RAAGEDY RIGHT-ANGLED COXETER GROUPS II: IN THE QUASIISOMETRY CLASS OF THE TREE RAAGS

CHRISTOPHER H. CASHEN

Abstract. We classify two-dimensional right-angled Coxeter groups that are quasiisometric to a right-angled Artin group defined by a tree, and show that when this is true the right-angled Coxeter group actually contains a visible finite index right-angled Artin subgroup.

1. Introduction

A *tree RAAG* is a right-angled Artin group whose presentation graph is a tree of diameter at least three. Such a group is the fundamental group of a compact 3–manifold with boundary whose JSJ decomposition, in the 3–manifold sense, contains only Seifert fibered pieces. A corollary of Behrstock and Neumann's [1] quasiisometry classification of graph manifolds is that all tree RAAGs belong to a single quasiisometry class. The diameter condition excludes the well-understood quasiisometry classes $\mathbb{Z}$, $\mathbb{Z}^2$, and $\mathbb{F}_2 \times \mathbb{Z}$.

Nguyen and Tran [20] classified planar, triangle-free graphs Γ such that the right-angled Coxeter group (RACG) W_Γ with presentation graph Γ is in the quasiisometry class of the tree RAAGs. Dani and Levcovitz [11] subsequently showed that when Nguyen and Tran's conditions are satisfied then W_Γ actually has a finite index tree RAAG subgroup that is *visible* in Γ, in a sense that will be made precise later.

Nguyen and Tran's proof uses planarity in an essential way twice: once to invoke the Jordan Curve Theorem to construct a certain suspension-decomposition of Γ, and again to say that W_Γ is virtually a 3–manifold group. Then their suspension-decomposition corresponds to the JSJ decomposition of the manifold and they apply Behrstock and Neumann's result. However, there are easy examples, such as the graph in Figure 1, of nonplanar graphs Γ such that W_Γ is quasiisometric to the tree RAAGs (see Example 3.2), so planarity is not a necessary condition.

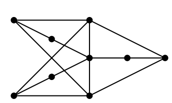

Figure 1. A nonplanar, triangle-free graph Γ such that W_Γ is in the quasiisometry class of the tree RAAGs.

Date: April 9, 2025.

2020 *Mathematics Subject Classification.* 20F65,20F55.

Key words and phrases. Right-angled Coxeter group, right-angled Artin group, RACG, RAAG, quasiisometry.

This research was funded by the Austrian Science Fund (FWF) 10.55776/P34214.

We give a complete classification of triangle-free graphs Γ such that W_Γ is in the quasiisometry class of the tree RAAGs, and show that in this case W_Γ contains a visible finite index tree RAAG subgroup. See Theorem 3.1. The proof uses the group theoretic version of JSJ decompositions (over 2–ended subgroups). Such decompositions of RAAGs and RACGs can be read off from their presentation graphs. One step in the proof is to describe what the JSJ decomposition of a group in the quasiisometry class of the tree RAAGs must look like. The second is to show that, for a 2–dimensional RACG, these conditions are already strong enough to produce a visible finite index RAAG subgroup.

There are infinitely many commensurability classes of tree RAAGs [2, 3], so the main theorem does not say that quasiisometry to a particular tree RAAG implies commensurability to that same tree RAAG. Nevertheless, we find the commensurability aspect of the theorem surprising. Usually when a geometric relation can be promoted to an algebraic relation it is because there is some strong rigidity phenomenon at play, but the class of tree RAAGs is quite flexible. For example, the results of Huang and Kleiner [17] on groups quasiisometric to RAAGs with finite outer automorphism group do not apply to tree RAAGs.

Acknowledgements. This paper is an spin-off of [6], where the question is when W_Γ is quasiisometric to any RAAG whatsoever.

Alexandra Edletzberger brought Nguyen and Tran's paper to my attention, and had already realized that their decomposition was related to the JSJ decomposition, and that cycles of cuts should be obstructions to being quasiisometric to a RAAG [15, Remark 1.38, Remark 4.29]. The latter became [6, Theorem 5.16].

2. Preliminaries

2.1. **Terminology.** The *right-angled Artin group with presentation graph* Δ, for a finite simplicial graph Δ, is the group A_Δ presented by a generator for each vertex of Δ and a commuting relation between two generators when they are connected by an edge of Δ. The *right-angled Coxeter group with presentation graph* Γ, W_Γ, is defined similarly, with additional relations saying that each generator has order 2. In both cases these groups have a geometric action on a CAT(0) cube complex, which has dimension at most 2 when the presentation graph is triangle-free.

$\mathbb{F}_n$ is the free group of rank n, and $\mathbb{F}$ is $\mathbb{F}_n$ for some unspecified $n \geqslant 2$ when the precise rank is not important. A *basis* of $\mathbb{F}$ means a free generating set.

An *induced* subgraph of a simplicial graph is a subgraph that contains all possible edges between its vertices. If Υ' is an induced subgraph of Υ then the inclusion at the graph level extends to inclusions $A_{\Upsilon'} < A_\Upsilon$ and $W_{\Upsilon'} < W_\Upsilon$. A subgroup of a RAAG or RACG defined in this way is called a *special subgroup*.

If Γ is a presentation graph of a RACG then a *thick join* means a join subgraph $A * B$ of Γ such that W_A and W_B are infinite groups, which happens when A and B are incomplete subgraphs. A *suspension* is a join in which A is a 2–anticlique. If A is a 2–anticlique and $|B| > 2$ then A is called the *pole* or the *suspension points* of $A * B$ and B is called the *suspended points*. A vertex is *essential* if it has valence at least 3. A *cone vertex* is everybody's neighbor. If v is a vertex, its *link* $\mathrm{lk}(v)$ is its set of neighbors and its *star* is $\mathrm{st}(v) := \mathrm{lk}(v) \cup \{v\}$. A graph is *biconnected* if it is connected with no cut vertex. Thus, a single edge is biconnected.

2.2. **Visible RAAG subgroups.** If Γ is a graph, its complement graph Γ^c is the graph on the same vertex set that has an edge precisely when Γ does not. We will have $\Lambda \subset \Gamma^c$ that is a disjoint union of trees Λ_0 and Λ_1. For a set of vertices V in Γ, if V is contained in one of the Λ_i then $\mathrm{hull}_\Lambda(V)$ is the set of vertices of Γ contained in the convex hull of V in Λ_i. The *commuting graph* Δ of Λ is the graph that has a vertex for each edge of Λ, with an edge between two of them if their support spans a square in Γ. There is a natural homomorphism $A_\Delta \to W_\Gamma$ defined by sending the generator of A_Δ corresponding to vertex $\{a,b\}$ of Δ to the element $ab \in W_\Gamma$. it is easy to come up with examples where this homomorphism is not injective, but Dani and Levcovitz [11] give conditions that imply the homomorphism is injective and the image has finite index.

Theorem 2.1 ([11, Theorem 4.8, Corollary 3.33, Remark 4.3])**.** *Let Γ be an incomplete, triangle-free graph without separating cliques. Suppose Λ is a disjoint union of trees Λ_0 and Λ_1 in Γ^c such that Λ spans Γ and for each i, no two vertices of Λ_i are adjacent in Γ. Let Δ be the commuting graph of Λ.*

If Λ satisfies the following conditions, then the homomorphism $A_\Delta \to W_\Gamma$ is injective and its image has index 4 in W_Γ.

- $\mathcal{R}_3$: *If $\{a,b\} * \{c,d\}$ is a square in Γ then $\mathrm{hull}_\Lambda\{a,b\} * \mathrm{hull}_\Lambda\{c,d\} \subset \Gamma$.*
- $\mathcal{R}_4$: *If $a - b$ is an edge in a cycle $\gamma \subset \Gamma$ then there is a square $\{a,a'\} * \{b,b'\}$ with $a', b' \in \mathrm{hull}_\Lambda(\gamma)$.*

Call a graph Λ satisfying the hypotheses of Theorem 2.1 a *finite index Dani-Levcovitz* Λ (FIDL–Λ). When we have a FIDL–Λ, the commuting graph Δ gives a 'visible' finite index RAAG subgroup of W_Γ. It is 'visible' in the sense that vertices of Δ correspond to diagonals of squares of Γ, and vertices of Δ span an edge precisely when their support spans a square in Γ.

2.3. **CFS graphs.** For W_Γ to be quasiisometric to a RAAG it must have at most quadratic divergence, which is characterized by Γ having property *CFS* [12].

Definition 2.2 (cf.[7, Definition 2.3])**.** The *diagonal graph* $\boxslash(\Gamma)$ of Γ is the graph with an edge for each induced square of Γ, with vertices representing diagonal pairs. The *support* of a vertex of $\boxslash(\Gamma)$ is the two vertices of Γ in the corresponding diagonal pair. The support of a subset of $\boxslash(\Gamma)$ is the union of supports of its vertices. Γ is *CFS* if $\boxslash(\Gamma)$ contains a component whose support is all non-cone vertices of Γ.

We will not invoke the CFS property directly, but $\boxslash(\Gamma)$ will figure prominently. Consider that A_Δ is 1–ended when Δ is connected and not a single vertex, and W_Γ is 1–ended when Γ is incomplete without separating cliques. Thus, in the 1–ended case, for a FIDL–Λ as in the previous section, we may identify Δ with a subgraph of $\boxslash(\Gamma)$, since Δ is connected and an edge in Δ corresponds to a square in Γ.

2.4. **JSJ decompositions.** We assume familiarity with Bass-Serre theory. In this paper we are interested in JSJ decompositions of finitely presented groups over 2–ended subgroups (equivalently, over virtually $\mathbb{Z}$ subgroups), which are graph of groups decompositions that encode all possible splittings of the group over 2–ended subgroups. The existence of non-trivial JSJ decompositions, for finitely presented groups not commensurable to a surface group, is a quasiisometry invariant [22], but the precise structure of a particular JSJ decomposition is not. One can make a canonical object by passing to the Bass-Serre tree of any JSJ decomposition and collapsing cylinders to make the *JSJ tree of cylinders* [16]. It is a bipartite tree where one part consists of *cylinder* vertices, which correspond to commensurability classes of universally elliptic 2–ended splitting subgroups, and one part is made up of *rigid* and *hanging* vertices that belong to more than one cylinder. Hanging vertices contain the non-(universally elliptic) 2–ended splitting subgroups; that is, they contain collections of 2–ended splitting subgroups that give incompatible group splittings. Rigid vertices correspond to subgroups that are not split by any 2–ended splitting of the ambient group.

It follows from Papasoglu's work [22] (cf. [9]) that a quasiisometry $\phi\colon G \to G'$ between two finitely presented groups induces an isomorphism of their JSJ trees of cylinders, $\phi_*\colon T \to T'$, that preserves vertex type: cylinder/rigid/hanging. Furthermore, if X is a tree of spaces for G over T, and similarly X' for G' over T', then restriction of ϕ to each vertex space X_v is uniformly close to a quasiisometry $\phi_v\colon X_v \to X'_{\phi_*(v)}$ that induces a bijection of incident edge spaces by taking $X_e \subset X_v$ to a set uniformly coarsely equivalent to $X'_{\phi_*(e)} \subset X'_{\phi_*(v)}$.

The quotient of the action of a group on its JSJ tree of cylinders gives the *graph of cylinders*. Since the JSJ tree of cylinders can be recovered by development, this graph of groups contains information on quasiisometry invariants for the group, such as the vertex types that occur, the quasiisometry classes the vertex groups, and the relative quasiisometry types of the vertices relative to the peripheral pattern of their incident edge groups, which will be discussed in Section 2.5.

Work of Clay [10] and Margolis [18], says that the graph of cylinders of a RAAG can be described 'visually' in terms of the presentation graph Δ as follows, where in each case the local group is the special subgroup defined by the relevant subgraph:

- Cylinders are stars of cut vertices of Δ.
- Rigid vertices are maximal biconnected subgraphs of Δ that are not contained in a single cylinder.
- Hanging vertices do not occur.
- There is an edge between a cylinder and rigid vertex when they intersect.

Edletzberger [14], building on work of Mihalik and Tschantz [19] and Dani and Thomas [13], gave a visual description of the graph of cylinders of a 1–ended RACG with a triangle-free presentation graph Γ that is not a cycle graph.

- Cylinders are suspensions $\{a, b\} * (\mathrm{lk}(a) \cap \mathrm{lk}(b))$ where *there is a cut* $\{a - b\}$, meaning either:
  - $\{a, b\}$ is a cut pair, that is, $\Gamma - \{a, b\}$ is not connected.
  - $\{a, b\}$ is not a cut pair, but there is a vertex $c \in \mathrm{lk}(a) \cap \mathrm{lk}(b)$ such that the 2–path $a \leftrightarrow c \leftrightarrow b$ disconnects Γ. In this case $a \leftrightarrow c \leftrightarrow b$ is a *cut 2–path.*
- Rigid vertices are sets of essential vertices of Γ of size at least four that cannot be separated by a cut and that are maximal with respect to inclusion among such sets.
- Hanging vertices consist of either collections of cut pairs that pairwise separate one another's points, or collections of cut 2–paths that share a common center and pairwise separate each other's endpoints. (Cut pairs cannot cross cut 2–paths, and vice versa [7, after Definition 2.8].)
- There is an edge between a cylinder $\{a, b\} * (\mathrm{lk}(a) \cap \mathrm{lk}(b))$ and a rigid/hanging vertex if latter contains a cut $\{a - b\}$. In this case the edge corresponds to intersection of subgraphs.

2.5. **Peripheral patterns.** A *peripheral pattern* $\mathcal{P}(G,\mathcal{H})$ in a group G relative to a collection $\mathcal{H} := \{H_0,\dots,H_{n-1}\}$ of subgroups of G is the set of coarse equivalence classes of left cosets of the H_i in G. When $\mathcal{H}$ consists of 2–ended subgroups the peripheral pattern it generates may be called a *line pattern.* A map between two groups with peripheral patterns is *pattern preserving* if it takes each coset of one pattern to within uniformly bounded Hausdorff distance of one in the other, and vice versa, inducing a bijection between patterns. The *relative quasiisometry type* of $\mathcal{P}(G,\mathcal{H})$ is the equivalence class of spaces with peripheral patterns up to pattern preserving quasiisometry.

Line patterns in free groups have been considered before [21, 8, 5]. Associated to a line pattern there is a *decomposition space*, which is the quotient of the boundary of the free group obtained by identifying the two endpoints of each line in the pattern. Pattern preserving quasiisometries induce homeomorphisms of decomposition spaces. Two cases are of particular interest for us. It turns out that each of these types defines a single relative quasiisometry type of line pattern.

Definition 2.3. A *surface type (line) pattern* in $\mathbb{F}$ is a peripheral pattern $\mathcal{P}(\mathbb{F},\mathcal{H})$ with $\mathcal{H} = \{\langle f_0\rangle,\dots,\langle f_{n-1}\rangle\}$ such that, up to passing to roots, conjugates, and inverses, and removing duplicate entries, there is a compact surface with fundamental group $\mathbb{F}$ such that $\{f_0,\dots,f_{n-1}\}$ are the elements of $\mathbb{F}$ represented by the boundary curves. Equivalently, its decomposition space is a circle.

Definition 2.4. A *basic (line) pattern* in $\mathbb{F}$ is a peripheral pattern $\mathcal{P}(\mathbb{F},\mathcal{H})$ with $\mathcal{H} = \{\langle f_0\rangle,\dots,\langle f_{n-1}\rangle\}$ such that, up to passing to roots, conjugates, and inverses, and removing duplicate entries, $\{f_0,\dots,f_{n-1}\}$ is a subset of a basis of $\mathbb{F}$. Equivalently, its decomposition space is totally disconnected.

In the present work we are interest not in line patterns in free groups, but in *plane patterns* in $\mathbb{F}\times\mathbb{Z}$. Let z be a generator of the center. By a *plane pattern* in $\mathbb{F}\times\mathbb{Z}$, we mean a peripheral pattern $\mathcal{P}(\mathbb{F}\times\mathbb{Z},\mathcal{H})$ where $\mathcal{H}$ consists of $\mathbb{Z}^2$ subgroups, which we might as well assume, up to coarse equivalence, are maximal $\mathbb{Z}^2$ subgroups. A maximal $\mathbb{Z}^2$ subgroup of $\mathbb{F}\times\mathbb{Z}$ is of the form $\langle f\rangle\times\langle z\rangle$, where $f\in\mathbb{F}$ and $z\in\mathbb{Z}$ are nontrivial elements that are not proper powers. Quasiisometries of $\mathbb{F}\times\mathbb{Z}$ coarsely preserve the center, so from a quasiisometry $\phi\colon \mathbb{F}\times\mathbb{Z}\to\mathbb{F}'\times\mathbb{Z}$ we get a quasiisometry $\pi_{\mathbb{F}'}\circ\phi|_{\mathbb{F}\times\{1\}}\colon \mathbb{F}\to\mathbb{F}'$, where $\pi_{\mathbb{F}'}\colon \mathbb{F}'\times\mathbb{Z}$ is projection to the first factor given by killing the center. The projections $\pi_{\mathbb{F}}\colon \mathbb{F}\times\mathbb{Z}\to\mathbb{F}$ and $\pi_{\mathbb{F}'}\colon \mathbb{F}'\times\mathbb{Z}\to\mathbb{F}'$ send peripheral plane patterns to peripheral line patterns, and the original quasiisometry ϕ preserves plane patterns if and only if $\pi_{\mathbb{F}'}\circ\phi|_{\mathbb{F}\times\{1\}}$ preserves the projected line patterns. We extend the basic/surface terminology to the plane patterns according to the type of the projection. More generally, extend the terminology to patterns of virtually $\mathbb{Z}^2$ subgroups in virtually $\mathbb{F}\times\mathbb{Z}$ groups, according to the type of the pattern obtained by pushing forward the pattern by the map induced by the restriction to a finite index $\mathbb{F}\times\mathbb{Z}$ subgroup.

Lemma 2.5. *Let Δ be a finite tree of diameter at least three, so that A_Δ is a 1–ended RAAG with a nontrivial graph of cylinders. For every cylinder in the graph of cylinders, the peripheral pattern coming from the incident edge groups is basic.*

Proof. Since Δ is a tree, every non-leaf is a cut vertex, so cylinders are stars of non-leaf vertices v of Δ. If v is a non-leaf then $\mathrm{lk}(v) = \{w_0,\dots,w_{n-1},\ell_0,\dots,\ell_{m-1}\}$, where the w_i are non-leaves and the ℓ_j are leaves. Since the diameter of Δ is greater than 2, $n > 0$ and $m + n \geqslant 2$. The cylinder group is $A_{\mathrm{st}(v)} = \langle v \rangle \times \langle w_0, \dots, w_{n-1}, \ell_0, \dots, \ell_{m-1} \rangle \cong \mathbb{Z} \times \mathbb{F}_{m+n}$. Since Δ is a tree, the maximal biconnected subgraphs are single edges, and those that are not contained in a single cylinder are the non-leaf edges. Thus, the neighbors of the cylinder $\mathrm{st}(v)$ in the graph of cylinders are the rigid vertices $v \,\text{—}\, w_i$ for $0 \leqslant i < n$, with $A_{\{v,w_i\}} \cong \mathbb{Z}^2$. These give a basic plane pattern, since $\{w_0, \dots, w_{n-1}\}$ is a subset of the basis $\{w_0, \dots, w_{n-1}, \ell_0, \dots, \ell_{m-1}\}$ of the $\mathbb{F}_{m+n}$ factor of $A_{\mathrm{st}(v)}$. □

Example 2.6. Let Γ be the nonplanar graph in Figure 2a. Let W_Γ be the RACG whose presentation graph is Γ, so that the generators of W_Γ are s_i where s_i corresponds to vertex i of Γ. (For compactness of notation, we omit the 's' in figures and subscripts.) The graph of cylinders of W_Γ is shown in Figure 2b, where the labels are the cylinder groups and each unlabelled vertex is a rigid vertex whose group is the intersection of its neighboring cylinder groups, which in each case is isomorphic to the product of two infinite dihedral groups. All edge groups are isomorphic to their neighboring rigid vertex group, and all edge maps are inclusion.

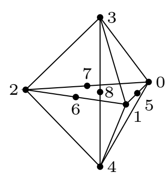

(a) Γ

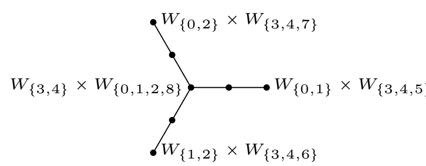

(b) Graph of cylinders of W_Γ.

FIGURE 2. Graph Γ and graph of cylinders of W_Γ in Example 2.6.

The middle cylinder vertex $\mathcal{C} := W_{\{3,4\}} \times W_{\{0,1,2,8\}}$ has peripheral pattern generated by $\mathcal{H} := \{W_{\{0,1\}} \times W_{\{3,4\}},\, W_{\{1,2\}} \times W_{\{3,4\}},\, W_{\{0,2\}} \times W_{\{3,4\}}\}$. The group $\mathcal{C}$ has an index 4 subgroup isomorphic to $\mathbb{F}_3 \times \mathbb{Z}$, namely $\langle w, x, y \rangle \times \langle z \rangle$, where $w := s_0 s_8$, $x := s_0 s_1$, $y := s_1 s_2$, and $z := s_3 s_4$. Restriction to the finite index subgroup induces a pattern preserving quasiisometry between $\mathcal{P}(\mathcal{C}, \mathcal{H})$ and $\mathbb{F}_3 \times \mathbb{Z}$ with plane pattern generated by $\{\langle x \rangle \times \langle z \rangle,\, \langle y \rangle \times \langle z \rangle,\, \langle xy \rangle \times \langle z \rangle\}$. Projection to the $\mathbb{F}_3$ factor gives line pattern generated by $\{x, y, xy\}$. Since the generators live in a free factor $\langle x, y \rangle$ of $\mathbb{F}_3$, it follows that the decomposition space for $\mathcal{P}(\langle x, y \rangle, \{x, y, xy\})$ embeds into the decomposition space for $\mathcal{P}(\mathbb{F}_3, \{x, y, xy\})$. The former is a surface type pattern: it matches the standard description of the fundamental group and boundary curves of a 3–holed sphere. Therefore, the decomposition space for $\mathcal{P}(\mathbb{F}_3, \{x, y, xy\})$ contains circles; in particular, it is not totally disconnected, so the line pattern is not basic. Thus, W_Γ is not quasiisometric to a tree RAAG, since the non-basic plane pattern in the cylinder $W_{\{0,1,2,8\}} \times W_{\{3,4\}}$ of the RACG cannot match the plane patterns in cylinder vertices of tree RAAGs, which by Lemma 2.5 are all basic. ◇

In the planar case, according to Nguyen and Tran's argument, Γ decomposes into a tree of maximal suspensions, with neighbors determined by two suspensions sharing a square. Playing the game as in the example by restricting to a finite index torsion-free subgroup and then projecting to the free factor, one can see that there are only two possibilities for a maximal suspension $\{a, b\} * \{c_0, \dots, c_{n-1}\}$: either there are fewer than n neighbors and the line pattern is basic, or there are exactly n neighbors and the line pattern is surface type. Nguyen and Tran use 'suspension of n vertices has exactly n neighboring suspensions' as an obstruction to being quasiisometric to a tree RAAG. Example 2.6 illustrates that in the nonplanar case there are more varied ways for a pattern to fail to be basic; $\mathcal{C}$ has a non-basic pattern, but it is a suspension of 4 vertices with only 3 neighboring suspensions.

3. RACGs quasiisometric to tree RAAGs

Theorem 3.1. *Let Γ be an incomplete, triangle-free graph without separating cliques. The following are equivalent:*

(1) W_Γ has an index four visible tree RAAG subgroup.
(2) W_Γ is in the quasiisometry class of the tree RAAGs.
(3) W_Γ has a nontrivial graph of cylinders with no hanging vertices and:
  (a) Cylinders are virtually $\mathbb{F} \times \mathbb{Z}$.
  (b) Rigid vertices are virtually $\mathbb{Z}^2$. Each is adjacent to two cylinders, with edge group equal to the rigid vertex group.
  (c) In each cylinder the incident edge groups form a basic plane pattern.
(4) Γ has the following structure:
  (a) Maximal thick joins are non-square suspensions. Every vertex and edge is contained in a maximal suspension. There is a cut coming from the pole of each maximal suspension, and all cuts are of this form.
  (b) Every set of at least four essential vertices that is not separated by a cut is a square, each of whose diagonals is the pole of a maximal suspension.
  (c) There does not exist a maximal suspension $\{a, b\} * \{c_0, \dots, c_{n-1}\}$ and $3 \leqslant m \leqslant n$ such that for all $0 \leqslant i < m$ the pair $\{c_i, c_{(i+1) \mod m}\}$ is the pole of a maximal suspension.

Proof. Inclusion of a finite index subgroup induces a quasiisometry, so (1) $\implies$ (2).

The graph of cylinders of a tree RAAG was described in the proof of Lemma 2.5. Quasiisometry invariance of the JSJ tree of cylinders implies that any group quasiisometric to a tree RAAG has JSJ tree of cylinders whose cylinders are quasiisometric to $\mathbb{F} \times \mathbb{Z}$, whose rigid vertices are quasiisometric to $\mathbb{Z}^2$, has no hanging vertices, has exactly two cylinders adjacent to each rigid vertex, and such that in each cylinder the plane pattern coming from incident edges is a basic plane pattern. 'Quasiisometric to' $\mathbb{F} \times \mathbb{Z}$ or $\mathbb{Z}^2$ can be promoted to 'virtually' $\mathbb{F} \times \mathbb{Z}$ or $\mathbb{Z}^2$, respectively, since these groups are quasiisometrically rigid.

Finally, if $\tilde{r}$ is a valence 2 rigid vertex of the JSJ tree of cylinders and r is its image in the graph of cylinders then r either has two adjacent cylinders connected to r via edge groups equal to the vertex group of r, or it has one adjacent cylinder connected to r via an edge group that is an index 2 subgroup of the vertex group of r. The latter is impossible, by visibility of the graph of cylinders in RACGs, since when Γ is triangle-free the only virtually $\mathbb{Z}^2$ special subgroups are defined by squares of Γ, so one cannot be properly contained in another. Thus, (2) $\implies$ (3).

Assume (3). Cylinders are suspensions $\{a, b\} * (\mathrm{lk}(a) \cap \mathrm{lk}(b))$ where there is an uncrossed cut $\{a - b\}$. Since cylinders are virtually $\mathbb{F} \times \mathbb{Z}$, the suspension is non-square. Since there are no hanging vertices, all cuts arise this way.

Rigid vertices are maximal subsets of Γ of at least four essential vertices that cannot be separated by a cut. Since they are virtually $\mathbb{Z}^2$, they correspond to squares in Γ. A square is contained in a maximal suspension if and only if one of its diagonals is the pole of the suspension. Since the rigid vertex must be adjacent to two cylinders, both diagonals of the square must be poles of maximal suspensions giving cylinders.

A $K_{3,3}$ subgraph of Γ cannot be separated by a cut, and a suspension $\{a,b\} * C$ with $|C| \geqslant 3$ can only be separated by a cut of the form $\{a - b\}$, so the existence of either a maximal thick join that is not a suspension or a non-square suspension whose suspension points do not make a cut implies there is a non-virtually–$\mathbb{Z}^2$ rigid vertex, which is not true, so maximal joins are suspensions whose pole gives a cut.

Since the graph of cylinders gives a visual decomposition, every vertex and edge of Γ is contained in a subgraph corresponding to a vertex of the graph of cylinders. Since there are no hanging vertices and rigid vertices are assumed to be contained in cylinders, every vertex and edge is contained in a cylinder.

Item (4c) describes a cylinder vertex of the graph of cylinders whose incident edges form a non-basic plane pattern, contrary to (3c).

It remains to show (4) $\implies$ (1). A graph Δ for which A_Δ is a finite index tree RAAG subgroup of W_Γ can be constructed essentially by taking the graph of cylinders of Γ and adding some leaves. We prove this by constructing a FIDL-Λ with Δ as its commuting graph, as in Theorem 2.1. This will take some work.

Assume (4). Any graph of groups decomposition of W_Γ has underlying graph a tree because W_Γ is generated by torsion elements, so it cannot surject onto $\mathbb{Z}$. Let T be the tree with one vertex for each cylinder in the graph of cylinders and an edge between cylinders if they share a square. This is just the underlying graph of the graph of cylinders, where we have 'forgotten' the rigid vertices, in the sense that they are all of valence two, so we think of them as the midpoint of an edge instead of as a valence two vertex. In fact, we will think of T as a subgraph of $\boxslash(\Gamma)$, whose edges were defined to be the squares of Γ.

We make some claims about the graph structure of Γ deduced from (4). Assuming the claims, we finish the proof of (1). After that we will prove the claims.

Claim 3.1.1. For every square $\{a,b\} * \{c,d\}$ of Γ, at least one of the diagonals $\{a,b\}$ or $\{c,d\}$ is the pole of a non-square suspension that is a maximal thick join.

Claim 3.1.2. $\boxslash(\Gamma)$ is a tree and T can be identified, via the map sending a maximal thick join $\{a,b\} * C$ to the vertex $\{a,b\} \in \boxslash(\Gamma)$, with the subtree obtained by removing all leaves of $\boxslash(\Gamma)$,

Examples illustrating Claim 3.1.3 are shown in Figure 3. In the claim, when $\{p_0,q_0\}$,$\{p_1,q_1\}$,...,$\{p_n,q_n\}$ is a geodesic in $T \subset \boxslash(\Gamma)$ then for a vertex $v \in \Gamma$ let $I_v := \{i \mid v \in \{p_i,q_i\}\}$. Also, let $\mathbb{E}$ denote the even numbers, and $\mathbb{O}$ the odd.

Claim 3.1.3. If $\{p_0,q_0\}$,$\{p_1,q_1\}$,...,$\{p_n,q_n\}$ is a geodesic in T with $n > 0$ and there exists a vertex $z \in \Gamma$ with $\{p_0,q_0,p_n,q_n\} \subset \mathrm{lk}(z)$ then:

- (i) $n \in \mathbb{E}$
- (ii) $I_z = [0,n] \cap \mathbb{O}$
- (iii) For all $0 < i < n$, there is a cut $\{p_i - q_i\}$ separating $\{p_j,q_j \mid j < i\} - \{p_i - q_i\}$ from $\{p_j,q_j \mid i < j\} - \{p_i - q_i\}$. If $i \in \mathbb{E}$ then the cut is $\{p_i - q_i\} = p_i \bullet\!\!-\!\!\bullet z \bullet\!\!-\!\!\bullet q_i$.

Claim 3.1.4. If $\{a,b\} \neq \{a,c\}$ are vertices in $\boxslash(\Gamma)$ then the distance between them is even and every second vertex on the geodesic between them contains a.

Claim 3.1.5. Every vertex of Γ is suspended in some maximal suspension. A vertex that is not a suspension point of any maximal suspension is contained in a unique maximal suspension.

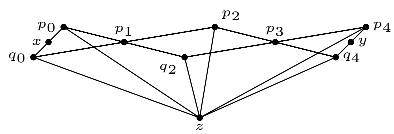

(a) Γ_1, with $z = q_1 = q_3$

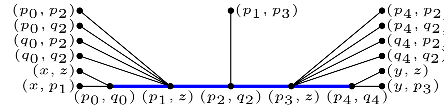

(b) $\boxslash(\Gamma_1)$

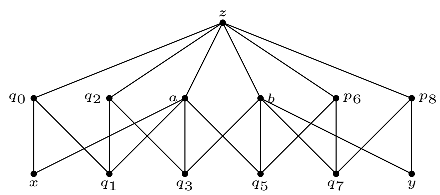

(c) Γ_2, with $z = p_1 = p_3 = p_5 = p_7$, $a = p_0 = p_2 = p_4$, and $b = q_4 = q_6 = q_8$

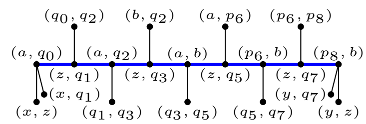

(d) $\boxslash(\Gamma_2)$

FIGURE 3. Two graphs as in Claim 3.1.3 with their diagonal graphs containing the geodesic $\{p_0, q_0\}, \{p_1, q_1\}, \ldots, \{p_n, q_n\}$.

Assuming the claims, we construct a FIDL–Λ.

By Claim 3.1.2, $\boxslash(\Gamma)$ is a tree, so it can be 2–colored 0/1 such that neighboring vertices have different colors. By Claim 3.1.4, if v is in the support of two different vertices of $\boxslash(\Gamma)$ then they have even distance in $\boxslash(\Gamma)$, so they have the same color. Thus, the 2–coloring on $\boxslash(\Gamma)$ induces a 2–coloring of Γ.

Let $\{a, b\}$ be a vertex of T, corresponding to a maximal suspension $\{a, b\} * \{c_0, \ldots, c_{n-1}\}$. Let $\Lambda_{a,b} \subset \Gamma^c$ be a graph on vertices $\{c_0, \ldots, c_{n-1}\}$, constructed as follows. Add an edge in $\Lambda_{a,b}$ between c_i and c_j if $\{c_i, c_j\}$ is a vertex of T. Call these 'mandatory edges'. Condition (4c) implies the mandatory edges form a forest. Add additional 'discretionary' edges to make $\Lambda_{a,b}$ a tree. Let $T_i \subset T$ be the vertices with color i. Let Λ_0 be the union of trees $\Lambda_{a,b}$ such that $\{a, b\} \in T_1$, so that $\mathrm{supp}(\Lambda_0)$ consists of all vertices of Γ colored 0. Define Λ_1 analogously for T_0. By Claim 3.1.5, $\Lambda := \Lambda_0 \sqcup \Lambda_1$ spans Γ. Vertices of Λ_i are colored i, so they are not adjacent in Γ.

We prove Λ_0 is a tree by induction. Pick a vertex $\{a_0, b_0\} \in T_1$. By construction, Λ_{a_0,b_0} is a tree. Now suppose that $\bigcup_{i=0}^{m-1} \Lambda_{a_i,b_i}$ is a tree, where the set of vertices $\{a_i, b_i\} \in T_1$ with $i < m$ is the intersection of a convex subset of T with T_1. Let $\{a_m, b_m\} \in T_1$ that is closest to the existing set. Since $\boxslash(\Gamma)$ is a tree, the convexity condition implies that there is a unique vertex $\{c, d\} \in T_0$ separating $\{a_m, b_m\}$ from all of the other $\{a_i, b_i\}$, and there is at least one index $j < m$ such that $\{c, d\}$ is a neighbor of $\{a_j, b_j\}$ in T. Thus, $c \leftrightarrow d$ is a mandatory edge in Λ_{a_m,b_m} and in any such Λ_{a_j,b_j}. We have $c \leftrightarrow d \subset \Lambda_{a_m,b_m} \cap \bigcup_{i=0}^{m-1} \Lambda_{a_i,b_i}$. Conversely, suppose, for some $j < m$, that $\Lambda_{a_j,b_j} \cap \Lambda_{a_m,b_m} \neq \varnothing$, so there is a vertex $z \in \Gamma$ with $\{a_j, b_j, a_m, b_m\} \subset \mathrm{lk}(z)$. Since $\{c, d\}$ is the penultimate vertex on the unique geodesic from $\{a_j, b_j\}$ to $\{a_m, b_m\}$ in $\boxslash(\Gamma)$, Claim 3.1.3 says $z \in \{c, d\}$. Thus, $\Lambda_{a_m,b_m} \cap \bigcup_{i=0}^{m-1} \Lambda_{a_i,b_i}$ is exactly the edge $c \leftrightarrow d$, but the union of two trees with exactly an edge in common is a tree, so Λ_0 is a tree by induction, and Λ_1 similarly.

Now we check that conditions $\mathcal{R}_3$ and $\mathcal{R}_4$ of Theorem 2.1 are satisfied.

Suppose $\{a, b\} * \{c, d\}$ is a square in Γ. Claim 3.1.1 says either $\{a, b\}$ or $\{c, d\}$ is the pole of a maximal suspension; suppose $\{a, b\} * C$ is a maximal suspension with

$\{c,d\} \subset C$. Then $\mathrm{hull}_\Lambda\{a,b\} = \{a,b\}$ and $\mathrm{hull}_\Lambda\{c,d\} \subset \Lambda_{a,b} \subset C$, so $\mathrm{hull}_\Lambda\{a,b\} * \mathrm{hull}_\Lambda\{c,d\} \subset \{a,b\} * C \subset \Gamma$. Condition $\mathcal{R}_3$ is satisfied.

Suppose condition $\mathcal{R}_4$ is not satisfied. Then there is a shortest simple cycle γ in Γ containing an edge e that is not in a square with vertices in $\mathrm{hull}_\Lambda(\gamma)$. Such a γ is induced, because if γ contains vertices v and w that are adjacent in Γ but not in γ then we could surger γ into two simple cycles γ' and γ'' that are each an arc of γ plus the edge $v \leftrightarrow w$. Both are strictly shorter than γ, and one of them, say γ', contains e, but $\mathrm{hull}_\Lambda(\gamma') \subset \mathrm{hull}_\Lambda(\gamma)$, so e is not contained in a square with vertices in $\mathrm{hull}_\Lambda(\gamma')$. This would contradict minimality of γ.

Let $u \leftrightarrow v \leftrightarrow w \leftrightarrow x$ be the subsegment of γ such that $e = v \leftrightarrow w$. If $u \leftrightarrow v$ and $w \leftrightarrow x$ are contained in a common suspension then so is $v \leftrightarrow x$, and $\{u,w\} * \{v,x\}$ is a square, contradicting the definition of e. Thus, we may assume that there is not a suspension containing all three of the edges. The argument will be to produce a cut through v or w separating u from x. Since γ is a loop, it must pass back through the cut again at some other vertex v', which will lead to a contradiction with the assumption that γ is induced.

Case 1: There is a maximal suspension σ containing $u \leftrightarrow v \leftrightarrow w$ and a different maximal suspension σ' containing $v \leftrightarrow w \leftrightarrow x$. Suppose the distance between σ and σ' in T is minimal among pairs with these properties.

Case 1a: Suppose σ and σ' are adjacent in T. Their intersection is the join of their poles. If $\{u,w\}$ is the pole of σ then all four of u, v, w, and x are contained in σ', contrary to hypothesis, and similarly for σ', so there are vertices $v' \neq w'$ such that $\sigma \cap \sigma' = \{v,v'\} * \{w,w'\}$. If no cut separates u from x then $\{u\} * \{v,v'\} * \{w,w'\} * \{x\}$ is a non-square rigid vertex, which is impossible, so there is at least one cut of the form $\{v - v'\}$ or $\{w - w'\}$ separating u from x. These cases are symmetric, so assume it is $\{v - v'\}$. Since v is adjacent to u and w, if γ passes through v' then γ fails to be induced. However, the only way for γ to pass back through the cut avoiding v' is if the cut is a cut 2–path $v \leftrightarrow u' \leftrightarrow v'$ with $u' \neq w$. In this case, γ passes back through u', which is adjacent to v, so γ is not induced.

Case 1b: Suppose σ and σ' are not adjacent in T. For $\{a,b\} \in T$, let $\sigma_{a,b} := \{a,b\} * (\mathrm{link}(a) \cap \mathrm{link}(b))$ be the maximal suspension of Γ with pole $\{a,b\}$. Let $\{p_0,q_0\}, \ldots, \{p_n,q_n\}$ be the geodesic in T between $\sigma = \sigma_{p_0,q_0}$ and $\sigma' = \sigma_{p_n,q_n}$. Vertices v and w cannot both be suspended in the same suspension since they are connected by an edge and Γ is triangle free, so, up to symmetry, there are two subcases: either w is a suspended and v is suspension in both σ and σ', or v is a suspension point of σ and suspended in σ' and the opposite is true for w.

In the first subcase, $\{p_0,q_0\} = \{v,v'\}$, for some v', and $\{p_n,q_n\} = \{v,x\}$. Claim 3.1.3 says $n \in \mathbb{E}$, $w \in \{p_1,q_1\}$, and there is a cut $\{p_1 - q_1\}$ separating $\{v,v'\} - \{p_1 - q_1\}$ from $\{v,x\} - \{p_1 - q_1\}$. Since $v \in \{v,v'\} \cap \{v,x\}$, the cut is the cut 2–path $p_1 \leftrightarrow v \leftrightarrow q_1$, and it separates v' from x. It also separates u from x, since u is adjacent to v'. So, γ crosses the cut once through v to get from u to x, and must cross back again later, but both of the other vertices of the cut are adjacent to v, contradicting that γ is induced.

In the second subcase we have that w is suspended in $\sigma_{p_{n-1}q_{n-1}}$, and, by minimality of n, $u \notin \sigma_{p_1,q_1}$ and $x \notin \sigma_{p_{n-1},q_{n-1}}$. Apply Claim 3.1.3 to $\{p_0,q_0\},\ldots,\{p_{n-1},q_{n-1}\}$ with $z = w$, which gives that $w \in \{p_1,q_1\}$ and there is a cut $\{p_1 - q_1\}$ in Γ separating $\{p_0,q_0\} - \{p_1 - q_1\}$ from $\{p_{n-1},q_{n-1}\} - \{p_1 - q_1\}$. Since $u \notin \sigma_{p_1,q_1}$ and $x \notin \sigma_{p_{n-1},q_{n-1}}$, $\{p_1 - q_1\}$ separates u from x. Since v is adjacent to u and v and x

are both suspended in σ' with two common neighbors $\{p_n, q_n\} \neq \{p_1, q_1\}$, the cut $\{p_1 - q_1\}$ is a cut 2–path $p_1 \leftrightarrow v \leftrightarrow q_1$. Thus, γ crosses $\{p_1 - q_1\}$ once through v and then back again through a vertex adjacent to v, contradicting that it is induced.

Case 2: The path $u \leftrightarrow v \leftrightarrow w$ is not contained in a maximal suspension. Let $\{p_0, q_0\}, \dots, \{p_n, q_n\}$ be a geodesic in T such that $u, v \in \sigma_{p_0,q_0}$ and $v, w \in \sigma_{p_n,q_n}$, and such that $n > 0$ is minimal among such geodesics. Edges sharing a common vertex cannot be contained in adjacent maximal suspensions without both being contained in at least one of the two, so $n \geqslant 2$. Minimality implies $|\{u, v\} \cap \{p_0, q_0\}| = 1$ and $\{u, v\} \cap \{p_1, q_1\} = \varnothing$ and $|\{v, w\} \cap \{p_n, q_n\}| = 1$ and $\{v, w\} \cap \{p_{n-1}, q_{n-1}\} = \varnothing$.

If v is suspended in both σ_{p_0,q_0} and σ_{p_n,q_n} then Claim 3.1.3 says $v \in \{p_1, q_1\}$, but this contradicts minimality of n.

If v is suspended in σ_{p_0,q_0} but not in σ_{p_n,q_n} then $v \in \{p_n, q_n\}$ is adjacent to both p_{n-1} and q_{n-1}. Apply Claim 3.1.3 to the geodesic subsegment from $\{p_0, q_0\}$ to $\{p_{n-1}, q_{n-1}\}$. Again, this implies $v \in \{p_1, q_1\}$, contradicting minimality of n. The symmetric argument works if v is suspended in σ_{p_n,q_n} but not in σ_{p_0,q_0}.

Finally, suppose v is a suspension point in both σ_{p_0,q_0} and σ_{p_n,q_n}. If $n = 2$ then apply Claim 3.1.3 with $z = p_1$. If $n > 2$ then apply Claim 3.1.3 to the geodesic subsegment $\{p_1, q_1\}, \dots, \{p_{n-1}, q_{n-1}\}$ with $z = v$. In both cases we get that $p_{n-1} \leftrightarrow v \leftrightarrow q_{n-1}$ is a cut 2–path, and the loop γ contains v and has vertices u and w on different sides of the cut, so it must cross the cut a second time in either p_{n-1} or q_{n-1}, both of which are adjacent to v, contradicting that γ is induced.

Up to symmetry, this accounts for all possibilities, so we have constructed a FIDL–Λ. By Theorem 2.1, the commuting graph Δ of Λ defines an index 4 visible RAAG subgroup of W_Γ. It follows from Behrstock and Neumann's [1] work that Δ must be a tree, because tree RAAGs are not quasiisometric to nontree RAAGs, but we can say explicitly what Δ is from the construction: it is a tree built by starting from T and at each vertex $\{a, b\} \in T$ adding a new leaf for each discretionary edge of $\Lambda_{a,b}$. This completes the proof of (4) $\implies$ (1), modulo the claims.

Proof of Claim 3.1.1: Any square $\{a, b\} * \{c, d\}$ in Γ is a thick join, so it is contained in some maximal thick join σ, which, by (4a) is a non-square suspension. Since Γ is triangle-free either $\{a, b\}$ or $\{c, d\}$ is the pole of σ. ◊

Proof of Claim 3.1.2: If $\{a, b\} * \{c, d\}$ in Γ is a square and $\{a, b\} * \{x, y\}$ is another square then $|\{c, d, x, y\}| \geqslant 3$. The thick join $\{a, b\} * \{c, d, x, y\}$ is contained in a maximal thick join, which is a suspension, by (4a), so $\{a, b\}$ is the pole of a maximal thick join that is a non-square suspension.

Every non-leaf of $\boxslash(\Gamma)$ is a pair of vertices of Γ that are the diagonal of more than one square, so, as above, is the pole of a non-square suspension. These correspond to vertices of T, so all of the non-leaf vertices of $\boxslash(\Gamma)$ are identified with vertices of T. Furthermore, adjacent vertices of $\boxslash(\Gamma)$ span a square, so T–vertices of $\boxslash(\Gamma)$ are adjacent in $\boxslash(\Gamma)$ if and only if they are adjacent in T. Thus, we identify T with a subgraph of $\boxslash(\Gamma)$ that includes all non-leaves of $\boxslash(\Gamma)$. By Claim 3.1.1, $\boxslash(\Gamma)$ has no isolated edges, so every leaf is attached to one of the T–vertices of $\boxslash(\Gamma)$. Thus, $\boxslash(\Gamma)$ is a tree and T is a subtree containing all non-leaves.

For $\{a, b\} \in T$ there is a maximal thick join $\{a, b\} * \{c_0, c_1, c_2, \dots\}$ in Γ. Since Γ is triangle free, there are induced squares $\{a, b\} * \{c_i, c_j\}$ for all $i \neq j$. Thus, $\{a, b\}$ has at least three neighbors in $\boxslash(\Gamma)$. In particular, it is not a leaf. ◊

Proof of Claim 3.1.3: The proof is by induction on n. Since Γ is triangle-free, $n > 1$. Suppose $n = 2$. Since $\{p_0, q_0\}$ and $\{p_2, q_2\}$ are distinct vertices in T, $|\{p_0, q_0, p_2, q_2\}| \geqslant 3$, so $\{p_0, q_0, p_2, q_2, z\}$ is a set of at least 4 essential vertices that is not a square, since z is adjacent to at least three of the others. By (4b), the set must be separated by a cut, since otherwise there would be a non-(virtually $\mathbb{Z}^2$) rigid vertex in the graph of cylinders. By (4a), for all i, vertices p_i and q_i have at least three common neighbors, so they cannot be separated by a cut. Now, $\{z, p_1, q_1\} * \{p_0, q_0, p_2, q_2\} \subset \Gamma$, but thick joins in Γ are suspensions, by (4a), so $|\{z, p_1, q_1\}| = 2$. Thus, $z \in \{p_1, q_1\}$, and there is a cut $\{p_1 - q_1\}$ separating $\{p_0, q_0\} - \{p_1 - q_1\}$ from $\{p_2, q_2\} - \{p_1 - q_1\}$. This proves the claim when $n = 2$.

Now consider a geodesic of length $n > 2$ as in the statement of the claim, and suppose the claim is true for all geodesics in T of length strictly less than n. As in the base case, for all i there is no cut separating p_i from q_i, but there must be a cut separating $\{p_0, q_0, p_n, q_n, z\}$, and the cut must contain z. Thus, there is some cut containing z and separating a subset of $\{p_0, q_0\}$ from a subset of $\{p_n, q_n\}$. The geodesic $\{p_0, q_0\},\dots,\{p_n, q_n\}$ in T gives a chain of squares in Γ with adjacent squares sharing a diagonal. Such a set is 2–connected, and the only potential cut pairs are $\{p_i, q_i\}$ for $0 < i < n$. Thus:

(†) There exists $0 < i_0 < n$ such that there is a cut $\{p_{i_0} - q_{i_0}\} \supset \{p_{i_0}, q_{i_0}, z\}$ separating the set $\{p_0, q_0, p_n, q_n, z\}$.

The goal will be to split the T–geodesic into two subsegments at index i_0 and apply the induction hypothesis to both sides. We will also use the following fact:

(‡) No cut of the form $\{p_i - q_i\}$ separates the set $\{p_j, q_j \mid j < i\} - \{p_i - q_i\}$ or separates the set $\{p_j, q_j \mid j > i\} - \{p_i - q_i\}$.

To see (‡), first note that for all $j < i$, $\{p_j, q_j\} \not\subset \{z, p_i, q_i\}$. This is true because $\{p_j, q_j\} \neq \{p_i, q_i\}$, and p_j and q_j are not adjacent, but z is adjacent to both of p_i and q_i. Thus, for all $j < i$ at least one of p_j or q_j survives in $\{p_j, q_j \mid j < i\} - \{p_i - q_i\}$ and any vertices with consecutive indices are adjacent.

The combination of (†) and (‡) implies there exists $0 < i_0 < n$ such that there is a cut $\{p_{i_0} - q_{i_0}\} \supset \{p_{i_0}, q_{i_0}, z\}$ separating $\{p_i, q_i \mid i < i_0\} - \{p_{i_0} - q_{i_0}\}$ from $\{p_i, q_i \mid i > i_0\} - \{p_{i_0} - q_{i_0}\}$. Consider two cases based on whether z is in $\{p_{i_0}, q_{i_0}\}$.

Case $z \notin \{p_{i_0}, q_{i_0}\}$**:** Then the cut $\{p_{i_0} - q_{i_0}\}$ is a cut 2–path $p_{i_0} \leftrightarrow z \leftrightarrow q_{i_0}$. Consider the geodesic subsegments $\{p_0, q_0\},\dots,\{p_{i_0}, q_{i_0}\}$ and $\{p_{i_0}, q_{i_0}\},\dots,\{p_n, q_n\}$. Each is strictly shorter than the one we started with and has all vertices in the supports of its two endpoints contained in $\mathrm{lk}(z)$. Apply the induction hypothesis. Then $i_0 \in \mathbb{E}$ and $n - i_0 \in \mathbb{E}$, so $n \in \mathbb{E}$. We have the even number i_0 is not in I_z, and by induction z occurs in the support of each odd index of each geodesic subsegment, so $I_z = [0, n] \cap \mathbb{O}$. Finally, the induction hypothesis for the first subsegment says that for every $i < i_0$ there is a cut $\{p_i - q_i\}$ separating $\{p_j, q_j \mid 0 \leqslant j < i\} - \{p_i - q_i\}$ from $\{p_j, q_j \mid i < j \leqslant i_0\} - \{p_i - q_i\}$, but (‡) says $\{p_i - q_i\}$ does not separate $\{p_j, q_j \mid i < j \leqslant n\} - \{p_i - q_i\}$, so it separates $\{p_j, q_j \mid 0 \leqslant j < i\} - \{p_i - q_i\}$ from $\{p_j, q_j \mid i < j \leqslant n\} - \{p_i - q_i\}$. A symmetric argument on the second geodesic subsegment finishes the proof of (iii).

Case $z \in \{p_{i_0}, q_{i_0}\}$**:** Without loss of generality, assume $z = p_{i_0}$. First suppose $1 < i_0 < n-1$, and consider the geodesic subsegments $\{p_0, q_0\},\dots,\{p_{i_0-1}, q_{i_0-1}\}$ and $\{p_{i_0+1}, q_{i_0+1}\},\dots,\{p_n, q_n\}$. Again, each is strictly shorter than the one we started with and has all vertices in the support of its two endpoints contained in $\mathrm{lk}(z)$, so induct. We get $i_0 - 1 \in \mathbb{E}$ and $n - (i_0 - 1) \in \mathbb{E}$, so $n \in \mathbb{E}$ and $i_0 \in \mathbb{O}$. We have that I_z contains the odd index i_0, since $z = p_{i_0}$, and the induction hypothesis for the two subsegments gives that I_z contains every other odd index as well. The proof of (iii) is the same as in the previous case.

If $i_0 = 1$ or $i_0 = n - 1$ the proof is the same, except one of the geodesic subsegments is a single vertex, so the induction is only necessary on the other. $\diamondsuit$

Proof of Claim 3.1.4: By Claim 3.1.2, $\boxslash(\Gamma)$ is a tree, so there is a unique geodesic $\{a, b\} = \{p_0, q_0\}, \ldots, \{p_n, q_n\} = \{a, c\}$. The support of a single edge has 4 vertices, so $n > 1$. If $n = 2$ the claim is true. We cannot have $n = 3$, because then $\{p_1, q_1\} * \{p_2, q_2\}$ is a square in Γ contained in $\mathrm{lk}(a)$, contradicting that Γ is triangle-free. If $n > 3$, apply Claim 3.1.3 to $\{p_1, q_1\}, \ldots, \{p_{n-1}, q_{n-1}\}$ with $z = a$. $\diamondsuit$

Proof of Claim 3.1.5: Since Γ is triangle-free, the intersection of suspensions corresponding to adjacent vertices in T is exactly the square formed by the join of the two poles. By (4a), every vertex of Γ is contained in some maximal suspension. If $v \in \Gamma$ is a suspension point of a maximal suspension σ, then it is suspended in each neighbor of σ in T, so every vertex of Γ is suspended in some suspension.

If $z \in \sigma \cap \sigma'$ is contained in two maximal suspension, but is not in the pole of either, then σ and σ' do not give adjacent vertices in T. Claim 3.1.3 says z occurs in the support of each odd index vertex on the T-geodesic from σ to σ'. Thus, every vertex that is contained in distinct maximal suspensions is in the pole of some maximal suspension. $\diamondsuit$ $\square$

Example 3.2. The graph of Figure 1 is shown again with its graph of cylinders in Figure 4. The hypotheses of Theorem 3.1 are satisfied, because the central cylinder vertex has basic plane pattern: $W_{\{3,4\}*\{0,1,2,8\}}$ has an index 4 subgroup $\langle w, x, y\rangle \times \langle z\rangle \cong \mathbb{F}_3 \times \mathbb{Z}$ given by $w := s_0s_8$, $x := s_1s_8$, $y := s_2s_8$, and $z := s_3s_4$, such that restriction to the subgroup induces a pattern preserving quasiisometry to $\mathbb{F}_3 \times \mathbb{Z}$ with plane pattern generated by $\{\langle w\rangle \times \langle z\rangle, \langle x\rangle \times \langle z\rangle, \langle y\rangle \times \langle z\rangle\}$, so the projected line pattern is generated by $\{w, x, y\}$ in $\langle w, x, y\rangle$, which is basic. $\diamond$

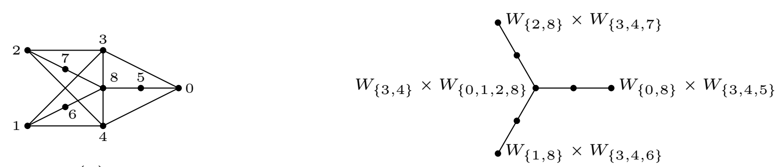

(a) Γ

(b) graph of cylinders of W_Γ.

FIGURE 4. Graph and graph of cylinders for Example 3.2.

4. Further remarks

The implication (3) $\implies$ (2) of Theorem 3.1 can be proven directly as in [4]. The description of the JSJ decomposition means these groups, like the tree RAAGs, are 'two-line tubular groups with bounded height change', which are all quasiisometric by [4, Example 5.1]. The implication (4) $\implies$ (3) is also easy. It should be possible to generalize Edletzberger's [14] visual description of the graph of cylinders in the case that Γ has triangles, in which case we should still get an equivalence between (2) and (3) and (4). The 'correct' hypothesis on Γ should be that it contains no icosahedron, so that W_Γ does not contain $\mathbb{Z}^3$. Theorem 2.1 is an application of conditions of Dani and Levcovitz [11] specialized to the triangle-free case. When there are triangles there are more conditions and they are only necessary, not sufficient, for Λ to spawn a visible RAAG subgroup. The 'sufficient' direction would be needed to extend Theorem 3.1 to handle triangles.

References

[1] J. A. Behrstock and W. D. Neumann, *Quasi-isometric classification of non-geometric 3-manifold groups*, J. Reine Angew. Math. **669** (2012), 101–120.

[2] M. Casals-Ruiz, I. Kazachkov, and A. Zakharov, *On commensurability of right-angled Artin groups I: RAAGs defined by trees of diameter 4*, Rev. Mat. Iberoam. **35** (2019), no. 2, 521–560.

[3] M. Casals-Ruiz, I. Kazachokov, and A. Zakharov, *On commensurability of right-angled Artin groups II: RAAGs defined by paths*, Math. Proc. Cambridge Philos. Soc. **170** (2021), no. 3, 559–608.

[4] C. H. Cashen, *Quasi-isometries between tubular groups*, Groups Geom. Dyn. **4** (2010), no. 3, 473–516.

[5] C. H. Cashen, *Splitting line patterns in free groups*, Algebr. Geom. Topol. **16** (2016), no. 2, 621–673.

[6] C. H. Cashen, P. Dani, A. Edletzberger, and A. Karrer, *RAAGedy right-angled Coxeter groups*, in preparation.

[7] C. H. Cashen and A. Edletzberger, *Visual right-angled Artin subgroups of two-dimensional right-angled Coxeter groups*, J. Group Theory (in press).

[8] C. H. Cashen and N. Macura, *Line patterns in free groups*, Geom. Topol. **15** (2011), no. 3, 1419–1475.

[9] C. H. Cashen and A. Martin, *Quasi-isometries between groups with two-ended splittings*, Math. Proc. Cambridge Philos. Soc. **162** (2017), no. 2, 249–291.

[10] M. Clay, *When does a right-angled Artin group split over $\mathbb{Z}$?*, Internat. J. Algebra Comput. **24** (2014), no. 6, 815–825.

[11] P. Dani and I. Levcovitz, *Right-angled Artin subgroups of right-angled Coxeter and Artin groups*, Algebr. Geom. Topol. **24** (2024), 755–802.

[12] P. Dani and A. Thomas, *Divergence in right-angled Coxeter groups*, Trans. Amer. Math. Soc. **367** (2015), no. 5, 3549–3577.

[13] P. Dani and A. Thomas, *Bowditch's JSJ tree and the quasi-isometry classification of certain Coxeter groups*, J. Topol. **10** (2017), no. 4, 1066–1106.

[14] A. Edletzberger, *Quasi-isometries for certain right-angled Coxeter groups*, Groups Geom. Dyn. **18** (2024), no. 3, 1037–1098.

[15] A. Edletzberger, *Quasi-isometries for two-dimensional right-angled Coxeter groups*, Ph.D. thesis, University of Vienna, 2024.

[16] V. Guirardel and G. Levitt, *JSJ decompositions of groups*, Astérisque (2017), no. 395.

[17] J. Huang and B. Kleiner, *Groups quasi-isometric to right-angled Artin groups*, Duke Math. J. **167** (2018), no. 3, 537–602.

[18] A. Margolis, *Quasi-isometry classification of RAAGs that split over cyclic subgroups*, Groups Geom. Dyn. **14** (2020), 1351–1417.

[19] M. Mihalik and S. Tschantz, *Visual decompositions of Coxeter groups*, Groups Geom. Dyn. **3** (2009), no. 1, 173–198.

[20] H. T. Nguyen and H. C. Tran, *On the coarse geometry of certain right-angled Coxeter groups*, Algebr. Geom. Topol. **19** (2019), no. 6, 3075–3118.

[21] J.-P. Otal, *Certaines relations d'equivalence sur l'ensemble des bouts d'un groupe libre*, J. London Math. Soc. **s2-46** (1992), no. 1, 123–139.

[22] P. Papasoglu, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2) **161** (2005), no. 2, 759–830.

Faculty of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, 0000-0002-6340-469X

Email address: `christopher.cashen@univie.ac.at`