \dots \hat{a}_l \dots \hat{a}_{l'} \dots a_m$ is an expression for w. □

References

[Bah05] Patrick Bahls, *The isomorphism problem in Coxeter groups*, Imperial College Press, London, 2005. MR 2162146

[Bas02] Şerban A. Basarab, *Partially commutative Artin-Coxeter groups and their arboreal structure*, J. Pure Appl. Algebra **176** (2002), no. 1, 1–25. MR 1930963

[BB05] Anders Björner and Francesco Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, vol. 231, Springer, New York, 2005. MR 2133266

[BC12] Jason Behrstock and Ruth Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356. MR 2874959

[Beh19] Jason Behrstock, *A counterexample to questions about boundaries, stability, and commensurability*, Beyond hyperbolicity, London Math. Soc. Lecture Note Ser., vol. 454, Cambridge Univ. Press, Cambridge, 2019, pp. 151–159. MR 3966609

[BKS08] Mladen Bestvina, Bruce Kleiner, and Michah Sageev, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol. **12** (2008), no. 3, 1653–1699. MR 2421136

[CH17] Matthew Cordes and David Hume, *Stability and the Morse boundary*, J. Lond. Math. Soc. (2) **95** (2017), no. 3, 963–988. MR 3664526

[Cha07] Ruth Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158. MR 2322545

[CS15] Ruth Charney and Harold Sultan, *Contracting boundaries of* CAT(0) *spaces*, J. Topol. **8** (2015), no. 1, 93–117. MR 3339446

[Dav15] Michael W. Davis, *The geometry and topology of Coxeter groups*, Introduction to modern mathematics, Adv. Lect. Math. (ALM), vol. 33, Int. Press, Somerville, MA, 2015, pp. 129–142. MR 3445448

[Deo89] Vinay V. Deodhar, *A note on subgroups generated by reflections in Coxeter groups*, Arch. Math. (Basel) **53** (1989), no. 6, 543–546. MR 1023969

[DJ00] Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235. MR 1783167

[DL19] Pallavi Dani and Ivan Levcovitz, *A study of subgroups of right-angled Coxeter groups via stallings-like techniques*, arXiv:1908.09046, 2019.

[DT15] Pallavi Dani and Anne Thomas, *Divergence in right-angled Coxeter groups*, Trans. Amer. Math. Soc. **367** (2015), no. 5, 3549–3577. MR 3314816

[Dye90] Matthew Dyer, *Reflection subgroups of Coxeter systems*, J. Algebra **135** (1990), no. 1, 57–73. MR 1076077

[KK13] Sang-hyun Kim and Thomas Koberda, *Embedability between right-angled Artin groups*, Geom. Topol. **17** (2013), no. 1, 493–530. MR 3039768

[LaF17] Garret LaForge, *Visible Artin Subgroups of Right-Angled Coxeter Groups*, ProQuest LLC, Ann Arbor, MI, 2017, Thesis (Ph.D.)–Tufts University. MR 3768473

[NT17] Hoang Thanh Nguyen and Hung Cong Tran, *On the coarse geometry of certain right-angled Coxeter groups*, arXiv:1712.01079, to appear in Algebraic & Geometric Topology, 2017.

[Sag95] Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. (3) **71** (1995), no. 3, 585–617. MR 1347406

[Tit69] Jacques Tits, *Le problème des mots dans les groupes de Coxeter*, Symposia Mathematica (INDAM, Rome, 1967/68), Vol. 1, Academic Press, London, 1969, pp. 175–185. MR 0254129

[Wis11] Daniel T. Wise, *The structure of groups with a quasiconvex hierarchy*, 2011.