RAAGEDY RIGHT-ANGLED COXETER GROUPS

CHRISTOPHER H. CASHEN, PALLAVI DANI, ALEXANDRA EDLETZBERGER, AND ANNETTE KARRER

Abstract. We give criteria for deciding whether or not a triangle-free simple graph is the presentation graph of a right-angled Coxeter group that is quasiisometric to some right-angled Artin group, and, if so, producing a presentation graph for such a right-angled Artin group.

We introduce two new graph modification operations, cloning and unfolding, to go along with an existing operation called link doubling. These operations change the presentation graph but not the quasiisometry type of the resulting group. We give criteria on the graph that imply it can be transformed by these operations into a graph that is recognizable as presenting a right-angled Coxeter group commensurable to a right-angled Artin group.

In the converse direction we derive coarse geometric obstructions to being quasiisometric to a right-angled Artin group, first by specializing existing results from the literature to this setting, then by developing new approaches using configurations of maximal product regions. In all cases we give sufficient graphical conditions that imply these geometric obstructions.

We implemented our criteria on a computer and applied them to an enumeration of small graphs. Our methods completely answer the motiving question when the graph has at most 10 vertices.

Contents

Date: June 20, 2025.

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

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

1. Introduction

We would like to understand the large scale geometry of right-angled Coxeter groups (RACGs). Previous work on this problem has focused on understanding a related collection of quasiisometry invariants (divergence, thickness, and hypergraph index) [29, 5, 64, 65, 66], understanding certain hyperbolic and relatively hyperbolic cases [25, 9, 30, 28, 54, 14], or certain other hyperbolic-like features [2, 48, 42, 62], and cases with nontrivial JSJ decompositions [30, 73, 40]. See also the survey [26].

The hyperbolic and relatively hyperbolic cases have exponential divergence, while linear divergence corresponds to being a product. This leaves the quadratic divergence case as the simplest 'interesting' case orthogonal to hyperbolicity, in the sense of having polynomial divergence. This class of RACGs is relatively unexplored and wide open for investigation. Furthermore, this class of groups contains, up to commensurability, all 1-ended right-angled Artin groups (RAAGs) [34], and there is extensive work on understanding the large scale geometry of RAAGs [3, 11, 12, 13, 55, 56, 58, 59, 69]. So, it is natural to ask when a given RACG with quadratic divergence is quasiisometric to some RAAG, in which case we say it is *RAAGedy*. By extension, we say that a simple graph Γ is *RAAGedy* if it is the presentation graph of a right-angled Coxeter group W_Γ that is RAAGedy.

Question 1.1. *Which RACGs are RAAGedy?*

Surprisingly little is known about this problem. There is a graph property *CFS* characterizing quadratic divergence, which is therefore necessary for a graph to be RAAGedy. This and other existing results are reviewed in Section 1.1. We develop multiple criteria to address Question 1.1. These are summarized in Section 1.2. We take a 'hands dirty' approach: large scale geometric features are interpreted in terms of the presentation graph of the Coxeter group, so that our criteria are effectively verifiable. In fact, we have computerized[1] all of the constructions, enumerated small triangle-free CFS graphs, and then checked all of them for RAAGediness. The results are shown in Figure 1, where the region labels are the list items from Section 1.2.

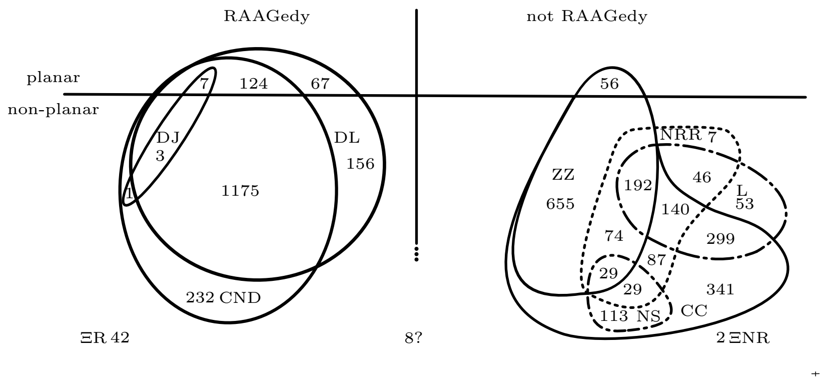

FIGURE 1. The 3938 isomorphism types of triangle-free CFS graphs with at most 11 vertices. There are exactly 8 graphs, all with 11 vertices, for which we do not know if they are RAAGedy/non-RAAGedy. They are listed in Section 8.

Preexisting work covers only the 'planar' part of the diagram, the regions labelled 'DL' and 'DJ', and a handful of isolated examples, so the figure shows that our new constructions vastly improve the state of knowledge about Question 1.1, and almost completely settle it for small triangle-free graphs. It also illustrates the inherent complexity of the problem— particularly on the non-RAAGedy side we see several overlapping but independent reasons that a graph can fail to be RAAGedy.

Practical criteria for answering Question 1.1 are an important outcome of this work, but the techniques developed to establish these criteria also support our broader aim of understanding the connections between the combinatorial properties of presentation graphs and the large scale geometry of the resulting right-angled Coxeter groups. By systematically identifying and exploiting these connections, we establish a foundation for exploring more general questions about the quasiisometric classification for RACGs.

1.1. **Background.** We summarize the current state of knowledge of Question 1.1:

There is a standard reduction of the quasiisometry problem to the case of 1–ended factors of the Grushko–Stallings–Dunwoody decomposition of a finitely presented group. In both the RAAG and RACG cases these are special subgroups, so they are again RAAGs or RACGs, respectively. Thus, for RAAGs we restrict to the case that Γ is connected with more than one vertex, and for RACGs we restrict to the case that Γ is incomplete without separating cliques.

A RAAG A_Δ has linear divergence when it is a product, which happens when Δ is a join, and it has quadratic divergence whenever Δ is not a join and A_Δ is 1–ended [3, Corollary 4.8]. One-ended RACGs have a richer divergence spectrum [29, 65, 66], and the divergence can be calculated from the structure of Γ. In particular, linear divergence corresponds to W_Γ being a product of infinite subgroups, which occurs when Γ is a thick join, and quadratic divergence occurs when Γ is not a join and has property CFS (component of full support/constructed from squares) [29, 5, 64].

[1]Code available at: https://github.com/cashenchris/RACG

Thus, for a 1–ended, irreducible RACG W_Γ, the CFS property for Γ is necessary for it to be RAAGedy.

Behrstock [2] mentions the 'Folk Question' of whether quadratic divergence for RACGs implies commensurability to a RAAG and gives a counterexample, that is not even RAAGedy, by constructing a RACG with quadratic divergence containing a stable subgroup (see Figure 2 and Section 5.2).

Nguyen and Tran [73] show that Behrstock's obstruction is not the only kind. They give a complete characterization of RAAGedy graphs that are triangle-free, have no separating cliques, and are planar. Their proof uses planarity in an essential way to port the problem to the realm of 3–manifolds.

In the RAAGedy direction there are constructions of Davis and Januszkiewicz [34] and Dani and Levcovitz [27] that give sufficient conditions on Γ for W_Γ to be commensurable to a RAAG.

1.2. **Summary of our main results.**

RAAGedy criteria: There are two existing criteria for showing W_Γ is commensurable to a RAAG:

- (DJ) The graph is a 'double' as in Definition 2.2, so W_Γ is commensurable to a RAAG by a result of Davis and Januszkiewicz.
- (DL) The graph satisfies the conditions of Dani and Levcovitz, so W_Γ has a finite-index RAAG subgroup.

There are also some graph operations that change a graph in such a way that the resulting RACG is a finite-index subgroup of the one we started with. These include *link doubling* (Section 2.1.3), as well as some others (Section 2.2.1). We call a graph a *near double* if it can be turned into a double via a sequence of link doubling operations. It turns out that there are concise criteria (Proposition 4.7) for recognizing when a graph Γ is a near double in terms of its *twin graph* $\amalg(\Gamma)$ (Definition 2.1), and when a graph is a near double it is always possible to change it into a double using only one or two link doubling steps.

We develop two more graph modification operations, *cloning* (Section 4.2) and *unfolding* (Section 4.3) that allow us to change the graph Γ without changing the quasiisometry type of W_Γ. Cloning, in particular, leads to a nice generalization of Proposition 4.7 as Theorem 4.16. We call graphs satisfying the hypotheses of Theorem 4.16 *coarse near doubles.* These include doubles and near doubles.

- (CND) If Γ is a coarse near double, then Γ is RAAGedy. A particularly simple case is that if every vertex of Γ has a twin then Γ is RAAGedy. See Theorem 4.16.

Finally, any graph that can be changed into one of our known RAAGedy types using a sequence of graph modification operations is also RAAGedy:

- (ΞR) If Γ can be changed into a coarse near double or a graph satisfying Dani and Levcovitz's conditions by a sequence of link doubles, clonings, and unfoldings, then Γ is RAAGedy. See Section 4.4.

Non-RAAGedy crietria: We also work in the opposite direction, detecting geometric obstructions to a graph being RAAGedy.

In Section 5.1 and Section 5.2 we describe such obstructions by exploring the Morse subgroups of 2–dimensional RACGs. The presence of 1–ended, infinite-index Morse subgroups is a known obstruction to being RAAGedy. Having a stable cycle in Γ is a sufficient condition for the existence of such a subgroup. We show by example that stable cycles can appear after iterated link doubling, and recount a condition for the Morse boundary of W_Γ to be totally disconnected, which rules out such behavior.

The results of Section 5.3 and Section 6 can be seen as two incarnations of a general strategy that has been used often in quasiisometric rigidity arguments (see the survey [36, Chapter 25]): find a geometrically distinguished feature and make a combinatorial object encoding interactions between subspaces so distinguished. One such geometrically distinguished feature is that of being a coarsely separating quasiline. Classes of parallel quasilines are grouped together into 'cylinders', and JSJ theory organizes these cylinders into a tree, the JSJ tree of cylinders. The second version of distinguished feature is top dimensional quasiflats. These are grouped together into maximal product regions and there is a maximal product region graph (MPRG) encoding their intersections. In both situations quasiisometries of the base groups induce isomorphisms of the corresponding combinatorial objects.

In Section 5.3 we compare JSJ decompositions of RAAGs and RACGs and derive several quasiisometry obstructions:

(NRR) If the JSJ graph of cylinders of W_Γ contains a rigid vertex whose group is not quasiisometric to a RAAG, then Γ is not RAAGedy, see Section 5.3.

(ZZ) If the JSJ graph of cylinders of W_Γ contains a virtually $\mathbb{Z}^2$ edge incident to a not virtually $\mathbb{Z}^2$ rigid vertex then Γ is not RAAGedy, by Corollary 5.14.

(CC') If the JSJ graph of cylinders of W_Γ contains a collection of cylinders forming a cycle of cuts, as in Theorem 5.16, then Γ is not RAAGedy.

In Section 6.1 we make precise the fact that the MPRG of a RAAG is a (connected) quasitree with a 1–bottleneck property. Then we give criteria on the presentation graph of a RACG that show this is not always the case for RACGs.

(NS) Γ is not RAAGedy if it is not strongly CFS, Theorem 6.1, because strongly CFS is equivalent to connectivity of the MPRG.

(L) We define a 'ladder' in the MPRG as a subgraph that is a coarse axis for the action of an infinite order element of W_Γ on its MPRG that is too wide to be compatible with the 1–bottleneck property. We give sufficient conditions for the existence of a ladder in Theorem 6.16, which therefore prevent Γ from being RAAGedy. A novel point here is that we are really discovering an invariant that has a fine dependence on the isometry type of the MPRG, not just its quasiisometry type. That is, an MPRG containing a ladder might still be a quasitree, it is just not a quasitree in precisely the same way that the MPRG of a RAAG is.

In Section 7 we construct a bespoke obstruction that is tailored to our specific problem comparing RACGS to RAAGs. To do this, we leverage a slight difference between the behavior of closest point projection to standard subcomplexes in RAAGs and RACGs: In RAAGs there is a dichotomy, the coarse intersection of two standard subcomplexes has diameter either infinte or 0. In RACGs it can be finite and nonzero, so many small diameter projections can add up to a large total.

It is a fact that maximal product regions are standard subcomplexes in RAAGs and RACGs, so quasiisometries between them coarsely preserve this particular family of standard subcomplexes. We bootstrap from this fact to inductively define (Definition 7.1) a class of *compliant subcomplexes* such that quasiisometries send them close to other compliant subcomplexes.

(CC) We deduce an obstruction to being RAAGedy if Σ_Γ contains a cycle $X_0, \dots, X_{n-1}$ of compliant subcomplexes such that consecutive X_i come close to one another and for each $i \neq 0$ the diameter of the projection of X_i to X_0 is finite but the diameter of the union of all the projections is large. Such a cycle can exist for RAAGedy Γ only if Γ satisfies some very restrictive conditions, as described in Theorem 7.5, so if we additionally rule these out then Γ is not RAAGedy.

It turns out that condition (CC) generalizes conditions (CC') and (ZZ).

Finally, we reconsider the graph modification operations:

(ΞNR) If Γ can be changed into one of the above non-RAAGedy types by a sequence of link doubles, clonings, and unfoldings, then Γ is not RAAGedy.

The practical criteria defining conditions (CC) and, especially, (L) are designed to take some link doubling into account, and what we actually computed for the purpose of Figure 1 is whether there is a sequence of at most 3 link doubles after which the presentation graph has the desired property. Even allowing this, we see that (CC), (L), and (NRR) appear to be independent.

1.3. **Comparison to Huang-Kleiner.** Huang and Kleiner [59] have results on groups quasiisometric to RAAGs, but their paper is very different. They start with a fixed RAAG A_Δ, with the additional hypothesis that it has finite outer automorphism group, and classify finitely generated groups quasiisometric to A_Δ. Our goal is start with a RACG W_Γ and decide whether or not it can be quasiisometric to any RAAG whatsoever, and, if so, produce Δ such that W_Γ and A_Δ are quasiisometric. Furthermore, many of the groups we consider do not have the rigid geometry associated with RAAGs with finite outer automorphism group. The extra flexibility makes it harder to find obstructions to being quasiisometric to some RAAG.

Acknowledgements. This paper contains material from the PhD thesis of A.E. [41], submitted in partial fulfilment of the requirements for the degree of Doktorin der Naturwissenschaft (Dr. rer. nat.) at the University of Vienna.

This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/P34214 and 10.55776/PAT7799924.

This material is based upon work supported by the National Science Foundation under Award No. 2407104 and Award No. 2407438.

Part of this work was conducted during the Erwin Schrödinger Institute Research in Teams program “Rigidity in Coxeter groups”, July 2023, C.C. and P.D. members.

We also thank Jingyin Huang and Sangrok Oh.

2. Preliminaries

In this section we fix terminology and notation, and recall some existing technology that can be used in quasiisometry results, which we will apply in subsequent sections in the special case of 2–dimensional RACGs and RAAGs. We also adapt some ideas that are present in the literature to forms that will be useful to us later on. This makes the paper more self-contained, and serves as a warm-up. We do not regard this material as new.

2.1. **Graphs.** A *simple graph*, also known as a simplicial graph, is a 1–dimensional simplicial complex, or, equivalently, it is a collection of vertices and edges such that every edge connects distinct vertices and two edges have at most one vertex in common. **All graphs in this paper are simple.** On the rare occasion that we want to entertain the possibility of self-loops or multiple edges between a pair of vertices we will call that structure a *multigraph*. A *full* or *induced* subgraph Γ' of a graph Γ is one that has an edge between two vertices if and only if Γ does. A *spanning* subgraph is one that contains every vertex of Γ. The graph is *complete* if it contains an edge between every pair of distinct vertices, and *incomplete* otherwise. A (possibly empty) set of vertices form a *clique* if they induce a complete subgraph, and an *anticlique* if there is no edge between any pair of vertices in the set. The *join* of nonempty Γ and Γ', written $\Gamma * \Gamma'$, is the graph made from the disjoint union of Γ and Γ' by adding an edge between every vertex of Γ and every vertex of Γ'. A graph is a *cone* if it can be written as the join of a subgraph with a singleton. The singleton is then called a *cone vertex*. We can also *cone-off a subgraph* Γ' of Γ

by adding to Γ a new vertex c connected to every vertex of Γ'. This is written as $\Gamma *_{\Gamma'} c$.

A graph is a *suspension* if it can be written as a join of a subgraph with a 2 point anticlique. The anticlique is called the *pole* or the *suspension points* for this particular realization of the graph as a suspension.

If Γ' is a subgraph of Γ, an *n–chord* is a path of length n in Γ whose endpoints are vertices of Γ' whose path distance in Γ' is strictly greater than n. Saying a subgraph has no 1–chords is equivalent to it being an induced subgraph.

If two graphs Γ and Γ' are isomorphic, we write $\Gamma \cong \Gamma'$.

When treating a connected graph as a metric object, we always use the path metric induced by metrizing edges as unit intervals.

2.1.1. *Twins and satellites.* A vertex $v \neq w$ is a *satellite* of w if $\mathrm{lk}(v) \subset \mathrm{lk}(w)$. In the context of RAAGs there is an existing notion of a relation between vertices called the *domination relation (of Servatius)*, where w dominates v if $\mathrm{lk}(v) \subset \mathrm{st}(w)$. In a triangle-free graph w dominates v if v is a satellite of w or if v is a leaf at w.

Vertices $v \neq w$ are *twins* if $\mathrm{lk}(v) = \mathrm{lk}(w)$.[2]

A *module* M in a graph Γ is a set of vertices such that for all $v \in \Gamma - M$, either every vertex in M is adjacent to v or no vertex in M is adjacent to v. A module is *non-trivial* if it is a proper subset containing more than one vertex.

There is a general fact that any partition of Γ into modules gives a quotient graph with one vertex for each module. For two modules M_0 and M_1 in the partition, either $M_0 * M_1 \subset \Gamma$ or there are no edges of Γ between M_0 and M_1. The quotient graph contains an edge between M_0 and M_1 when $M_0 * M_1 \subset \Gamma$.

Given a vertex v of Γ, the set consisting of v and all its twins forms a module; call this a *twin module*. The set of twin modules partitions Γ into disjoint anticliques.

Definition 2.1. The *twin graph* $\amalg(\Gamma)$ of Γ is the quotient graph coming from the decomposition of Γ into twin modules.

2.1.2. *RACGs and RAAGs.* Given a finite graph Γ the *right-angled Artin group with presentation graph* Γ is the group A_Γ whose generators are the vertices of Γ, and whose defining relations are commutation relations corresponding to pairs of vertices that are connected by an edge in Γ. The *right-angled Coxeter group with presentation graph* Γ is the group W_Γ that includes the same generators and relations as A_Γ and additional defining relations saying that each generator has order 2. In both cases, an induced subgraph Γ' of Γ determines a *special subgroup* $A_{\Gamma'} < A_\Gamma$ or $W_{\Gamma'} < W_\Gamma$. Tits' solution to the word problem for Coxeter groups [32, Sec. 3.4] implies that for $w \in W_{\Gamma'} < W_\Gamma$, every minimal length word representing w in W_Γ uses generators only from Γ'. The same is true in RAAGs [82].

We will call a RAAG or RACG *irreducible* if the given presentation graph is not a join. This is equivalent to saying that the group does not split as a direct product of nontrivial subgroups [70, 77].

Both families of groups admit cocompact actions on CAT(0) cube complexes, which have the additional property that the 1–skeleton is the Cayley graph corresponding to the generating set defined by the presentation graph. For a RAAG defined by Δ, this cube complex is the universal cover of the Salvetti complex, denoted $\Sigma(A_\Delta)$ or Σ_Δ. For a RACG defined by Γ, the cube complex is the Davis complex, denoted $\Sigma(W_\Gamma)$ or Σ_Γ. It will be clear from context whether Σ_Γ refers to

[2]In the literature the definition given here is sometimes called 'open twins' or 'false twins' and there is another definition in which v and w are twins ('closed twins' or 'true twins') if $\{v\} \cup \mathrm{lk}(v) = \{w\} \cup \mathrm{lk}(w)$, which requires twins to be adjacent. In our setting of connected, incomplete, triangle-free graphs, if v and w are adjacent then $\mathrm{lk}(v) - \{w\} \neq \mathrm{lk}(w) - \{v\}$, so only the definition with v and w nonadjacent is meaningful.

a Salvetti complex or a Davis complex. We will usually use Γ for the presentation graph of a RACG and Δ for the presentation graph of a RAAG. A consequence of word problem result above is that for $\Gamma' < \Gamma$ there is a convex subcomplex $\Sigma_{\Gamma'} \subset \Sigma_\Gamma$, and, again, the analogue is also true in RAAGs.

For both RAAGs and RACGs, there is a bijection between join subgraphs of the presentation graph and special subgroups that split as direct products. We will be interested in the case where both factors of the direct product are infinite groups. For RAAGs this is always true, but for RACGs we need an additional condition that both factors of the join subgraph are incomplete. In the context of RACGs we use the term *thick join* to mean $A * B \subset \Gamma$ with A and B incomplete.

For further background on RAAGs and RACGs, see [18, 32, 26, 33].

2.1.3. *Doubles.*

Definition 2.2. Let Γ be a graph.

- The *double* $\mathfrak{D}(\Gamma)$ of Γ is the graph with vertex set $\Gamma \times \{0, 1\}$ such that for every edge $v \bullet\!\!-\!\!\bullet w$ in Γ there are four edges $(v, \epsilon) \bullet\!\!-\!\!\bullet (w, \delta)$ for $\epsilon, \delta \in \{0, 1\}$.
- The *link double of* Γ *at a vertex* v is the graph:

  $$\mathfrak{D}^\circ_v(\Gamma) := (\Gamma \times \{0, 1\}/\{(w, 0) \sim (w, 1) \mid w \in \mathrm{lk}(v)\}) - (\{v\} \times \{0, 1\})$$

- The *star double of* Γ *at a vertex* v is the graph:

  $$\mathfrak{D}^*_v(\Gamma) := \Gamma \times \{0, 1\}/\{(w, 0) \sim (w, 1) \mid w \in \mathrm{st}(v)\}$$

These doubles have the following, well known, algebraic significance:

Lemma 2.3. *Let* Γ *and* Δ *be graphs.*

- *There is an exact sequence:*

  $$1 \to W_{\mathfrak{D}^\circ_v(\Gamma)} \to W_\Gamma \to \mathbb{Z}_2 \to 1$$

  The first map sends $(w, 0) \mapsto w$ *and* $(w, 1) \mapsto v^{-1}wv$ *for all* $w \in \Gamma - \{v\}$. *The second map sends* v *to 1 and all other generators to 0.*
- *There is an exact sequence:*

  $$1 \to A_{\mathfrak{D}^*_v(\Delta)} \to A_\Delta \to \mathbb{Z}_2 \to 1$$

  The maps are as in the previous case, with the addition that $(v, 0) \mapsto v^2$.

Theorem 2.4 (Davis and Januszkiewicz [34])**.** A_Γ *is commensurable to* $W_{\mathfrak{D}(\Gamma)}$.

2.2. **More finite-index subgroup constructions.** Lemma 2.3 gave us one way to pass to a finite-index RACG subgroup of a RACG, and Theorem 2.4 gave us one way to find a commensurable RAAG. In this section we give an additional construction of each of these types.

2.2.1. *Some other finite-index RACG subgroups.* In the RAAG case of Lemma 2.3 the '2' is not special; one can take the kernel of the map $A_\Delta \to \mathbb{Z}_n$ obtained by killing the n–th power of v and all of the other generators, and express it as the RAAG on the graph obtained by taking n–copies of Δ and identifying them along the star of v. Obviously, we cannot make such a short exact sequence with a RACG, but actually we can do a similar graph construction to get an index-n subgroup if Γ contains a pair of twins, as follows:

Proposition 2.5. *Fix* $n > 2$ *and a pair of twins* v *and* w *in* Γ. *Let* $\{a_0, \ldots, a_{\ell-1}\}$ *be the vertices of* $\mathrm{lk}(v) = \mathrm{lk}(w)$. *Consider*

$$\Gamma' := \Gamma \times \{0, \ldots, n-1\}/\{(x, i) \sim (x, 0) \mid x \in \{v, w\} * \{a_0, \ldots, a_{\ell-1}\}, 1 \leqslant i < n\}.$$

Then there is an injective homomorphism $\iota : W_{\Gamma'} \to W_\Gamma$ *such that* $\iota(W_{\Gamma'})$ *is an index-*n *subgroup of* W_Γ.

See Section 2.7 for a refresher on walls in cube complexes.

Proof. Let $\{b_0, \dots, b_{m-1}\}$ be the vertices of Γ not in $\{v, w\} * \{a_0, \dots, a_{\ell-1}\}$. Define the homomorphism ι as follows:

$$\begin{aligned}
(v,0) &\mapsto v \\
(w,0) &\mapsto \begin{cases} (wv)^{(n-1)/2} w (vw)^{(n-1)/2} & \text{if } n \text{ is odd} \\ (wv)^{(n-2)/2} wvw (vw)^{(n-2)/2} & \text{if } n \text{ is even} \end{cases} \\
(a_i,0) &\mapsto a_i \\
(b_j,k) &\mapsto \begin{cases} (wv)^{k/2} b_j (vw)^{k/2} & \text{if } k \text{ is even} \\ (wv)^{(k-1)/2} w b_j w (vw)^{(k-1)/2} & \text{if } k \text{ is odd} \end{cases}
\end{aligned}$$

To see that the image is really an index-n subgroup, consider the vw–bicolored geodesic $\Sigma_{\{v,w\}}$ through the identity vertex 1 of Σ_Γ. Let $\mathcal{H}_v$ be the wall dual to the $1 \bullet\!\!-\!\!\bullet v$ edge, and similarly for the other generators. Consider the subsegment X of $\Sigma_{\{v,w\}}$ containing vertices 1, w, wv, wvw,…, x, where x is the alternating product of w and v of length $n-1$, starting with w. This is a finite convex subcomplex of Σ_Γ containing n vertices. It is a consequence of a theorem of Dyer [39] and Deodhar [35] (in not necessarily right-angled Coxeter groups) that the reflection subgroup of W_Γ generated by reflections through walls closest to X that do not cross X is a right-angled Coxeter group with those reflections as fundamental generators. We work through edges incident to X and see that the walls dual to such edges correspond to the elements described above, and commuting relations between them correspond to the edges described for Γ'. At vertex 1 we have edges corresponding to each of the generators except w, since the edge $1 \bullet\!\!-\!\!\bullet w$ is contained in X. The reflections in these walls correspond to the $(y, 0)$ for $y \neq w$. For each i, a_i commutes with both v and w, so the single wall $\mathcal{H}_{a_i}$ runs through every a_i–colored edge along $\Sigma_{\{v,w\}}$. Thus, we only need one reflection, given by the image of $(a_i, 0)$. For each j, since v and w are twins and b_j does not commute with both of v and w, it commutes with neither of them, so every b_j–colored edge incident to $\Sigma_{\{v,w\}}$ is dual to a different wall, and we get n different walls for each j. The corresponding reflections are b_j, wb_jw, wvb_jvw,…, which are the images of $(b_j, 0)$, $(b_j, 1)$, $(b_j, 2)$,… under the homomorphism ι. Finally, at vertex x there is one additional edge outside of X that continues along $\Sigma_{\{v,w\}}$, colored either v or w according to whether n is even or odd. The reflection in the wall dual to this edge is the image of $(w, 0)$ under ι.

It remains to argue that Γ' is actually the presentation graph for the subgroup. It is clear that the edges that are present belong there, as the corresponding elements of W_Γ commute. In the other direction, the only case of any interest is to show that the commutator of the images of (b_i, j) and (b_k, ℓ) is a nontrivial element of W_Γ when $j \neq \ell$. However, such a commutator has incomplete cancellation of the alternating v, w conjugating words if $j \neq \ell$, so it is a word of the form $z_0 b_i z_1 b_k z_2 b_i z_3 b_k z_4$, where each z_p is a nontrivial alternating word in v and w. Since neither b_i nor b_k commute with either of v or w, and v and w do not commute with each other, there are no elementary operations (see [32, Section 3.4]) that rearrange or shorten that word, so it represents a nontrivial element of W_Γ. □

Using the theorem of Deodhar and Dyer cited above, one can find other finite-index RACG subgroups of W_Γ by considering various finite convex subcomplexes of Σ_Γ. The challenge, from the point of view of this paper, would then be to understand the presentation graph of the subgroup in terms of operations on Γ.

We will mention one more case in which we understand the resulting graph: suppose v and w are adjacent in Γ, and take X to be the square $\Sigma_{\{v,w\}} \subset \Sigma_\Gamma$. We claim the presentation graph of the corresponding index-4 subgroup is $\mathfrak{D}^\circ_{(w,0)} \circ \mathfrak{D}^\circ_v(\Gamma) \cong$

$\mathfrak{D}^\circ_{(v,0)} \circ \mathfrak{D}^\circ_w(\Gamma)$. Explicitly, vertices of those two graphs are of the form (γ, i, j) for $\gamma \in \Gamma$ and $i, j \in \{0,1\}$, and the isomorphism exchanges the second and third coordinates. When v and w are not adjacent it can happen that the graphs $\mathfrak{D}^\circ_{(w,0)} \circ \mathfrak{D}^\circ_v(\Gamma)$ and $\mathfrak{D}^\circ_{(v,0)} \circ \mathfrak{D}^\circ_w(\Gamma)$ are not isomorphic. More generally, in higher dimensional RACGs one can also double over larger cliques, cf. [14, Section 5].

2.2.2. *Visible RAAG subgroups.* Let Γ be a finite graph, and let Γ^c be its complement: the graph with the same vertex set as Γ that has an edge whenever Γ does not. Let Λ be a subgraph of Γ^c, and let Δ be its commuting graph: an edge $a \bullet\!\!-\!\!\bullet b$ of Λ is a vertex $\{a,b\}$ of Δ and vertices $\{a,b\}$ and $\{c,d\}$ of Δ span an edge if they commute, which is to say, $\{a,b\} * \{c,d\}$ is an induced square in Γ. Obtain a homomorphism $A_\Delta \to W_\Gamma$ by sending a vertex $\{a,b\}$ of Δ to the element ab of W_Γ. It is easy to see that this homomorphism need not be injective. Dani and Levcovitz [27] give criteria on Λ that imply that the corresponding homomorphism $A_\Delta \to W_\Gamma$ is injective and has finite-index image. We refer to a graph Λ satisfying their conditions as a *finite-index Dani-Levcovitz* Λ (FIDL–Λ). The corresponding subgroup A_Δ is a *visible RAAG subgroup*. It is 'visible' (or 'visual') in the sense that we see generators of A_Δ as (products of the entries of) diagonals of squares of Γ, and we see relations of A_Δ as squares in Γ.

The simplest example is to take Γ to be a square and $\Lambda = \Gamma^c$, so that Δ is a single edge whose vertices are the two diagonals of Γ. Then $A_\Delta \cong \mathbb{Z}^2$ includes into $W_\Gamma \cong D_\infty \times D_\infty$ as an index 4 subgroup.

2.3. **Coarse geometry.** Let X and Y be metric spaces.

For $Z \subset X$, let $\mathcal{N}_r(Z) := \{x \in X \mid d(x,Z) < r\}$, and let $\bar{\mathcal{N}}_r(Z) := \{x \in X \mid d(x,Z) \leqslant r\}$. Two subsets Z, Z' of X are *coarsely equivalent* if they are at finite Hausdorff distance, where Hausdorff distance is defined as:

$$d_{Haus}(Z,Z') := \inf\{r \mid Z \subset \bar{\mathcal{N}}_r(Z'),\ Z' \subset \bar{\mathcal{N}}_r(Z)\}$$

A subset $Z \subset X$ is *coarsely dense* if it is coarsely equivalent to X.

A *coarse map* $\phi\colon X \to Y$ is a map from X to uniformly bounded diameter subsets of Y. Two coarse maps $\phi, \psi\colon X \to Y$ are *coarsely equivalent* if there exists C such that $\operatorname{diam}(\phi(x) \cup \psi(x)) \leqslant C$ for all $x \in X$. A coarse map $\phi\colon X \to Y$ is (L,A)*–coarse Lipschitz* if $\operatorname{diam}(\phi(x) \cup \phi(x')) \leqslant Ld(x,x') + A$ for all x, $x' \in X$. It is (L,A)*–coarse biLipschitz*, or an (L,A)*–quasiisometric embedding*, if, in addition, $(1/L)d(x,x') - A \leqslant \operatorname{diam}(\phi(x) \cup \phi(x'))$ for all x, $x' \in X$. It is an (L,A)*–quasiisometry* if it is (L,A)–biLipschitz with A–coarsely dense image.

When we do coarse geometry and X and Y are graphs we take the point of view that 'the spaces' are discrete metric spaces given by the vertex sets of X and Y. The edges serve to help visualize the discrete metrics. Thus, a coarse map $\phi\colon X \to Y$ only need be defined on vertices of X. Of course, it is possible to adjust ϕ to be an actual map and then extend it to interior points of edges of the geometric realization of X, but such choices are non-canonical, and all reasonable choices yield coarsely equivalent maps.

The following says that coarsely inverse coarse Lipschitz maps are actually quasiisometries. It can be proved by using the inverse map to establish the quasiisometry lower bounds.

Lemma 2.6. *If $\phi\colon X \to Y$ and $\bar{\phi}\colon Y \to X$ are coarse Lipschitz and $\bar{\phi} \circ \phi$ is coarsely equivalent to the identity map on X and $\phi \circ \bar{\phi}$ is coarsely equivalent to the identity map on Y, then ϕ and $\bar{\phi}$ are inverse quasiisometries.*

A *geodesic* is an isometric embedding of an interval, and a *quasigeodesic* is a quasiisometric embedding of an interval.

Given subsets Y and Z of X, if for all sufficiently large r the sets $\bar{\mathcal{N}}_r(Y)\cap\bar{\mathcal{N}}_r(Z)$ are coarsely equivalent to one another, then we call that coarse equivalence class the *coarse intersection* of Y and Z, and denote it $Y \overset{c}{\cap} Z$.

Lemma 2.7. *If $\phi\colon X \to X'$ is a quasiisometry and $Y, Z < X$ such that $Y \overset{c}{\cap} Z$ and $\phi(Y) \overset{c}{\cap} \phi(Z)$ exist, then $\phi(Y \overset{c}{\cap} Z)$ is coarsely equivalent to $\phi(Y) \overset{c}{\cap} \phi(Z)$.*

Proof. Let ϕ be an (L, A)–quasiisometry. For any $r \geqslant LA$:

$$\bar{\mathcal{N}}_{\frac{r}{L}-A}(\phi(Y))\cap\bar{\mathcal{N}}_{\frac{r}{L}-A}(\phi(Z)) \subset \phi(\bar{\mathcal{N}}_r(Y)\cap\bar{\mathcal{N}}_r(Z)) \subset \bar{\mathcal{N}}_{Lr+A}(\phi(Y))\cap\bar{\mathcal{N}}_{Lr+A}(\phi(Z))$$

□

Lemma 2.8 (Coarse intersections of cosets exist)**.** *If X is the Cayley graph of a finitely generated group, a and b are elements of the group, and G and H are subgroups, then $aG \overset{c}{\cap} bH$ exists and is represented by $aGa^{-1} \cap bHb^{-1}$.*

Proof. This is an immediate corollary of [72, Lemma 2.2], which says the intersection of subgroups represents their coarse intersection. Then observe that aG is coarsely equivalent to aGa^{-1}, and bH is coarsely equivalent to bHb^{-1}. □

A tree is connected graph with no cycles. It is *bushy* if the set of vertices with more than two unbounded complementary components is coarsely dense. A graph is a *quasitree* if it is quasiisometric to a tree.

In a tree, if γ is a geodesic segment then every path between its endpoints touches every point of γ. There are various ways to coarsen this property to characterize quasitrees in terms of 'bottleneck constants' [67, Theorem 4.6], [68, Lemma 2.16]. The formulation we will use is that X is a quasitree if and only if there exist A, B, C such that for every geodesic γ of length at least A there is a vertex v within distance B of the midpoint of γ such that the C–ball about v separates the endpoints of the geodesic in X. We refer to C as 'the' bottleneck constant.

The following example shows why some additional constraint on the location of the separating ball is necessary; it is not sufficient to say a space is a quasitree if every sufficiently long geodesic has its endpoints separated by a C–ball.

Example 2.9. Let Γ be any connected graph that is not a quasitree— the $\mathbb{Z}^2$ grid, for example. Its double $\mathfrak{D}(\Gamma)$ is quasiisometric to Γ, so it is also not a quasitree. However, points at distance at least 3 in $\mathfrak{D}(\Gamma)$ are separated by balls of radius 1: For any x, y with $d(x, y) \geqslant 3$ let $z \neq x$ be the vertex that has the same image as x under the projection of $\mathfrak{D}(\Gamma)$ onto Γ. Then $d(x, z) = 2$ and $\mathrm{lk}(x) = \mathrm{lk}(z)$, so $x, y \notin \bar{\mathcal{N}}_1(z)$ but every path from x to y goes through $\bar{\mathcal{N}}_1(z)$. ◇

2.4. **Group decompositions and model spaces.** A *graph of groups* consists of a finite connected multigraph, *local groups* associated to each vertex and edge, and injections of each edge group into each of its two incident vertex groups. A graph of groups has an associated *fundamental group* obtained by amalgamating the vertex groups along the edge groups. Conversely, a graph of groups is said to give a *decomposition* of G if G is isomorphic to its fundamental group.

An equivalent formulation comes from letting a group G act cocompactly and without edge inversions on a tree T. One gets a graph of groups decomposition of G by considering the quotient multigraph of $G \curvearrowright T$, and setting the local groups to be stabilizers. In this setting, a subgroup is said to be *elliptic* if it fixes a vertex of T. In the other direction, given a graph of groups with fundamental group G there is an associated *Bass-Serre tree* T with a G action such that the graph of groups coming from the action of G on T recovers the original graph of groups.

Associated to a graph of groups decomposition of G with Bass-Serre tree T we can make a geometric model for G as a *tree of spaces* X over T: For each vertex and edge of T choose a metric space quasiisometric to the stabilizer group. The space X is built by taking each edge space X_e crossed with a unit interval $[0,1]$, and gluing $X_e \times \{0\}$ to the vertex space X_v of the initial vertex $v = \iota(e)$ of e compatibly with the group inclusion, and similarly for the terminal vertex $\tau(e)$. See [17] for details. The resulting metric space X is quasiisometric to the group G.

The *Grushko-Stallings-Dunwoody (GSD) decomposition* of a finitely presented group G is a graph of groups decomposition of G whose vertex groups are finite or 1–ended, and whose edge groups are finite. It turns out that this is well-defined, and the quasiisometry type of G is determined by the quasiisometry types of the 1–ended vertex groups in its GSD decomposition [76]. Thus, for quasiisometry questions we restrict to 1–ended groups.

The next higher complexity for splittings of 1–ended, finitely presented groups are the *Jaco-Shalen-Johannson (JSJ) decompositions*. In this paper we always use "JSJ decomposition" to mean "JSJ decomposition over 2–ended subgroups". This is a graph of groups decomposition of G that, in a sense, captures all possible splittings over 2–ended subgroups. The edge groups are all 2–ended. Unlike with finite edge groups, there can be incompatible splittings over 2–ended subgroups, and these are combined into *hanging vertices* in the graph of groups. The other vertices are either 2–ended or they are *rigid*, in the sense that they do not split further over 2–ended subgroups *relative to the existing incident 2–ended edge groups*, meaning a splitting in which all of these subgroups are elliptic.

The JSJ decomposition of a group G is not a well-defined graph of groups; instead, it is a deformation space of graphs of groups (see [50]). However, the existence of a non-trivial JSJ decomposition is a quasiisometry invariant [75]. One can make a canonical graph of groups decomposition from (any) JSJ decomposition of G as follows. Let T be the Bass-Serre tree of some nontrivial JSJ decomposition of G. A *cylinder* is a maximal collection of edges of T whose stabilizers are commensurable in G. One makes a *JSJ tree of cylinders* by collapsing all of the cylinders to single vertices. These are known as *cylinder vertices*. Vertices of T that are contained in a single cylinder are absorbed into that cylinder vertex; vertices that are not, which are necessarily rigid or hanging, survive in the JSJ tree of cylinders. The quotient of the JSJ tree of cylinders by the G action gives a graph of groups decomposition of G known as the *JSJ graph of cylinders* of G. It is a bipartite multigraph where one part consists of cylinder vertices, and the other part is rigid and hanging vertices that are not contained in a single cylinder. It turns out that the JSJ tree of cylinders and the JSJ graph of cylinders are well-defined, independent of the starting JSJ decomposition of G [50] (but the edges are not necessarily 2–ended anymore). Furthermore, it follows from [75] (see, eg, [17]) that a quasiisometry $\phi\colon G \to G'$ between two groups induces an isomorphism $\phi_*\colon T \to T'$ between their JSJ trees of cylinders, and if X and X' are trees of spaces of G and G' over T and T', respectively, then the restriction of ϕ to each vertex space X_v is uniformly coarsely equivalent to a quasiisometry $\phi_v\colon X_v \to X'_{\phi_*(v)}$ such that for each edge e of T incident to v, the set $\phi_v(X_e)$ is uniformly coarsely equivalent to $X'_{\phi_*(e)}$.

In [17] the construction is carried out in the other direction: If X and X' are trees of spaces of G and G' over some trees T and T' (not necessarily related to JSJ decompositions) and if there is an isomorphism $\chi\colon T \to T'$ and a collection of uniform quasiisometries $\phi_v\colon X_v \to X'_{\chi(v)}$ that uniformly coarsely agree on common edge spaces, then X and X' are quasiisometric by a map ϕ such that $\phi|_{X_v}$ coarsely agrees with ϕ_v for each vertex $v \in T$. The collection of quasiisometries is called a *tree of quasiisometries over* χ.

Proposition 2.10. *Suppose that G and G' are two groups that are defined by graphs of groups on the same underlying multigraph Γ, which is a tree. Suppose for* *each vertex v of Γ there is a quasiisometry ψ_v from the local group G_v of G to the local group G'_v of G'. If the collection of quasiisometries $\{\psi_v \mid v \in \Gamma\}$ satisfies the following conditions, then G and G' are quasiisometric.*

(1) For each edge e of Γ, $\psi_{\iota(e)}$ and $\psi_{\tau(e)}$ coarsely agree on G_e.
(2) For each edge e of Γ, $\psi_{\iota(e)}$ induces a bijection $\psi^*_{\iota(e)}$ between $G_{\iota(e)}/G_e$ and $G'_{\iota(e)}/G'_e$ by taking each coset of G_e in $G_{\iota(e)}$ to within uniformly bounded Hausdorff distance of a unique coset of G'_e in $G'_{\iota(e)}$, and vice versa.
(3) For each edge e of Γ and each coset gG_e of G_e in $G_{\iota(e)}$ there exists $h \in gG_e$ and $h' \in \psi^*_{\iota(e)}(gG_e)$ such that $h'\psi_{\iota(e)}h^{-1}$ and $\psi_{\iota(e)}$ coarsely agree on gG_e, and such that the coarseness constants are bounded uniformly over all cosets.

Proof. We build a tree of quasiisometries as described above. Let T be the Bass-Serre tree corresponding to the given splitting of G, and let X be a tree of spaces for G over T. Since Γ is a tree it admits a lift $\widetilde{\Gamma}$ to T, such that if $v \overset{e}{\bullet\!\!-\!\!\bullet} w$ is an edge of Γ then $\tilde{v} \overset{\tilde{e}}{\bullet\!\!-\!\!\bullet} \tilde{w}$ is an edge of $\widetilde{\Gamma}$. We use all the same notation with $'$s for the corresponding concepts with respect to G'.

For each vertex $v \in \Gamma$ define $\chi(\tilde{v}) := \tilde{v}'$ and define $\phi_{\tilde{v}} \colon X_{\tilde{v}} \to X'_{\tilde{v}'}$ to be $\psi_v \colon G_v \to G'_v$, where we implicitly identify $X_{\tilde{v}}$ with G_v and $X'_{\tilde{v}'}$ with G'_v. Condition (1) implies that for each edge $e \in \Gamma$, $\phi_{\iota(\tilde{e})}$ and $\phi_{\tau(\tilde{e})}$ coarsely agree on $X_{\tilde{e}}$, which is the intersection of their domains.

Now we inductively expand the domain of χ, specifying local quasiisometries between the corresponding vertex spaces as we go, in such a way that the maps on adjacent vertex spaces coarsely agree on their intersection. Once this is done, [17, Corollary 2.16] says the local quasiisometries patch together to give a quasiisometry.

As we will go, it will be useful for bookkeeping to define $\kappa \colon T \to G$ and $\kappa' \colon T' \to G'$ such that:

- $\kappa(g\tilde{v}) \in gG_v$
- $\kappa'(\chi(g\tilde{v})) \in g'G'_v$, where $g'\tilde{v}' = \chi(g\tilde{v})$
- $\phi_{g\tilde{v}} = \kappa'(\chi(g\tilde{v}))\psi_v\kappa(g\tilde{v})^{-1}$

For the base cases of vertices in $\widetilde{\Gamma}$ and $\widetilde{\Gamma}'$ define $\kappa(\tilde{v}) = 1$ and $\kappa'(\tilde{v}') = 1$. For the initial inductive step, consider an edge e of Γ with $\iota(e) = v$. The edges of T incident to $\tilde{v}$ and covering e are of the form $g_i\tilde{e}$, where $\{g_i\}$ is a set of coset representatives of G_v/G_e. Similarly, the edges of T' incident to $\tilde{v}'$ and covering e are translates of $\tilde{e}'$ by coset representatives of G'_v/G'_e. For each i let $g'_iG'_e = \psi^*_{\iota(e)}(g_iG_e)$ and define $\chi(g_i\tilde{e}) := g'_i\tilde{e}'$. By Condition (2), this defines a bijection between edges of T incident to $\tilde{v}$ covering e and edges of T' incident to $\tilde{v}' = \chi(\tilde{v})$ covering e. Now choose h_i and h'_i as in Condition (3) with respect to g_iG_e, and define:

$$\phi_{g_i\tau(\tilde{e})} := h'_i\psi_{\tau(e)}h_i^{-1}$$

We must check that $\phi_{\tau(g_i\tilde{e})}$ and $\phi_{\iota(g_i\tilde{e})}$ coarsely agree on $X_{g_i\tilde{e}} = g_iG_e$. For $k \in G_e$:

$$\begin{aligned}
\phi_{\tau(g_i\tilde{e})}(g_ik) &= \phi_{g_i\tau(\tilde{e})}(g_ik) \\
&= h'_i\psi_{\tau(e)}h_i^{-1}(g_ik) \\
&= h'_i\psi_{\tau(e)}(h_i^{-1}g_ik) \\
&\approx h'_i\psi_{\iota(e)}(h_i^{-1}g_ik) && \text{by Condition (1), since } h_i^{-1}g_ik \in G_e \\
&= h'_i\psi_{\iota(e)}h_i^{-1}(g_ik) \\
&\approx \psi_{\iota(e)}(g_ik) && \text{by Condition (3), since } g_ik \in g_iG_e \\
&= \phi_{\iota(\tilde{e})}(g_ik) = \phi_{\iota(g_i\tilde{e})}(g_ik)
\end{aligned}$$

Define $\kappa(g_i\tau(\tilde{e})) := h_i$ and $\kappa'(\chi(g_i\tau(\tilde{e}))) := h'_i$.

For the general inductive step, suppose χ is defined on a subtree of T containing $\widetilde{\Gamma}$, and for each vertex $g\tilde{v}$ in the subtree we have defined a quasiisometry $\phi_{g\tilde{v}}\colon X_{g\tilde{v}} \to X'_{\chi(g\tilde{v})}$ and κ and κ' as above. Suppose $g\tilde{v} \notin \widetilde{\Gamma}$ is a leaf of the subtree and $g\tilde{v} \overset{g\tilde{e}}{\text{—}} g\tilde{w}$ is an edge of T such that $g\tilde{w}$ is farther from $\widetilde{\Gamma}$ than $g\tilde{v}$. Pull $g\tilde{v}$ back to $\tilde{v}$ by applying $\kappa(g\tilde{v})^{-1}$, which takes $g\tilde{e}$ to some edge $\kappa(g\tilde{v})^{-1}g\tilde{e}$ incident to $\tilde{v}$. Then ψ_v^* identifies this with some edge incident to $\tilde{v}'$ in the same orbit as $\tilde{e}'$. Push this edge forward by $\kappa'(\chi(g\tilde{v}))$ to get an outgoing edge of T' at $\chi(g\tilde{v})$. Define this edge to be $\chi(g\tilde{e})$, and define $\chi(g\tilde{w})$ to be its other endpoint. Similarly, to define $\phi_{g\tilde{w}}$, pull back via $\kappa(g\tilde{v})^{-1}$, do the construction of the previous paragraph, and then push forward the result by $\kappa'(\chi(g\tilde{v}))$.

The key points are that all quasiisometries between vertex spaces are one of the base quasiisometries ψ_v, pre- and post- composed by multiplication in the groups, which are isometries, so the quasiisometry constants are bounded by the maxima of the constants for the base maps. Similarly, the coarse agreement between maps of neighboring vertex spaces on their common intersection is, up to isometry, the same as that coarse agreement described in Condition (3), which is assumed to be uniform for each vertex/edge pair in Γ, of which there are finitely many. □

2.4.1. *JSJ decompositions of RAAGs.* A graph is *biconnected* if it connected with no cut vertex. By this definition, a biconnected graph either has at most two vertices, or every vertex has valence at least 2.

It is not hard to see that a RAAG on at least two generators is 1–ended if and only if its presentation graph is connected. According to Clay [22] and Margolis [69, Proposition 3.6], the JSJ graph of cylinders of a RAAG A_Δ can be described "visually" in Δ: cylinders are stars of cut vertices, rigid vertices are maximal biconnected subgraphs that either contain two cut vertices or are not contained in any cylinder, and edges between them are intersections of the corresponding subgraphs. From this fact and quasiisometry invariance of the JSJ tree of cylinders, several quasiisometry invariants appear:

Lemma 2.11. *RAAGs have no hanging vertices in their JSJ decompositions.*

Lemma 2.12. *Let Δ be a connected graph with more than one vertex. The rigid vertex groups of the JSJ graph of cylinders of A_Δ are one-ended special subgroups $A_{\Delta'}$ of A_Δ that do not split further over 2–ended subgroups.*

Proof. A_Δ is finitely presented and one-ended, so it has a JSJ decomposition. Any rigid vertices correspond to maximal biconnected subgraphs Δ' of Δ that either contain at least two cut vertices or are not contained in the star of any cut vertex. In particular, rigid vertex groups are special subgroups. Since Δ is connected with more than one vertex, every vertex is contained in an edge. Edges are biconnected, so no single vertex is a maximal biconnected subgraph. Thus, Δ' is connected with more than one vertex, so $A_{\Delta'}$ is one-ended. Furthermore, since Δ' is biconnected it contains no cut vertex, so $A_{\Delta'}$ has no two-ended splittings. □

2.4.2. *JSJ decompositions of RACGs.* Mihalik and Tschantz [71] show that in some sense all decompositions of Coxeter groups are visual.

For RACGs it is not hard to see that splittings over finite subgroups correspond to separating cliques in the presentation graph. We will also exclude the well understood cases that the presentation graph Γ is complete (W_Γ is finite) and Γ is a cycle (W_Γ is virtually a surface group). Assuming that Γ is triangle-free, incomplete, with no separating clique, and is not a cycle, Edletzberger [40], extending work of Dani and Thomas [30] from the hyperbolic case, gives a description of the JSJ graph of cylinders in terms of subgraphs of the defining graph. In particular, 2–ended splittings arise in two ways:

- $\{a, b\}$ is a cut pair of Γ, meaning that $\Gamma - \{a, b\}$ is not connected.
- $\{a, b\}$ is not a cut pair, but there is a common neighbor $c \in \mathrm{lk}(a) \cap \mathrm{lk}(b)$ such that $\Gamma - \{a, b, c\}$ is not connected. In this case $a \multimap c \multimap b$ is called a *cut 2–path.*

We combine the two by saying $\{a - b\}$ *is a cut* to mean that either $\{a, b\}$ is a cut pair or that there exists c such that $a \multimap c \multimap b$ is a cut 2–path. In the first case $\langle ab \rangle \cong \mathbb{Z}$ is an index-2 subgroup of $W_{\{a,b\}}$, and in the second case it is an index-4 subgroup of $W_{\{a,b,c\}}$. A cut is *crossed* if there is another cut containing vertices in multiple of its complementary components, and *uncrossed* otherwise. Crossed cuts group together to make hanging vertices of the JSJ graph of cylinders. Cylinder vertices of the JSJ graph of cylinders are commensurators of uncrossed cuts, which can be described explicitly: if $\{a - b\}$ is an uncrossed cut of Γ then there is a corresponding cylinder vertex with vertex group $W_{\{a,b\}*(\mathrm{lk}(a)\cap\mathrm{lk}(b))}$. Such a group is commensurable to one of $\mathbb{Z}$, $\mathbb{Z}^2$, or $F_2 \times \mathbb{Z}$, according to whether $|\mathrm{lk}(a) \cap \mathrm{lk}(b)|$ is less than 2, equal to 2, or greater than 2, respectively.

Rigid vertices can also be described explicitly ([40, Proposiiton 3.8]), they are subsets of size at least 4 of vertices of Γ, each with valence at least 3, that cannot be separated by any pair of vertices or 2–path in Γ, and that are maximal with respect to inclusion among sets with these properties.

2.4.3. *Further splittings.* In this paper we are always using 'JSJ decomposition' in the sense of decomposition over 2–ended subgroups. JSJ theory also exists for other classes of splittings [37, 43, 50]. In particular, Groves and Hull [49] describe the structure of splittings of RAAGs over Abelian groups. In the 2–dimensional case, it would therefore be interesting to consider whether the kind of invariants we develop for 2–ended JSJ decompositions of RACGs vs RAAGs can be extended to splittings over virtually $\mathbb{Z}^2$ subgroups. We leave this line of inquiry for future work.

2.5. Morse property and stability.

A subspace Z of a metric space X is *μ–Morse* if for every L and A we have that every (L, A)–quasigeodesic segment γ of X with endpoints in Z is contained in the $\mu(L, A)$–neighborhood of Z.

A subspace is *Morse* if there is some μ for which it is μ–Morse.

Morseness is a quasigeodesic quasiconvexity condition on Z that describes how it sits in X, but says nothing about the intrinsic geometry of Z. There is a further property, *stability*, that, when X is a geodesic metric space, is equivalent to Z being Morse and Z itself being a hyperbolic space.

The concept of subgroup stability was introduced by Durham and Taylor [38] as a geometric group-theoretic interpretation of convex-cocompactness for subgroups mapping class groups of surfaces. Another characterization of such subgroups is that the subgroup is convex cocompact if and only if its orbit map into the curve graph of the surface is a quasiisometric embedding [63, 53]. These results inspired work to characterize stability in other families of groups, and then to search for a *stability recognizing space* to play the role of the curve graph, in the sense that a subgroup of the given group is stable if and only if its orbit map into the stability recognizing space is a quasiisometric embedding.

The theory of Morse and stable subgroups of RAAGs and RACGs is well developed, as we will briefly recall. We will return to the topic of stability recognizing spaces in Section 6.2.

Theorem 2.13 ([24, Theorem F],[80, Corollary 7.4][3])**.** *In a 1–ended RAAG, every Morse subset is coarsely dense or quasiisometric to a finite valence tree.*

[3] This result also says that in a strongly CFS RACG every Morse subset is either hyperbolic or coarsely dense.

The Morse property is easily identifiable for special subgroups of RACGs. A subgraph Γ' of Γ is *square complete* if whenever Γ' contains a diagonal of an induced square it contains the whole square. An induced subgraph is *minsquare* if it contains an induced square, is square complete, and is minimal with respect to inclusion among all subgraphs that satisfy the first two conditions.

Theorem 2.14 ([81, Theorem 1.11],[44, Proposition 4.9])**.** *A special subgroup of a RACG is Morse if and only if its presentation subgraph is square complete.*

Corollary 2.15. *A special subgroup of a RACG is stable if and only if its presentation subgraph is square complete and contains no induced square.*

Proposition 2.16. *A 1–ended RAAGedy RACG has presentation graph that is a join of a clique and a minsquare subgraph.*

Proof. Removal of the clique factor of the join is passage to a finite-index subgroup. The resulting RACG is still 1–ended and quasiisometric to a RAAG, so we may assume the presentation graph has no cone vertices.

A 1–ended RAAG contains $\mathbb{Z}^2$, so it is not hyperbolic. Thus, the presentation graph of the RACG contains an induced square. Given a square, the smallest square complete subgraph containing it is a minsquare subgraph, so such subgraphs exist. Suppose the presentation graph contains a proper subgraph that is minsquare. Its complement does not consist of cone vertices, since there are none, so the special subgroup defined by the subgraph is of infinite index and is Morse and 1–ended. Its image under the quasiisometry to the RAAG is a Morse subset that is neither coarsely dense nor quasiisometric to a tree, contradicting Theorem 2.13. Thus, the only minsquare subgraph is the presentation graph itself. □

2.6. **Maximal product regions.** This section recapitulates the setup for a theorem of Oh [74], who builds on work of Haglund and Wise [52] and Huang [57].

Definition 2.17. A *weakly special square complex* is a non-positively curved square complex X such that no wall of X self-osculates or self-intersects.

Huang [57, Section 5.3] considers two families of examples of compact weakly special cube complexes: the Salvetti complex of a RAAG and the *commutator complex*[4] P_Γ of a RACG W_Γ. The commutator complex is a standard construction [31, 32, 33] of a cube complex whose fundamental group is the commutator subgroup of W_Γ, so is a finite-index normal subgroup of W_Γ, and whose universal cover is the Davis complex Σ_Γ.

Huang asserts that a compact weakly special cube complex has a finite cover in which walls are 2–sided. In particular, the universal cover has 2–sided walls. Oh makes a standing assumption that we have already passed to such a finite cover. The examples that we care about in this paper, the Salvetti complex and commutator complex, already have 2–sided walls, but to accurately quote Oh, we make an auxiliary definition:

Definition 2.18. A compact cube complex is *weakly special** if it is weakly special with 2–sided walls.

Remark. Compared to *special* [52, Def. 3.1], Definition 2.18 legalizes inter-osculation.

Definition 2.19. A *standard graph* is a topologically nontrivial graph without leaves. A *standard product subcomplex* of a compact weakly special* square complex X is the image of a local isometry from the product of two standard graphs into X. A *standard product subcomplex/region* of $\tilde{X}$ is a product subcomplex of $\tilde{X}$ that is a lift to $\tilde{X}$ by the universal covering map of a standard product subcomplex of X.

[4]Davis [33] calls the complex P_Γ the *polyhedral product of intervals.*

The following lemma of Oh allows us to work directly with standard product regions in $\tilde{X}$ without explicitly demonstrating that they are lifts of standard product subcomplexes of X.

Lemma 2.20. [74, Lemma 2.8] *Every product subcomplex of $\tilde{X}$ that is a product of infinite trees without leaves is a standard product region.*

Definition 2.21. The *intersection graph*, or *maximal standard product region graph (MPRG)*, $\Pi(X)$ of a compact, weakly special* square complex X is the graph whose vertices are maximal standard product regions in $\tilde{X}$ such that two vertices are connected by an edge if the corresponding product regions intersect in a standard product region.

Theorem 2.22 ([74, Corollary 3.3])**.** *For each L and A there is C such that if X and Y are compact, weakly special* square complexes and $\phi\colon \tilde{X} \to \tilde{Y}$ is an (L, A)–quasiisometry between their universal covers then ϕ induces a bijection ϕ_* between maximal standard product subcomplexes of $\tilde{X}$ and $\tilde{Y}$ such that for each maximal standard product subcomplex P of $\tilde{X}$, $d_{Haus}(\phi(P), \phi_*(P)) \leqslant C$. It follows that P and $\phi_*(P)$ are quasiisometric.*

Corollary 2.23. *The maximal standard product region graph of a compact weakly special* cube complex can be decorated by adding to each vertex v the quasiisometry type of the maximal product region corresponding to v. A quasiisometry between universal covers of compact, weakly special* square complexes induces an isomorphism of their maximal standard product regions graphs that respects these decorations.*

Definition 2.24. The *reduced intersection graph* $\mathrm{Ric}(X)$ of a compact, weakly special* square complex X is the graph whose vertices are maximal standard product subcomplexes of X such that two vertices are connected by an edge if the corresponding product subcomplexes intersect in a standard product subcomplex.

We will work with $\Pi(X)$ and $\mathrm{Ric}(X)$ as graphs, but actually these graphs are the 1–skeleta of higher dimensional complexes that Oh calls the *intersection complex* and *reduced intersection complex*, respectively. Higher dimensional cells are made by filling in a simplex whenever there is a clique of maximal standard product subcomplexes all containing a common standard product subcomplex.

Theorem 2.25 ([74, Theorem 3.9])**.** *If X is a compact, weakly special* square complex, $\Pi(X)$ is its intersection complex, and $\mathrm{Ric}(X)$ is its reduced intersection complex, then the $\pi_1(X)$ action on $\tilde{X}$ by deck transformations induces an action on $\Pi(X)$ with fundamental domain isomorphic to $\mathrm{Ric}(X)$. Furthermore, this isomorphism is compatible with the definitions of $\Pi(X)$ and $\mathrm{Ric}(X)$ in the sense that if $v \in \mathrm{Ric}(X)$ and $\tilde{v}$ is the corresponding vertex in the fundamental domain of $\Pi(X)$ then the maximal standard product region of $\tilde{X}$ corresponding to $\tilde{v}$ is a lift to $\tilde{X}$ of the maximal standard product subcomplex corresponding to v in X.*

Lemma 2.26. *With notation as in the previous theorem, if $\gamma = e_0, \ldots, e_{n-1}$ is an edge path in $\Pi(X)$ then there are $g_i \in \pi_1(X)$ such that $\gamma' = g_0e_0, \ldots, g_{n-1}e_{n-1}$ is a path in $\mathrm{Ric}(X)$. Furthermore, vertices and edges of γ contained in $\mathrm{Ric}(X)$ are invariant under this projection.*

Proof. By definition, an edge $\tilde{e}$ in $\Pi(X)$ represents a standard product region $\tilde{P}_{\tilde{e}}$ of $\tilde{X}$ that is the intersection of two maximal standard product regions. Standard product regions of $\tilde{X}$ are, by definition, lifts to $\tilde{X}$ of standard product subcomplexes of X, so there is a standard product subcomplex P_e of X and corresponding edge e of $\mathrm{Ric}(X)$ such that the quotient by the $\pi_1(X)$–action sends $\tilde{P}_{\tilde{e}}$ to P_e. Thus, identifying $\mathrm{Ric}(X)$ with the fundamental domain of $\pi_1(X)\backslash\Pi(X)$ as in Theorem 2.25, for each of the edges $e_i = v_i \multimap v_{i+1}$ of γ there exists $g_i \in \pi_1(X)$ such that $g_i e_i \in \mathrm{Ric}(X)$. Then $g_i v_i$ and $g_{i+1} v_i$ are vertices of $\mathrm{Ric}(X)$ in the same $\pi_1(X)$–orbit, so they are the same vertex, so γ' is a path. □

Proposition 2.27 (Visibility of reduced intersection graphs for RAAGs)**.** *Let Δ be a connected, triangle-free graph. The reduced intersection graph* Ric_Δ *of the Salvetti complex of A_Δ is isomorphic to the graph obtained from Δ by taking a vertex for each maximal join subgraph and connecting them by an edge if the join subgraphs have an edge in common.*

Proof. Edges of the universal cover Σ_Δ of the Salvetti complex are colored by the corresponding vertex of Δ. Let $\Theta = \Theta_1 \times \Theta_2$ be a product region of Σ_Δ. Let Δ_i be the set of vertices $v \in \Delta$ such that some edge of Θ_i has color v. Since Θ is a product, every edge in Θ_1 spans a square with every edge in Θ_2, which means that $\Delta_1 * \Delta_2$ is a join subgraph of Δ. Conversely, if $\Delta_1 * \Delta_2$ is a subgraph of Δ, then $\Sigma_{\Delta_1} \times \Sigma_{\Delta_2}$ is a product region of Σ_Δ, some translate of which contains Θ. If $\Delta_1' * \Delta_2'$ is a strictly larger join subgraph containing $\Delta_1 * \Delta_2$ then $\Sigma_{\Delta_1'} \times \Sigma_{\Delta_2'}$ is a product region strictly containing $\Sigma_{\Delta_1} \times \Sigma_{\Delta_2}$, hence, up to translation, Θ. Thus, Θ is a maximal product region if and only if it is a translate of $\Sigma_{\Delta_1} \times \Sigma_{\Delta_2}$, such that $\Delta_1 * \Delta_2$ is a maximal join in Δ.

Edges in the product region graph correspond to having a common product sub-region, which equates to two maximal join subgraphs of Δ having a common join subgraph. Every join subgraph contains an edge, and edges are joins, so it suffices to consider edges. □

Proposition 2.28 (Visibility of reduced intersection graphs for RACGs)**.** *Let Γ be a triangle-free graph without separating cliques. The reduced intersection graph* Ric_Γ *of the commutator complex P_Γ of Γ is isomorphic to the graph obtained from Γ by taking a vertex for each maximal thick join subgraph, and connecting them by an edge if the join subgraphs have a square in common.*

Furthermore, the action of W_Γ on the Davis complex $\Sigma_\Gamma = \widetilde{P_\Gamma}$ induces an action of W_Γ on the maximal standard product region graph $\Pi(W_\Gamma) := \Pi(P_\Gamma)$ that has a fundamental domain isomorphic to Ric_Γ*.*

Proof. The argument for identifying Ric_Γ with the graph of thick joins in Γ is the same as for Proposition 2.27, except that product regions are supposed to be infinite in both factors, so we add the requirement that the join subgraphs are thick.

The content of the second part is that Theorem 2.25 tells us that Ric_Γ is the fundamental domain for the action of $\pi_1(P_\Gamma)$ on Σ_Γ, and W_Γ is a supergroup of $\pi_1(P_\Gamma)$, so a priori might have had smaller fundamental domain, but it does not, because $W_\Gamma \curvearrowright \Sigma_\Gamma$ preserves the edge coloring. □

Definition 2.29. If Υ is the presentation graph of a RAAG/RACG, let Ric_Υ be a choice of lift of the reduced intersection graph of the Salvetti complex/commutator complex to the maximal standard product region graph Π_Υ, which gives a fundamental domain for the group action.

For $v \in \mathrm{Ric}_\Upsilon$ let J_v be the corresponding maximal (thick, in the RACG case) join subgraph of Υ, as described in Proposition 2.27 and Proposition 2.28.

2.7. **Convex sets and projections in CAT(0) cube complexes.** Huang calls a subcomplex of the universal cover Σ_Δ of a RAAG A_Δ *standard* if it is a translate of $\Sigma_{\Delta'}$ for some $\Delta' \subset \Delta$. After fixing an identity vertex, thus identifying the 1–skeleton of Σ_Δ with the Cayley graph of A_Δ, vertex sets of standard subcomplexes are precisely cosets of special subgroups of A_Δ. This description applies equally well to RACGs, so we can also speak of standard subcomplexes of Davis complexes of RACGs. By Proposition 2.27 and Proposition 2.28, this is consistent with the 'standard product subcomplex' terminology of Oh, in the sense that standard product subcomplexes are examples of standard subcomplexes.

In Section 7 we need some results about coarse intersections of standard subcomplexes. These follow from some properties of projections to convex sets in (finite dimensional) CAT(0) cube complexes. We will work in the combinatorial metric: the metric on the vertex set obtained by restricting the path metric on the 1–skeleton.

Lemma 2.30. *Let X be a CAT(0) cube complex with its combinatorial metric, and let Y be a convex subcomplex. There is a Lipschitz map $\pi_Y\colon X \to Y$ sending each vertex $x \in X$ to the vertex of Y that is closest to x. The map π_Y is a* gate projection*: for all $x \in X$ and $y \in Y$ the vertex $\pi_Y(x)$ is on a geodesic from x to y.*

In a CAT(0) cube complex, consider the equivalence relation on edges generated by the condition that two edges are equivalent if they are opposite edges of some square in the complex. Dual to each equivalence class is a *wall* (also called a *hyperplane*), which can be thought of geometrically as the union of midcube hyperplanes dual to these edges in the cubes containing them. A wall $\mathcal{H}$ separates the cube complex into two complementary sets of vertices, each of which spans a convex subcomplex, called *halfspaces* and usually denoted $\mathcal{H}^+$ and $\mathcal{H}^-$, such that every combinatorial geodesic from $\mathcal{H}^-$ to $\mathcal{H}^+$ contains an edge in the equivalence class defining the wall. The number of walls separating two vertices is equal to the combinatorial distance between them. Combinatorial geodesics cross each wall at most once, and two combinatorial geodesics between the same points cross the same set of walls. Two walls *cross* if there is a cube containing an edge dual to each of them, or, equivalently, if all four possible intersections of their halfspaces are nonempty.

The following result is well known to experts. Similar results have been proved and reproved several times in different settings. In the CAT(0) metric, compare [13, Lemma 2.3], [57, Lemma 2.10], also [3]. Chatterji, Fernós, and Iozzi [20, Lemma 2.18] prove a version of Proposition 2.31 in the combinatorial metric, but only state it for halfspaces Y and Z. Full proofs of the more general statement appear in unpublished sources [5] [51, Theorem 1.22], [46, Proposition 1.5.2]. A proof can also be deduced from results of Isbell on parallelism between gate projections in discrete median algebras [60, cf. Proclamation 2.5 and the Corollary to Theorem 5.12], together with the well known equivalence between discrete median algebras and 0-skeleta of CAT(0) cube complexes [47, 79, 21].

Proposition 2.31. *[Bridge Lemma] Let X be a CAT(0) cube complex with its combinatorial metric and let Y and Z be convex subcomplexes. The* bridge $Y \asymp Z$*, consisting of the subcomplex spanned by vertices that lie on some minimal length geodesic between Y and Z, is a combinatorially convex subcomplex isomorphic to $\pi_Y(Z) \times [y, \pi_Z(y)]$, where y is any vertex in $\pi_Y(Z)$ and $[y, \pi_Z(y)]$ is the subcomplex spanned by vertices that lie on a combinatorial geodesic from y to $\pi_Z(y)$.*

The walls that meet $Y \asymp Z$ are partitioned into the set that traverse the bridge and the set that transect the bridge. A wall traverses the bridge if and only if meets both $\pi_Z(Y)$ and $\pi_Y(Z)$, which is true if and only if it meets both Y and Z. A wall transects the bridge if and only if it separates Y and Z, which is true if and only if its intersection with $Y \asymp Z$ separates $\pi_Y(Z)$ from $\pi_Z(Y)$ in $Y \asymp Z$. Every wall that traverses the bridge crosses every wall that transects the bridge, and vice versa.

Corollary 2.32. *For X, Y, Z as in Proposition 2.31, $\pi_Y(Z)$, $\pi_Z(Y)$, and $Y \asymp Z$ are coarsely equivalent and represent $Y \overset{c}{\cap} Z$.*

---

[5] We thank Anthony Genevois for the references.

Proof. Coarse equivalence is the observation $\pi_Y(Z) \subset Y \asymp Z \subset \bar{\mathcal{N}}_{d(Y,Z)}(\pi_Y(Z))$.

For $r \geqslant \lceil d(Y,Z)/2 \rceil$, consider $x \in \bar{\mathcal{N}}_r(Y) \cap \bar{\mathcal{N}}_r(Z) \neq \emptyset$, and let $y := \pi_Y(x)$ and $z := \pi_Z(x)$. By Lemma 2.30, $\pi_Z(y)$ is on a geodesic from y to z, so $2r \geqslant d(y,z) \geqslant d(y, Y \asymp Z) + d(z, Y \asymp Z)$, so at least one of them is bounded above by r, which implies $d(x, Y \asymp Z) \leqslant 2r$. Thus, $Y \asymp Z \subset \bar{\mathcal{N}}_r(Y) \cap \bar{\mathcal{N}}_r(Z) \subset \bar{\mathcal{N}}_{2r}(Y \asymp Z)$. ☐

2.8. **Coarse geometry of standard subcomplexes in RAAGs and RACGs.** The Davis complex of a RACG and the universal cover of the Salvetti complex of a RAAG are CAT(0) cube complexes in which every edge is labelled/colored by a generator of the group. Furthermore, since all squares correspond to commutation relations in the group, opposite sides of any square have the same label.

Lemma 2.33. *If Σ is the Davis complex of a RACG or the universal cover of the Salvetti complex of a RAAG, X_0 and X_1 are convex subcomplexes, e is an edge of $\pi_{X_0}(X_1)$ labelled a, and $\gamma\colon [0,L] \to \Sigma$ is a geodesic from $\iota(e)$ to $\pi_{X_1}(\iota(e))$, then Σ contains a subcomplex $\gamma \times [0,1]$ such that $\gamma(0) \times [0,1] = e$, $\gamma(L) \times [0,1]$ is an edge of X_1, every edge $\gamma(i) \times [0,1]$ has label a, and for every $0 < i \leqslant L$, the two edges $\gamma([i-1,i]) \times \{0,1\}$ have the same label b_i, which commutes with a.*

Proof. If X_0 and X_1 intersect there is nothing to prove: $L = 0$ and there is a single edge $e \in X_0 \cap X_1$ labelled a and the set of commuting b_i's is empty. Otherwise, Proposition 2.31 implies that Σ contains a subcomplex $\gamma \times [0,1]$ with $\gamma(0) \times [0,1] = e$ and $\gamma(L) \times [0,1] \subset X_1$. Since squares are labelled by commutators, $\gamma(1) \times [0,1]$ has label a, and $\gamma([0,1]) \times \{0\}$ and $\gamma([0,1]) \times \{1\}$ are edges with the same label, which commutes with a. Repeat for each edge of γ. ☐

Proposition 2.34. *Let Υ be a graph, let G_Υ be either a RACG or RAAG with presentation graph Υ, and let Σ_Υ be its Davis complex or the universal cover of its Salvetti complex, respectively. Consider $S_0, S_1 \subset \Upsilon$ and $g_0, g_1 \in G_\Upsilon$. Let w be a minimal length representative of $G_{S_0}(g_0^{-1}g_1)G_{S_1}$, and let T be the set of generators that appear in w. Then either $T = \emptyset$, which occurs when $g_0\Sigma_{S_0} \cap g_1\Sigma_{S_1} \neq \emptyset$, or T contains an element not in S_0 and an element not in S_1. Furthermore:*

$$\pi_{g_0\Sigma_{S_0}}(g_1\Sigma_{S_1}) = g_0\Sigma_{S_0 \cap S_1 \cap \bigcap_{t \in T} \operatorname{lk}(t)}$$

In particular, projection of a standard subcomplex to a standard subcomplex is a standard subcomplex.

Proof. The description of w is equivalent to that of a word read on the edges of a minimal length geodesic connecting $g_0\Sigma_{S_0}$ to $g_1\Sigma_{S_1}$. Proposition 2.31 implies that the set T does not depend on the choice of such a geodesic. Using the group action, take $h \in \pi_{g_0\Sigma_{S_0}}(g_1\Sigma_{S_1}) \subset g_0\Sigma_{S_0}$ so that $1 \in h^{-1}\pi_{g_0\Sigma_{S_0}}(g_1\Sigma_{S_1}) = \pi_{h^{-1}g_0\Sigma_{S_0}}(h^{-1}g_1\Sigma_{S_1}) = \pi_{\Sigma_{S_0}}(h^{-1}g_1\Sigma_{S_1})$. Again by Proposition 2.31, every vertex in $\pi_{\Sigma_{S_0}}(h^{-1}g_1\Sigma_{S_1})$ is connected to a vertex in $\pi_{h^{-1}g_1\Sigma_{S_1}}(\Sigma_{S_0})$ by a geodesic labelled w, so the geodesic starting at 1 labelled w ends at $w \in \pi_{h^{-1}g_1\Sigma_{S_1}}(\Sigma_{S_0}) \subset h^{-1}g_1\Sigma_{S_1}$. Thus, it suffices to consider the case that $g_0 = 1$ and $g_1 = w$.

The set T is empty, and w is the empty word, precisely when $\Sigma_{S_0} \cap \Sigma_{S_1} \neq \emptyset$. In this case $S_0 \cap S_1 \cap \bigcap_{t \in T} \operatorname{lk}(t) = S_0 \cap S_1$, and $\pi_{\Sigma_{S_0}}(\Sigma_{S_1}) = \Sigma_{S_0 \cap S_1}$.

Suppose that T is not empty and w is not the empty word. Minimality of w implies that it does not start with an element of S_0 or end with an element of S_1, so T contains an element not in S_0 and an element not in S_1.

Let S' be the set of labels that occur on edges of $\pi_{\Sigma_{S_0}}(w\Sigma_{S_1})$. By Lemma 2.33 every $s \in S'$ commutes with every letter of w, and if $e \in \pi_{\Sigma_{S_0}}(w\Sigma_{S_1})$ is an edge labelled s then the edge $we \in \pi_{w\Sigma_{S_1}}(\Sigma_{S_0})$ is also labelled s, so $S' \subset S_0 \cap S_1 \cap \bigcap_{t \in T} \operatorname{lk}(t)$. Conversely, if $s \in S_0 \cap S_1 \cap \bigcap_{t \in T} \operatorname{lk}(t)$ then $1 \bullet\!\!-\!\!\bullet s$ is an edge of Σ_{S_0} that is parallel via w to the edge $w\bullet\!\!-\!\!\bullet ws$ in $w\Sigma_{S_1}$, so $S' = S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)$. Furthermore, $\Sigma_{S'} \subset \Sigma_{S_0}$ is parallel to $w\Sigma_{S'} \subset w\Sigma_{S_1}$ via w, so $\Sigma_{S'} = \pi_{\Sigma_{S_0}}(w\Sigma_{S_1})$. □

Corollary 2.35. *If X_0 and X_1 are standard subcomplexes in either the Davis complex of a RACG or the universal cover of the Salvetti complex of a RAAG, then X_0 and X_1 have unbounded coarse intersection if and only if there is an unbounded standard subcomplex X_2 such that X_0 and X_1 contain parallel copies of X_2.*

Proof. Corollary 2.32 and Proposition 2.34 say $\pi_{X_0}(X_1)$ and $\pi_{X_1}(X_0)$ are parallel standard subcomplexes representing $X_0 \overset{c}{\cap} X_1$. □

Corollary 2.36. *Two standard subcomplexes X_0 and X_1 of the universal cover of the Salvetti complex of a RAAG have bounded coarse intersection if and only if the combinatorial closest point projection map from X_0 to X_1 is constant.*

Proof. In a RAAG, once a standard subcomplex contains an edge it contains the entire bi-infinite monochrome geodesic containing that edge, so $\pi_{X_0}(X_1)$ and $\pi_{X_1}(X_0)$ are either single vertices or unbounded. □

Lemma 2.37. *If W_Γ is a RACG and $A, B \subset \Gamma$ then $\Sigma_B \subset \Sigma_A \overset{c}{=} \Sigma_B$ if and only if $A = B * C$, where C is a clique.*

Proof. If C is a clique then $d_{Haus}(\Sigma_B, \Sigma_{B*C}) = |C|$. Conversely, if $\Sigma_B \subsetneq \Sigma_A \overset{c}{=} \Sigma_B$ and $A \neq B * C$ for a clique C, then either there is an $a \in A - B$ not adjacent to some $b \in B$, or there are a and b in $A - B$ not adjacent to each other. In either case $\Sigma_{\{a,b\}}$ is a line. Since $\pi_{\Sigma_{\{a,b\}}}$ is Lipschitz, $d_{Haus}(\Sigma_A, \Sigma_B) \geqslant d_{Haus}(\pi_{\Sigma_{\{a,b\}}}(\Sigma_A), \pi_{\Sigma_{\{a,b\}}}(\Sigma_B))$, the latter of which is infinite, since $\Sigma_{\{a,b\}} \subset \Sigma_A$, but $\pi_{\Sigma_{\{a,b\}}}(\Sigma_B)$ is finite. □

Corollary 2.38. *If W_Γ is a RACG and $S_0, S_1 \subset \Gamma$ then $\Sigma_{S_0} \overset{c}{=} \Sigma_{S_1}$ implies $d_{Haus}(\Sigma_{S_0}, \Sigma_{S_1})$ is bounded above by the size of the largest clique in Γ.*

Proof. Let C be the largest clique of Γ. Let S_i' be S_i minus its cone vertices. By Lemma 2.37, $d_{Haus}(\Sigma_{S_i}, \Sigma_{S_i'}) \leqslant |C|$, so $\Sigma_{S_0} \overset{c}{=} \Sigma_{S_1} \implies \Sigma_{S_0'} \overset{c}{=} \Sigma_{S_1'}$. But then $\Sigma_{S_0' \cap S_1'} \subset \Sigma_{S_0'} \overset{c}{=} \Sigma_{S_0' \cap S_1'}$, so Lemma 2.37 says $S_0' \cap S_1' = S_0'$. Repeat the argument for S_1', and conclude $S := S_0' = S_1'$. Now every vertex of Σ_{S_i} is distance at most $|C|$ from $\Sigma_S \subset \Sigma_{S_0} \cap \Sigma_{S_1}$. □

Corollary 2.39. *If W_Γ is a RACG, $g \in W_\Gamma$, and $S_0, S_1 \subset \Gamma$ then $\Sigma_{S_0} \overset{c}{=} g\Sigma_{S_1}$ implies there exists S and cliques C_0 and C_1 such that $S_0 = S * C_0$, $S_1 = S * C_1$ and g centralizes S.*

Proof. By Corollary 2.38, we may assume S_0 and S_1 have no cone vertices. By Corollary 2.32, if if T is the set of generators appearing in minimal length elements of $W_{S_0} g W_{S_1}$ then:

$$\Sigma_{S_0} \overset{c}{\cap} g\Sigma_{S_1} \overset{c}{=} \pi_{\Sigma_{S_0}}(g\Sigma_{S_1}) = \Sigma_{S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)}$$

So if $\Sigma_{S_0} \overset{c}{=} g\Sigma_{S_1}$ then:

$$\Sigma_{S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)} \subset \Sigma_{S_0} \overset{c}{=} (\Sigma_{S_0} \overset{c}{\cap} g\Sigma_{S_1}) \overset{c}{=} \Sigma_{S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)}$$

Lemma 2.37 says $S_0 = (S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)) * C$ for a clique C, but S_0 has no cone vertex, so $C = \varnothing$, $S_0 \subset S_1$, and $S_0 \subset \bigcap_{t\in T} \mathrm{lk}(t)$. The same argument applies to S_1, so we conclude that $S := S_0 = S_1$ and T centralizes S. □

Lemma 2.40. *If A_Δ is a RAAG, $g \in A_\Delta$, and $S_0, S_1 \subset \Delta$ then $\Sigma_{S_0} \overset{c}{=} g\Sigma_{S_1}$ implies $S := S_0 = S_1$ and g centralizes S.*

Proof. Let $\Sigma_S := \pi_{\Sigma_{S_0}}(g\Sigma_{S_1})$. Since $\Sigma_{S_0} \overset{c}{=} g\Sigma_{S_1}$, $\Sigma_{S_0} \overset{c}{=} \Sigma_{S_0} \overset{c}{\cap} g\Sigma_{S_1}$, which, by Corollary 2.32, is coarsely equivalent to Σ_S. By Proposition 2.34, $S = S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t)$, where T is the letters appearing in a minimal word in $A_{S_0} g A_{S_1}$. But in a RAAG no standard subcomplex is coarsely equivalent to one of its proper standard subcomplexes, since generators have infinite order, so $S = S_0$, which implies $S_0 = S_1$ and g centralizes S. ☐

3. CFS graphs

We know from Dani and Thomas [29] that an incomplete, triangle-free graph without separating cliques is RAAGedy only if it is CFS. In this section we establish some basic results about the structure of CFS graphs.

3.1. The diagonal graph.

Definition 3.1. The *diagonal graph* $\boxslash(\Gamma)$ of Γ is the graph with a vertex $\{a, b\}$ if a and b are vertices of Γ that are the diagonal vertices of some induced square. There is an edge $\{a, b\} \bullet\!\!-\!\!\bullet \{c, d\}$ in $\boxslash(\Gamma)$ when $\{a, b\} * \{c, d\}$ is an induced square of Γ.

The *support* of a vertex $\{a, b\}$ of $\boxslash(\Gamma)$ is the set $\{a, b\} \subset \Gamma$. The support of a subset of $\boxslash(\Gamma)$ is the union of supports of its vertices.

Definition 3.2. [6,7] Γ is *CFS* if $\boxslash(\Gamma)$ contains a connected component whose support is all non-cone vertices of Γ.

Γ is *strongly CFS* if it is CFS and $\boxslash(\Gamma)$ is connected.

Lemma 3.3. *A triangle in* $\boxslash(\Gamma)$ *corresponds to an induced octahedron in* Γ.

Corollary 3.4. *A triangle-free graph has a triangle-free diagonal graph.*

If A is a set, let $\binom{A}{2}$ be the collection of 2–element subsets of A. The following relates joins in Γ and joins in $\boxslash(\Gamma)$, and will be used later.

Lemma 3.5. *Let* Γ *be triangle-free. If* $A * B$ *is a thick join in* Γ *then* $\binom{A}{2} * \binom{B}{2}$ *is a join in* $\boxslash(\Gamma)$. *If* $A * B$ *is a join in* $\boxslash(\Gamma)$ *then* $\mathrm{supp}(A) * \mathrm{supp}(B)$ *is a thick join in* Γ *and* $\binom{\mathrm{supp}(A)}{2} * \binom{\mathrm{supp}(B)}{2}$ *is a join in* $\boxslash(\Gamma)$.

Proof. Suppose $A * B$ is a thick join in Γ. Let $\{a_0, a_1\}$ be any 2–element subset of A, and let $\{b_0, b_1\}$ be any 2–element subset of B. Since $A * B$ is a join, there is a square $\{a_0, a_1\} * \{b_0, b_1\}$, which is induced, since Γ is triangle-free, so there is an edge $\{a_0, a_1\} \bullet\!\!-\!\!\bullet \{b_0, b_1\}$ in $\boxslash(\Gamma)$.

Conversely, suppose $A * B$ is a join in $\boxslash(\Gamma)$. Let $a \in \mathrm{supp}(A)$ and $b \in \mathrm{supp}(B)$. Then there exists a' and b' such that $\{a, a'\} \in A$ and $\{b, b'\} \in B$, so $\{a, a'\} \bullet\!\!-\!\!\bullet \{b, b'\}$ is an edge in $A * B$, hence in $\boxslash(\Gamma)$, so $\{a, a'\} * \{b, b'\}$ is an induced square in Γ. In particular, a and b are adjacent in Γ. Since this was true for arbitrary elements of $\mathrm{supp}(A)$ and $\mathrm{supp}(B)$, Γ contains $\mathrm{supp}(A) * \mathrm{supp}(B)$. Moreover, each factor has size at least two, since supports of single vertices of $\boxslash(\Gamma)$ have size two, and since Γ is triangle-free this implies each factor is incomplete, so this is a thick join in Γ. Finally, by the first part of this lemma, $\binom{\mathrm{supp}(A)}{2} * \binom{\mathrm{supp}(B)}{2}$ is a join in $\boxslash(\Gamma)$. ☐

Next we state two easy lemmas to formalize constructions that have appeared in examples elsewhere in the literature. We use these results without further comment.

[6]This is equivalent to other definitions of CFS [29, 5, 4] and strongly CFS [80]. These sources make definitions in terms of a 'square graph' of Γ that is slightly different than $\boxslash(\Gamma)$ but ultimately contains the same information.

[7]It is possible [4] to produce CFS graphs with multiple components of full support.

Lemma 3.6. *Let Γ' be a CFS graph and let V be a subset of the vertices of Γ'. Let Γ be obtained from Γ' by adding some new edges between vertices of V. If no new edge connects the diagonals of a square of Γ' then Γ is CFS.*

In particular, if V is a set of vertices of Γ' that are pairwise at distance at least 3 from one another, then no squares of Γ' are killed, so Γ is CFS. If we space V even wider, the new edges do not interact with squares of Γ' at all:

Lemma 3.7. *Let Γ' be a CFS graph and let V be a subset of the vertices of Γ' that are pairwise at distance at least 4 apart. Let Γ be obtained from Γ' by adding some new edges between vertices of V, and let Γ'' be the subgraph induced by the new edges. Then Γ is CFS and $\boxslash(\Gamma)$ is a disjoint union of $\boxslash(\Gamma')$ and $\boxslash(\Gamma'')$.*

In this way one can build all kinds of exotic examples of CFS graphs starting from ones with sufficiently large diameter. This is the content of an example of Behrstock [2], which can be interpreted as taking Γ' to be the large diameter CFS graph obtained by doubling a path graph of length at least 12 and taking V to be 5 vertices at pairwise distance at least 3 in Γ' whose induced graph is a pentagon, see Figure 2. Russell, Spriano, and Tran do something similar [80, Proposition 7.6].

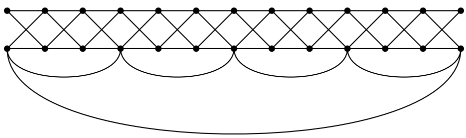

FIGURE 2. The example of Behrstock of a strongly CFS graph containing a stable cycle.

3.2. **First examples.**

3.2.1. *Large diameter contructions.* Here are some simple examples of strongly CFS graphs that can be constructed to have large diameter. They serve as basic building blocks for further examples.

Example 3.8. Consider the path graph P_n of length $n > 0$. Doubling gives a strongly CFS graph of diameter n.

For example $\mathfrak{D}(P_7) :=$ 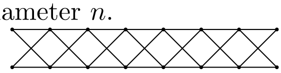

Deleting a vertex from one or both ends does not change the strong CFS property: the graphs 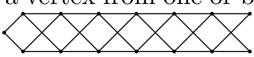 and 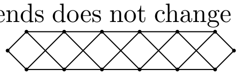 are both strongly CFS.

Consider a graph Γ that is a star with $m > 2$ legs of length $n \geqslant 2$, so that there are m leaves that are pairwise at distance $2n$ from one another. Take $\mathfrak{D}(\Gamma)$. This is the (m, n)–*spider*. Optionally, from each foot we can either leave both valence 2 vertices, and call this a pincer foot, or delete one of them. The result is commensurable to a RAAG tree, and has the property that vertices on different feet have pairwise distance $2n$.

Here is the $(4, 3)$–spider with one pincer foot: 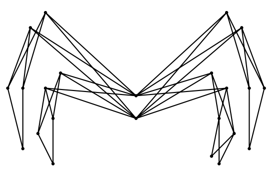 ◇

Proposition 3.9. *Every finite connected graph* Γ *can be isometrically embedded in a strongly CFS graph* Γ'*. Furthermore, if* Γ *is triangle-free then so is* Γ'*, and the construction can be arranged so that* $v, w \in \Gamma$ *are the diagonal vertices of an induced square of* Γ' *if and only if they are diagonal vertices of an induced square of* Γ*.*

Proof. Let Γ be a connected finite graph. Let Δ be a spider with legs of length $\max\{3, \lceil \frac{\mathrm{diam}(\Gamma)}{2} \rceil\}$, with one single-vertex foot for each vertex of Γ that is not contained in an induced square of Γ and one pincer foot for each component of $\boxslash(\Gamma)$. Identify each single vertex foot with its corresponding vertex of Γ. For each component of $\boxslash(\Gamma)$, choose a vertex, which is a pair of vertices of Γ that are the diagonal of some induced square, and identify these two vertices with the two vertices of the corresponding pincer foot of the spider. The resulting graph Γ' contains Γ as an isometrically embedded subgraph because the legs of the spider are long, so we have not created any shortcuts between vertices of Γ. In particular, all of the squares in Γ survive as squares in Γ'. Some vertices of Δ may get identified, but those that do come from different feet, so they have distance at least 6 in Δ; thus, all of the squares of Δ survive as squares in Γ', and we did not create any triangles.

By construction, every vertex of Γ' lies in a square, so to establish that Γ' is strongly CFS, it is enough to show that $\boxslash(\Gamma')$ is connected. Since Δ is strongly CFS, we conclude that vertices of $\boxslash(\Gamma')$ that are the diagonal of a square with edges only from Γ or only from Δ are contained in a common connected component of $\boxslash(\Gamma')$. This leaves us to consider vertices of $\boxslash(\Gamma')$ coming from squares that use edges from both Γ and Δ. The only vertices of Δ that attach to Γ and are at distance less than 6 from one another inside Δ are the pincers of a foot, so we may assume a and b are the pincers, and take d_0 and d_1 to be their two common neighbors in Δ. By construction, $\{a, b\}$ is also the diagonal of a square in Γ, so a and b have common neighbors $c_0, c_1, \ldots$ in Γ. For any i and j, there is a square $\{a, b\} * \{c_i, d_j\}$ with Γ–edges $c_i * \{a, b\}$ and Δ edges $d_j * \{a, b\}$. We have that $\{c_i, d_j\}$ is a vertex of $\boxslash(\Gamma')$ that does not occur in either $\boxslash(\Delta)$ or $\boxslash(\Gamma)$, but it is adjacent to the vertex $\{a, b\}$, which occurs in both.

The further claim about squares follows from the construction, as the 'mixed' squares with edges from Γ and Δ only occur when a pincer foot attached to the diagonal of an existing square of Γ. ☐

We remark that Γ embeds isometrically in $\mathfrak{D}(\Gamma)$, which is strongly CFS, but $\mathfrak{D}(\Gamma)$ does not satisfy the further claim about squares.

[80, Proposition 7.6] says any graph Γ can be isometrically embedded into a CFS graph Γ' in such a way that Γ is square complete in Γ'. If Γ contains a square then the Γ' produced by that construction will not be strongly CFS, whereas the Γ' produced by Proposition 3.9 will not contain Γ as a square complete subgraph. This tradeoff is unavoidable; it is not possible to embed a subgraph containing a square into a strongly CFS graph and have it be square complete, see Corollary 5.5.

3.2.2. *Blow-up graphs.*

Definition 3.10. Let (Δ, ω) be a *weighted graph*, consisting of a graph Δ and a weight function $\omega \colon \mathrm{Vertices}(\Delta) \to \mathbb{N}$. The *blow-up graph* Δ^ω is the graph with $\omega(v)$–many vertices $(v, 0), \ldots, (v, \omega(v) - 1)$ for each vertex $v \in \Delta$, and such that (v, i) and (w, j) are connected by an edge in Δ^ω if and only if v and w are connected by an edge in Δ.

Remark. $\Gamma \cong \mathbb{I}(\Gamma)^\omega$, where $\omega(M)$ is the number of vertices in the twin module M.

Remark 3.11. Without changing the graph structure of Δ^ω, we may assume that Δ is twin-free by replacing Δ with $\mathbb{I}(\Delta)$ and labelling each twin module by the sum of the labels of its vertices. This only changes the labelling of vertices of Δ^ω.

Assuming Δ is twin-free implies, in particular, that every vertex of Δ has at most one adjacent leaf. Furthermore, if Δ is twin-free then for $\Gamma := \Delta^\omega$ we have that $\amalg(\Gamma) \cong \Delta$ and that ω gives the cardinality of each twin module of Γ. Thus, if ω takes any odd values then Γ is not a graph double. On the other hand, if ω takes only even values then $\Gamma \cong \mathfrak{D}(\Delta^{\omega/2})$.

Proposition 3.12. *Let $\Gamma := \Delta^\omega$ be a blow-up graph, with the conditions that Δ is connected, triangle-free, twin-free, has at least two vertices, and ω takes the value 1 only on leaves of Δ, and if Δ is a single edge then ω does not take the value 1. Then Γ is triangle-free and strongly CFS.*

Proof. Projection to the first coordinate gives a map $\Gamma \to \Delta$ that sends edges to edges, and the preimage of a vertex is an anticlique. It follows that, when Δ is not a single vertex, Γ is connected and triangle-free if and only if Δ is.

If Δ is a single edge then by hypothesis $\Gamma \cong K_{m,n}$ with $m, n \geqslant 2$. It is easy to see that this graph is strongly CFS, so assume Δ is not a single edge.

For each edge $v \leftrightarrow w$ in Δ such that both of $\omega(v)$ and $\omega(w)$ are greater than one there are induced squares $\{(v,i),(v,j)\} * \{(w,k),(w,\ell)\}$ in Γ for all $i \neq j$ and $k \neq \ell$. Also, since weight 1 only occurs on leaves, for every embedded segment $v \leftrightarrow w \leftrightarrow x$ in Δ there are squares $\{(v,i),(x,j)\} * \{(w,k),(w,\ell)\}$ in Γ for all $k \neq \ell$. Finally, if $\{u,v\} * \{w,x\}$ is an induced square in Δ then there are induced squares $\{(u,i),(v,j)\} * \{(w,k),(x,\ell)\}$ in Γ. Conversely, each induced square in Γ projects to either a square or a path of length 1 or 2 in Δ, so the squares constructed above account for all of the induced squares of Γ, and this, in turn, accounts for all vertices and edges of $\boxslash(\Gamma)$. Since Δ is not a single edge, every vertex lies on some embedded segment of length 2, so every vertex belongs to some square. Thus, the support of $\boxslash(\Gamma)$ is all of Γ. We will show that $\boxslash(\Gamma)$ is connected, which then implies Γ is strongly CFS.

Since ω takes the value 1 only on leaves and Δ is twin-free, each vertex is adjacent to at most one vertex with weight 1. Furthermore, since Δ is twin-free and not a single edge, every vertex is adjacent to a non-leaf, so every vertex is adjacent to a vertex with weight greater than 1.

Define $\phi\colon\ \Delta \to \boxslash(\Gamma)$ by $\phi(v) = \{(v,0),(v,1)\}$ if $\omega(v) > 1$. If $\omega(v) = 1$ then the hypotheses on Δ imply there is a unique neighbor w of v, $\omega(w) > 1$, and there exist vertices at distance 2 from v. Choose any x at distance 2 from v, and define $\phi(v) = \{(v,0),(x,0)\}$, which is a diagonal of the induced square $\{(v,0),(x,0)\} * \{(w,0),(w,1)\}$ of Γ. Then ϕ sends edges of Δ to edges of $\boxslash(\Gamma)$ and is injective on vertices; the only potential source of collisions would be if v and x are both leaves weighted 1 at distance 2 from each other such that $\phi(v) = \{(v,0),(x,0)\} = \phi(x)$, but this is ruled out by the assumption that Δ is twin-free. Thus, $\phi(\Delta)$ is connected.

We claim that every vertex of $\boxslash(\Gamma)$ is in $\phi(\Delta)$ or is adjacent to a vertex of $\phi(\Delta)$. Since vertices of $\boxslash(\Gamma)$ come from one of the three types of induced squares in Γ enumerated above, we only need to consider the following three cases. If $\omega(v) > 1$ then, since v has some neighbor w not weighted 1, $\{(v,i),(v,j)\} \leftrightarrow \{(w,0),(w,1)\} = \phi(w)$ in $\boxslash(\Gamma)$ for all $i \neq j$. Similarly, if $d(v,x) = 2$ then they have a common neighbor w, which is not a leaf, so not weighted 1, and $\{(v,i),(x,j)\} \leftrightarrow \{(w,0),(w,1)\} = \phi(w)$ in $\boxslash(\Gamma)$ for all $i \neq j$. Finally, if $\{(u,i),(v,j)\} * \{(w,k),(x,\ell)\}$ is an induced square in Γ coming from an induced square $\{u,v\} * \{w,x\}$ of Δ then none of the vertices are leaves in Δ, so none are weighted 1, so $\{(v,i),(x,j)\} \leftrightarrow \{(w,0),(w,1)\} = \phi(w)$. □

In Section 4.2 we upgrade the conclusion of Proposition 3.12 to 'RAAGedy'.

3.3. Enumeration of small triangle-free CFS graphs.

We enumerated triangle-free CFS graphs of small order. Table 1 gives the number of isomorphism types of graph by number of vertices (V) and edges (E). Blank entries are 0.

E\V	4	5	6	7	8	9	10	11	12
4	1								
5									
6		1							
7									
8			2						
9			1						
10				3					
11				1					
12				1	8				
13					6				
14					3	19			
15					2	21			
16					1	17	61		
17						7	115		
18						4	119	207	
19						1	71	616	
20						1	37	950	828
21							17	782	3820
22							7	461	8722
23							3	212	10863
24							2	103	8492
25							1	42	4856
26								19	2385
27								7	1082
28								4	477
29								1	204
30								1	89
31									38
32									17
33									7
34									3
35									2
36									1

TABLE 1. Number of isomorphism types of triangle-free CFS graphs by numbers of vertices and edges. In each column the horizontal line designates the point below which all graphs are forced to be bipartite.

Mantel's Theorem says that a triangle-free graph with n vertices has at most $\lfloor \frac{n^2}{4} \rfloor$ edges, with equality if and only if the graph is $K_{\frac{n}{2},\frac{n}{2}}$ when n is even or $K_{\frac{n+1}{2},\frac{n-1}{2}}$ when n is odd. A complete bipartite graph is strongly CFS when both parts are non-singletons, so this gives us the unique largest triangle-free CFS graph for each fixed number of vertices. Similarly, it is an exercise in extremal graph theory to show that a triangle-free graph with n vertices and strictly greater than $\frac{(n-1)^2}{4}+1$ edges is bipartite. Further, it can be shown that these conditions also imply the graph is strongly CFS.

3.4. **Inductive construction.** [9, Theorem II] says that the class of thick RACGs is the class whose defining graphs belong to the smallest class of graphs containing the square and closed under the following operations:

- Cone off a subgraph that is not a clique.
- Amalgamate two graphs Γ_1, Γ_2 in the class over a common subgraph Γ' that is not a clique.
- Add additional edges to the previous item between vertices of $\Gamma_1 - \Gamma'$ and vertices of $\Gamma_2 - \Gamma'$.

CFS graphs are thick, so they can be built iteratively as above. However:

- Not all thick graphs are CFS.
- The 'amalgamate and add edges' operation is useful for constructing examples, but makes inductive arguments more complicated.

We show that for CFS graphs only the cone-off operation is needed.

We say that a property $\mathcal{P}$ is *constructible by coning from a square* if for every graph Γ with property $\mathcal{P}$ there exists a sequence $\Gamma_0 \subset \cdots \subset \Gamma_n$ where each Γ_i has property $\mathcal{P}$, Γ_0 is a square, $\Gamma_n = \Gamma$, and $\Gamma_{i+1} = \Gamma_i *_{N_i} v_i$ is obtained from Γ_i by coning off some subgraph N_i.

Proposition 3.13. *The property of being an incomplete CFS graph is constructible by coning from a square: An incomplete CFS graph can be built from a square as a sequence of CFS graphs such that each successive step is a cone-off of an incomplete subgraph of the previous graph.*

Proof. Let Γ be an incomplete CFS graph. Since cone vertices can be added last and not every vertex is a cone vertex, it suffices to assume that Γ has no cone vertex. Let C be a component of $\boxslash(\Gamma)$ with $\mathrm{supp}(C) = \Gamma$.

Pick an edge $\Delta_0 := \{a,b\} \bullet\!\!-\!\!\bullet \{c,d\}$ in C. By definition, this corresponds to an induced square $\Gamma_0 := \{a,b\} * \{c,d\}$ in Γ. If Γ is a square there is nothing more to prove, so assume Δ_0 is not all of C.

We inductively construct a sequence of connected subgraphs $\Delta_0 \subset \cdots \subset \Delta_i \subset C$ and induced CFS subgraphs $\Gamma_0 \subset \cdots \subset \Gamma_i \subset \Gamma$, such that each Γ_i is obtained from Γ_{i-1} by coning off an incomplete subgraph, and with the property that for each i one of the following is true:

- Δ_i is a full support connected subgraph of $\boxslash(\Gamma_i)$ (so Γ_i has no cone vertex).
- Γ_i has exactly one cone vertex, v, and Δ_i is a connected subgraph of $\boxslash(\Gamma_i)$ with support $\Gamma_i - v$.

Suppose Δ_i and Γ_i have been constructed. Suppose Δ_i is not all of C. Extend Δ_i to the largest connected subgraph Δ_i' of $\boxslash(\Gamma)$ that contains Δ_i and whose support is contained in Γ_i. Suppose Δ_i' is not all of C. We extend to Δ_{i+1} and Γ_{i+1} according to two (exhaustive) cases:

Case A: There is a vertex $\{v,w\}$ adjacent to some $\{e,f\} \in \Delta_i'$ in $\boxslash(\Gamma)$ such that $v \notin \Gamma_i$ and $w \in \Gamma_i$.

Case B: There is no such vertex as in Case A, but there is a vertex $\{v,w\}$ adjacent to some $\{e,f\} \in \Delta_i'$ in $\boxslash(\Gamma)$ such that $v \notin \Gamma_i$ and $w \notin \Gamma_i$.

In Case A, let $\Delta_{i+1} := \Delta_i' \cup \{v,w\}$, and let Γ_{i+1} be the induced subgraph of Γ containing Γ_i and v. We have $\Gamma_{i+1} = \Gamma_i *_{\mathrm{lk}_\Gamma(v)\cap\Gamma_i} v$. Since $\{e,f\} * \{v,w\}$ is an induced square in Γ, the subset $\mathrm{lk}_\Gamma(v) \cap \Gamma_i \subset \Gamma_i$ contains the non-adjacent vertices e and f, but does not contain w. So it is a proper, incomplete subgraph of Γ_i.

In Case B, supposing Γ_i has no cone vertex, we claim that both v and w are adjacent to every vertex of Γ_i and define $\Delta_{i+1} := \Delta_i'$ and $\Gamma_{i+1} := \Gamma_i * v$. By construction, Γ_i is not a clique.

To see the claim, suppose, to the contrary, that there is a vertex x of Γ_i that is not adjacent to v. Since x is not a cone vertex of Γ_i, there exist vertices of Δ_i' with support containing x. We may assume x and $\{x,y\} \in \Delta_i'$ have been chosen such that that $\{x,y\}$ is a closest vertex to $\{e,f\}$ in Δ_i' among those whose support contains a vertex not adjacent to v or w. Take a geodesic $\{e,f\} = \{e_0,f_0\}$, $\{e_1,f_1\}, \dots,$ $\{e_n,f_n\} = \{x,y\}$ in $\boxslash(\Gamma)$. The 'closest' hypothesis implies that for $j < n$ both of e_j and f_j are adjacent to both of v and w, so $\{v,w\} \bullet\!\!-\!\!\bullet \{e_j,f_j\}$ is an edge in $\boxslash(\Gamma)$. Thus, by replacing $\{e,f\}$ with $\{e_{n-1},f_{n-1}\}$, we may assume $\{x,y\}$ and $\{e,f\}$ are adjacent in $\boxslash(\Gamma)$. But this gives a join $\{v,w,x,y\} * \{e,f\}$ in Γ with v not adjacent to x, yielding an induced square $\{v,x\} * \{e,f\}$ in Γ, which is a contradiction, because in that case $\{v,x\}$ is available for Case A.

Notice that after performing a Case B extension, the vertex $\{v,w\}$ has $v \in \Gamma_{i+1}$ and $w \notin \Gamma_{i+1}$, so the subsequent inductive step will be a Case A extension. This justifies the supposition that Γ_i has no cone vertex when Case B is applied.

We now check that if the induction hypotheses are true up to stage i then they are true at $i+1$. Since the Γ_i are induced subgraphs of Γ, their diagonal graphs are subgraphs of $\boxslash(\Gamma)$.

Suppose Γ_i was obtained from Γ_{i-1} by a Case B extension. Since we never perform consecutive Case B extensions, the support of Δ_{i-1} is all of Γ_{i-1}, so the support of Δ_i is all of Γ_i except the cone vertex added in the last extension.

Suppose Γ_i is obtained by a Case A extension and Γ_{i-1} has no cone vertex. Let $\{v, w\}$ and $\{e, f\}$ be as in the description of Case A. Since Γ_{i-1} has no cone vertex, Δ_{i-1} has support all of Γ_{i-1}. Since Δ_i contains Δ_{i-1} and $\{v, w\}$, its support contains all vertices of Γ_i.

Finally, suppose Γ_{i-1} has a cone vertex. This occurs when Γ_{i-1} is constructed from Γ_{i-2} by a Case B extension. Let $\{v, w\}$ and $\{e, f\}$ be as in the description of Case B, so that $\Gamma_{i-1} = \Gamma_{i-2} * v$ and Γ_i is obtained from Γ_{i-1} by coning-off the proper subgraph $\Gamma_{i-2} \subset \Gamma_{i-1}$ to w. Since we never perform consecutive Case B extensions, the support of Δ_{i-2} contains all vertices of Γ_{i-2}. Since Δ_i contains Δ_{i-2} and $\{v, w\}$, its support is all vertices of Γ_i. □

Corollary 3.14. *An incomplete CFS graph with $n \geqslant 4$ vertices has at least $2n - 4$ edges.*

Proof. The claim is true for a square. Proceed by induction on the construction by cone-offs, each of which adds one vertex and at least two edges. □

Corollary 3.15. *A triangle-free CFS graph with $n \geqslant 4$ vertices has between $2n - 4$ and $\lfloor \frac{n^2}{4} \rfloor$ edges, and both extremes are realized.*

Proof. The minimum number of edges occurs for the suspension $K_{2,n-2}$, and the maximum occurs for either $K_{\frac{n}{2},\frac{n}{2}}$ or $K_{\frac{n-1}{2},\frac{n+1}{2}}$, according to the parity of n. □

Corollary 3.16. *The property of being a triangle-free CFS graph is constructible by coning from a square.*

Proof. If Γ is triangle-free then so are all subgraphs, so the proof of Proposition 3.13 produces a triangle-free graph at each stage. □

Example 3.17. The property of being a cone-vertex-free CFS graph is not constructible by coning from a square: The 1–skeleton of the octahedron is a CFS graph with no cone vertex, but removing any vertex leaves a cone on a square. ◇

Example 3.18. The property of being strongly CFS is not constructible by coning from a square: The graph $\Gamma := \texttt{circulant}(11, \{1, 3\}) = Cay(\mathbb{Z}_{11}, \{1, 3\})$ shown in Figure 3 is the smallest example of a triangle-free strongly CFS graph such that for every vertex v the graph $\Gamma - \{v\}$ is not strongly CFS. Thus, Γ cannot be built by coning from a square while remaining strongly CFS at each step. ◇

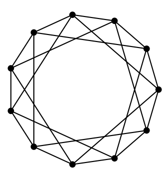

FIGURE 3. $\texttt{circulant}(11, \{1, 3\})$

The graph of Figure 3 turns out not to be RAAGedy. In fact, we do not know an example of a triangle-free graph without separating cliques that is RAAGedy but cannot be constructed by coning from a square through RAAGedy graphs.

Question 3.19. *Is the property of being triangle-free and RAAGedy constructible by coning from a square?*

4. Establishing that a graph is RAAGedy

In Section 4.1 we consider the possibility that a graph that is not a graph double might become a graph double after applying a sequence of link doubling operations. At the group level this corresponds to a sequence of passages to index-two subgroups, ending with one that is commensurable to a RAAG, so the original group is also commensurable to a RAAG. It turns out, Proposition 4.7, that if such a sequence exists, then at most two link doubling steps are necessary, and the existence of the sequence and identification of which vertices to double over the links of are recognizable in the presentation graph.

In Section 4.2 and Section 4.3 we introduce two new operations, *cloning* and *unfolding*, that change a graph Γ without changing the quasiisometry type of W_Γ. This gives us three such operations, the third being link doubling. Unlike link doubling, for cloning and unfolding we only know that they produce quasiisometric groups, not whether the resulting group is commensurable to the one we started with. In Section 4.4 we explore the connections within our enumeration of small CFS graphs given by these three graph modification operations.

4.1. Near doubles.

Recall the graph constructions of doubling, link doubling, and star doubling of Section 2.1 and that A_Δ is commensurable to $W_{\mathfrak{D}(\Delta)}$, which is Theorem 2.4. The proof of the following is elementary and is left to the reader. We will not use the lemma directly, but its interpretation as a statement that the property of being a graph double is stable under link doubling inspired Definition 4.2.

Lemma 4.1. $\mathfrak{D}^\circ_{(v,1)}(\mathfrak{D}(\Gamma)) \cong \mathfrak{D}(\mathfrak{D}^*_v(\Gamma))$*, that is, they are isomorphic graphs.*

Definition 4.2. A *near double* is a graph Γ such that there exists a sequence of link doubles of Γ such that the result is isomorphic to some $\mathfrak{D}(\Delta)$.

Proposition 4.3. *If Γ is a near double then W_Γ is commensurable to a RAAG.*

Proof. Apply Lemma 2.3 and Theorem 2.4. □

We first give a characterization of graph doubles, which then enables us to give an example of a near double that is not a double. After that we characterize near doubles in Proposition 4.7. It turns out that the length of the necessary sequence of link doubles is bounded by 2, but Proposition 4.7 is a more concrete description than just testing all such bounded sequences.

Recall the definition of a twin module in Section 2.1.1.

Proposition 4.4. *Let Γ be a graph. The following are equivalent:*

(1) Γ is a double, ie, there exists Δ such that $\Gamma \cong \mathfrak{D}(\Delta)$.
(2) Every twin module in Γ has even order.
(3) Γ admits a fixed-point-free involution that fixes each twin module.

Proof. (1) $\implies$ (2), since vertices of a double come in pairs with equal links.

(2) $\implies$ (3) by choosing a pairing of the vertices within each twin module. The map that exchanges the vertices of each such pair satisfies (3).

(3) $\implies$ (1) by defining a graph Δ by taking a vertex for each orbit of the involution ϕ, and connecting two vertices v and w by an edge if the twin module of v is adjacent to the twin module of w in the twin graph $\amalg(\Gamma)$. By construction, there is a bijection between vertices of $\mathfrak{D}(\Delta)$ and vertices of Γ; it is $(v,0) \mapsto v$ and $(v,1) \mapsto \phi(v)$. Every edge of $\mathfrak{D}(\Delta)$ belongs to an induced square of the form $\{(v,0),(v,1)\} * \{(w,0),(w,1)\}$, where $v \bullet\!\!-\!\!\bullet w$ is an edge of Δ, and corresponds to an induced square $\{v,\phi(v)\} * \{w,\phi(w)\}$ of Γ. Conversely, every edge of Γ belongs to such an induced square, since twin modules are anticliques and $v \bullet\!\!-\!\!\bullet w$ implies $v \bullet\!\!-\!\!\bullet \phi(w)$, since w and $\phi(w)$ are twins. □

Example 4.5. The graph Γ on the left in Figure 4 is not a double; vertices 6 and 7 have no twin. Two link doubles turn Γ into a doubled octagon, as shown in Figure 4. We compactify notation by writing (v,α) as v_α, (v,α,β) as $v_{\alpha\beta}$, etc. $\diamond$

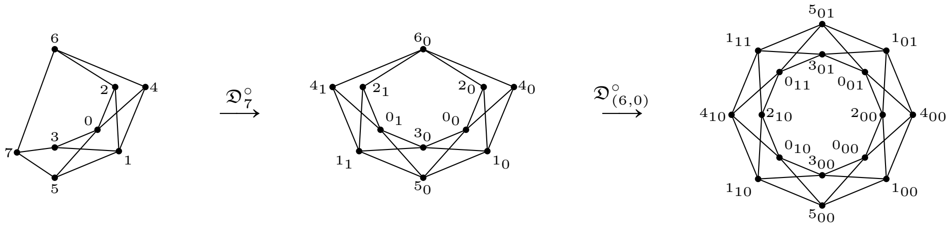

FIGURE 4. First link doubling example.

In Proposition 4.7 below, we describe twin modules as *even/odd* according to their orders. We refer to the *link* $\mathrm{lk}(M)$ of a twin module M, which we take to mean $\mathrm{lk}(M) := \mathrm{lk}_{\amalg(\Gamma)}(M)$, that is, the link of the vertex M in the twin graph $\amalg(\Gamma)$ (recall Definition 2.1). Let $M_\Gamma(v)$ denote the twin module of Γ containing v.

Lemma 4.6. *Let Γ be an incomplete triangle-free graph without separating cliques. For $v \in \Gamma$, the twin modules of $\mathfrak{D}^\circ_v(\Gamma)$ are as follows:*

- *If $A = M_\Gamma(v)$ and $|A| > 1$ then there is a single twin module corresponding to A in $\mathfrak{D}^\circ_v(\Gamma)$ of size $2(|A|-1)$, given by $(A - \{v\}) \times \{0,1\}$.*
- *If A is a satellite of $M_\Gamma(v)$ in $\amalg(\Gamma)$ then there is a single twin module corresponding to A in $\mathfrak{D}^\circ_v(\Gamma)$ of size $2|A|$, given by $A \times \{0,1\}$.*
- *If $A \in \mathrm{lk}_{\amalg(\Gamma)} M_\Gamma(v)$ then there is a single twin module corresponding to A in $\mathfrak{D}^\circ_v(\Gamma)$ of size $|A|$, given by $A \times \{0\} = A \times \{1\}$.*
- *Otherwise, there are two distinct twin modules of size $|A|$ corresponding to A in $\mathfrak{D}^\circ_v(\Gamma)$, given by $A \times \{0\}$ and $A \times \{1\}$.*

Proof. The link of every vertex of Γ is an anticlique of size at least 2.

If $w \in M_\Gamma(v) - \{v\}$, i.e., $w \neq v$ and $\mathrm{lk}_\Gamma(v) = \mathrm{lk}_\Gamma(w)$, then $w \notin \mathrm{lk}_\Gamma(v)$, so w has distinct preimages $(w,0)$ and $(w,1)$ in $\mathfrak{D}^\circ_v(\Gamma)$, but they have the same link, $\mathrm{lk}_{\mathfrak{D}^\circ_v(\Gamma)}(w,0) = \mathrm{lk}_{\mathfrak{D}^\circ_v(\Gamma)}(w,1) = \mathrm{lk}_\Gamma(v) \times \{0\}$, so the twin module $M_{\mathfrak{D}^\circ_v(\Gamma)}(w,0)$ contains 2 copies of each vertex in $M_\Gamma(v) - \{v\}$.

If $\mathrm{lk}_\Gamma(w) \subsetneq \mathrm{lk}_\Gamma(v)$, then $w \notin \mathrm{lk}_\Gamma(v)$, but all of its neighbors are, so there are two copies $(w,0)$ and $(w,1)$ of w whose links are both $\mathrm{lk}_\Gamma(w) \times \{0\}$. Suppose u is another vertex of $M_\Gamma(w)$, so $(u,0)$ has the same link as $(w,0)$ in $\mathfrak{D}^\circ_v(\Gamma)$. Then $u \neq v$, since $\mathrm{lk}_\Gamma(w) \subsetneq \mathrm{lk}_\Gamma(v)$ is proper containment, and $u \notin \mathrm{lk}_\Gamma(v)$ by our hypotheses on Γ, since if it were then u would be adjacent only to v, or to something also in $\mathrm{lk}_\Gamma(v)$, contradicting either that Γ is connected without cut vertices, or that it is triangle-free. Therefore, $u \notin \mathrm{lk}_\Gamma(v)$, so u contributes vertices $(u,0)$ and $(u,1)$ to $M_{\mathfrak{D}^\circ_v(\Gamma)}(w,0)$. Thus, $M_{\mathfrak{D}^\circ_v(\Gamma)}(w,0)$ contains 2 copies of each vertex in $M_\Gamma(w)$.

If $w \in \mathrm{lk}_\Gamma(v)$ then there is only one vertex $(w,0)$ in the preimage of w. The fact that Γ is triangle-free and connected without cut vertices implies that $\mathrm{lk}_\Gamma(w)$ contains only v and a nonempty set of vertices from outside $\mathrm{lk}_\Gamma(v)$. Choose $u \in \mathrm{lk}(w) - \{v\}$, so $(w,0)$ is adjacent to a distinct pair $(u,0)$ and $(u,1)$. Now, a vertex $(x,0)$ in $\mathfrak{D}^\circ_v(\Gamma)$ is adjacent to $(u,0)$ if and only if x and u are adjacent in Γ, but $(x,0)$ and $(u,1)$ are adjacent if and only if x and u are adjacent in Γ and $x \in \mathrm{lk}_\Gamma(v)$. So, the vertices adjacent to both $(u,0)$ and $(u,1)$ are $(\mathrm{lk}_\Gamma(u) \cap \mathrm{lk}_\Gamma(v)) \times \{0\}$. Thus, vertices with the same link as $(w,0)$ are of the form $(x,0)$ with $x \in \mathrm{lk}_\Gamma(v)$, where $\mathrm{lk}_\Gamma(w) = \mathrm{lk}_\Gamma(x)$. That is, $M_{\mathfrak{D}^\circ_v(\Gamma)}(w,0) = M_\Gamma(w) \times \{0\}$.

The remaining case is that $w \notin \mathrm{lk}_\Gamma(v)$ and w is adjacent to at least one vertex $x \notin \mathrm{lk}_\Gamma(v)$. Then $(w,0)$ is adjacent to $(x,0)$ but not to $(x,1)$. Suppose $\mathrm{lk}_{\mathfrak{D}_v^\circ(\Gamma)}(w,0) = \mathrm{lk}_{\mathfrak{D}_v^\circ(\Gamma)}(u,\epsilon)$. This implies $\mathrm{lk}_\Gamma(w) = \mathrm{lk}_\Gamma(u) - \{v\}$. However, it also implies (u,ϵ) is adjacent to $(x,0)$ but not $(x,1)$, which does not happen if $u \in \mathrm{lk}_\Gamma(v)$, so, actually, $\mathrm{lk}_\Gamma(w) = \mathrm{lk}_\Gamma(u)$. Furthermore, $\epsilon = 0$, and $(u,1)$ is adjacent to $(x,1)$ but not $(x,0)$, so $(u,1) \notin M_{\mathfrak{D}_v^\circ(\Gamma)}(w,0)$; we have that $M_{\mathfrak{D}_v^\circ(\Gamma)}(w,0) = M_\Gamma(w) \times \{0\}$ and $M_{\mathfrak{D}_v^\circ(\Gamma)}(w,1) = M_\Gamma(w) \times \{1\}$ are distinct twin modules. □

Proposition 4.7 (Recognizing near doubles)**.** *A triangle-free graph Γ without separating cliques is a near double if and only if one of the following is true:*

- *(0) There are no odd twin modules.*
- *(1) There exists a twin module A of Γ such that every odd twin module of Γ is either A or a satellite of A in $\amalg(\Gamma)$.*
- *(2) There exist twin modules A and B of Γ such that A and B are adjacent in $\amalg(\Gamma)$ and such that every odd twin module of Γ is either A or B or a satellite of A or B in $\amalg(\Gamma)$.*

Proof. Clearly $(0) \implies (1) \implies (2)$, but the conditions provide case distinctions.

If Γ is complete then it is either empty, a single vertex, or a single edge, and it is reduced to the empty graph by $|\Gamma|$–many link doublings. The empty graph is the double of itself, so the proposition is true in this case. Now assume Γ is incomplete so that we can apply Lemma 4.6.

In case (2), pick $v \in A$ and consider $\mathfrak{D}_v^\circ(\Gamma)$. From Lemma 4.6 we see that the only possible odd modules of $\mathfrak{D}_v^\circ(\Gamma)$ come from odd modules of Γ that are not A or one of its satellites in $\amalg(\Gamma)$. By (2) the only possible choices are B or a satellite of B. Since B is adjacent to A in $\amalg(\Gamma)$ there is a unique module $B \times \{0\} = B \times \{1\}$ in $\mathfrak{D}_v^\circ(\Gamma)$ corresponding to B. If C is a satellite of B in $\amalg(\Gamma)$ then for every $c \in C$ and every $b \in B$ we have $\mathrm{lk}(c) \subset \mathrm{lk}(b)$, so for every $\ell \in \mathrm{lk}(c)$ there are edges $c \multimap \ell \multimap b$ in Γ. If $\ell = v$ then it does not appear in $\mathfrak{D}_v^\circ(\Gamma)$. Otherwise, $\mathfrak{D}_v^\circ(\Gamma)$ contains edges:

$$(c,0) \multimap (\ell,0) \multimap (b,0) = (b,1) \multimap (\ell,1) \multimap (c,1)$$

Thus any module of $\mathfrak{D}_v^\circ(\Gamma)$ corresponding to a satellite of B in $\amalg(\Gamma)$ is a satellite of $B \times \{0\}$ in $\amalg(\mathfrak{D}_v^\circ(\Gamma))$. In particular, every odd module of $\mathfrak{D}_v^\circ(\Gamma)$ is either $B \times \{0\}$ or a satellite of $B \times \{0\}$, so $\mathfrak{D}_v^\circ(\Gamma)$ belongs to case (1).

In case (1), pick $v \in A$, and consider $\mathfrak{D}_v^\circ(\Gamma)$. Again, by Lemma 4.6 the possible odd modules of $\mathfrak{D}_v^\circ(\Gamma)$ come from odd modules of Γ that are not A or one of its satellites in $\amalg(\Gamma)$. But (1) says there are none of these, so $\mathfrak{D}_v^\circ(\Gamma)$ belongs to case (0)

In case (0) there are no odd modules, so Γ is a double by Proposition 4.4.

In the other direction we show that if there exists $v \in \Gamma$ such that (2) is true for $\mathfrak{D}_v^\circ(\Gamma)$ then (2) was already true in Γ. Thus, by induction on the link doubling sequence, if (2) is false in Γ then it remains false in any iterated link double of Γ. Therefore, every iterated link double of Γ contains odd twin modules, so is not a graph double, so Γ is not a near double.

Suppose there is $v \in \Gamma$ such that (2) is true for $\mathfrak{D}_v^\circ(\Gamma)$. Let $\sigma\colon\ \mathfrak{D}_v^\circ(\Gamma) \to \mathfrak{D}_v^\circ(\Gamma)$ be the involution that fixes the first coordinate and exchanges 0 and 1 in the second coordinate. Let $\pi\colon\ \mathfrak{D}_v^\circ(\Gamma) \to \Gamma$ be projection to the first coordinate. By construction of $\mathfrak{D}_v^\circ(\Gamma)$, the map π sends edges to edges. It follows that π sends twin modules to twin modules, so it induces $\pi\colon\ \amalg(\mathfrak{D}_v^\circ(\Gamma)) \to \amalg(\Gamma)$ that sends vertices to vertices and edges to edges.

Furthermore, π preserves the satellite relationship, as follows. Suppose B is a satellite of A in $\amalg(\mathfrak{D}_v^\circ(\Gamma))$. Then $\mathrm{lk}(B) \subset \mathrm{lk}(A)$, so $\pi(\mathrm{lk}(B))) \subset \pi(\mathrm{lk}(A)) \subset \mathrm{lk}(\pi(A))$. However, it is possible that $\pi(\mathrm{lk}(B)) \subsetneq \mathrm{lk}(\pi(B))$, which happens when $\{v\} = M_\Gamma(v) \in \mathrm{lk}(\pi(B))$, so that $\pi^{-1}(M_\Gamma(v)) = \varnothing$. To show $\pi(B)$ is a satellite of $\pi(A)$

we must rule out the possibility that $\{v\} = M_\Gamma(v) \in \mathrm{lk}(\pi(B)) - \mathrm{lk}(\pi(A))$. Suppose this is the case. Clearly $\pi(A) \notin \mathrm{st}(M_\Gamma(v))$. We cannot have $\pi(A)$ as a satellite of $M_\Gamma(v)$, because that would mean that $\mathrm{lk}(\pi(B)) - \{v\} \subset \mathrm{lk}(\pi(A)) \subset \mathrm{lk}(M_\Gamma(v))$, but then triangle-freeness implies $\mathrm{lk}(\pi(B)) = \{v\}$, which would make v a cut vertex of Γ, which is a contradiction. By Lemma 4.6, the remaining option is that A and $\sigma(A)$ are distinct. The condition $\pi(B) \in \mathrm{lk}(M_\Gamma(v))$ implies σ fixes the vertices in B. For any $C \in \mathrm{lk}(B) \subset \mathrm{lk}(A)$ we have $\pi(C) \notin \mathrm{st}(M_\Gamma(v))$ by triangle-freeness, so for all $a \in A$, $b \in B$, and $c \in C$ there are segments $a \leftrightarrow c \leftrightarrow b$ and $\sigma(b) \leftrightarrow \sigma(c) \leftrightarrow \sigma(a)$, with $b = \sigma(b)$, $\sigma(c) \neq c$, and $\sigma(a) \neq a$. By construction of $\mathfrak{D}^\circ_v(\Gamma)$, if $a \neq \sigma(a)$ and $c \neq \sigma(c)$ and there exist edges $a \leftrightarrow c$ and $\sigma(a) \leftrightarrow \sigma(c)$ then there do not exist edges $a \leftrightarrow \sigma(c)$ and $c \leftrightarrow \sigma(a)$. This contradicts $\sigma(c) \in \mathrm{lk}(b) \subset \mathrm{lk}(a)$. Thus, $\pi(B)$ is a satellite of $\pi(A)$.

Now we argue that (2) for $\mathfrak{D}^\circ_v(\Gamma)$ implies (2) for Γ. Since (2) is true for $\mathfrak{D}^\circ_v(\Gamma)$, there are modules A and B of $\mathfrak{D}^\circ_v(\Gamma)$ that are neighbors in $\amalg(\mathfrak{D}^\circ_v(\Gamma))$, and such that every odd module of $\mathfrak{D}^\circ_v(\Gamma)$ is either A or B or a satellite of A or B in $\amalg(\mathfrak{D}^\circ_v(\Gamma))$. Since $\pi\colon \amalg(\mathfrak{D}^\circ_v(\Gamma)) \to \amalg(\Gamma)$ preserves satellites, the potential odd modules of Γ are, according to Lemma 4.6, either $M_\Gamma(v)$ or one of its satellites, or the projection of an odd module of $\mathfrak{D}^\circ_v(\Gamma)$, which by hypothesis must be $\pi(A)$ or $\pi(B)$ or a satellite of one of these two.

Since A and B are adjacent, so are $\pi(A)$ and $\pi(B)$. By triangle-freeness, they cannot both be in $\mathrm{lk}(M_\Gamma(v))$. Suppose $\pi(B)$ is not in $\mathrm{lk}(M_\Gamma(v))$. The alternatives are $\pi(B) = M_\Gamma(v)$, $\pi(B)$ is a satellite of $M_\Gamma(v)$, or $\pi(B)$ is neither in $\mathrm{st}(M_\Gamma(v))$ nor a satellite of $M_\Gamma(v)$.

If $\pi(B) = M_\Gamma(v)$ or $\pi(B)$ is a satellite of $M_\Gamma(v)$ then since A and B are adjacent, $\pi(A) \in \mathrm{lk}(M_\Gamma(v))$. If C is a satellite of B then $\mathrm{lk}(\pi(C)) \subset \mathrm{lk}(\pi(B)) \subset \mathrm{lk}(M_\Gamma(v))$, so $\pi(C)$ is a satellite of $M_\Gamma(v)$, so every odd module of Γ is either $M_\Gamma(v)$ or $\pi(A)$ or a satellite of one of these. Since $\pi(A) \in \mathrm{lk}(M_\Gamma(v))$, vertices $\pi(A)$ and $M_\Gamma(v)$ are adjacent in $\amalg(\Gamma)$ so (2) is true for Γ.

Suppose $\pi(B)$ is neither in $\mathrm{st}(M_\Gamma(v))$ nor a satellite of $M_\Gamma(v)$. Then $\pi(B) \times \{0\} \neq \pi(B) \times \{1\}$ are distinct and symmetric, so we may assume $B = \pi(B) \times \{0\}$. We will show that B and its satellites are even. Suppose C is an odd satellite of B. Since C is odd, $\pi(C)$ is not $M_\Gamma(v)$ or a satellite of $M_\Gamma(v)$, by Lemma 4.6. Since C is a satellite of B and π preserves satellites, $M_\Gamma(v) \notin \mathrm{lk}(\pi(B)) \supset \mathrm{lk}(\pi(C))$, so $\pi(C) \times \{0\} \neq \pi(C) \times \{1\}$ are distinct odd modules. Triangle-freeness implies that $\pi(B) \times \{0\}$, $\pi(B) \times \{1\}$, and satellites of either of these two are distinct from A and satellites of A. So, by property (2) for $\mathfrak{D}^\circ_v(\Gamma)$, both of $\pi(C) \times \{0\}$ and $\pi(C) \times \{1\}$ are satellites of B. However, for $\pi(C) \times \{1\}$ to be a satellite of $B = \pi(B) \times \{0\}$ implies $\mathrm{lk}(\pi(C) \times \{1\}) \subset \pi^{-1}(\mathrm{lk}(M_\Gamma(v)))$, so $\pi(C)$ is a satellite of $M_\Gamma(v)$ in $\amalg(\Gamma)$, a contradiction, as we have already ruled this out. Similarly, if B is odd then $\pi(B) \times \{1\}$ is a satellite of $B = \pi(B) \times \{0\}$, so the same argument implies $\pi(B)$ is a satellite of $M_\Gamma(v)$ in $\amalg(\Gamma)$, which contradicts that B is odd. Thus, B and its satellites are even.

The same arguments apply after swapping the roles of A and B.

In conclusion, if one of $\pi(A)$ or $\pi(B)$ is equal to or a satellite of $M_\Gamma(v)$ then (2) is true for Γ and we are done. Otherwise, since $\pi(A)$ and $\pi(B)$ cannot both be in $\mathrm{lk}(M_\Gamma(v))$, we may assume that $\pi(B)$ is not, which implies that B and its satellites are even. If $\pi(A)$ is also not in $\mathrm{lk}(M_\Gamma(v))$ then A and its satellites are also even, leaving only $M_\Gamma(v)$ and its satellites as possible odd modules of Γ, so (1) is true for Γ. On the other hand, if $\pi(A) \in \mathrm{lk}(M_\Gamma(v))$ then the potential odd modules of Γ are $M_\Gamma(v)$, $\pi(A)$, or a satellite of one of these adjacent vertices, so (2) is true for Γ. □

4.2. **Cloning.**

Definition 4.8. Let v be a vertex of a graph Γ. To *clone* v means to add a new vertex v' and connect it by an edge to each vertex of $\mathrm{lk}(v)$ to make a new graph $\Gamma' := \Gamma *_{\mathrm{lk}(v)} v'$ in which v and v' are twins.

Definition 4.9. A vertex of a graph is a *singleton* if it has no twins. A vertex is *clonable* if it is a satellite of at least two other vertices, and *unclonable* otherwise.

Proposition 4.10. *Let v be a clonable vertex of a connected, triangle-free graph Γ. Cloning v produces a quasiisometric RACG: the Davis complex Σ_Γ is quasiisometric to $\Sigma_{\Gamma *_{\mathrm{lk}(v)} v'}$ by a quasiisometry that restricts to a color preserving isomorphism on each translate of $\Sigma_{\Gamma-\{v\}}$.*

The proof of Proposition 4.10 has two parts: we define base quasiisometries in Lemma 4.11 that are designed to facilitate an application of Proposition 2.10.

Lemma 4.11. *For every $n \geqslant 2$ there is a quasiisometry ϕ between the $(n+1)$–valent tree T_{n+1} with edges colored a, b, c, $x_1,\dots,x_{n-2}$ (with exactly one edge of each color at each vertex) and the $(n+2)$–valent tree T_{n+2} with edges colored a, b, c, d, $x_1,\dots,x_{n-2}$ with the following properties:*

- *ϕ is bijective on vertices.*
- *ϕ induces a bijection between the set of S–colored components of T_{n+1} and the set of S–colored components of T_{n+2}, where $S := \{a, b, x_1, \dots, x_{n-2}\}$ is the set of 'static' colors.*
- *ϕ restricts to a color-preserving isomorphism on each S–colored component.*

In the proofs of Lemma 4.11 and Lemma 4.22 we use without further explanation the following well-known construction. If p, q, r are vertices of a tree such that $q \in \mathrm{lk}(r)$ is between p and r then the map that replaces the edge $q \leftrightarrow r$ with an edge $p \leftrightarrow r$ is $(1 + d(p, q))$–biLipschitz. To see this, consider a geodesic $v_0, \dots, v_n$ in the original tree structure such that $v_m = q$ and $v_{m+1} = r$, for some $0 \leqslant m < n$. In the new tree structure there is a path from v_0 to v_n given by $v_0, \dots, v_m$ followed by a geodesic from $v_m = q$ to p, then the new edge $p \leftrightarrow r = v_{m+1}$, then continuing with $v_{m+2},\dots,v_n$. Thus, the distance from v_0 to v_n increased by at most an additive factor of $d(p, q)$. Since vertex distances are integral, distances increase by a multiplicative factor of at most $1 + d(p, q)$. The other inequality can be proved by making the same argument for the inverse map.

In the literature this map is sometimes described as an 'edge-slide' move, where it is imagined that the target of the initial incidence map for $q \leftrightarrow r$ continuously 'slides' along a specified edge path, which in this case is the geodesic from q to p. More generally, we can make infinitely many such edge-slides simultaneously without compounding the biLipschitz constant, provided that no edge-to-be-slid is contained in the slide path of another.

Proof of Lemma 4.11. Pick a base vertex w_0 in T_{n+1}. Given a vertex v we speak of 'incoming' and 'outgoing' edges with respect to the chosen basepoint; ie, an outgoing edge at v leads farther from w_0. Since we are working in a tree, for every vertex v and every word in letters a, b, c, $x_1,\dots$ there is a unique edge path in the tree starting from v whose edges are colored by successive letters of the given word. Words without repeated successive letters are geodesics, and if we identify w_0 with the empty word, then every vertex can be described uniquely by a geodesic word.

Each S–colored component has a unique vertex w_i closest to w_0. We assume that the vertices w_i are ordered by increasing distance from w_0; that is, if $i < j$, then the distance from w_j to w_0 is greater than or equal to the distance from w_i to w_0. For phase 1 of the construction, consider the following map that fixes all of the vertices and rearranges the c–colored edges. All of the vertices $w_1, w_2, \dots$ have an incoming c–edge, and every other vertex has an outgoing c–edge. Working inductively with increasing distance from w_0, for every vertex v with an outgoing c–edge, remove the c–edge between vca and $vcac$ and connect v and $vcac$ by a c–edge. This map is 3–biLipschitz on the vertex set. In the new tree structure, $w_1, w_2, \dots$ still have one incoming c–edge, the vertices $w_1a, w_2a, \dots$ have either zero or two c–edges, and every other vertex has two outgoing c–edges.

For phase 2 of the construction, for each $i \geqslant 1$ consider the two geodesic rays $w_i, w_ib, w_iba, w_ibab, \dots$ and $w_i, w_ia, w_iab, w_iaba, \dots$ based at w_i. On the first ray, let every vertex after w_i pass one of its c–edges to its predecessor. On the second ray, if w_ia has two c–edges, do nothing, and if it has no c–edges, then let every vertex after w_ia pass both of its c–edges its predecessor. Thus, w_i receives one additional c–edge and passes none, w_ia already has or receives two c–edges and passes none, and every other vertex receives the same number of c–edges from its successor as it passes to its predecessor. Hence, every vertex now has two incident c–edges. The phase 2 map is 2–biLipschitz on the vertex set.

The composition of the two phases is a 6–biLipschitz map on the vertex set taking T_{n+1} to an $(n+2)$–valent tree such that every vertex in the image has exactly two incident c–edges and one of each color from S. By construction, the map is a color preserving isometry on each S–colored component. The $\{c\}$–colored components of the image are biinfinite geodesics. On each such geodesic, recolor the edges so that they alternate c and d. The result is isomorphic to T_{n+2} as colored trees. $\square$

Proof of Proposition 4.10. Since Γ is connected, $\varnothing \neq L := \mathrm{lk}(v)$. Since v is clonable, $C := \{c \in \Gamma - \{v\} \mid L \subset \mathrm{lk}(c)\}$ contains at least two vertices. Since Γ is triangle-free and $L \neq \varnothing$, $C \cup \{v\}$ is an anticlique. Since $|C| \geqslant 2$, there is a quasiisometry $\psi\colon \Sigma_{C\cup\{v\}} \to \Sigma_{C\cup\{v,v'\}}$ given by Lemma 4.11 that restricts to color preserving isomorphisms on each copy of Σ_C. We may assume, up to pre-composing by translation in $W_{C\cup\{v\}}$, that ψ fixes Σ_C. Crossing with the identity on Σ_L gives a quasiisometry $\mathrm{Id}_{\Sigma_L} \times \psi$ from $\Sigma_L \times \Sigma_{C\cup\{v\}} = \Sigma_{L\cup C\cup\{v\}}$ to $\Sigma_L \times \Sigma_{C\cup\{v,v'\}} = \Sigma_{L\cup C\cup\{v,v'\}}$ that preserves edge colors for all edges colored by $L \cup C$. In particular, the map matches up translates of $\Sigma_{L\cup C}$ bijectively, and restricts to a color-preserving isomorphism on each such translate.

If the vertex set of Γ is $\{v\} \cup C \cup L$, we are done. Otherwise, we can write W_Γ and $W_{\Gamma *_L v'}$ as amalgams:

$$W_\Gamma = W_{\Gamma-\{v\}} *_{W_L\times W_C} (W_L \times (W_C * W_{\{v\}}))$$
$$W_{\Gamma *_L v'} = W_{\Gamma-\{v\}} *_{W_L\times W_C} (W_L \times (W_C * W_{\{v,v'\}}))$$

Apply Proposition 2.10 with base maps $\mathrm{Id}_{W_{\Gamma-\{v\}}}$ and $\mathrm{Id}_{\Sigma_L} \times \psi$ on the vertex groups of the given splittings. We check that the hypotheses of Proposition 2.10 are satisfied for $\mathrm{Id}_{\Sigma_L} \times \psi$. The check for $\mathrm{Id}_{W_{\Gamma-\{v\}}}$ is similar. The construction has been arranged so that Conditions (1) and (2) Proposition 2.10 are satisfied. For (3), note that $W_{L\cup C}$ acts by color preserving isomorphisms of $\Sigma_{L\cup C}$, freely and transitively on vertices. Thus, if $(\mathrm{Id}_{\Sigma_L} \times \psi)((1,g)\Sigma_{L\cup C}) = (1,g')\Sigma_{L\cup C}$ then let $h' = (\mathrm{Id}_{\Sigma_L} \times \psi)((1,g)1) \in (1,g')W_{L\cup C}$ so that $\mathrm{Id}_{\Sigma_L} \times \psi$ and $h'(\mathrm{Id}_{\Sigma_L} \times \psi)(1,g)^{-1}$ are color preserving isomorphisms of $(1,g)\Sigma_{L\cup C}$ that agree on one vertex, hence on all. $\square$

Proposition 4.12. *Let Γ be a triangle-free graph without separating cliques. Then W_Γ is quasiisometric to $W_{\Gamma'}$, where Γ' is the blow-up graph $\amalg(\Gamma)^{\omega'}$ for the weight function ω' derived from the weight function $\omega(M) := |M|$ of $\amalg(\Gamma)$ as follows:*

- *$\omega'(M) = 1$ if M is an unclonable singleton.*
- *$\omega'(M) = 2$ if M:*
  - *is a clonable singleton, or*

  – $\omega(M) = 2$, *or*
  – $\omega(M) > 2$ *and* M *is a satellite in* $\amalg(\Gamma)$.
- $\omega'(M) = 4$ *otherwise.*

Proof. We will change Γ into Γ' by cloning or retiring clones in such a way that W_Γ and W'_Γ are quasiisometric by Proposition 4.10.

If M is an unclonable singleton or if $|M| = 2$, then no change is necessary.

If $M = \{v\}$ is a clonable singleton then cloning v creates a graph in which v belongs to a twin module of size 2.

If M is a twin module of Γ containing a vertex v where $|M| > 2$ and M is a satellite of N in $\amalg(\Gamma)$ then for any vertex $w \in N$, v is a satellite of w and has at least two twins v' and v''. Delete v''. In the resulting graph v is clonable, since it is a satellite of v' and w. Cloning v is therefore a quasiisometry on the level of Coxeter groups, and results in a graph isomorphic to Γ, so deleting v'' induces the inverse quasiisometry on the level of Coxeter groups. Iterating this process, we can arrange $\omega'(M) = 2$. Similarly, if M is not a satellite but $|M| > 4$ then deleting a vertex from M induces a quasiisometry of Coxeter groups, so we can arrange $\omega'(M) = 4$.

Finally, if $|M| = 3$ but M is not a satellite in $\amalg(\Gamma)$ then we can clone a vertex of M without changing the quasiisometry type of W_Γ, so we make $\omega'(M) = 4$. □

Corollary 4.13. *If* Γ_i *and* $(\amalg(\Gamma_i), \omega'_i)$ *for* $i = 1, 2$ *are as in Proposition 4.12 and* $(\amalg(\Gamma_1), \omega'_1)$ *and* $(\amalg(\Gamma_2), \omega'_2)$ *are isomorphic as weighted graphs then* W_{Γ_1} *and* W_{Γ_2} *are quasiisometric.*

Corollary 4.14. *A triangle-free graph with no separating cliques and no unclonable singletons is RAAGedy.*

Proof. Let $\Gamma' := \amalg(\Gamma)^{\omega'}$ as in Proposition 4.12, so that W_Γ is quasiisometric to $W_{\Gamma'}$ and ω' does not take the value 1. Since ω' takes only even values, Γ' is isomorphic to the graph double of the blow-up graph $\amalg(\Gamma)^{\omega'/2}$. Apply Theorem 2.4. □

As an application, we upgrade the conclusion of Proposition 3.12:

Corollary 4.15. *Let* (Δ, ω) *be a weighted graph such that* Δ *is connected, has more than one vertex, is triangle-free and twin-free, and such that* ω *takes the value 1 only on leaves of* Δ*, and if* Δ *is a single edge then* ω *does not take the value 1. Then the blow-up graph* Δ^ω *is triangle-free without separating cliques and is RAAGedy.*

Proof. Apply Corollary 4.14. This requires showing that Δ^ω is triangle-free without separating cliques or unclonable singletons. By Proposition 3.12, Δ^ω is triangle-free and strongly CFS, and the latter implies there are no separating cliques.

Singletons of Δ^ω come from vertices of Δ with weight 1. By hypothesis, any such vertex is a leaf of Δ. Suppose u is a leaf in Δ with weight 1 and v is the vertex of Δ adjacent to u. By hypothesis, $u \bullet\!\!-\!\!\bullet v$ is not all of Δ. Since Δ is connected, v has some neighbor $w \neq u$. The vertex w is not a leaf, since if it were then u and w would be twins, but Δ is twin-free. Hence $\omega(w) > 1$. In Δ^ω, $(u, 0)$ is a satellite of all the vertices $(w, 0)$, …, $(w, \omega(w) - 1)$, of which there are at least two, so $(u, 0)$ is clonable. Thus, Δ^ω has no unclonable singletons. □

If we assume the Δ is twin-free then a blow-up graph $\Gamma = \Delta^\omega$ is a graph double precisely when ω takes only even values. In [16] it is shown that when Γ contains an induced cycle of length greater than 6 with no 2–chord, then Γ does not admit a FIDL–Λ. Thus, by taking Δ to be twin-free containing a long cycle without 2–chords and taking ω to be uneven we get many examples of blow-up graphs Δ^ω to which neither the Davis-Januszkiewicz nor Dani-Levcovitz conditions apply, but that are RAAGedy by Corollary 4.15.

Finally, we give an analogue of the 'near-double' construction to describe when passing to link doubles can eliminate unclonable singletons:

Theorem 4.16. *A triangle-free graph Γ without separating cliques is RAAGedy if any of the following are true:*

- *There are no unclonable singletons.*
- *There is a vertex v such that the set of unclonable singletons is contained in the set consisting of v and its satellites.*
- *There are adjacent vertices v and w such that the set of unclonable singletons is contained in the set consisting of v and w and their satellites.*

Definition 4.17. We call a graph satisfying Theorem 4.16 a *coarse near double*.

Proof. As in Proposition 4.12, without changing the quasiisometry type of W_Γ we may replace Γ by a graph Γ' in which all twin modules are even except those coming from unclonable singletons of Γ. The hypotheses control the relative arrangement of those unclonable singletons and match the hypothesis of Proposition 4.7, so Γ' is a near double. Thus, W_Γ is quasiisometric to $W_{\Gamma'}$, by Proposition 4.10, and $W_{\Gamma'}$ is commensurable to a RAAG, by Proposition 4.3 □

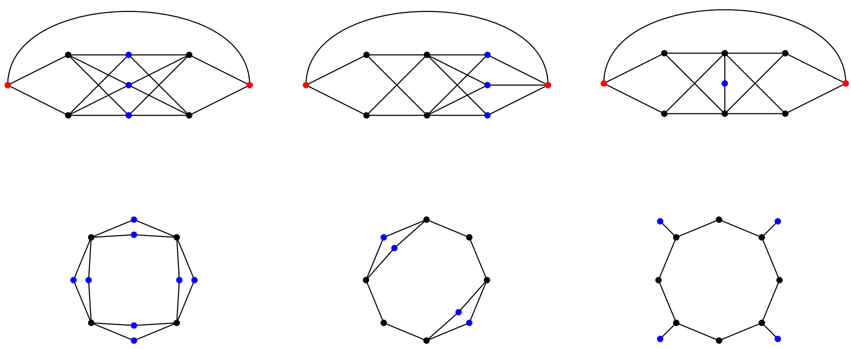

FIGURE 5. Some coarse near doubles Γ (top row) and graphs Δ (bottom row) such that W_Γ is quasiisometric to A_Δ, respectively.

Example 4.18. The top row of Figure 5 shows the only three triangle-free CFS graphs with at most 9 vertices that are not near doubles and for which Dani-Levcovitz does not produce a finite-index RAAG subgroup, but to which Theorem 4.16 applies. These are the smallest examples that we know for which W_Γ is quasiisometric to a RAAG, but we do not know if W_Γ is commensurable to some RAAG.

In each case there is an adjacent pair of unclonable singletons (the extreme left and right vertices, in red) and one other odd module, which consists of clonable vertices (blue). Clone a vertex from this module to make it even. The remaining odd modules are a pair of adjacent singletons. Doubling over these two vertices produces a graph that is a double of the corresponding graph in the bottom row. ◇

4.3. **Unfolding.** We introduce a new operation on graphs that induces a quasi-isometry of their respective RACGs. We first state a result, then give motivating examples, then prove the result.

Proposition 4.19. *Suppose Γ is triangle-free with a separating join $E * F$. Let G be a union of connected components of $\Gamma - (E * F)$, and let $H := (\Gamma - E * F) - G$.*

Let $\bar{G}$ be the union of F and G and the neighbors of G in E. Let $\bar{H}$ be the union of F and H and the neighbors of H in E. Partition E as follows:

$$A := E \cap \bar{G} \cap \bar{H}^c = \{a_0, a_1, \ldots, a_\ell\}$$
$$B := E \cap \bar{G}^c = \{b_0, b_1, \ldots, b_m\}$$
$$C := E \cap \bar{G} \cap \bar{H}$$

Suppose that A and B are nonempty and $C = \{c\}$ is a single vertex.

Then W_Γ is quasiisometric to $W_{\Gamma'}$, where Γ' is a graph constructed as follows. The vertex set is the vertex set of Γ plus one new vertex d, with $D := \{d\}$. Add edges so that $(D \sqcup E) * F \subset \Gamma'$. If $x \bullet\!\!-\!\!\bullet y$ is an edge in $\bar{G}$ then add an edge from x to y in Γ'. If $x \bullet\!\!-\!\!\bullet y$ is an edge of $\bar{H}$ such that $x \neq c \neq y$ then add an edge from x to y in Γ'. If $x \bullet\!\!-\!\!\bullet c$ is an edge of $\bar{H}$ then add an edge from x to d in Γ'.

The quasiisometry can be constructed so that each copy of $\Sigma_{\bar{G}}$ in Σ_Γ is sent to within uniformly bounded Hausdorff distance of a copy of $\Sigma_{\bar{G}}$ in $\Sigma_{\Gamma'}$, and each copy of $\Sigma_{\bar{G}}$ in $\Sigma_{\Gamma'}$ is Hausdorff close to the image of a copy from Σ_Γ.

Furthermore, on each copy of $\Sigma_{\bar{G}}$ in Σ_Γ the quasiisometry restricts to a map that is uniformly bounded distance from a color preserving cubical isomorphism.

Similar statements are true for $\Sigma_{\bar{H}}$, except that c–edges are sent to d–edges.

Example 4.20. Consider the graphs Γ and Γ' of Figure 6.

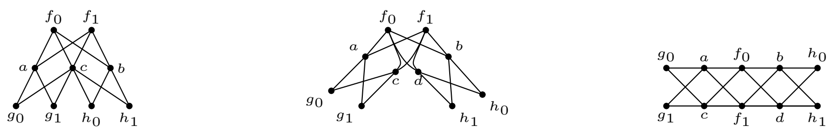

(a) Γ (b) Intermediate (c) Γ'

FIGURE 6. Unfolding visualized as a continuous deformation.

Γ' is obtained from Γ by unfolding as in Proposition 4.19 with $A := \{a\}$, $B := \{b\}$, $C := \{c\}$, $E := A \sqcup B \sqcup C$, $F := \{f_0, f_1\}$, $G := \{g_0, g_1\}$, and $H := \{h_0, h_1\}$, so W_Γ and $W_{\Gamma'}$ are quasiisometric. ◇

In the previous example Γ is a near double and admits a FIDL–Λ, so we already knew it was RAAGedy. In the next example unfolding is the only way that we know to say that the graph is RAAGedy.

Example 4.21. Consider the graphs Γ and Γ' of Figure 7.

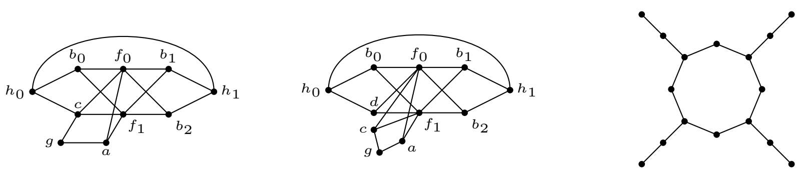

(a) Γ (b) Γ' (c) Δ

FIGURE 7. Unfolding example.

As in the previous example, the sets for Proposition 4.19 are indicated by their lowercase vertex labels, and we conclude that W_Γ is quasiisometric to $W_{\Gamma'}$.

In Γ, vertices h_0, h_1, b_0, c, and a are all unclonable singletons, since a and b_0 are only satellites of c, and vertices h_0, h_1, and c are not satellites at all. The graph is not a near double. It also contains an odd cycle, so it does not admit a FIDL–Λ.

In Γ', d and b_0 are twins and c and a are twins. The only unclonable singletons are h_0 and h_1, which are adjacent. So, clone g and then link double over h_0 and h_1. Conclude that W_Γ and $W_{\Gamma'}$ are quasiisometric to A_Δ, for the Δ of Figure 7c. $\diamond$

For the proof of Proposition 4.19 we need a variation of Lemma 4.11.

Lemma 4.22. *For every $m, n \geqslant 0$ there is a quasiisometry ϕ between the $(m+n+3)$–valent tree T_{m+n+3} with edges colored $a_0,\dots, a_m$, $b_0,\dots,b_n$, c (with exactly one edge of each color at each vertex) and the $(m+n+4)$–valent tree T_{m+n+4} with edges colored $a_0,\dots, a_m$, $b_0,\dots,b_n$, c, d with the following properties:*

- *ϕ sends each $\{a_0, \dots, a_m, c\}$–colored component of T_{m+n+3} within uniformly bounded Hausdorff distance of a unique $\{a_0, \dots, a_m, c\}$–colored component of T_{m+n+4}, and every such component of T_{m+n+4} is the coarse image of a unique component of T_{m+n+3}. Furthermore, for each such component the quasiisometry restricts to a color-preserving isomorphism except at a single vertex.*
- *ϕ takes every $\{b_0, \dots, b_n, c\}$–colored component of T_{m+n+3} within uniformly bounded Hausdorff distance of a unique $\{b_0, \dots, b_n, d\}$–colored component of T_{m+n+4}, and every such component of T_{m+n+4} is the coarse image of a unique component of T_{m+n+3}. Furthermore, for each $\{b_0, \dots, b_n, c\}$–colored component of T_{m+n+3} the quasiisometry restricts, except at a single vertex, to an isomorphism that preserves b_j–edges for each j and takes c–edges to d–edges.*

Proof. As in Lemma 4.11, choose a base vertex w_0 of $T_{m+n+3}$and describe vertices according to colored paths.

For phase 1 of the construction, consider every vertex v with an outgoing c–edge. Add a new edge labelled d at v, and call its opposite vertex vd. For each j, delete the edge between vc and vcb_j and instead connect vcb_j to vd by an edge labelled b_j. The effect is to 'unfold' all of the c–edges, leaving the a_i–edges in place at the end of the c–edge and moving the b_j–edges to the end of the new d–edge. See Figure 8. Call the result ϕ_1. It is a coarse map, in the sense that some vertices are sent to a set of 2 vertices at uniformly bounded distance 2 from one another. ϕ_1 is a $(3, 2)$–quasiisometry. It may be easier to visualize the inverse of the (coarse) map in Figure 8: it 'folds' co-incident edges $v \multimap vc$ and $v \multimap vd$ together to make a single edge $v \multimap vc$. This is actually a map, but is not 1 to 1 on vertices.

The coarse map ϕ_1 has the additional property that for any $\{a_0, \dots, a_m, c\}$–component C there is a unique color preserving isomorphism ϕ_1^C such that for all vertices $v \in C$ we have $\phi_1^C(v) \in \phi_1(v)$. The same is true for $\{b_0, \dots, b_n, c\}$–components, except that edges colored c are sent to edges colored d. This is to say that at the level of trees there is not a canonically nice way to make choices of image points such that ϕ_1 is an honest map instead of a coarse map, but at the level of $\{a_0, \dots, a_m, c\}$–component and $\{b_0, \dots, b_n, c\}$–component there is, and it is unique, so we will not further belabor the point, and simply speak of *the restriction of ϕ_1 to C* as a well defined map. So in Figure 8, we would say that ϕ_1 sends the unique a_0c (red-green) colored geodesic through the basepoint on the left isomorphically to the unique a_0c (red-green) colored geodesic through the basepoint on the right, and sends the unique bc (blue-green) colored geodesic through the basepoint on the left isomorphically to the unique bd (blue-olive) colored geodesic through the base point on the right.

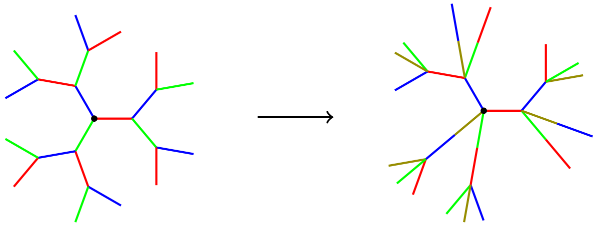

FIGURE 8. Phase 1 coarse map ϕ_1, with a_0 red, b_0 blue, c green, d olive

Vertices with an incoming c–edge have outgoing $\{a_0, \ldots, a_m\}$–edges, but no incident $\{b_0, \ldots, b_n, d\}$–edges. Vertices with an incoming d–edge have outgoing $\{b_0, \ldots, b_n\}$–edges, but no incident $\{a_0, \ldots, a_m, c\}$–edges. All other vertices have one incident edge of each color.

For phase 2 of the construction, define a map ϕ_2 inductively with increasing distance to w_0, as follows. See also Figure 9. Our map ϕ_2 will be injective on vertices, but rearrange the placement of some edges.

Suppose v is a vertex of less than full valence (of valence less than $m + n + 4$) and its incoming edge is colored c. Then v has no $\{b_0, \ldots, b_n, d\}$–edges. The vertex $va_0db_0a_0$ had full valence at the end of phase 1, since it has incoming edge colored a_0. We claim it still has full valence now, so we can take all of its $\{b_0, \ldots, b_n, d\}$–edges and donate them to v. We further claim that none of the other phase 2 moves affect an edge on the geodesic between v and $va_0db_0a_0$.

Now suppose v has less than full valence with incoming a_0–edge. This could happen if it had full valence at the end of phase 1, but donated its $\{b_0, \ldots, b_n, d\}$–edges to a predecessor earlier in phase 2. The vertex $vca_0b_0a_0$ has full valence, because it did at the end of phase 1 and its last four edges are not $a_0db_0a_0$, so it was not called upon to donate to a vertex with incoming c–edge. Take its $\{b_0, \ldots, b_n, d\}$–edges and donate them to v.

The construction for v with incoming d– or b_0–edges is similar, swapping the roles of a_0 and b_0 and those of c and d.

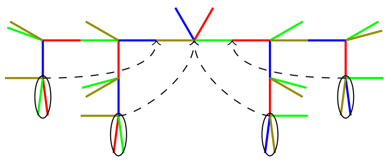

FIGURE 9. Choices of donors for phase 2 map.

Now we confirm the claim that these four types of donations do not interfere with one another. Given a vertex u, let $\delta(u)$ denote its donor, if any. Suppose v is a predecessor of w and let w' be the immediate predecessor of w. For the geodesic $[v, \delta(v)]$ to end on $[w, \delta(w)]$, and thus potentially have the donation to v interfere with the donation to w, would require a terminal segment of $[v, \delta(v)]$ to coincide with an initial segment of $[w', \delta(w)]$. The four possible labels of $[v, \delta(v)]$ are:

$$a_0db_0a_0 \quad ca_0b_0a_0 \quad b_0ca_0b_0 \quad db_0a_0b_0$$

The four possible labels of $[w', \delta(w)]$ are:

$$ca_0db_0a_0 \quad a_0ca_0b_0a_0 \quad db_0ca_0b_0 \quad b_0db_0a_0b_0$$

The maximal overlap between a suffix of the former and a prefix of the latter is a single letter, which means that $[v, \delta(v)] \cap [w, \delta(w)]$ is empty unless $w = \delta(v)$. Even in this case we have arranged that the edges donated to v by w are not one of the first edges on $[w, \delta(w)]$, since when w donates $\{b_0, \dots, b_n, d\}$–edges, $[w, \delta(w)]$ starts with c, and when w donates $\{a_0, \dots, a_n, d\}$–edges, $[w, \delta(w)]$ starts with d.

The new edges added connect vertices that were at distance at most 8 after phase 1, so the map ϕ_2 of phase 2 is 8–biLipschitz on vertices.

Consider a $\{b_0, \dots, b_n, d\}$–component in the tree before phase 2, and suppose w is its closest vertex to w_0. Suppose that during phase 2 a predecessor v of w took the $\{b_0, \dots, b_n, d\}$–edges from w. By construction, all predecessors of v are already full valence, so no vertex of $\{b_0, \dots, b_n, d\}$–component at w other than w will be called upon to donate $\{b_0, \dots, b_n, d\}$–edges. Thus, for any nontrivial geodesic word u in letters $\{b_0, \dots, b_n, d\}$, $\phi_2(wu) = vu$. However, $\phi_2(w) = w$. So ϕ_2 restricted to the $\{b_0, \dots, b_n, d\}$–component based at w is a color preserving isomorphism except at the vertex w, since an isomorphism of this component would have sent w to v.

Similarly, each $\{a_0, \dots, a_m, c\}$–component is sent to within bounded Hausdorff distance of an $\{a_0, \dots, a_m, c\}$–component, and the map restricts to a color preserving isomorphism except possibly at the unique vertex closest to w_0. □

Proof of Proposition 4.19. Since Γ is triangle-free, $E = A \cup B \cup C$ is an anticlique, so Σ_E is a tree, as is $\Sigma_{E'}$ for $E' := A \cup B \cup C \cup D$. By Lemma 4.22 there is a quasiisometry $\psi\colon \Sigma_{A\cup B\cup C} \to \Sigma_{A\cup B\cup C\cup D}$ that coarsely takes $\Sigma_{A\cup C}$–components to $\Sigma_{A\cup C}$–components, $\Sigma_{B\cup C}$–components to $\Sigma_{B\cup D}$–components, and on each such component restricts to a color preserving isomorphism except at a single vertex, and except for the fact that in the second case c–edges are sent to d–edges. Define:

$$\begin{aligned}\psi' := \mathrm{Id}_{\Sigma_F} \times \psi\colon\ \Sigma_{F*E} &= \Sigma_F \times \Sigma_{A\cup B\cup C} \\ &\to \Sigma_F \times \Sigma_{A\cup B\cup C\cup D} = \Sigma_{F*E'}\end{aligned}$$

Because this map is just a product, the good behavior of ψ on colored components carries over. Specifically, if the restriction of ψ to the $(A \cup C)$–component of Σ_E based at vertex w sends it to the $(A \cup C)$–component of $\Sigma_{E'}$ based at v then the restriction of ψ is a color preserving isomorphism on that component, except at w if $\psi(w) \neq v$, and the distance from v to $\psi(w)$ is uniformly bounded. Then ψ' restricted to the $\Sigma_{F*(A\cup C)}$–component of Σ_{F*E} based at $(1, w)$ sends it to the $\Sigma_{F*(A\cup C)}$–component of $\Sigma_{F*E'}$ based at $(1, v)$, and is a color preserving cubical isomorphism except along $\Sigma_F \times \{w\}$, which is sent to $\Sigma_F \times \{\psi(w)\}$ instead of $\Sigma_F \times \{v\}$. But these two sets are parallel at distance $d(v, \psi(w))$, which was uniformly bounded, so the restriction of ψ' to an $(A \cup C)$–component is uniformly bounded distance from a color preserving cubical isomorphism.

Similar statements hold for copies of $\Sigma_{F*(B\cup C)}$, except that c–edges change to d–edges.

According to our setup, $\Gamma = \bar{G} \cup \bar{H}$ and $\bar{G} \cap \bar{H} = F \cup C \subset F * E$. This gives the following splittings of W_Γ and $W_{\Gamma'}$ as graphs of groups, where each edge group is simply the intersection of its two vertex groups.

$$W_\Gamma = W_{\bar{G}} \overset{W_{F*(A\cup C)}}{\overline{\hspace{3em}}} W_{F*E} \overset{W_{F*(B\cup C)}}{\overline{\hspace{3em}}} W_{\bar{H}}$$

$$W_{\Gamma'} = W_{\bar{G}} \overset{W_{F*(A\cup C)}}{\overline{\hspace{3em}}} W_{F*E'} \overset{W_{F*(B\cup D)}}{\overline{\hspace{3em}}} W_{\bar{H}'}$$

Now apply Proposition 2.10 to get a quasiisometry between W_Γ and $W_{\Gamma'}$ as a tree of quasiisometries with respect to the given splittings, with base maps:

- The identity on $W_{\bar{G}}$.
- $\psi'\colon W_{F*E} \to W_{F*E'}$
- The isomorphism $W_{\bar{H}} \to W_{\bar{H}'}$ fixing each b_j and sending c to d. □

The check that the hypotheses of Proposition 2.10 are satisfied is similar to the one in the proof of Proposition 4.10: Condition (2) has been established explicitly, and

Conditions (1) and (3) follow from the fact that ψ' is uniformly bounded distance from a color preserving isomorphism on each coset.

Here is an application of Proposition 4.19.

Proposition 4.23. *For every triangle-free CFS graph Γ there is a triangle-free CFS graph Γ' with no cut 2–paths, such that W_Γ and $W_{\Gamma'}$ are quasiisometric.*

The proof will require the following lemma.

Lemma 4.24. *Suppose Γ is incomplete, triangle-free, and CFS. Suppose C is a cut, either an anticlique $\{a, b\}$ or a 2–path $a \bullet\!\!-\!\!\bullet\, c \bullet\!\!-\!\!\bullet\, b$, such that $\Gamma - C$ is not connected. Then every component of $\Gamma - C$ contains a vertex in $\mathrm{lk}(a) \cap \mathrm{lk}(b)$.*

Proof. Let X be a component of $\Gamma - C$, and let $\bar{X} := X \cup C$.

Suppose there is a vertex $\{x, y\}$ of $\boxslash(\Gamma)$ such that $x \in X$ and $y \in \Gamma - \bar{X}$. Then x and y are the diagonals of a square. The other diagonal of that square consists of two nonadjacent vertices in $\mathrm{lk}(x) \cap \mathrm{lk}(y)$, but every path from x to y passes through C, and the only nonadjacent points of C are a and b, so the square is $\{a, b\} * \{x, y\}$. In this case we are done: $x \in X \cap \mathrm{lk}(a) \cap \mathrm{lk}(b)$.

Now we show there must be such a vertex of $\boxslash(\Gamma)$ containing a point each from X and $\Gamma - \bar{X}$. Choose any $x \in X$ and $y \in \Gamma - \bar{X}$. Since Γ is incomplete, triangle-free, and CFS, it has no cone vertices and $\boxslash(\Gamma)$ contains a component whose support is all of Γ, so there is a nontrivial path in $\boxslash(\Gamma)$ from a vertex with x in its support to a vertex with y in its support. The path corresponds to a chain of squares $S_0, \ldots, S_n$ in Γ such that consecutive squares share a diagonal and $x \in S_0$ and $y \in S_n$. Let m be the least index such that S_m contains a vertex of $\Gamma - \bar{X}$. Call that vertex y'. If S_m contains a vertex from X we are done, so assume not. This implies $m > 0$ Consider S_{m-1}. It is contained in $\bar{X}$, so it contains some $x' \in X$, since $|S_{m-1}| = 4$ and $|\bar{X} - X| \leqslant 3$. We have that S_{m-1} does not contain y' and S_m does not contain x'. The shared diagonal of S_{m-1} and S_m consists of two nonadjacent vertices in $\mathrm{lk}(x') \cap \mathrm{lk}(y') \subset C$, which must be $\{a, b\}$. Conversely, x' and y' are then nonadjacent vertices in $\mathrm{lk}(a) \cap \mathrm{lk}(b)$, so $\{a, b\} * \{x', y'\}$ is a square with a diagonal containing vertices from X and $\Gamma - \bar{X}$. ☐

Proof of Proposition 4.23. Suppose Γ has a cut 2–path $x \bullet\!\!-\!\!\bullet\, c \bullet\!\!-\!\!\bullet\, y$. Take a component $\mathcal{C}$ of $\Gamma - \{x, y, c\}$. Since Γ is CFS it does not have separating cliques, and, by definition of cut 2–paths, $\{x, y\}$ is not a cut pair, so $\mathcal{C}$ contains vertices adjacent to x and y and c and to no other vertices of $\Gamma - \mathcal{C}$. Let $\bar{\mathcal{C}} := \mathcal{C} \cup \{x, y, c\}$. By the same reasoning $\bar{\mathcal{C}}^c$ contains vertices adjacent to x and y and c. By Lemma 4.24, $A := \mathcal{C} \cap \mathrm{lk}(x) \cap \mathrm{lk}(y) \neq \varnothing$ and $B := \bar{\mathcal{C}}^c \cap \mathrm{lk}(x) \cap \mathrm{lk}(y) \neq \varnothing$. Apply Proposition 4.19 with A, B, $F = \{x, y\}$ and $C := \{c\}$. In the resulting graph Γ' the set $\{x, y\}$ becomes a cut pair. If Γ' has other cut 2–paths repeat the argument until they have all been unfolded into cut pairs. By construction, Γ' is triangle-free, as a triangle in Γ' would give a triangle in Γ under the natural map collapsing Γ' to Γ. Finally, $W_{\Gamma'}$ has quadratic divergence, since it is quasiisometric to W_Γ, and Γ is CFS, so Γ' is CFS as well. ☐

Illustrating Proposition 4.23, in both Example 4.20 and Example 4.21, the graph Γ contains a cut 2–path $f_0 - c - f_1$ that becomes a cut pair $\{f_0, f_1\}$ in Γ' after unfolding.

Example 4.25. One might wonder whether the restriction $|C| = 1$ in Proposition 4.19 is a limitation of the proof or of the concept. Here is an example with $|C| > 1$ that shows the analogue of Proposition 4.19 is not true. Consider the graphs Γ and Γ' in Figure 10.

FIGURE 10. Graphs showing that unfolding is not a quasiisometry when $|C| > 1$.

As in Example 4.21, the labelling of vertices in Figure 10 suggests the sets to which they belong in the statement of Proposition 4.19. It appears that Γ' is obtained from Γ by an unfolding-like operation on $F * C$, but W_Γ and $W_{\Gamma'}$ are not quasiisometric, because their JSJ decompositions are incompatible, according to Section 2.4.2: W_Γ does not split over a 2–ended subgroup, since Γ has no cut pairs or cut 2–paths, while $W_{\Gamma'}$ splits over $W_{\{f_0,f_1\}}$. ◇

4.4. **Equivalence classes under the graph modification operations.** Consider the graph Ξ whose vertices represent the triangle-free CFS graphs with at most 12 vertices, with an edge between vertices if the graph of one can be obtained from the other by either link doubling, cloning, or unfolding. If two graphs Γ and Γ' are in the same connected component of Ξ, then W_Γ and $W_{\Gamma'}$ are quasiisometric.

Example 4.26. Consider the graphs Γ and Γ' shown in Figure 11. They are contained in the same component of Ξ, and each is a local minimum in Ξ with respect to number of vertices.

FIGURE 11. Two graphs in the same Ξ component.

Γ is not a coarse near double: 5, 6, 7, and 8 are all unclonable singletons. It does not admit a FIDL–Λ since it is not bipartite.

Γ' is a near double: it has a pair of adjacent singletons, 4 and 5, and one additional singleton, 8, that is a satellite of 4, so $\mathfrak{D}^\circ_{(5,0)} \circ \mathfrak{D}^\circ_4(\Gamma')$ is a graph double.

Γ and Γ' are connected in Ξ by a cloning edge and a link doubling edge; specifically, the graph obtained by cloning vertex X of Γ is isomorphic to $\mathfrak{D}^\circ_7(\Gamma')$. ◇

4.5. **Inductive construction revisited.** Recall that Question 3.19 asks if every RAAGedy, triangle-free graph without separating cliques is constructible by coning from a square. It seems implausible that this should be true in light of the fact that the unfolding and cloning operations add vertices, but in fact we can computationally verify an affirmative answer for triangle-free CFS graph with at most 12 vertices that are known to be RAAGedy by the results of this paper so far. For graphs of any finite cardinality that are RAAGedy by virtue of satisfying Dani and Levcovitz's or Davis and Januszkiewicz's conditions, the question has a positive answer, even if we strengthen the condition 'RAAGedy' to 'commensurable to a RAAG':

Proposition 4.27. *If Γ is a triangle-free graph without separating cliques that is either a graph double or admits a FIDL–Λ, then Γ is constructible by coning from a square through graphs that define RACGs commensurable to RAAGs.*

Proof. For graphs admitting a FIDL–Λ, this is the main result of [16]; the idea is to use the fact that Λ is a forest and show that there is a Λ–leaf that can be removed first, and then induct on the number of vertices. Here we prove the graph double case, also by induction on the number of vertices.

Suppose Γ is a graph double. If it is a square we are done, so suppose not. Then $\Gamma = \mathfrak{D}(\Theta)$ where Θ is a connected, triangle-free, incomplete graph. Recall that the vertex set of Γ is identified with $\Theta \times \{0,1\}$.

Since Θ is finite it contains some vertex v that is not a cut vertex. Let $\Theta' := \Theta - \{v\}$ and $\Gamma' = \mathfrak{D}(\Theta') \subset \Gamma$. Let $L := \mathrm{lk}_\Gamma((v,0)) = \mathrm{lk}_\Gamma((v,1)) = \mathrm{lk}_\Theta(v) \times \{0,1\}$. Let $\Gamma'' := \Gamma - \{(v,1)\}$, so that $\Gamma = \Gamma'' *_L (v,1)$ and $\Gamma'' = \Gamma' *_L (v,0)$. Since Θ' is connected and triangle-free, it follows that Γ' and Γ'' are triangle-free, incomplete graphs without separating cliques. Since $(v,0)$ is the only vertex of Γ'' with no twin, apply Proposition 4.7 to see Γ'' is a near double, so defines a RACG commensurable to a RAAG, by Proposition 4.3.

Now Γ is constructible from a square by coning through graphs defining RACGs commensurable to RAAGs, first by applying the induction hypothesis to Γ', which is a graph double with fewer vertices than Γ, then coning once to get Γ'' and once more to get Γ. □

5. Obstructions to being RAAGedy

We now work from the other direction to find reasons that a right-angled Coxeter group cannot be quasiisometric to any right-angled Artin group.

5.1. **Minsquare versus CFS.** Recall that for incomplete triangle-free graphs Γ without separating cliques, an obstruction to W_Γ being RAAGedy is provided by Γ not being minsquare (see Proposition 2.16) or not being CFS (see Section 3). In fact, we show in Theorem 6.1 that Γ must be strongly CFS if W_Γ is RAAGedy. In light of these results, we now explore the relationships between these properties.

Lemma 5.1. *Let Γ be a graph. Minsquare subgraphs of Γ are in bijection with collections $\mathcal{C}$ of connected components of $\boxslash(\Gamma)$ satisfying:*

(1) $\mathcal{C}$ is nonempty.
(2) If $\{v,w\} \subset \bigcup_{C\in\mathcal{C}} \mathrm{supp}(C)$ and $\{v,w\}$ is a vertex of $\boxslash(\Gamma)$ contained in connected component C_0, then $C_0 \in \mathcal{C}$.
(3) $\mathcal{C}$ is minimal with respect to inclusion among collections of connected components of $\boxslash(\Gamma)$ satisfying the previous two conditions.

Proof. If $\mathcal{C}$ is a collection of connected components of $\boxslash(\Gamma)$ and Δ is a subgraph of Γ, let $\Phi(\mathcal{C})$ be the span of $\bigcup_{C\in\mathcal{C}} \mathrm{supp}(C)$ in Γ and let $\Psi(\Delta)$ be the collection of connected components C of $\boxslash(\Gamma)$ such that Δ contains an induced square with vertices in $\mathrm{supp}(C)$. We claim that Φ and Ψ give inverse bijections between the set of minsquare subgraphs of Γ and collections of connected components of $\boxslash(\Gamma)$ satisfying the conditions in the statement of the lemma.

Since $\boxslash(\Gamma)$ has no isolated vertices, $\mathcal{C}$ is nonempty if and only if $\Phi(\mathcal{C})$ contains an induced square of Γ, and Δ contains an induced square of Γ if and only if $\Psi(\Delta)$ is a nonempty collection.

Now we want to show that Condition (2) describes square completeness. If $\mathcal{C}$ satisfies Condition (2) then whenever $\{v,w\} \subset \Phi(\mathcal{C})$ is a vertex of $\boxslash(\Gamma)$ contained in a connected component C_0, we have $C_0 \in \mathcal{C}$. Since C_0 is a connected component of $\boxslash(\Gamma)$, every vertex adjacent to $\{v,w\}$ is also contained in C_0, so every induced square of Γ with diagonal $\{v,w\}$ is contained in $\mathrm{supp}(C_0) \subset \Phi(\mathcal{C})$. Thus, $\Phi(\mathcal{C})$ is square complete.

Conversely, if Δ is a square complete subgraph of Γ and if $\{v,w\} \subset \Delta$ is a vertex of $\boxslash(\Gamma)$ contained in connected component C_0, then every neighbor $\{x,y\}$ of $\{v,w\}$, of which there is at least one, corresponds to an induced square $\{v, w\} * \{x, y\}$ of Γ. By square completeness, $\{x, y\} \subset \Delta$, so $C_0 \in \Psi(\Delta)$. Furthermore, by induction on distance to $\{v, w\}$ in C_0, $\operatorname{supp}(C_0) \subset \Delta$. Applying this reasoning to each component in $\Psi(\Delta)$, we see that $\Phi(\Psi(\Delta)) \subset \Delta$. This shows that $\Psi(\Delta)$ satisfies Condition (2), as follows. If $\{v, w\} \subset \bigcup_{C \in \Psi(\Delta)} \operatorname{supp}(C)$ is a vertex of $\boxslash(\Gamma)$ contained in a component C_0, then $\{v, w\} \subset \Delta$, since $\bigcup_{C \in \Psi(\Delta)} \operatorname{supp}(C)$ induces the graph $\Phi(\Psi(\Delta)) \subset \Delta$, so by the first part of this paragraph, $C_0 \in \Psi(\Delta)$.

Now observe that every collection $\mathcal{C}$ satisfying (2) fulfills $\Psi(\Phi(\mathcal{C})) = \mathcal{C}$. Indeed, $C_0 \in \Psi(\Phi(\mathcal{C}))$ means that $\Phi(\mathcal{C})$ contains an induced square with vertices in $\operatorname{supp}(C_0)$, so Condition (2) gives $C_0 \in \mathcal{C}$, ie. $\Psi(\Phi(\mathcal{C})) \subseteq \mathcal{C}$. Conversely, if $C_0 \in \mathcal{C}$ then $C_0 \in \Phi(\operatorname{supp}(C_0))$, hence, $C_0 \in \Psi(\operatorname{supp}(C_0)) \subseteq \Psi(\Phi(\mathcal{C}))$.

Finally, we show that Condition (3) is the analogue of the minimality condition in the definition of minsquare. Suppose $\mathcal{C}$ satisfies all three conditions, and define $\Delta := \Phi(\mathcal{C})$. By the previous two steps of the argument, Δ is square complete and contains a square. Suppose that Δ' is a square complete subgraph of Δ that contains a square. Then $\Psi(\Delta')$ is a subcollection of $\Psi(\Delta)$ that satisfies the first two conditions, so by minimality $\Psi(\Delta') = \Psi(\Delta)$. Now:

$$\Delta' \subset \Delta = \Phi(\mathcal{C}) = \Phi(\Psi(\Delta)) = \Phi(\Psi(\Delta')) \subset \Delta'$$

Thus, Δ is square complete and contains a square, and there is no proper subgraph of Δ with both properties, so Δ is minsquare.

Conversely, if Δ is minsquare then it is square complete and contains a square, so $\Psi(\Delta)$ is a collection of connected components of $\boxslash(\Gamma)$ that satisfies the first two conditions. Suppose $\mathcal{C}$ is a subcollection of $\Psi(\Delta)$ satisfying the first two conditions. Then $\Phi(\mathcal{C})$ is a square complete subgraph of Δ that contains a square. By minimality of Δ, we have $\Phi(\mathcal{C}) = \Delta$. But then $\mathcal{C} \subset \Psi(\Delta) = \Psi(\Phi(\mathcal{C})) = \mathcal{C}$. Thus, $\Psi(\Delta)$ is minimal with respect to inclusion among collections of components of $\boxslash(\Gamma)$ satisfying the first two conditions. □

Genevois [45, Example 7.3] showed by example that the minsquare and CFS properties are independent. Using Lemma 5.1, we give triangle-free examples.

Example 5.2 (minsquare does not imply CFS)**.** Consider the strongly CFS graph

$\Gamma_0 :=$ 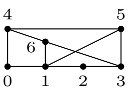

, which contains an isometrically embedded path $P := (0, 1, 2, 3)$ that contains a pair $\{0, 2\}$ of vertices that is not the diagonal of any square and a pair $\{1, 3\}$ that is. Construct Γ by taking two copies of Γ_0 and identifying the two copies of P with opposite orientations:

$\Gamma :=$ 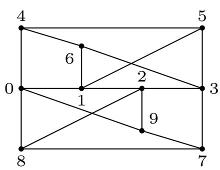

$\boxslash(\Gamma) =$ 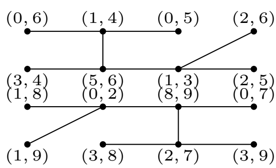

No induced square of Γ enters both copies of $\Gamma_0 - P$, so $\boxslash(\Gamma)$ consists of two disjoint copies of $\boxslash(\Gamma_0)$, as shown. Thus, Γ is not CFS. On the other hand, $\{0, 2\}$ is in the support of the top component of $\boxslash(\Gamma)$ and also appears as a vertex in the bottom component. Similarly, $\{1, 3\}$ is in the support of the bottom component and appears as a vertex in the top component. Since neither component has square complete support, Lemma 5.1 says Γ is minsquare. ◇

Example 5.3 (CFS does not imply minsquare)**.** Consider the graphs:

$\Gamma_0 =$ 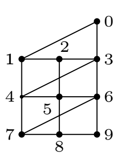

$\Gamma :=$ 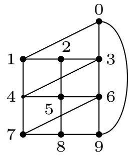

Γ_0 is (strongly) CFS and contains a path $\gamma := (0,3,6,9)$ that is isometrically embedded and does not contain a diagonal of any square. Adding the edge $0 \leftrightarrow 9$ to make Γ does not disturb the CFS property, since the new edge does not kill any square. The square spanned by $\{0,3,6,9\}$ in Γ does not share a diagonal with any other square, so it is a proper minsquare subgraph of a CFS graph. ◇

Example 5.4 (minsquare and CFS does not imply strongly CFS)**.** Let Γ_0 be as in Example 5.2. Let Γ_1 be a $(6,2)$–spider with one pincer foot (recall Example 3.8). Identify the non-pincer feet of the spider with vertices $\{1,3,4,5,6\}$, and the two vertices of the pincer foot with vertices 0 and 2. See Figure 12. All edges of the resulting graph Γ come from either Γ_0 or Γ_1. The spider Γ_1 is strongly CFS and contains all vertices of Γ, so Γ is CFS.

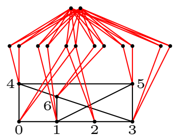

FIGURE 12. A spider attacking a small graph.

We claim that $\boxslash(\Gamma)$ consists of one component isomorphic to $\boxslash(\Gamma_0)$, and one component consisting of $\boxslash(\Gamma_1)$ with two additional leaf vertices attached, coming from the two possible squares that use the segment $(0,1,2)$ of Γ_0 and one of the common neighbors of 0 and 2 in Γ_1. This uses the fact that $\{0,2\}$ is not the diagonal of any square in Γ_0, so we do not create any connection between the copies of $\boxslash(\Gamma_0)$ and $\boxslash(\Gamma_1)$ in $\boxslash(\Gamma)$. Since $\boxslash(\Gamma)$ has two components, Γ is not strongly CFS.

Γ is minsquare because neither component of $\boxslash(\Gamma)$ satisfies Lemma 5.1. ◇

Corollary 5.5. *Strongly CFS with no cone vertices implies minsquare.*

Proof. If Γ is strongly CFS with no cone vertices then $\boxslash(\Gamma)$ is connected, so the lone component $C = \boxslash(\Gamma)$ has $\operatorname{supp}(C) = \Gamma$ and the collection $\{C\}$ satisfies the conditions of Lemma 5.1, so $\Gamma = \Phi(\{C\})$ is minsquare. □

5.2. **Stable subspaces and Morse boundaries.** By a *stable cycle* we mean a simple cycle in Γ of length at least 5 that is square complete. The corresponding special subgroup is stable and 1–ended; it is virtually a hyperbolic surface group. If Γ is RAAGedy it cannot have stable cycles, by Theorem 2.13. Indeed, this is exactly the construction used by Behrstock [2] to give the first example of a non-RAAGedy CFS graph (recall Figure 2). Since passing to iterated link double of Γ induces a quasiisometry on the level of RACGs, if Γ is RAAGedy then no link double of Γ can contain a stable cycle. In this subsection we address the possibility of finding stable cycles in iterated link doubles.

A conjecture proposed in [81] suggested that if a graph Γ contains no stable cycles, then none would appear in any of its link doubles. This was disproven by a counterexample provided by Graeber et al. [48]. Their example was initially constructed by link-doubling a graph with triangles to produce a triangle-free graph without stable cycles [61][Sec. 5.5]. A second link-doubling of this graph, however, resulted in the emergence of a stable cycle. Notably, none of the three graphs in this construction are CFS, prompting the question of whether a similar phenomenon could occur in CFS graphs. Furthermore, the example raised the guess that one doubling might always suffice to produce stable cycles in the triangle-free case. As the following example shows, both assumptions are false—even within the class of triangle-free, strongly CFS graphs.

Example 5.6 (Deeply buried stable cycle)**.** Figure 13a gives a triangle-free, strongly CFS graph Γ that has a deeply buried stable cycle; specifically:

1. Γ contains no stable cycle.
2. $\forall v \in \Gamma$, $\mathfrak{D}^\circ_v(\Gamma)$ contains no stable cycle.
3. $\exists v, w \in \Gamma$, $\mathfrak{D}^\circ_{(w,0)}(\mathfrak{D}^\circ_v(\Gamma))$ contains a stable cycle.

The first two claims are verified by enumerating and checking the possibilities. The third is achieved by doubling first over vertex 0 and then over vertex $(2, 0)$. Consider the cycle $((9, 0), 0)$, $((8, 0), 0)$, $((4, 0), 1)$, $((\mathrm{x}, 0), 1)$, $((4, 1), 1)$, $((3, 1), 0)$, $((5, 0), 0)$, shown in red in Figure 13b. It is induced and has a 2–chord $((4, 0), 1) \multimap ((3, 0), 0) \multimap ((5, 0), 0)$, but there is no other 2–path between vertices of the cycle that is not a subsegment of the cycle, so it is square complete. ◇

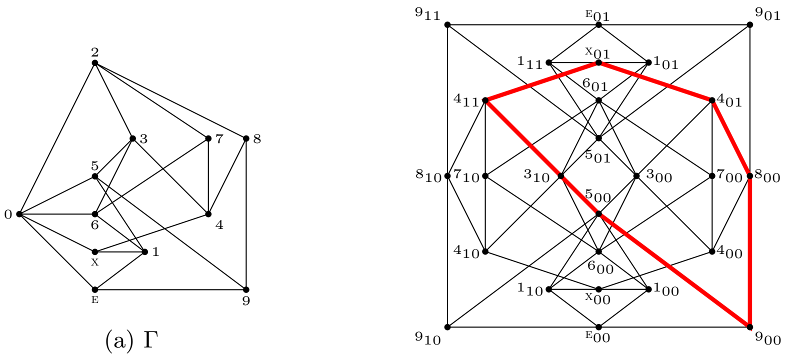

(a) Γ

(b) $\mathfrak{D}^\circ_{(2,0)}(\mathfrak{D}^\circ_0(\Gamma))$

Figure 13. A graph with a deeply buried stable cycle for Example 5.6

We also know an example where a stable cycle appears only after performing three link doubles, so it seems unlikely that there should be any universal bound on how deeply stable cycles can be buried. These examples were found by computer in our enumeration of small graphs; we do not know how to build them on demand.

Question 5.7. *Can stable cycles be buried arbitrarily deeply? That is, given n, does there exist a (triangle-free, strongly) CFS graph Γ such that for all $m < n$, no m–fold iterated link double of Γ that contains a stable cycle, but there is an n–fold iterated link double of Γ that does?*

There is a notion of *Morse boundary* of a group [23], which is a boundary at infinity made of equivalence classes of Morse geodesic rays. Its topological type is quasiisometry invariant.

Theorem 5.8 ([19])**.** *The Morse boundary of a RAAG is totally disconnected.*

A stable cycle appearing in some iterated link double of Γ implies that the Morse boundary of W_Γ contains a circle, so is an obstruction to W_Γ being quasiisometric to a RAAG. However, as discussed in [48], the 1–skeleton of a 3–cube does not have stable cycles in any link double, but it does have circles in its Morse boundary. Thus, connectivity in the Morse boundary is a better obstruction to being RAAGedy than the existence of buried stable cycles.

Fioravanti and Karrer [42] show that a group has totally disconnected Morse boundary if it splits as amalgam of groups with totally disconnected boundaries (including, possibly, the empty set) over a subgroup with empty relative Morse boundary, generalizing the results in [62]. For our purposes it is enough to say that if Γ is a thick join then W_Γ has empty Morse boundary, and if $\Delta < \Gamma$ is a subgraph contained in a thick join subgraph of Γ then the relative Morse boundary of W_Δ in W_Γ is empty.

Proposition 5.9 (Base sufficient criteria for t.d. Morse boundary, cf [42, 62])**.**

(1) If Γ is a clique or a thick join then W_Γ has empty Morse boundary.
(2) If Γ contains a subgraph Δ such that all of the following are true:
  (a) $\Gamma - \Delta$ has more than one component.
  (b) Δ is a clique or is contained in a thick join subgraph of Γ.
  (c) For each component C of $\Gamma - \Delta$, $W_{C\cup\Delta}$ has totally disconnected Morse boundary.

Then W_Γ has totally disconnected Morse boundary.

Definition 5.10. A *decomposition-sequence* of Γ is a rooted tree T of graphs where:

- The graph associated to the root is Γ.
- The graph associated to a vertex equals the union of the graphs of its descendants, and the intersection of its descendants is contained in a thick join subgraph or a clique of Γ.
- Every non-leaf vertex of T has at least 2 descendants.

Corollary 5.11 (Inductive sufficient criteria for t.d. Morse boundary)**.** *If Γ admits a decomposition-sequence such that the graphs associated to the leaves are either cliques or thick joins, then W_Γ has totally disconnected Morse boundary.*

If Corollary 5.11 is satisfied then we save ourselves the work of iterating link doubles of Γ and searching for stable cycles; none exist.

Corollary 5.12. *If Γ satisfies Corollary 5.11 then no iterated link double of Γ contains a stable cycle.*

Forthcoming work of Cordes, Karrer, and Ruane will show that the conditions of Corollary 5.11 are also necessary. This will imply, in the 2–dimensional case, that a RAAGedy Γ must admit a decomposition-sequence.

5.3. **Obstructions from JSJ Decompositions.** Recall from Section 2.4.1 that we know some properties of the JSJ graph of cylinders of a RAAG. For instance, it has no hanging vertices (Lemma 2.11), and rigid vertices are special subgroups that admit no finite or 2–ended splittings (Lemma 2.12). From the quasiisometry invariance of JSJ trees of cylinders we can conclude that Γ is not RAAGedy if the JSJ graph of cylinders of W_Γ has any of the following:

- A hanging vertex group.
- A rigid vertex group that splits over a finite or 2–ended subgroup.
- A rigid vertex group that is not quasiisometric to any RAAG.

Here is another obstruction that requires a little more work to justify.

Lemma 5.13 (Cf [41, Proposition 4.44])**.** *Let Δ be a connected, triangle-free graph that contains a cut vertex v, and let A_Δ be the RAAG defined by Δ. Let $\tilde{v}$ be the cylinder vertex in the JSJ tree of cylinders $\mathcal{T}$ of A_Δ whose stabilizer is $A_{\mathrm{st}(v)}$. Let $\tilde{r}$*

be a rigid vertex adjacent to $\tilde{v}$ in $\mathcal{T}$, such that $\tilde{r}$ is stabilized by A_{Δ_r}, where Δ_r is some maximal biconnected subgraph of Δ containing v. Then:

- *$\tilde{v}$ has infinite valence in $\mathcal{T}$.*
- *If Δ_r is a single edge $v \bullet\!\!-\!\!\bullet w$ then $\tilde{r}$ has valence 2 in $\mathcal{T}$, and its stabilizer and that of its two incident edges are $A_{\{v,w\}} \cong \mathbb{Z}^2$.*
- *If Δ_r is not a single edge then $\tilde{r}$ has infinite valence in $\mathcal{T}$ and the stabilizers of $\tilde{r}$ and all of its incident edges are RAAGs of rank at least 3.*

Proof. Partition $\mathrm{lk}(v)$ into subsets P_1, P_2,... according to which component of $\Delta - \{v\}$ each vertex belongs. There are at least two such parts, since v is a cut vertex. For each P_i there is a maximal biconnected subgraph Δ_i containing P_i as well as v, and containing no vertex from any other P_j. For each i such that Δ_i is either not an edge or is an edge $v \bullet\!\!-\!\!\bullet w$ where w is a cut vertex of Γ, there is a rigid vertex $\tilde{r}_i$ in $\mathcal{T}$ adjacent to $\tilde{v}$ whose stabilizer is A_{Δ_i}. The stabilizer of the edge $\tilde{e}_i$ between $\tilde{r}_i$ and $\tilde{v}$ is $A_{\Delta_i} \cap A_{\mathrm{st}(v)} = A_{\{v\}\cup P_i}$. By assumption, there is at least one rigid vertex $\tilde{r}_i$. For this i, since $\mathrm{lk}(v) - P_i$ is nonempty, $A_{\{v\}\cup P_i}$ has infinite index in $A_{\mathrm{st}(v)}$, so there are infinitely many distinct translates of $\tilde{e}_i$ in $\mathcal{T}$ incident to $\tilde{v}$.

Now consider a rigid vertex $\tilde{r}_i$ adjacent to $\tilde{v}$ in $\mathcal{T}$. If $P_i = \{w\}$ is a singleton, then w is necessarily also a cut vertex and Δ_i is the single edge $w \bullet\!\!-\!\!\bullet v$.

Then $\tilde{r}_i$ has incident edges in $\mathcal{T}$, connecting it to $\tilde{v}$ and $\tilde{w}$, and $\tilde{r}_i$ as well as both of these edges are stabilized by $A_{v,w} \cong \mathbb{Z}^2$.

If P_i is not a singleton then $\Delta_i \cap \mathrm{st}(v)$ contains at least three vertices, since it contains P_i and v, so the stabilizer of $\tilde{r}_i$ and the edge $\tilde{e}_i$ connecting it to $\tilde{v}$ are RAAGs of rank at least 3.

Finally, since Δ is triangle-free, P_i is an anticlique. This implies that the biconnected subgraph Δ_i contains an additional vertex that is not in $\mathrm{st}(v)$, so $A_{\mathrm{st}(v)\cap\Delta_i}$ has infinite index in A_{Δ_i}, which implies there are infinitely many edges in the orbit of $\tilde{e}_i$ incident to $\tilde{r}_i$. □

Corollary 5.14. *If Γ is an incomplete, triangle-free graph with no separating clique and the JSJ graph of cylinders for W_Γ contains a rigid vertex group that is not virtually $\mathbb{Z}^2$, but has an incident edge that is virtually $\mathbb{Z}^2$, then Γ is not RAAGedy.*

Example 5.15. Consider the graph Γ of Figure 14. The JSJ graph of cylinders is:

$$W_\Gamma = W_{\{0,1\}*\{2,7,8\}} \overset{W_{\{0,1\}*\{2,7\}}}{\overline{\qquad\qquad}} W_{\{0,1,2,3,4,5,6,7\}}$$

The vertex on the left is a cylinder; the vertex on the right is rigid. The graph Γ is not RAAGedy, by Corollary 5.14, since the JSJ graph of cylinders of W_Γ has a non-virtually–$\mathbb{Z}^2$ rigid vertex group with an incident virtually $\mathbb{Z}^2$ edge. ⋄

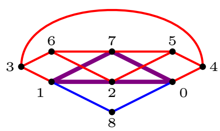

FIGURE 14. The blue/violet subgraph corresponds to a cylinder in the JSJ graph of cylinders of the RACG. The red/violet subgraph corresponds to a non-virtually–$\mathbb{Z}^2$ rigid vertex. Their intersection (violet) corresponds to a virtually–$\mathbb{Z}^2$ edge.

We state another kind of obstruction to being RAAGedy that one can derive from the JSJ decomposition. We will not give a proof here, as Theorem 5.16 is a special case of Theorem 7.5.

Theorem 5.16 (No cycles of cuts)**.** *Let Γ be a triangle-free graph without separating cliques. Suppose for some $n \geqslant 3$ there is an anticlique $\{a_0, \dots, a_{n-1}\}$ such that for all i there is a cut $\{a_i - a_{(i+1)}\}$, for subscripts modulo n. Then Γ is not RAAGedy.*

6. The maximal product region graph

In this section we use the maximal product region graph to distinguish some RACGs from RAAGs. Recall Section 2.6. In Section 6.1 we show that RAAGedy graphs are strongly CFS and establish connectivity properties of the MPRG. In Section 6.2 we show that in the strongly CFS case the MPRG is equivariantly quasiisometric to the stability recognizing space of Abbott, Behrstock, and Durham. In Section 6.3 we define ladders in the MPRG as an obstruction to being RAAGedy, and give sufficient conditions for their presence. In Section 6.4 we give examples and non-examples.

6.1. **Connectivity properties of the MPRG.** In this section we establish connectivity properties of the maximal product region graph. As a corollary:

Theorem 6.1. *If Γ is a triangle-free, RAAGedy graph with no separating clique then it is strongly CFS.*

The proof of Theorem 6.1 is to apply quasiisometry invariance of the MPRG (Corollary 2.23) and compare the following result (Theorem 6.2, which says the MPRG of a 1–ended, 2–dimensional RAAG is connected), to the subsequent Proposition 6.3 about connectivity of the MPRG for RACGs.

Theorem 6.1 appeared as [41, Proposition 4.52]. The proof here is similar. [4] asserts that it is possible to deduce the same result without the triangle-free hypothesis using results from [10].

Theorem 6.2 ([74, Corollary 4.9])**.** *The MPRG of a 1–ended, irreducible, 2–dimensional RAAG is an unbounded quasitree. In particular, it is connected.*

Recall Definition 2.29: Denote the MPRG by Π_Γ and let $\mathrm{Ric}_\Gamma = W_\Gamma \backslash \Pi_\Gamma$. Each vertex v of Ric_Γ corresponds to a maximal thick join J_v of Γ, and W_{J_v} is the stabilizer of v for $W_\Gamma \curvearrowright \Pi_\Gamma$, since it is the stabilizer of Σ_{J_v} for $W_\Gamma \curvearrowright \Sigma_\Gamma$.

Proposition 6.3. *Let Γ be a triangle-free CFS graph. The following are equivalent:*

- Π_Γ *is connected.*
- Ric_Γ *is connected.*
- $\boxslash(\Gamma)$ *is connected.*
- Γ *is strongly CFS.*

Proof. Since Γ is CFS, being strongly CFS is equivalent to $\boxslash(\Gamma)$ being connected. In addition, Γ is triangle-free, so every vertex belongs to some square, hence to some maximal thick join subgraph. Thus, every generator s of W_Γ fixes at least one vertex of Ric_Γ, so $\mathrm{Ric}_\Gamma \cap s.\,\mathrm{Ric}_\Gamma \neq \varnothing$. By induction on word length, for every word w whose i–th prefix is w_i there is a chain $\mathrm{Ric}_\Gamma = w_0.\,\mathrm{Ric}_\Gamma, \dots, w_n.\,\mathrm{Ric}_\Gamma = w.\,\mathrm{Ric}_\Gamma$ such that $w_i.\,\mathrm{Ric}_\Gamma \cap w_{i+1}.\,\mathrm{Ric}_\Gamma \neq \varnothing$. If Ric_Γ is connected this implies Π_Γ is connected. Conversely, Lemma 2.26 implies that if Π_Γ is connected then Ric_Γ is too.

The proof is completed by establishing a bijection between connected components of Ric_Γ and connected components of $\boxslash(\Gamma)$.

By Lemma 3.5, we can define maps ϕ from thick joins of Γ to joins of $\boxslash(\Gamma)$ by $\phi(A * B) := \binom{A}{2} * \binom{B}{2}$ and ψ from joins of $\boxslash(\Gamma)$ to thick joins of Γ by $\psi(A * B) := \mathrm{supp}(A) * \mathrm{supp}(B)$. Consider the two compositions: $\psi \circ \phi(A * B) = \psi(\binom{A}{2} * \binom{B}{2}) = A * B$ and $\phi \circ \psi(A * B) = \phi(\mathrm{supp}(A) * \mathrm{supp}(B)) = \binom{\mathrm{supp}(A)}{2} * \binom{\mathrm{supp}(B)}{2}$, which is a join subgraph of $\boxslash(\Gamma)$ containing $A * B$.

By Proposition 2.28, a vertex $v \in \mathrm{Ric}_\Gamma$ corresponds to a maximal thick join $J_v \subset \Gamma$, and $v \bullet\!\!-\!\!\bullet\, w$ is an edge of Ric_Γ when J_v and J_w contain a common square. If $v \bullet\!\!-\!\!\bullet\, w$ is an edge of Ric_Γ then $\phi(J_v)$ and $\phi(J_w)$ are joins in $\boxslash(\Gamma)$ with at least one edge in common, so ϕ takes components of Ric_Γ into components of $\boxslash(\Gamma)$.

Conversely, if $\{a,b\} \bullet\!\!-\!\!\bullet \{c,d\}$ is an edge in $\boxslash(\Gamma)$ then $\psi(\{a,b\} * \{c,d\}) = \{a,b\} * \{c,d\}$ is a square in Γ, and the set of maximal thick joins containing $\{a,b\} * \{c,d\}$ is a nonempty clique in Ric_Γ. Moreover, if $\{a,b\} \bullet\!\!-\!\!\bullet \{e,f\}$ is another edge in $\boxslash(\Gamma)$ then $\psi(\{a,b\} * \{\{c,d\},\{e,f\}\}) = \{a,b\} * \binom{\{c,d,e,f\}}{2}$ is a thick join containing $\{a,b\} * \{c,d\}$ and $\{a,b\} * \{e,f\}$, so the set of maximal thick joins containing it is a nonempty clique in the intersection of the clique of those containing $\{a,b\} * \{c,d\}$ and the clique of those containing $\{a,b\} * \{e,f\}$. Thus, ψ takes connected components of $\boxslash(\Gamma)$ into connected components of Ric_Γ.

Finally, by the observation on compositions of ϕ and ψ, we have that ϕ and ψ induce inverse bijections between connected components of Ric_Γ and connected components of $\boxslash(\Gamma)$. $\square$

In the next results we will need to go between paths in the maximal product region graph and paths in the square complex. Let Υ be a triangle-free graph, let G_Υ be the RACG or RAAG presented by Υ, let Σ_Υ be its Davis complex or universal cover of its Salvetti complex, respectively. Let $\mathrm{Ric}_\Upsilon = G_\Upsilon \backslash \Pi_\Upsilon$ be the fundamental domain for the action on the MPRG. Let $\gamma\colon [0,L] \to \Pi_\Upsilon$ be a combinatorial path. We construct a path γ' in Σ_Γ *shadowing* it as follows. For each $i \in [0,L)$, the maximal standard product regions $\gamma(i)$ and $\gamma(i+1)$ intersect in a standard product region, by definition of Π_Υ. Further, Ric_Υ corresponds to maximal standard product regions of Σ_Υ containing the vertex 1. Choose vertices a in the maximal standard product region $\gamma(0)$ and b in the maximal standard product region $\gamma(L)$. Let γ'_0 be a path in Σ_Υ starting at a, contained in the maximal standard product region $\gamma(0)$, and ending in $\gamma(0) \cap \gamma(1)$. For $i+1 \in (0,L)$, let γ'_{i+1} be a path in Σ_Υ that starts where γ'_i ended, is contained in $\gamma(i+1)$, and ends in $\gamma(i+1) \cap \gamma(i+2)$. Let γ'_L be a path in Σ_Υ contained in $\gamma(L)$ that starts at the end of γ'_{L-1} and ends at b. Then the concatenation γ' of the γ'_i is a path from a to b in Σ_Υ composed of subsegments that are contained in product regions corresponding to successive vertices of γ.

The next results Lemma 6.4 and Corollary 6.5 are implicit in [74], but we need more precise statements than what appear there explicitly.

Lemma 6.4 (cf. [74, Section 4.1],[41, Theorem 2.35])**.** *Let Δ be a connected, triangle-free, incomplete graph without a cut vertex, and that is not a join. Let Σ_Δ be the universal cover of the Salvetti complex of the RAAG A_Δ, and let $\mathrm{Ric}_\Delta = A_\Delta \backslash \Pi_\Delta$ be the fundamental domain of its action on its maximal standard product region graph as in Definition 2.29. Let $\mathcal{H}_a$ be the wall in Σ_Δ dual to the edge $1 \bullet\!\!-\!\!\bullet\, a$ for $a \in \Delta$. Let $\sigma(\mathcal{H}_a) \in \mathrm{Ric}_\Delta \cap a.\mathrm{Ric}_\Delta \subset \Pi_\Delta$ be the vertex corresponding to the maximal standard product region $\Sigma_{\mathrm{lk}(a)} \times \Sigma_{\{b \in \Delta \mid \mathrm{lk}(a) \subset \mathrm{lk}(b)\}}$. If g and h are separated by $\mathcal{H}_a$ in Σ_Δ then $g.\mathrm{Ric}_\Delta$ and $h.\mathrm{Ric}_\Delta$ are separated by $\mathrm{st}(\sigma(\mathcal{H}_a))$ in Π_Δ.*

Proof. There is a unique maximal join subgraph $\mathrm{lk}(a) * \{b \in \Delta \mid \mathrm{lk}(a) \subset \mathrm{lk}(b)\}$ of Δ containing $\mathrm{st}(a)$, so $\sigma(\mathcal{H}_a) \in \mathrm{Ric}_\Delta$ is well-defined. Furthermore, since $a \in \mathrm{st}(a)$, this vertex is a–invariant, so $\sigma(\mathcal{H}_a) \in \mathrm{Ric}_\Delta \cap a.\mathrm{Ric}_\Delta$.

The set of edges dual to the wall $\mathcal{H}_a$ is $A_{\mathrm{lk}(a)}(1 \bullet\!\!-\!\!\bullet\, a)$. For $\ell \in A_{\mathrm{lk}(a)}$, any standard product region containing an edge $\ell(1 \bullet\!\!-\!\!\bullet\, a)$ contains a square $\ell(1 \bullet\!\!-\!\!\bullet\, a \times 1 \bullet\!\!-\!\!\bullet\, b)$ for some $b \in \mathrm{lk}(a)$, thus contains the flat $\ell\Sigma_{\{a,b\}} \subset \Sigma_{\mathrm{st}(a)}$. Thus, any maximal standard product region distinct from $\sigma(\mathcal{H}_a)$ that contains an edge dual to $\mathcal{H}_a$ is adjacent to $\sigma(\mathcal{H}_a)$ in Π_Δ.

A path γ in Π_Δ from $g.\mathrm{Ric}_\Delta$ to $h.\mathrm{Ric}_\Delta$ can be shadowed by a path γ' in Σ_Δ from g to h, which must cross $\mathcal{H}_a$, by hypothesis, so γ contains a vertex corresponding to a maximal standard product region that contains an edge dual to $\mathcal{H}_a$. These are all in $\mathrm{st}(\sigma(\mathcal{H}_a))$, so $\mathrm{st}(\sigma(\mathcal{H}_a))$ separates $g.\,\mathrm{Ric}_\Delta$ from $h.\,\mathrm{Ric}_\Delta$ in Π_Δ. □

As a corollary we make precise the way in which Π_Δ is a quasitree. If it were actually a tree we would have the statement that for every combinatorial geodesic $\gamma\colon [0,L] \to \Pi_\Delta$ and every $t \in [0,L]$, the vertices $\gamma(0)$ and $\gamma(L)$ are not in a common component of $\Pi_\Delta - \{\gamma(t)\}$: either $\gamma(t)$ coincides with one of the other two or it separates them. The actual situation is weaker in two ways:

- We will need stars of vertices to separate Π_Δ, not just single vertices.
- $\mathrm{Ric}_\Delta \subset \Pi_\Delta$ could be highly connected, so if $\gamma \cap \mathrm{st}(\gamma(t)) \subset g.\,\mathrm{Ric}_\Delta$ then it might be possible to detour around $\mathrm{st}(\gamma(t))$ in $g.\,\mathrm{Ric}_\Delta$. It will turn out that the bottlenecks appear at transition points between different translates of Ric_Δ, so instead of looking at $\gamma(t)$ we should shift to such a transition point on γ near to $\gamma(t)$, where 'near' $\approx \mathrm{diam}(\mathrm{Ric}_\Delta)$.

Corollary 6.5. *With notation as in Lemma 6.4, let $\gamma\colon [0,L] \to \Pi_\Delta$ be a combinatorial geodesic. For every $t \in [0,L]$ there is a vertex $v \in \Pi_\Delta$ such that $\gamma(0)$ and $\gamma(L)$ are not contained in the same connected component of $\Pi_\Delta - \mathrm{st}(v)$, and such that $d_{\Pi_\Delta}(v, \gamma(t)) \leqslant \mathrm{diam}(\mathrm{Ric}_\Delta) + 2$.*

Proof. We may assume that no translate of Ric_Δ contains both $\gamma(0)$ and $\gamma(L)$, since otherwise choosing $v := \gamma(0)$ satisfies the corollary vacuously.

Given t, by translating by the A_Δ action, if necessary, we may assume $\gamma(t) \in \mathrm{Ric}_\Delta$. Take any $g, h \in A_\Delta$ with $\gamma(0) \in g.\,\mathrm{Ric}_\Delta$ and $\gamma(L) \in h.\,\mathrm{Ric}_\Delta$. By assumption, $g \neq h$. The 1–skeleton of a CAT(0) cube complex is a median graph, so there exists a unique vertex m of Σ_Δ that is the median of $\{1, g, h\}$. This implies that, with respect to m, every wall $\mathcal{H}$ has a 'majority side' containing m and a complementary 'minority side' containing at most one of 1, g, and h, counted with multiplicity.

We will use the fact that the map σ can be extended A_Δ–equivariantly to all walls of Σ_Δ. Furthermore, for distinct $b, c \in A_\Delta$ let $\mathcal{H}_{b,c}$ be any wall dual to the first edge of any geodesic from b to c in Σ_Δ. Then $\sigma(\mathcal{H}_{b,c}) \in b.\,\mathrm{Ric}_\Delta$, which follows from equivariance and the fact that $\sigma(\mathcal{H}_a) \in \mathrm{Ric}_\Delta \cap a.\,\mathrm{Ric}_\Delta$ for all $a \in \Delta$.

We estimate that $d_{\Pi_\Delta}(\gamma(t), m.\,\mathrm{Ric}_\Delta) \leqslant 2$. If $m = 1$ then the distance is 0, so the estimate holds. Otherwise, let $\mathcal{H}_{m,1}$ be a wall dual to the first edge on a geodesic from m to 1 in Σ_Δ. We have $\sigma(\mathcal{H}_{m,1}) \in m.\,\mathrm{Ric}_\Delta$. Since 1 is on the minority side of $\mathcal{H}_{m,1}$, vertices g and h are on the majority side, so $\mathcal{H}_{m,1}$ separates 1 from $\{g, h, m\}$ in Σ_Δ. By Lemma 6.4, $\mathrm{st}(\sigma(\mathcal{H}_{m,1}))$ separates Ric_Δ from both $g.\,\mathrm{Ric}_\Delta$ and $h.\,\mathrm{Ric}_\Delta$. Thus, $t \geqslant t_0 := \min\{t' \mid \gamma(t') \in \mathrm{st}(\sigma(\mathcal{H}_{m,1}))\}$ and $t \leqslant t_1 := \max\{t' \mid \gamma(t') \in \mathrm{st}(\sigma(\mathcal{H}_{m,1}))\}$. Since γ is geodesic, $t_1 - t_0 \leqslant 2$, so $d_{\Pi_\Delta}(\gamma(t), \{\gamma(t_0), \gamma(t_1)\}) \leqslant 1$. This gives $d_{\Pi_\Delta}(\gamma(t), m.\,\mathrm{Ric}_\Delta) \leqslant d_{\Pi_\Delta}(\gamma(t), \sigma(\mathcal{H}_{m,1})) \leqslant 2$.

Suppose $m \neq g$ and let $\mathcal{H}_{m,g}$ be a wall in Σ_Δ dual to the first edge of some geodesic from m to g. Then it suffices to take $v := \sigma(\mathcal{H}_{m,g}) \in m.\,\mathrm{Ric}_\Delta$, as follows. The desired distance bound is satisfied, since:

$$d_{\Pi_\Delta}(v, \gamma(t)) \leqslant \mathrm{diam}(m.\,\mathrm{Ric}_\Delta) + d(m.\,\mathrm{Ric}_\Delta, \gamma(t)) \leqslant \mathrm{diam}(\mathrm{Ric}_\Delta) + 2$$

By definition, g is on the minority side of $\mathcal{H}_{m,g}$, so $\mathcal{H}_{m,g}$ separates g from h in Σ_Δ, which, by Lemma 6.4, implies $\mathrm{st}(\sigma(\mathcal{H}_{m,g}))$ separates $g.\,\mathrm{Ric}_\Delta$ from $h.\,\mathrm{Ric}_\Delta$. Since $\gamma(0) \in g.\,\mathrm{Ric}_\Delta$ and $\gamma(L) \in h.\,\mathrm{Ric}_\Delta$, vertices $\gamma(0)$ and $\gamma(L)$ are not contained in a common component of $\Pi_\Delta - \mathrm{st}(v)$.

If $m = g \neq h$ apply the same argument for $v := \sigma(\mathcal{H}_{m,h})$. □

Lemma 6.6. *With notation as in Lemma 6.4,* $\mathrm{diam}(\mathrm{Ric}_\Delta) \leqslant \mathrm{diam}(\Delta) + 2$.

Proof. Pick vertices r and s in Ric_Δ with $d(r, s) = \mathrm{diam}(\mathrm{Ric}_\Delta)$. They correspond to maximal joins J_r and J_s of Δ. Take a shortest geodesic $\gamma\colon [0,L] \to \Delta$ from a vertex in J_r to a vertex in J_s. For each integer $i \in [0, L]$ there is a maximal join J_i containing $\mathrm{st}(\gamma(i))$, and $J_i \cap J_{i+1}$ contains the edge $\gamma(i) \bullet\!\!-\!\!\bullet \gamma(i+1)$, so J_i and J_{i+1} correspond to vertices of Ric_Δ whose distance is at most 1. Furthermore, J_r contains $\gamma(0)$, so it has edges in common with J_0, so r is at distance at most 1 in Ric_Δ from the vertex corresponding to J_0. Similarly, s is distance at most 1 from the vertex of Ric_Δ corresponding to J_L. Thus, $\mathrm{diam}(\mathrm{Ric}_\Delta) \leqslant L + 2 \leqslant \mathrm{diam}(\Delta) + 2$. □

Corollary 6.7. *If Π is a graph containing sequences of vertices (x_i) and (y_i) such that $d(x_i, y_i) \xrightarrow{i\to\infty} \infty$ and for all sufficiently large i there does not exist a vertex v at distance at least 3 from each of x_i and y_i whose star separates x_i and y_i, then Π is not the MPRG of an irreducible, 1–ended, 2–dimensional RAAG.*

Proof. Suppose Δ is a finite, connected, triangle-free graph and x and y are vertices in Π_Δ that are not separated by the star of any vertex that is not in the 2–neighborhood of one of them. By Corollary 6.5, if $d(x, y)$ is large enough then there is an approximate midpoint m of x and y whose star separates x from y, and together with Lemma 6.6 it follows that $d(m, \{x, y\}) \geqslant d(x, y)/2 - \mathrm{diam}(\Delta) - 4$. But $d(m, \{x, y\}) \leqslant 2$ by hypothesis, so $d(x, y) \leqslant 2\,\mathrm{diam}(\Delta) + 12$. For any fixed Δ, eventually $d(x_i, y_i) \xrightarrow{i\to\infty} \infty$ exceeds this bound, so $\Pi \ncong \Pi_\Delta$. □

Oh [74, Lemma 4.13] characterizes cut vertices of the MPRG of a RAAG A_Δ: a cut vertex of Π_Δ contained in Ric_Δ is either a cut vertex of Ric_Δ or is fixed by an element of A_Δ that does not fix any of its neighbors in Ric_Δ. Lemma 6.8 is the analogous result for a RACG W_Γ. It is more complicated because $W_\Gamma \curvearrowright \Pi_\Gamma$ is more complicated than $A_\Delta \curvearrowright \Pi_\Delta$. Specifically, nonadjacent vertices in Ric_Γ can have common elements in their stabilizers, and Γ may contain edges that do not belong to any maximal thick join. Both of these phenomenon give rise to loops in Π_Γ that are not visible in Ric_Γ. This observation will be key in Section 6.3.

Lemma 6.8. *Let Γ be a triangle-free strongly CFS graph that is not a join. One of the following conditions hold if and only if $v \in \mathrm{Ric}_\Gamma$ is a cut vertex of Π_Γ.*

1. *There are induced subgraphs R_0 and R_1 of Ric_Γ properly containing $\{v\}$ with $R_0 \cap R_1 = \{v\}$ and $\mathrm{Edges}(\mathrm{Ric}_\Gamma) = \mathrm{Edges}(R_0) \cup \mathrm{Edges}(R_1)$, and such that for $i \in \{0, 1\}$ and Γ_i defined to be the subgraph of Γ spanned by $\bigcup_{u\in R_i} J_u$, we have $\Gamma_0 \cap \Gamma_1 = J_v$ and $\mathrm{Edges}(\Gamma) = \mathrm{Edges}(\Gamma_0) \cup \mathrm{Edges}(\Gamma_1)$.*
2. *There is a vertex s of Γ such that $\mathrm{st}(s) \subset J_v$ and s is not contained in any other maximal thick join.*

Proof. Since Γ is a triangle-free strongly CFS graph that is not a join, Ric_Γ and Π_Γ are connected and are not single vertices, by Proposition 6.3.

Suppose $v \in \mathrm{Ric}_\Gamma$ is a cut vertex of Π_Γ. If v separates Ric_Γ in Π_Γ then choose a complementary component and take R_0 to be the subgraph of Ric_Γ spanned by the union of $\{v\}$ and vertices of Ric_Γ in that connected component of $\Pi_\Gamma - \{v\}$. Let R_1 be the subgraph of Ric_Γ spanned by v and $\mathrm{Ric}_\Gamma - R_0$. By construction, $\{v\} \subsetneq R_0$, $\{v\} \subsetneq R_1$, $R_0 \cap R_1 = \{v\}$ and $\mathrm{Edges}(\mathrm{Ric}_\Gamma) = \mathrm{Edges}(R_0) \cup \mathrm{Edges}(R_1)$.

There are two ways for (1) to fail, and we show that from either of them we can produce a contradictory path that connects R_0 to R_1 in Π_Γ while avoiding v. We conclude that v separating Ric_Γ in Π_Γ implies (1).

The first way for (1) to fail is if there exists $s \in \Gamma_0 \cap \Gamma_1 - J_v$. Then $s.v \neq v$ but s fixes vertices in $R_0 - \{v\}$ and $R_1 - \{v\}$. Let γ be a shortest path in Ric_Γ from the fixed set of s in R_0 to the fixed set of s in R_1. Since v separates R_0 from R_1 in Π_Γ, γ goes through v. The path $s.\gamma$ does not go through v, since $s.v \neq v$, and it has the same endpoints as γ, so it connects R_0 to R_1 in Π_Γ and avoids v.

The other way for (1) to fail is if there is an edge of Γ from a vertex $s_0 \in \Gamma_0 - \Gamma_1$ to a vertex $s_1 \in \Gamma_1 - \Gamma_0$. By construction, this means there are vertices $u_i \in R_i - \{v\}$

such that $s_i \in J_{u_i} - J_v$. Furthermore, $s_0, s_1 \notin J_v$ implies $W_{\{s_0,s_1\}} \cap W_{J_v} = \{1\}$. Let γ be a minimal length path in Ric_Γ from u_0 to u_1. The concatenation of $s_0.\gamma$, $s_1 s_0.\gamma = s_0 s_1.\gamma$, and $s_1.\gamma$ connects R_0 and R_1 in Π_Γ and avoids v.

Suppose v does not separate Ric_Γ in Π_Γ. Then there is some translate $w.\mathrm{Ric}_\Gamma$ such that $\mathrm{Ric}_\Gamma \cap w.\mathrm{Ric}_\Gamma = \{v\}$ and v separates Ric_Γ from $w.\mathrm{Ric}_\Gamma$ in Π_Γ. We induct on the word length of w after considering what happens for generators. Consider $s \in J_v$ not satisfying (2), so either s fixes a vertex $u_s \in \mathrm{Ric}_\Gamma - \{v\}$, or s only fixes v but is adjacent in Γ to a vertex t that fixes $u_t \in \mathrm{Ric}_\Gamma - \{v\}$ but does not fix v. We claim in both cases that v does not separate Ric_Γ and $s.\mathrm{Ric}_\Gamma$ in Π_Γ. In the first case $v \neq u_s \in \mathrm{Ric}_\Gamma \cap s.\mathrm{Ric}_\Gamma$. In the second case take a shortest path γ in Ric_Γ from v to u_t. Then $t.\gamma$ contains u_t and $t.v$ but not v, and $st.\gamma$ contains $s.u_t \in s.\mathrm{Ric}_\Gamma - \{v\}$ and $st.v = ts.v = t.v$ but not $s.v = v$. Thus, there is a path in Π_Γ from $s.\mathrm{Ric}_\Gamma - \{v\}$ to $\mathrm{Ric}_\Gamma - \{v\}$ that avoids $\{v\}$. Now write w as a minimal length word $s_1 \cdots s_n$ for $s_i \in J_v$. If every $s_i \in J_v$ fails to satisfy (2) then every pair Ric_Γ and $s_i.\mathrm{Ric}_\Gamma$ are not separated by v in Π_Γ. But then $s_1 \cdots s_i.\mathrm{Ric}_\Gamma$ is not separated from $s_1 \cdots s_i s_{i+1}.\mathrm{Ric}_\Gamma$ by v in Π_Γ, so all of $\mathrm{Ric}_\Gamma - \{v\}$, $s_1.\mathrm{Ric}_\Gamma - \{v\}$,…,$s_1 \cdots s_n.\mathrm{Ric}_\Gamma - \{v\} = w.\mathrm{Ric}_\Gamma - \{v\}$ are in the same component of $\Pi_\Gamma - \{v\}$, contradicting the choice of w. Thus, there is some $s \in J_v$ satisfying (2).

In the other direction we suppose (1) or (2) and produce a cut vertex v of Π_Γ.

First suppose (1). Consider the splitting $W_\Gamma = W_{\Gamma_0} *_{W_{J_v}} W_{\Gamma_1}$. Let T be its Bass-Serre tree. Edges of T belong to a single orbit and they are in bijection with cosets of W_{J_v}. There are two orbits of vertices in T corresponding to cosets of W_{Γ_0} and W_{Γ_1}. We define a W_Γ–equivariant map $\phi\colon W_\Gamma \times \mathrm{Ric}_\Gamma \to \mathrm{Vertices}(T) \cup \mathrm{Edges}(T)$ and then check it actually defines a map on $\Pi_\Gamma = W_\Gamma.\mathrm{Ric}_\Gamma$.

$$\phi(w.u) := \begin{cases} \text{the vertex } wW_{\Gamma_0} & \text{if } u \in R_0 - \{v\} \\ \text{the vertex } wW_{\Gamma_1} & \text{if } u \in R_1 - \{v\} \\ \text{the edge } wW_{J_v} & \text{if } u = v \end{cases}$$

Suppose $u \in R_0 - \{v\}$ and $w'.u = w.u$. Then $w' \in wW_{J_u} \subset wW_{\Gamma_0}$, so $\phi(w.u) = wW_{\Gamma_0} = w'W_{\Gamma_0} = \phi(w'.u)$. Similar arguments hold for $R_1 - \{v\}$ and v, so ϕ is well-defined on vertices of Π_Γ.

Every edge in Π_Γ is a translate of one in Ric_Γ (recall Lemma 2.26), so it can be written $w.u_0 \bullet\!\!-\!\!\bullet w.u_1$ where $u_0 \bullet\!\!-\!\!\bullet u_1 \subset \mathrm{Ric}_\Gamma$. If u_0, u_1 are both in $R_0 - \{v\}$ or both in $R_1 - \{v\}$ then they map to the same vertex of T. If $u \in R_i$ is adjacent to v then $\phi(v) = 1W_{J_v}$ is an edge incident to the vertex $\phi(u) = 1W_{\Gamma_i}$. By hypothesis, there are no edges between a vertex of $R_0 - \{v\}$ and a vertex of $R_1 - \{v\}$. Conclude that edge paths in Π_Γ are sent by ϕ to connected subsets of T.

Consider the edge of T corresponding to the coset $1W_{\Gamma_0 \cap \Gamma_1}$. Its two vertices are $1W_{\Gamma_0}$ and $1W_{\Gamma_1}$. Consider any path in Π_Γ from a vertex of $R_0 - \{v\}$ to a vertex of $R_1 - \{v\}$. Its ϕ–image is a connected set in T that contains $1W_{\Gamma_0}$ and $1W_{\Gamma_1}$, so it contains a vertex $w.u$ such that $\phi(w.u)$ is the edge $1W_{J_v}$. The definition of ϕ requires $u = v$ and $wW_{J_v} = 1W_{J_v}$, so $w \in W_{J_v}$, which gives $w.v = v$. Thus, every path in Π_Γ from $R_0 - \{v\}$ to $R_1 - \{v\}$ passes through v, so v is a cut vertex of Π_Γ.

Now suppose (2). Let $\gamma\colon [0, L] \to \Pi_\Gamma$ be a combinatorial path from $\mathrm{Ric}_\Gamma - \{v\}$ to $s.\mathrm{Ric}_\Gamma - \{v\}$, and let γ' be a path from 1 to s in Σ_Γ shadowing it. Since γ' starts and ends on opposite sides of the wall $\mathcal{H}_s$ dual to the edge $1 \bullet\!\!-\!\!\bullet s$, it contains some edge crossing $\mathcal{H}_s$. The condition of (2) saying s is not contained in any other maximal thick join implies that Σ_{J_v} is the only maximal standard product subcomplex containing the edge $1 \bullet\!\!-\!\!\bullet s$. The condition that $\mathrm{lk}(s) \subset J_v$ implies that every square containing the edge $1 \bullet\!\!-\!\!\bullet s$ is contained in Σ_{J_v}. It follows that every edge crossing $\mathcal{H}_s$ is contained in Σ_{J_v} and not in any other maximal standard product subcomplex, so γ' can only cross $\mathcal{H}_s$ if $v \in \gamma$. Thus, v is a cut vertex of Π_Γ. $\square$

6.2. **Relation to hierarchical hyperbolic structures.** This section is predominantly to relate the maximal product region graph to other results in the literature, but, having done this, Corollary 6.10 gives us an easy way to guarantee that the orbit map gives a quasiisometric embedding of certain subgroups of W_Γ into Π_Γ, which will be useful in the next subsection.

The definition of hierarchically hyperbolic spaces and groups (HHS/HHG) is complicated, and we will not repeat it; see [8, 7, 6, 10]. There is a hyperbolic graph coming from the hierarchical hyperbolic structure, and quasiisometries between HHSs induce quasiisometries of these hyperbolic graphs. We will show that if Γ is strongly CFS then Π_Γ is W_Γ–equivariantly quasiisometric to the HHS graph for W_Γ. For RAAGs the HHS graph is a quasitree and the MPRG is a quasitree with bottleneck constant 1. In the next subsection we construct wide ladders in the MPRGs of certain RACGs. This allows the possibility that they are still quasitrees, but the bottleneck constant is at least half the width of the ladder, so must be larger than 1. To conclude that these groups are not quasiisometric to a RAAG we really need the finer control that quasiisometries induce *isomorphisms* of MPRGs, not just that they induce quasiisometries between the HHS graphs.

For RAAGs the relevant hyperbolic graph from the standard HHS structure is just the contact graph. For RACGs some modifications must be made. Abbott, Behrstock, and Durham [1, Theorems A,B] show that a HHG G admits an action on a hyperbolic space $ABD(G)$ with the following properties. The structure of $ABD(W_\Gamma)$ will be described in the course of the proof of Theorem 6.9.

- $G \curvearrowright ABD(G)$ is a largest acylindrical action.
- In certain cases, including when G is RACG, $G \curvearrowright ABD(G)$ is universal, in the sense that every generalized loxodromic element of G acts loxodromically in this particular action.
- $ABD(G)$ is a stability recognizing space: a finitely generated subgroup H of G is stable if and only if any orbit map of H into $ABD(G)$ is a quasiisometric embedding.

Theorem 6.9. *Let Γ be a triangle-free, strongly CFS graph. Then Π_Γ is W_Γ–equivariantly quasiisometric to $ABD(W_\Gamma)$.*

Proof. Let $\mathcal{X}_0 := \Sigma_\Gamma$, $\mathcal{X}_1 := ABD(W_\Gamma)$, $\mathcal{X}_2$ be the graph obtained from $\mathcal{X}_0$ by coning off each maximal standard product region, and $\mathcal{X}_3 := \Pi_\Gamma$. We will construct W_Γ–equivariant quasiisometries:

$$\mathcal{X}_3 \xrightarrow{\phi} \mathcal{X}_2 \xrightarrow{\iota} \mathcal{X}_1$$

The standard HHS structure on W_Γ is $(\mathcal{X}_0, \mathfrak{S})$, where $\mathfrak{S}$ is the projection closure of hyperplane carriers, and for $S \in \mathfrak{S}$, the hyperbolic space CS associated to S is its contact graph. To construct $ABD(G)$, Abbott, Behrstock, and Durham modify this HHS structure [1, Theorem 3.7] by replacing the top level hyperbolic space as follows. Start with $\mathcal{X}_0$. For $S \in \mathfrak{S}$ not the maximal element, if there exists $T, U \in \mathfrak{S}$ with $S \subset T$ and $T \perp U$ with both CT and CU of infinite diameter, then cone off S by adding a new vertex c_S attached to each vertex of S. The resulting space is $\mathcal{X}_1 := ABD(W_\Gamma)$.

The inclusion map $\iota\colon \mathcal{X}_2 \to \mathcal{X}_1$ is W_Γ–equivariant and Lipschitz. The idea for showing it is a quasiisometry is that while there may be more cone vertices in $\mathcal{X}_1$, the extra ones are coning off subsets of $\mathcal{X}_0$ that were already coned off in $\mathcal{X}_2$, so they are not making much difference. To see this precisely, we will define a coarse inverse $\bar{\iota}\colon \mathcal{X}_1 \to \mathcal{X}_2$ to be the identity on $\mathcal{X}_0$ and extend it to $\mathcal{X}_1 - \mathcal{X}_0$.

First we characterize S such that there is a cone vertex $c_S \in \mathcal{X}_1 - \mathcal{X}_0$. Such a cone vertex comes from a convex subcomplex S of some hyperplane carrier, for a hyperplane dual to edges labelled by some $s \in \Gamma$. Up to the W_Γ–action, we may assume $1 \in S \subset \Sigma_{\mathrm{st}(s)}$ Any T containing S is a convex subcomplex of the full hyperplane carrier $\Sigma_{\mathrm{lk}(s)} \times \Sigma_s = \Sigma_{\mathrm{st}(s)}$ with $T \subset \Sigma_{\tau \cup \{s\}} \subset \Sigma_{\mathrm{st}(s)}$, where $\tau \cup \{s\}$ is the set of edge labels that occur in T. Now, $\Sigma_{\tau \cup \{s\}}$ has unbounded associated hyperbolic space when W_τ is infinite and not a product, which, since Γ is triangle-free, is simply the case that τ has at least two vertices. In the other direction, if $U \perp T$ then $U \perp \Sigma_{\tau \cup \{s\}}$, but, by triangle-freeness, the biggest subcomplex perpendicular to $\Sigma_{\tau \cup \{s\}}$ is Σ_υ, where υ is the set of common neighbors of τ. This has infinite associated hyperbolic space when υ has more than one vertex. Thus, we have a cone vertex $c_S \in \mathcal{X}_1$ when there exists a subset of $\mathrm{lk}(s)$ that contains the non-s edge labels in S, has at least 2 vertices, and has at least one common neighbor other than s. This is equivalent to saying that S is contained in a standard product region.

For $c_S \in \mathcal{X}_1 - \mathcal{X}_0$ define $\bar{\iota}(c_S)$ to be the set of cone vertices $c_{S'}$ of $\mathcal{X}_2$ such that S' is a maximal standard product region containing S. Then $\bar{\iota} \circ \iota = Id_{\mathcal{X}_2}$, and we have $\varnothing \neq \mathrm{lk}(c_S) \subset \bigcap_{c_{S'} \in \iota \circ \bar{\iota}(c_S)} \mathrm{lk}(c_{S'})$, which implies that $\iota \circ \bar{\iota}$ is at distance at most 2 from $Id_{\mathcal{X}_1}$ and that distances between points in $\mathcal{X}_0$ are the same in $\mathcal{X}_1$ as in $\mathcal{X}_2$. Now apply Lemma 2.6 to see that ι and $\bar{\iota}$ are inverse quasiisometries.

Define:

$$\phi\colon \mathcal{X}_3 \to \mathcal{X}_2\colon S \mapsto c_S$$

$$\bar{\phi}\colon \mathcal{X}_2 \to \mathcal{X}_3 : \begin{cases} c_S \mapsto S & \text{for } c_S \in \mathcal{X}_2 - \mathcal{X}_0 \\ x \mapsto \{S \mid x \in S\} & \text{for } x \in \mathcal{X}_0 \end{cases}$$

By construction , ϕ is W_Γ–equivariant. If distinct maximal standard product regions S_1 and S_2 intersect in a standard flat S_0 then in $\mathcal{X}_1$ there are cone vertices c_{S_i} with $\varnothing \neq \mathrm{lk}(c_{S_0}) \subset \mathrm{lk}(c_{S_1}) \cap \mathrm{lk}(c_{S_2})$, so $d_{\mathcal{X}_2}(\phi(S_1), \phi(S_2)) = d_{\mathcal{X}_2}(c_{S_1}, c_{S_2}) = 2$. Since $\mathcal{X}_3$ is connected, by Proposition 6.3, this implies ϕ is 2–Lipschitz.

The maximal standard product regions containing 1 correspond to maximal nontrivial joins in Γ. There are finitely many of these, and there do exist some, since the CFS property implies Γ contains a square. Thus, $\bar{\phi}(1)$ is a non-empty set that is finite. Since $\mathcal{X}_3$ is connected, finite sets have finite diameter. By definition, $\bar{\phi}$ is W_Γ–equivariant, so $\bar{\phi}$ is a well defined map from $\mathcal{X}_2$ to nonempty subsets of $\mathcal{X}_3$ of uniformly bounded diameter. When $x \in \mathcal{X}_0$ and $c_S \in \mathcal{X}_2 - \mathcal{X}_0$ is an adjacent cone vertex then $x \in S$, so $\bar{\phi}(c_S) \in \bar{\phi}(x)$. Similarly, the CFS property implies that for adjacent vertices in $\mathcal{X}_0$, say across an edge labelled s, there is a standard 2–flat containing both, since the vertex $s \in \Gamma$ is in the support of a square of Γ. Thus, adjacent vertices in $\mathcal{X}_0 \subset \mathcal{X}_2$ have intersecting $\bar{\phi}$–images. It follows that $\bar{\phi}$ is $(\mathrm{diam}\, \bar{\phi}(1), \mathrm{diam}\, \bar{\phi}(1))$–coarsely Lipschitz. Clearly, $\bar{\phi} \circ \phi = Id_{\mathcal{X}_3}$. The map $\phi \circ \bar{\phi}$ agrees with $Id_{\mathcal{X}_2}$ on $\mathcal{X}_2 - \mathcal{X}_0$ and sends $x \in \mathcal{X}_0$ to the set of its cone vertex neighbors, which are all adjacent to x. Apply Lemma 2.6. □

Corollary 6.10. *Let Γ be triangle-free and strongly CFS. Then Π_Γ is a hyperbolic graph, and it is a stability recognizing space for W_Γ.*

Corollary 6.11. *Let Γ be triangle-free and strongly CFS. If W_Γ contains a 1–ended stable subgroup then Π_Γ is not a quasitree.*

6.3. **Ladders.** In this subsection we introduce an obstruction to a graph being the MPRG of a RAAG by the presence of a 'ladder', which will be a 2–ended subgraph whose ends are not separated by stars of vertices, allowing us to apply Corollary 6.7. While this formulation is conceptually clear, it is a large-scale geometric condition in a locally infinite graph, so it is not tangible. We therefore develop explicit, computer verifiable conditions, described purely in terms of the presentation graph Γ, that are sufficient to imply the geometric condition.

The idea for building a ladder is to find a highly connected subset Q of the fundamental domain $\mathrm{Ric}_\Gamma = W_\Gamma \backslash \Pi_\Gamma$ and a pair of generators r and s of W_Γ such that $s.Q \cap Q$ and $r.Q \cap Q$ are large enough. Then the 'ladder' will be $\langle r, s\rangle .Q$. The precise conditions for r, s, and Q appear in Theorem 6.16.

Lemma 6.12. *Suppose* Π *is a graph that contains a sequence of connected subgraphs* $\Theta_i \subset \Pi$ *such that:*

- $\mathrm{diam}_\Pi(\Theta_i) \geqslant i$
- *For all* i *and all* $v \in \Pi$, $\Theta_i - \mathrm{st}(v)$ *has exactly one non-singleton component.*

Then Π *is not the MPRG of an irreducible, one-ended, 2–dimensional RAAG.*

Proof. Pick $x_i, y_i \in \Theta_i$ realizing the diameter. Since Θ_i is connected, x_i has neighbors, and if $d(x_i, v) > 2$ then no neighbor of x_i is in $\mathrm{st}(v)$, so x_i is contained in the unique non-singleton component of $\Theta_i - \mathrm{st}(v)$. The same is true for y_i, so x_i and y_i are not separated by $\mathrm{st}(v)$ except possibly if $d(v, \{x_i, y_i\}) \leqslant 2$. Apply Corollary 6.7. □

Our prototypical example for Lemma 6.12 is to construct an infinite graph Θ by taking a square Q of side length at least three contained in $\mathrm{Ric}_\Gamma = W_\Gamma \backslash \Pi_\Gamma$, such that there are nonadjacent r and s in Γ that act as reflections fixing opposite sides of Q. The Θ_i of Lemma 6.12 are an increasing nested sequence of consecutive translates of Q, with Θ as their infinite union. See Figure 15.

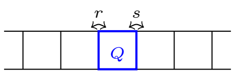

FIGURE 15. Prototypical ladder.

A variant is shown in Figure 16. The variant highlights the reason that in Lemma 6.12 we say that the graph minus a star has one non-singleton component, rather than saying the graph minus a star is connected. The ladder of Figure 16 does obstruct Π from being the MPRG of a RAAG, because even though it is separated by the star of a vertex v (red), its ends are not.

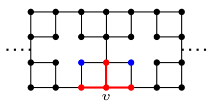

FIGURE 16. A non-prototypical example of a ladder

Definition 6.13 (wide ladder)**.** A *wide ladder* L in a graph Π is a graph satisfying:

(a) There is a connected graph $\mathcal{C}$, each vertex of which is associated to a connected subgraph Q of Π, and the subgraphs associated with adjacent vertices of $\mathcal{C}$ intersect. L is a graph of infinite diameter that is the union of these subgraphs.
(b) The ladder *has wide rungs*, in the sense that if Q and Q' are adjacent in $\mathcal{C}$, then $\mathrm{diam}_\Pi(Q \cap Q') \geqslant 3$.
(c) For all $Q \in \mathcal{C}$ and all vertices $v \in \Pi$, there is exactly one non-singleton component of $Q - \mathrm{st}(v)$.
(d) There do not exist adjacent Q and Q' in $\mathcal{C}$ and $v \in \Pi$ such that every vertex of $(Q \cap Q') - \mathrm{st}(v)$ is a singleton component of either $Q - \mathrm{st}(v)$ or $Q' - \mathrm{st}(v)$. See Figure 17.

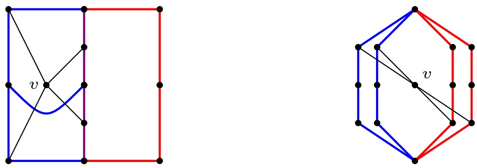

FIGURE 17. In each picture Q is blue, Q' is red, any edges in their intersection are violet. These satisfy Definition 6.13 (b) and (c) but not (d), and there is a vertex v whose star separates their union into multiple non-singleton components.

Lemma 6.14. *A graph containing a wide ladder is not the MPRG of a RAAG.*

Proof. Let Π be a graph containing a wide ladder L defined by $\mathcal{C}$ as in Definition 6.13. The goal is to show that L satisfies the hypotheses of Lemma 6.12. If there is a sequence of vertices c_i of $\mathcal{C}$ such that the corresponding graphs Q_{c_i} have diameters that grow without bound then we can apply Lemma 6.12 with $\Theta_i = Q_{c_i}$. Otherwise, we choose nested connected subgraphs $\mathcal{C}_i$ of $\mathcal{C}$ such that the graphs $\Theta_i := \bigcup_{c\in\mathcal{C}_i} Q_c$ have diameters growing without bound. This is possible since $L = \bigcup_{c\in\mathcal{C}} Q_c$ has infinite diameter. We show by induction on the size of the subgraph of $\mathcal{C}$ that such a Θ_i satisfies the condition that it has a unique non-singleton complementary component for every star.

Suppose Q and Q' are adjacent elements in $\mathcal{C}$ and v is some vertex of Π. Condition (d) says for every $v \in \Pi$ there is a $u \in (Q \cap Q') - \mathrm{st}(v)$ such that u is not a singleton component of $Q - \mathrm{st}(v)$ and u is not a singleton component of $Q' - \mathrm{st}(v)$. Condition (c) says $Q - \mathrm{st}(v)$ and $Q' - \mathrm{st}(v)$ have unique non-singleton components U and U', respectively, so $u \in U \cap U'$. Thus, $U \cup U'$ is connected. Every vertex of $(Q \cup Q') - (U \cup U' \cup \mathrm{st}(v))$ is a singleton component of $(Q \cup Q') - \mathrm{st}(v)$, so $U \cup U'$ is the unique non-singleton component of $(Q \cup Q') - \mathrm{st}(v)$.

Now, suppose that every connected subset $\mathcal{A} \subset \mathcal{C}$ of size at most n has the property that for all $v \in \Pi$, $\bigcup_{a\in\mathcal{A}} Q_a - \mathrm{st}(v)$ has a unique non-singleton component containing the non-singleton component of $Q_a - \mathrm{st}(v)$ for all $a \in \mathcal{A}$. Add one more graph Q such that the corresponding vertex of $\mathcal{C}$ is adjacent to $a_0 \in \mathcal{A}$. For any $v \in \Pi$, $Q \cup Q_{a_0} - \mathrm{st}(v)$ and $\bigcup_{a\in\mathcal{A}} Q_a - \mathrm{st}(v)$ have unique non-singleton components, and they both contain the non-singleton component of $Q_{a_0} - \mathrm{st}(v)$, so $Q \cup \bigcup_{a\in\mathcal{A}} Q_a - \mathrm{st}(v)$ has a unique non-singleton component. □

Next we would like to give sufficient conditions in terms of the graph structure of Γ for Π_Γ to contain a wide ladder. In the proof we will need the following lemma.

Lemma 6.15. *Let Γ be an incomplete, triangle-free graph without a separating clique. Suppose $Z = \{z_1, \ldots, z_k\} \subset \mathrm{st}(v') \cap \mathrm{Ric}_\Gamma$ for some $v' \in \Pi_\Gamma - \mathrm{Ric}_\Gamma$. Let v be the unique vertex in $W_\Gamma.v' \cap \mathrm{Ric}_\Gamma$. Then $Z \subset \mathrm{st}(v)$ and $\left(\bigcap_{i=1}^k J_{z_i}\right) - J_v \neq \varnothing$.*

Proof. By Lemma 2.26, for each i there exists $g_i \in W_\Gamma$ such that $g_i^{-1}.z_i = z_i$ and $g_i^{-1}.v' = v$. This shows $Z \subset \mathrm{st}(v)$. Furthermore, $g_i v = v'$ for each i, so $g_i\Sigma_{J_v}$ is the maximal standard product region of Σ_Γ corresponding to $v' \in \Pi_\Gamma$. In particular, $g_i\Sigma_{J_v} = g_1\Sigma_{J_v}$ for all i.

Since Ric_Γ corresponds to the maximal standard product regions of Σ_Γ containing the identity vertex, $\varnothing \neq \Sigma_{J_v} \cap \bigcap_i \Sigma_{J_{z_i}}$. Thus, for all i, $\varnothing \neq g_i(\Sigma_{J_{z_i}} \cap \Sigma_{J_v}) = \Sigma_{J_{z_i}} \cap g_1\Sigma_{J_v}$. The Helly Property for CAT(0) cube complexes says that since the convex subcomplexes $\{g_1\Sigma_{J_v}, \Sigma_{J_{z_1}}, \ldots, \Sigma_{J_{z_k}}\}$ pairwise have nonempty intersection, their mutual intersection contains a vertex $w \in \Sigma_\Gamma$. Now, $w \in \Sigma_{J_{z_i}}$ implies that every minimal expression of w as a word in W_Γ uses only generators from J_{z_i}. On the other hand, $g_1\Sigma_{J_v}$ and Σ_{J_v} are disjoint, so $w \notin \Sigma_{J_v}$, which means that w cannot be written using only generators from J_v. Thus, every minimal expression of w uses only generators from $\bigcap_i J_{z_i}$, and uses at least one that is not in J_v. $\square$

Theorem 6.16. *Let Γ be triangle-free and strongly CFS. Suppose there exist $r, s \in \Gamma$ and maximal thick joins $J_0, \ldots, J_{n-1}$ of Γ satisfying the following conditions:*

(1) The graph Q with a vertex q_i for each J_i, and such that q_i and q_j span an edge when $J_i \cap J_j$ contains a square, is connected.
(2) Let J be any maximal thick join of Γ. Let I be a subset of $\{q_i \mid J \cap J_{q_i}$ contains a square$\}$ that is either the whole set or is a subset for which $\left(\bigcap_{q_i \in I} J_{q_i}\right) - J \neq \varnothing$. Then $Q - I$ has exactly one non-singleton component, B, and $\mathrm{diam}_Q(Q - B) \leqslant 2$.
(3) The vertices r and s are not adjacent and not contained in a common thick join of Γ.
(4) There are indices a_r, b_r, a_s, $b_s \in \{0, \ldots, n-1\}$, necessarily distinct, with:
   - *$r \in J_{a_r} \cap J_{b_r}$*
   - *J_{a_r} and J_{b_r} do not share a square, and no maximal thick join of Γ shares a square with both of them.*
   - *$s \in J_{a_s} \cap J_{b_s}$*
   - *J_{a_s} and J_{b_s} do not share a square, and no maximal thick join of Γ shares a square with both of them.*

Then Π_Γ contains a wide ladder. Consequently, Γ is not RAAGedy.

Proof. By the description in Item (1), we may regard Q as an induced subgraph of Ric_Γ, which itself is induced in Π_Γ, by Lemma 2.26.

Item (3) says $W_{\{r,s\}}$ is an infinite dihedral group that does not act elliptically on Π_Γ. By Corollary 2.15 it is stable, so by Corollary 6.10 its orbit map into Π_Γ is a quasiisometric embedding. We will take $\mathcal{C}$ to be the Cayley graph of $W_{\{r,s\}}$ with respect to $\{r, s\}$, which is a line, and associate to vertex g the subgraph $g.Q$ in Π_Γ. The ladder $L := \bigcup_{g \in W_{\{r,s\}}} g.Q$ is unbounded, since the orbit map of $W_{\{r,s\}}$ is a quasiisometric embedding. So far, we have shown that condition (a) of Definition 6.13 is satisfied. Item (4) gives condition (b), since it implies $3 \leqslant d_{\mathrm{Ric}_\Gamma}(q_{a_r}, q_{b_r})$ and Lemma 2.26 implies $d_{\mathrm{Ric}_\Gamma}(q_{a_r}, q_{b_r}) = d_\Pi(q_{a_r}, q_{b_r})$, and analogously for s.

To show that condition (c) is satisfied, we claim that the intersection of the star of a vertex of Π_Γ with Q is one of the sets I as described in Item (2). When the vertex is $v \in \mathrm{Ric}_\Gamma$, then $\mathrm{st}(v) \cap Q = \{q_i \mid J_v \cap J_i$ contains a square$\}$. When the vertex is $v' \notin \mathrm{Ric}_\Gamma$, let $\{v\} = W_\Gamma.v' \cap \mathrm{Ric}_\Gamma$; then Lemma 6.15 says $I := \mathrm{st}(v') \cap Q \subset \{q_i \mid J_v \cap J_i$ contains a square$\}$ and that $\left(\bigcap_{i \in I} J_{q_i}\right) - J_v \neq \varnothing$.

Finally, we show that condition (d) of Definition 6.13 is satisfied. Neighboring elements of $\mathcal{C}$ differ by the action of r or s, so, up to a symmetric argument and the group action, it suffices to consider that the adjacent translates of condition (d) are Q and $r.Q$. We know that $Q \cap r.Q$ contains q_{a_r} and q_{b_r}, so the negation of condition (d) requires, up to reversing the roles of q_{a_r} and q_{b_r}, that there is a vertex $v \in \Pi_\Gamma$ such that either q_{a_r} and q_{b_r} are both singleton components of $Q - \mathrm{st}(v)$, or q_{a_r} is a singleton component of $Q - \mathrm{st}(v)$ and q_{b_r} is a singleton component of $r.Q - \mathrm{st}(v)$. Suppose the latter is true. Let v' be the unique vertex in $W_\Gamma.v \cap \mathrm{Ric}_\Gamma$. For $w \in \mathrm{lk}_Q(q_{a_r}) \subset \mathrm{st}(v)$, either $w = v$ or there is an element of W_Γ taking the edge $w \bullet\!\!-\!\!\bullet v$ to $w \bullet\!\!-\!\!\bullet v'$ in Ric_Γ by Lemma 2.26. Similarly, for $w' \in \mathrm{lk}_{r.Q}(q_{b_r}) \subset \mathrm{st}(v)$, either $w' = v$ and $r.w' = r.v = v'$ or there is an element of W_Γ taking the edge $w' \bullet\!\!-\!\!\bullet v$ to $r.w' \bullet\!\!-\!\!\bullet v'$ in Ric_Γ. Thus, $\mathrm{lk}_Q(q_{a_r}) \cup \mathrm{lk}_Q(q_{b_r}) \subset \mathrm{st}(v')$, so q_{a_r} and q_{b_r} are singleton components of $Q - \mathrm{st}(v')$. The second claim of Item (2) forces a contradiction:

$$d_Q(q_{a_r}, q_{b_r}) \leqslant 2 < 3 \leqslant d_\Pi(q_{a_r}, q_{b_r}) \leqslant d_Q(q_{a_r}, q_{b_r})$$ $\square$

The original ladder example in [41] was constructed by hand directly for the graph Γ of Example 6.17. Theorem 6.16 applies to an iterated link double of Γ.

6.4. Examples.

Example 6.17. Figure 18a depicts a graph Γ with its maximal thick join subgraphs shadowed in different colors. The Ric_Γ for $W_\Gamma \rightharpoonup \Pi_\Gamma$ is shown in Figure 18b.

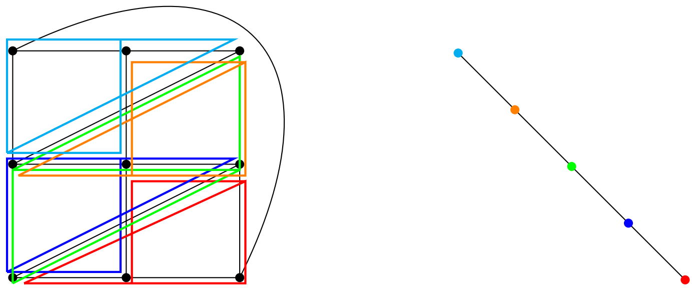

(a) Γ (b) Ric_Γ

FIGURE 18. The graph Γ of Example 6.17 with its maximal thick joins colored and $\mathrm{Ric}_\Gamma = W_\Gamma \backslash \Pi_\Gamma$ with vertices of matching colors.

There is one edge of Γ that does not belong to any thick join, and each of its endpoints belongs to a unique maximal join, corresponding to the two ends of Ric_Γ. The orbit of Ric_Γ by the action of the order 4 subgroup represented by that edge makes a 16–cycle in Π_Γ. We can see this more clearly by passing to the finite-index subgroup obtained by link doubling over the two vertices of the extraordinary edge. The 16–cycle Q is the fundamental domain for the action of the subgroup on Π_Γ, as seen in Figure 20. Notice that all vertices with first coordinate 2 belong to five consecutive maximal join subgraphs; pick two on opposite sides, say $r := 2_{00}$ and $s := 2_{11}$. Then r, s, and Q satisfy Theorem 6.16.

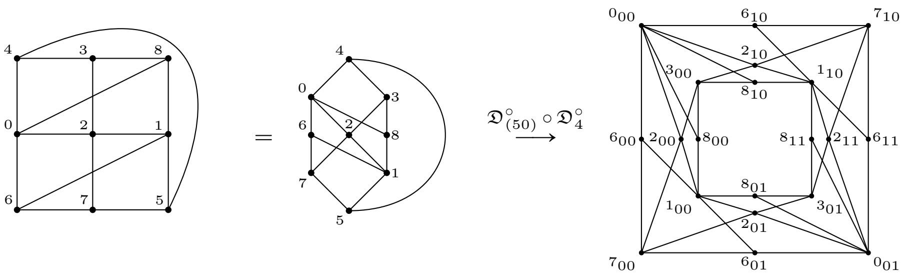

FIGURE 19. Link doubling of Γ in example Example 6.17.

Note also that W_Γ is strongly CFS, has no 2–ended splittings, contains no compliant cycle as in Theorem 7.5, and has totally disconnected Morse boundary, by Proposition 5.9, so contains no 1–ended stable subgroups. The ladder is the only way we know to say Γ is not RAAGedy. ◇

Example 6.18. Figure 21 shows the other 9–vertex graph that is strongly CFS, has no 2–ended splittings, contains no compliant cycle as in Theorem 7.5, and has totally disconnected Morse boundary, by Proposition 5.9, so contains no 1–ended

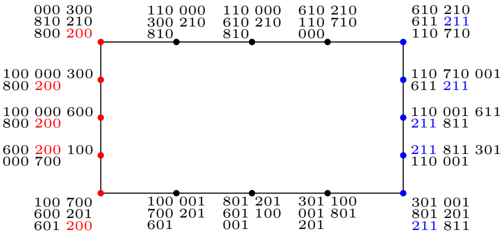

FIGURE 20. Fundamental domain for the action of the subgroup on the MPRG of Example 6.17.

stable subgroups. It has an edge not contained in a thick join. Link double over the two vertices of this edge, as in the previous example.

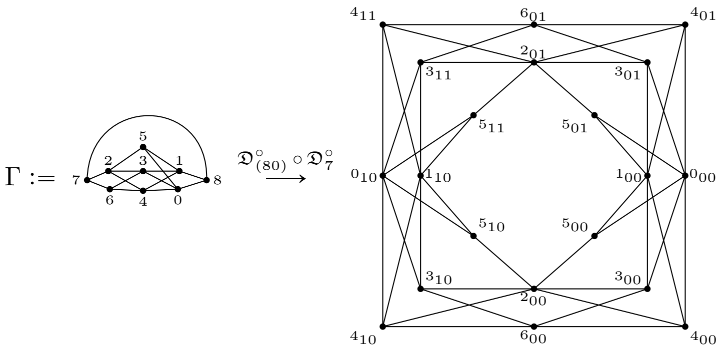

FIGURE 21. The Γ of Example 6.18 and an iterated link double.

The fundamental domain for the action of the finite-index subgroup is shown in Figure 22. Take Q to be the entire fundamental domain. There are choices $r = 2_{01}$ and $s = 2_{00}$ that satisfy Theorem 6.16 for this Q, so Γ is not RAAGedy.

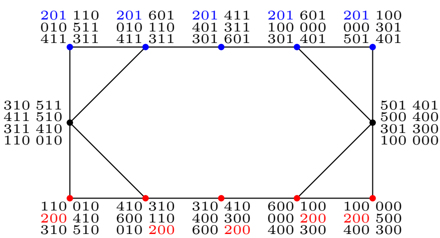

FIGURE 22. Fundamental domain of the MPRG in Example 6.18.

In this example we cannot take Q to be just the outer boundary of the fundamental domain as drawn in Figure 22 because that would not be an induced subgraph, and we also cannot take just the interior octagon to be Q, since then neighboring translates of Q would only intersect in paths of length 2. ◇

7. COMPLIANT CYCLES

In this section we will show that certain cycles of subgraphs in Γ give an obstruction to being RAAGedy. Recall Corollary 2.36: In RAAGs there is a dichotomy, the closest point projection of one standard subcomplex to another has diameter either zero or infinite, while in RACGs it is possible to have finite, nonzero projection diameter. We show that for some RACGs it is possible to build cycles $X_0,\dots,X_{m-1}$ of standard subcomplexes such that consecutive pairs are close, the projection of any single X_i to X_0 is small, and the projection of $\cup_{i=1}^{m-1} X_i$ to X_0 is large. Then we would like to say that if such a RACG were quasiisometric to a RAAG we could derive a contradiction, based on the dichotomy for projection diameters in RAAGs. This is accomplished in Theorem 7.5. For the plan to make sense we need to know that these particular standard subcomplexes X_i in the RACG are sent by quasiisometry close to standard subcomplexes of the RAAG. This additional hypothesis is formalized in the first subsection, where we bootstrap from the quasiisometry invariance of maximal standard product subcomplexes to define a class of 'compliant' subcomplexes.

7.1. **Compliant sets and subcomplexes.** The reader should refresh their memory of the background material on the coarse geometry of standard subcomplexes of RAAGs and RACGs established in Section 2.7 and Section 2.8. We will use it now.

The next statement defines compliant sets of the presentation graph, and corresponding compliant subcomplexes of Σ_Υ. In terms of subcomplexes, the first two bullet points say the collection is closed under adding or removing a finite direct factor. Recalling Proposition 2.34, (a) says the collection is closed under projection, which includes intersection in the special case that $T = \varnothing$. Finally, (b) says the collection is closed under passing to the $\mathbb{R}$ factor of a tree $\times$ $\mathbb{R}$ subcomplex.

Definition 7.1. Let G_Υ be a 2–dimensional, 1–ended RACG or RAAG, and let Σ_Υ be its Davis complex or the universal cover of its Salvetti complex, respectively. We recursively define strata $\mathcal{C}^n_\Upsilon$ of subsets of the vertex set of Υ, inductively extend to higher strata, and finally take $\mathcal{C}_\Upsilon := \bigcup_{n=0}^\infty \mathcal{C}^n_\Upsilon$.

Seed $\mathcal{C}^0_\Upsilon$ by taking the collection of subsets of vertices of Υ that consists of the empty set and the vertex set of each maximal (thick, in the RACG case) join of Υ.

Recursive Phase: Having defined some subsets of $\mathcal{C}^n_\Upsilon$, join and 'unjoin' spherical[8] factors; that is, if $S_0 * S_1 \subset \Upsilon$ and S_0 is spherical then:

- If $S_1 \in \mathcal{C}^n_\Upsilon$ set $\mathcal{C}^n_\Upsilon := \mathcal{C}^n_\Upsilon \cup \{S_0 \sqcup S_1\}$.
- If $S_0 \sqcup S_1 \in \mathcal{C}^n_\Upsilon$ set $\mathcal{C}^n_\Upsilon := \mathcal{C}^n_\Upsilon \cup \{S_1\}$.

Repeat until $\mathcal{C}^n_\Upsilon$ stabilizes, which happens in finitely many steps since Υ only has finitely many subgraphs.

Inductive Phase: Seed $\mathcal{C}^{n+1}_\Upsilon$ by taking $\mathcal{C}^n_\Upsilon$ and all subsets of the following forms:

(a) If $S_0, S_1 \in \mathcal{C}^n_\Upsilon$ and T is a set of vertices of Υ such that either $T = \varnothing$ or $T \not\subset S_0$ and $T \not\subset S_1$, then $S_0 \cap S_1 \cap \bigcap_{t\in T} \mathrm{lk}(t) \in \mathcal{C}^{n+1}_\Upsilon$.
(b) If $S_0 \sqcup S_1 \in \mathcal{C}^n_\Upsilon$ with $S_0 * S_1 \subset \Upsilon$ such that Σ_{S_0} is a line and Σ_{S_1} is a bushy tree then $S_0 \in \mathcal{C}^{n+1}_\Upsilon$.

Perform the Recursive Phase, joining and unjoining spherical factors to elements of $\mathcal{C}^{n+1}_\Upsilon$ until $\mathcal{C}^{n+1}_\Upsilon$ stabilizes.

For $S \in \mathcal{C}_\Upsilon$, let its *index* be $\mathrm{ind}(S) := \min\{n \mid S \in \mathcal{C}^n_\Upsilon\}$.

We call $\mathcal{C}_\Upsilon$ the *compliant subsets* of Υ and call the standard subcomplexes $g\Sigma_S$, for $g \in G_\Upsilon$ and $S \in \mathcal{C}_\Upsilon$, the *compliant subcomplexes* of Σ_Υ.

Proposition 7.2. *For all L, A, N, and n there exists C_n with the following property. Let G_Υ be a 2–dimensional, 1–ended RACG or RAAG, and let Σ_Υ be its Davis complex or the universal cover of its Salvetti complex, respectively, and define $G_{\Upsilon'}$ and $\Sigma_{\Upsilon'}$ similarly. Suppose $|\Upsilon| \leqslant N$. For every $g \in G_\Upsilon$, $S \in \mathcal{C}^n_\Upsilon$, and*

[8]A subset of vertices of the presentation graph of a Coxeter group is *spherical* if the special subgroup they generate is finite. For RACGs this happens if and only if the vertex set is a clique. For RAAGs no nontrivial special subgroup is finite, so this phase does not apply.

(L, A)–quasiisometry $\phi\colon \Sigma_\Upsilon \to \Sigma_{\Upsilon'}$ there exists $g' \in G_{\Upsilon'}$ and $S' \in \mathcal{C}^n_{\Upsilon'}$ such that $d_{Haus}(\phi(g\Sigma_S), g'\Sigma_{S'}) \leqslant C_n$.

Furthermore, $\mathcal{C}_\Upsilon = \mathcal{C}^{2N}_\Upsilon$, so we can take $C := \max\{C_n \mid 0 \leqslant n \leqslant 2N\}$ as a uniform Hausdorff distance bound for all of $\mathcal{C}_\Upsilon$.

Proof. Let C_0' be the constant of Theorem 2.22 with respect to L and A. Let $C_0 := \max\{C_0', 2L + A\}$. Consider $g\Sigma_S$ for $g \in G_\Upsilon$ and $S \in \mathcal{C}^0_\Upsilon$. If S is the vertex set of a maximal (thick) join of Υ then Theorem 2.22 and Lemma 2.20 combine to say there exist $g' \in G'$ and S' the vertex set of a maximal (thick) join of Υ' such that $d_{Haus}(\phi(g\Sigma_S), g'\Sigma_{S'}) \leqslant C_0' \leqslant C_0$. By definition $S' \in \mathcal{C}^0_{\Upsilon'}$. If S is such that G_S is finite then the diameter of Σ_S is 0 if G_Υ is a RAAG or at most 2 if G_Υ is a 2–dimensional RACG. Thus, for any $g' \in G_{\Upsilon'}$ such that $g'\Sigma_\varnothing \in \phi(g\Sigma_S)$ we have $\varnothing \in \mathcal{C}^0_{\Upsilon'}$ and $d_{Haus}(\phi(g\Sigma_S), g'\Sigma_\varnothing) \leqslant 2L + A \leqslant C_0$. In the RAAG case this is all of $\mathcal{C}^0_\Upsilon$. In the RACG case it is also all of $\mathcal{C}^0_\Upsilon$, since the triangle-free hypothesis implies that we cannot add a cone vertex to a maximal thick join.

Now induct on index: supposing the proposition is true for all indices up to and including n, we extend it to $n + 1$.

Induction step (a) corresponds to projection between subcomplexes. Suppose $S := S_0 \cap S_1 \cap \bigcap_{t \in T} \mathrm{lk}(t) \in \mathcal{C}^{n+1}_\Upsilon$ for some $S_0, S_1 \in \mathcal{C}^n_\Upsilon$ and T either empty or not contained in either of S_0 or S_1. If $T = \varnothing$ define $h = 1 \in G_\Upsilon$. Then $S = S_0 \cap S_1$, so $\Sigma_S = \Sigma_{S_0} \cap \Sigma_{S_1} = \pi_{\Sigma_{S_0}}(h\Sigma_{S_1})$. If $T \neq \varnothing$ then there exists $t_0 \in T - S_0$ and $t_1 \in T - S_1$. Let $h \in G_\Upsilon$ be any shortest word that begins with t_0 and ends with t_1 and uses every letter in T at least once. So $|h| = |T|$ if $|T| = 1$ or $t_0 \neq t_1$, and $|h| = |T| + 1$ if $|T| > 1$ but $t_0 = t_1$. Again we have $\Sigma_S = \pi_{\Sigma_{S_0}}(h\Sigma_{S_1})$. Define $h_0 := 1 \in G_\Upsilon$ and $h_1 := h$.

By the induction hypothesis, for every $g \in G_\Upsilon$ and every (L, A)–quasiisometry $\phi\colon \Sigma_\Upsilon \to \Sigma_{\Upsilon'}$ there are $S_0', S_1' \in \mathcal{C}^n_{\Upsilon'}$ and $g_0', g_1' \in G_{\Upsilon'}$ such that:

$$d_{Haus}(\phi(gh_i\Sigma_{S_i}), g_i'\Sigma_{S_i'}) \leqslant C_n$$

Let g' be a shortest element of $G_{S_0'}((g_0')^{-1}g_1')G_{S_1'}$, so $|g'| = d_{\Sigma_{\Upsilon'}}(g_0'\Sigma_{S_0'}, g_1'\Sigma_{S_1'})$ and if g' is nontrivial then it begins with a letter not in S_0' and ends with a letter not in S_1'. Let T' be the letters appearing in g', so that for $S' := S_0' \cap S_1' \cap \bigcap_{t \in T'} \mathrm{lk}(t) \in \mathcal{C}^{n+1}_{\Upsilon'}$ we have $g_0'\Sigma_{S'} = \pi_{g_0'\Sigma_{S_0'}}(g_1'\Sigma_{S_1'})$.

Corollary 2.32 says $g\Sigma_S = g\pi_{\Sigma_{S_0}}(h\Sigma_{S_1}) \stackrel{c}{=} (g\Sigma_{S_0} \stackrel{c}{\cap} gh\Sigma_{S_1})$, and Lemma 2.7 says ϕ sends this coarse intersection to within bounded Hausdorff distance of $\phi(g\Sigma_{S_0}) \stackrel{c}{\cap} \phi(gh\Sigma_{S_1})$. But using the induction hypothesis and Corollary 2.32 gives:

$$\begin{aligned} \phi(g\Sigma_{S_0}) \stackrel{c}{\cap} \phi(gh\Sigma_{S_1}) &\stackrel{c}{=} g_0'\Sigma_{S_0'} \stackrel{c}{\cap} g_1'\Sigma_{S_1'} \\ &\stackrel{c}{=} \pi_{g_0'\Sigma_{S_0'}}(g_1'\Sigma_{S_1'}) \\ &= g_0'\Sigma_{S'} \end{aligned}$$

Moreover, the coarse equivalences provided by Corollary 2.32 and Lemma 2.7 depend on L and A and the distances between the sets of the coarse intersections, so on $d(g\Sigma_{S_0}, gh\Sigma_{S_1}) = |h| \leqslant |T| + 1 \leqslant |\Upsilon| + 1$ and on:

$$d(g_0'\Sigma_{S_0'}, g_1'\Sigma_{S_1'}) \leqslant 2C_n + d(\phi(g\Sigma_{S_0}), \phi(gh\Sigma_{S_1})) \leqslant 2C_n + A + L(|\Upsilon| + 1)$$

Thus, $d_{Haus}(\phi(g\Sigma_S), g'\Sigma_{S'})$ is bounded by a constant C_{n+1}' depending on L, A, C_n, and $|\Upsilon|$, hence on L, A, N, and n.

For induction step (b), suppose that $S_0 \in \mathcal{C}^{n+1}_\Upsilon$ is obtained from $S \in \mathcal{C}^n_\Upsilon$ by virtue of S decomposing as $S = S_0 \sqcup S_1$, with $S_0 * S_1 \subset \Upsilon$ where Σ_{S_0} is line and Σ_{S_1} is a bushy tree. By the induction hypothesis, for every $g \in G_\Upsilon$ and every (L, A)–quasiisometry $\phi\colon \Sigma_\Upsilon \to \Sigma_{\Upsilon'}$ there are $g' \in G_{\Upsilon'}$ and $S' \in \mathcal{C}^n_{\Upsilon'}$ such that $d_{Haus}(\phi(g\Sigma_S), g'\Sigma_{S'}) \leqslant C_n$. Then $\pi_{g'\Sigma_{S'}} \circ \phi|_{g\Sigma_S}$ is a quasiisometry from a (bounded valence bushy tree) $\times\mathbb{R}$ to $g'\Sigma_{S'}$ whose quasiisometry constants depend only on L, A, and C_n. Thus, $\Sigma_{S'}$ is also a (bounded valence bushy tree) $\times\mathbb{R}$, which implies $S'=S'_0\sqcup S'_1$, where $S'_0*S'_1\subset\Upsilon'$, with $\Sigma_{S'_0}$ a line and $\Sigma_{S'_1}$ a bushy tree. Thus, $S'_0\in\mathcal{C}^{n+1}_{\Upsilon'}$. Furthermore, since the $\mathbb{R}$–factor defines the only coarse equivalence class of separating quasiline in a (bounded valence bushy tree) $\times\mathbb{R}$, it follows from [75] that up to post-composition by an element of $G_{S'_1}$, we have that $g\Sigma_{S_0}$ is sent C''_{n+1}–Hausdorff close to $g'\Sigma_{S'_0}$, with C''_{n+1} depending on L, A, and C_n, hence on L, A, N, and n.

In the RAAG case this is all of $\mathcal{C}^{n+1}_{\Upsilon}$. In the RACG case the Recursive Phase allows for the possibility of taking joins with spherical subsets, which, since $n+1>0$ and Υ is triangle-free, simply means adding or removing a cone vertex. Corollary 2.38 says that if $S_0*S_1\subset\Upsilon$ and $S'_0*S_1\subset\Upsilon$ with Σ_{S_0} and $\Sigma_{S'_0}$ finite then Σ_{S_1}, $\Sigma_{S_0*S_1}$, $\Sigma_{S'_0*S_1}$ are pairwise at Hausdorff distance at most 2. Suppose we know that one of these sets is in $\mathcal{C}^{n+1}_{\Upsilon}$ by one of the previous constructions, so its ϕ–image is $\max\{C'_{n+1},C''_{n+1}\}$ close to some $g'\Sigma_{S'}$ for $g'\in G_{\Upsilon'}$ and $S'\in\mathcal{C}^n_{\Upsilon'}$. Then the ϕ–images of the other two are within Hausdorff distance $2L+A+\max\{C'_{n+1},C''_{n+1}\}$ of the same $g'\Sigma_{S'}$.

We have shown it suffices to take $C_{n+1}:=2L+A+\max\{C'_{n+1},C''_{n+1}\}$.

For the further statement, in a RAAG there are no nontrivial spherical subsets so there is no Recursive Phase, only operations (a) and (b) of the Inductive Phase of Definition 7.1 apply, and these both decrease the size of the sets involved, so the sequence $\mathcal{C}^0_{\Upsilon}\subset\mathcal{C}^1_{\Upsilon}\subset\cdots$ stabilizes after a number of steps bounded above by the size of largest set in $\mathcal{C}^0_{\Upsilon}$, which is the largest join in Υ, whose size is at most $|\Upsilon|\leqslant N$. In a RACG we can add a cone vertex, increasing the size of a compliant set without changing its index. For operation (b), it takes at least three generators to make a bushy tree, so if Σ_{S_0} is a line and Σ_{S_1} is a bushy tree and $\Sigma_{S_0*S_1}$ is compliant then S_0 is compliant and contains fewer elements than $S_0\sqcup S_1$, and a cone on S_0 is also compliant and contains fewer elements than $S_0\sqcup S_1$.

Now consider the case that $S_0,S_1\in\mathcal{C}^n_{\Upsilon}$ and $S:=S_0\cap S_1\cap\bigcap_{t\in T}\mathrm{lk}(t)\in\mathcal{C}^{n+1}_{\Upsilon}-\mathcal{C}^n_{\Upsilon}$ and $S'=S*\{c\}$ is a cone on S. Since Υ is triangle-free, S is an anticlique. Suppose $|S'|\geqslant|S_0|$ and $|S'|\geqslant|S_1|$. Neither of these can be strict, since $S\notin\mathcal{C}^n_{\Upsilon}$ implies $S\subsetneq S_0$ and $S\subsetneq S_1$. So $S_0=S\cup\{a_0\}$ and $S_1=S\cup\{a_1\}$. Furthermore, $S\notin\mathcal{C}^n_{\Upsilon}$ implies a_i is not a cone vertex of S_i, which implies S_i contains no cone vertex, since S is an anticlique. Thus, we entertain the notion that $|S'|=|S_0|=|S_1|$, but such an inconvenience does not propagate any further, since this S' having a cone vertex means it cannot be the S_0 or S_1 in an iterate of this paragraph. Thus, we conclude that the size of a set in $\mathcal{C}^{n+2}_{\Upsilon}$ is strictly less than the sets of $\mathcal{C}^n_{\Upsilon}$ from which it is derived. Thus, $\mathcal{C}_{\Upsilon}=\mathcal{C}^{2N}_{\Upsilon}$. □

Remark. One might guess a more general definition of 'compliant subcomplex' as one that is sent uniformly Hausdorff close to a standard subcomplex by any quasiisometry. Definition 7.1 represents the only examples we know that satisfy this property, but the conclusion of Proposition 7.2 is stronger: images are not just close to standard, they are close to the specific standard subcomplexes coming from $\mathcal{C}_{\Upsilon'}$.

Remark. There is a sense in which the class of products of trees is quasiisometrically rigid [78], but this involves projection to the factors. If $S_0\sqcup S_1\in\mathcal{C}_{\Upsilon}$ with $S_0*S_1\subset\Upsilon$, Σ_{S_0} a bushy tree, and Σ_{S_1} a tree, then we cannot automatically conclude $S_0\in\mathcal{C}_{\Upsilon}$. For example, the automorphism $x\mapsto xz$, $y\mapsto y$, $z\mapsto z$ of the RAAG $F_2\times\mathbb{Z}=\langle x,y\rangle\times\langle z\rangle$ does not send the standard subcomplex $\Sigma_{\{x,y\}}$ close to a standard subcomplex.

7.2. **Coarse geometry of compliant cycles.** We would like the next result to say that if a configuration of compliant subcomplexes exists in a RACG then it is not RAAGedy, but the actual outcome is more subtle: it says if the configuration exists and if there is a quasiisometry to a RAAG then there exists a subcomplex X' with a particular set of properties. The trailing corollary says if there is no such subcomplex then the group was not RAAGedy. In Theorem 7.5 we will translate these conditions to Γ, with one set of conditions describing the configuration of compliant sets, and a second set of conditions implying that the mystery subcomplex X' does not exist. If both sets of conditions are true then the group is not RAAGedy.

Theorem 7.3 (Compliant cycles)**.** *Let Γ be an incomplete triangle-free graph without separating cliques. Suppose there exists $B \geqslant 0$ such that for every sufficiently large r there are compliant subcomplexes $X_0, X_1, \dots, X_{n-1}$ of the Davis complex Σ_Γ, for some $n \geqslant 3$, such that for all $0 \leqslant i < n$ there is a vertex $b_i \in X_i \cap \bar{N}_B(X_{(i+1)\%n})$ and the following conditions are satisfied:*

(i) $d(b_0, b_{n-1}) \geqslant r$.
(ii) X_0 has bounded coarse intersection with every other X_i.

If Γ is RAAGedy, then for all sufficiently large r there is a quasiisometry $\psi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ taking X_0 to within bounded Hausdorff distance of a compliant subcomplex X' that contains a quasigeodesic edge path $\bar{\gamma}'$ such that:

(1) $\bar{\gamma}'$ is contained in a bounded neighborhood of X_0.
(2) $\bar{\gamma}'$ comes close to b_{i_0} for some $1 \leqslant i_0 \leqslant n-2$.

Furthermore, the quasiisometry constants of ψ, the quasigeodesic constants of $\bar{\gamma}'$, $d_{Haus}(X', \psi(X_0))$, and $d(\bar{\gamma}', b_{i_0})$ are independent of r.

In applications we will arrange that for all $1 \leqslant i \leqslant n-2$, $d(b_i, X_0) \geqslant r$, since if some b_i is close to X_0 we could take $\psi = \mathrm{Id}_{\Sigma_\Gamma}$, and the theorem would be vacuous.

Corollary 7.4. *Suppose the hypotheses of Theorem 7.3 are satisfied and there does not exist a quasiisometry $\psi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ taking X_0 to within bounded Hausdorff distance of a compliant subcomplex that has unbounded coarse intersection with X_0 and comes within the required distance of some b_i. Then Γ is not RAAGedy.*

Proof of Theorem 7.3. Suppose that $\phi\colon W_\Gamma \to A_\Delta$ and $\bar{\phi}\colon A_\Delta \to W_\Gamma$ are coarse inverse (L, A)–quasiisometries between W_Γ and some RAAG A_Δ. Let C be the constant of Proposition 7.2 for this L and A and $N = \max\{|\Gamma|, |\Delta|\}$.

Assume r is large compared to A, B, C, and L; specifically, $r > L(6C+4A+3LB)$ is the estimate that will be needed later. Choose X_i and b_i with respect to this r.

Since each X_i is compliant, by Proposition 7.2 there is a compliant subcomplex Y_i of Σ_Δ at Hausdorff distance at most C from $\phi(X_i)$. Let δ_i be a path in Y_i from a point δ_i^- of Y_i closest to $\phi(b_{(i-1)\%n})$ to a point δ_i^+ of Y_i closest to $\phi(b_i)$. Let ϵ_i be a geodesic from δ_i^+ to $\delta_{(i+1)\%n}^-$. Note that $|\epsilon_i| \leqslant 2C + LB + A$ for all i, and that δ_0 and the concatenation $\epsilon_0 + \delta_1 + \epsilon_1 + \delta_2 + \epsilon_2 + \cdots + \delta_{n-1} + \epsilon_{n-1}$ are paths with the same endpoints δ_0^+ and δ_0^-. From the assumptions $d(b_0, b_{n-1}) \geqslant r$ and $d(b_{n-1}, X_0) \leqslant B$ we get an estimate:

$$d(\delta_0^+, \delta_0^-) \geqslant d(\phi(b_0), \phi(b_{n-1})) - LB - A - 2C \geqslant r/L - LB - 2A - 2C$$

Combinatorial closest point projection to Y_0 in Σ_Δ is a Lipschitz, hence combinatorial, map, so its sends the concatenation of δ and ϵ paths to an edge path in Y_0 from δ_0^+ to δ_0^-. Since $X_i \overset{c}{\cap} X_0$ is bounded, so is $Y_i \overset{c}{\cap} Y_0$, but for standard subcomplexes of a RAAG this means that $\pi_{Y_0}(Y_i)$ is a single vertex. Thus, for all $i \neq 0$, $\pi_{Y_0}(\delta_i)$ is a single vertex.

Since we assumed $r > L(6C + 4A + 3LB)$, we have:

$$d(\delta_0^-, \delta_0^+) > 2(2C + LB + A) \geqslant |\epsilon_0| + |\epsilon_{n-1}| \geqslant |\pi_{Y_0}(\epsilon_0)| + |\pi_{Y_0}(\epsilon_{n-1})|$$

This means that $\pi_{Y_0}(\epsilon_0)$ and $\pi_{Y_0}(\epsilon_{n-1})$ alone are not long enough to reach from δ_0^+ to δ_0^-, so $\pi_{Y_0}(\delta_1 + \epsilon_1 + \cdots + \epsilon_{n-2} + \delta_{n-1})$ is a nontrivial edge path in Y_0. All of the

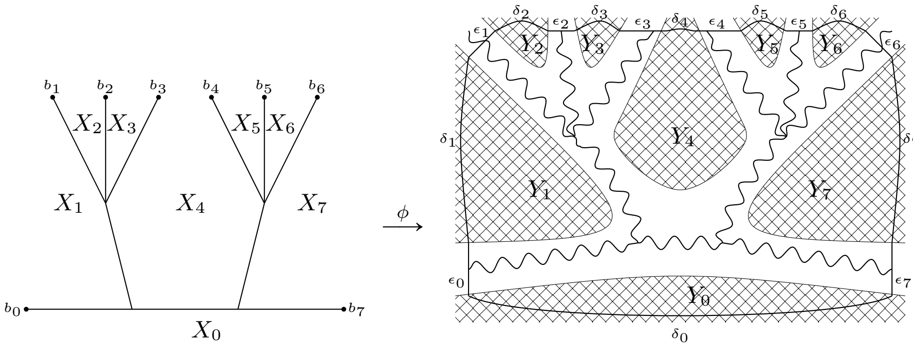

FIGURE 23. Quasiisometry carrying compliant cycle into a RAAG. Each compliant set X_i on the left, visualized as complementary regions of a tree in the plane, is sent by ϕ close to a compliant subcomplex Y_i. Consecutive Y_i, Y_{i+1} may no longer come B–close to each other, but they both come close to $\phi(b_i)$, so the path ϵ_i crossing between them is uniformly short.

δ_i project to single vertices, so there exists some $1 \leqslant i_0 \leqslant n-2$ such that ϵ_{i_0} has nontrivial projection to Y_0. Thus, there are parallel edges $e \in Y_0$ and $e' \in \epsilon_{i_0}$.

Recalling Lemma 2.33, let $b \in A_\Delta$ be represented by a word labelling a geodesic from e to e'. The fact that e and e' are parallel means that that they have the same label a, and that a and b commute. Up to translation, we may assume that e is the edge between the vertices labelled 1 and a in Σ_Δ, in which case, $e' = be$. Let $Y' := bY_0$. Since it is standard, Y' contains the entire standard geodesic γ' containing e', just as Y_0 contains the entire standard geodesic γ containing e, and these geodesics are parallel, since every edge in both geodesics is labelled a and we have an element b commuting with a realizing the parallel translation.

Compliance of subcomplexes is preserved by group translation, by definition, so applying Proposition 7.2 to Y' gives a compliant subcomplex X' of Σ_Γ at Hausdorff distance at most C from $\bar{\phi}(Y')$. Define $\psi := \bar{\phi} \circ (b\cdot) \circ \phi$, where $b\cdot$ is left-multiplication by b, so that $\psi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ is a quasiisometry taking X_0 within bounded Hausdorff distance of X'. The quasiisometry constants of ψ only depend on L and A, while $d_{Haus}(\psi(X_0), X')$ depends only on L, A, and C.

Push the vertices of γ' first to Σ_Γ via $\bar{\phi}$, and then into X' via $\pi_{X'}$. Since X' is convex and $d_{Haus}(X', \bar{\phi}(Y')) \leqslant C$, a standard connect-the-dots argument says there is a quasigeodesic edge path $\bar{\gamma}'$ contained in X' whose quasigeodesic constants depend only on L, A, and C, and that is bounded Hausdorff distance from $\bar{\phi}(\gamma')$ and $\bar{\phi}(\gamma)$, hence contained in a bounded neighborhood of $X_0 \stackrel{c}{=} \bar{\phi}(Y_0)$.

The following estimate shows that $d(\bar{\gamma}', b_{i_0})$ is bounded above, independent of r:

$$
\begin{aligned}
d(\bar{\gamma}', b_{i_0}) &\leqslant d(\pi_{X'}(\bar{\phi}(\gamma')), b_{i_0}) \\
&\leqslant C + d(\bar{\phi}(\gamma'), b_{i_0}) \\
&\leqslant C + A + d(\bar{\phi}(\gamma'), \bar{\phi}\phi(b_{i_0})) \\
&\leqslant C + 2A + Ld(\gamma', \phi(b_{i_0})) \\
&\leqslant C + 2A + LC + Ld(\gamma', \delta_{i_0}^+) \\
&\leqslant C + 2A + LC + L|\epsilon_{i_0}| \\
&\leqslant C + 2A + LC + L(2C + LB + A) \qquad \square
\end{aligned}
$$

7.3. **Graphical criteria.** The next result gives practical, graphical criteria for applying Theorem 7.3. The point is that conditions (1)–(5) imply the hypotheses of Theorem 7.3, but conditions (a)–(b) say that there is no possible target subcomplex for a quasiisometry as in Corollary 7.4, so Γ cannot be RAAGedy.

Theorem 7.5. *Let Γ be an incomplete triangle-free graph without separating cliques, and let $\mathcal{C}_\Gamma$ be the compliant subsets of Γ, as in Definition 7.1. Suppose there exist, for $q \geqslant 1$ and all $0 \leqslant i \leqslant q-1$, sets $S_i \in \mathcal{C}_\Gamma$ and paths $P_i := (a_{i,0}, a_{i,1}, \dots, a_{i,\ell(i)-1})$ in Γ, containing $\ell(i) \geqslant 1$ vertices, such that all of the following hold (for S_j, P_j, and $a_{j,k}$ we always implicitly take j mod q and k mod $\ell(j)$):*

(1) The paths P_i are disjoint, and $P := \bigcup_j P_j$ is an induced subgraph of Γ.
(2) $\forall i$, $S_i \cap P = \{a_{i-1,\ell(i-1)-1}, a_{i,0}\}$, which are the last vertex of P_{i-1} and the first vertex of P_i.
(3) If $q = 1$ then $P_0 \not\subset S_0$.
(4) $\forall i \neq 0$, $S_i \cap S_0$ is a clique.
(5) $\forall (i,j) \notin \{(0,0), (q-1, \ell(q-1)-1)\}$, $\mathrm{lk}(a_{i,j}) \cap S_0$ is a clique.

Assume that no proper subset of S_0 belongs to $\mathcal{C}_\Gamma$ and contains the first and last vertices $a_{0,0}$ and $a_{q-1,\ell(q-1)-1}$ of P. Let Γ_0 be the subgraph of Γ induced by S_0. If either of the following are true then Γ is not RAAGedy:

(a) Γ_0 is square-free.
(b) No $S' \in \mathcal{C}_\Gamma$ spans a subgraph Γ' satisfying all of the following conditions:
  (I) $\Gamma_0 \cap \Gamma'$ is incomplete.
  (II) $\Gamma' \cap P$ is incomplete.
  (III) $\{a_{0,0}, a_{q-1,\ell(q-1)-1}\} \not\subset \Gamma'$.
  (IV) There exists a quasiisometry $\psi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ with $\psi(\Sigma_{\Gamma_0}) \stackrel{c}{=} \Sigma_{\Gamma'}$.

The minimality condition on S_0 is justified because a proper subset of S_0 belonging to $\mathcal{C}_\Gamma$ and containing $\{a_{0,0}, a_{q-1,\ell(q-1)-1}\}$ satisfies hypotheses (1)–(5) for the same choices of S_i and P_i.

In Section 7.4 we give some ways to rule out condition (IV).

Proof. Let C_m be the cycle graph of length $m \geqslant 3$. Its commutator complex is a closed surface, since it is a connected square complex such that the link of every vertex is a circle. Its Euler characteristic is $2^m - m2^{m-1} + m2^{m-2} = 2^{m-2}(4-m)$. Think of $P = \sqcup_i P_i$ as a subgraph of C_m for $m := \sum_{i=0}^{q-1} \ell(i)$, where the vertices $a_{i,j}$ are ordered lexicographically, so that for each i there is an edge $a_{i,\ell(i)-1} \multimap\!\!\bullet\, a_{i+1,0}$ of C_m that is not an edge of P. Thus, the commutator complex of P is homotopy equivalent to the commutator complex of C_m after puncturing each square labelled by a commutator $[a_{i,\ell(i)-1}, a_{i+1,0}]$. There are $q2^{m-2}$ such squares, so the Euler characteristic of the commutator complex of P is $2^{m-2}(4-m-q)$. We claim this is a negative number, so W_P is a virtually a nonAbelian free group. To see this, consider that the alternative is that either $q = 1$ and $\ell(0) \leqslant 3$ or $q = 2$ and $\ell(0) = \ell(1) = 1$. If $q = 1$ then Hypothesis (3) says $\ell(0) > 2$, but it cannot be that $q = 1$ and $\ell(0) = 3$ because this would either give a triangle in Γ or contradict Hypothesis (5). We cannot have $q = 2$ and $\ell(0) = \ell(1) = 1$ because this would either violate Hypothesis (1), if $a_{0,0}$ and $a_{1,0}$ are adjacent, or Hypothesis (4) if not.

Σ_P is the universal cover of a closed surface with some open faces removed, so Σ_P admits a planar embedding in which all of its vertex links are copies of P ordered as in C_m or its reverse, and whose boundary components are the bicolored geodesics whose colors are $a_{i,\ell(i)-1}a_{i+1,0}$, for each i; that is, the boundary components are the lifts of the boundaries of the missing faces.

If $\{a\} = P_i$ is an isolated vertex of P then edges of Σ_P labelled a belong to two different components of $\partial\Sigma_P$, one bicolored with a and the last vertex of P_{i-1} and one bicolored with a and the first vertex of P_{i+1}. If $a_{i,0}$ is the first vertex of a non-singleton component P_i of P then an edge e of Σ_P labelled $a_{i,0}$ is contained in a unique component of $\partial\Sigma_P$ and is a face of a unique square whose sides are colored $a_{i,0}$ and $a_{i,1}$. The opposite face of this square is the unique edge of Σ_P parallel to e.

Fix an identity vertex 1 of Σ_P, let γ_0 be the $a_{0,0}a_{q-1,\ell(q-1)-1}$ bicolored geodesic through 1, parameterized by arclength with $\gamma_0(0) = 1$. For any $r > 0$, consider the following set, where the overbar means closure in the 1–skeleton:

$$\hat{\Sigma}_P := \bar{\mathcal{N}}_r(\gamma_0) \cap \overline{\pi_{\gamma_0}^{-1}(\gamma_0(0,r))}$$

This should be imagined as an r–tubular neighborhood in Σ_P of the subsegment of γ_0 of length r starting at 1. We specify a collection of components of $\partial\Sigma_P$ that contains all of the vertices of $\hat{\Sigma}_P$. First, include every boundary component of Σ_P that contains an edge in $\hat{\Sigma}_P$. Then, for each vertex x of $\hat{\Sigma}_P$ that is not contained in one of these boundary components, choose any one of the components of $\partial\Sigma_P$ containing x. Clockwise with respect to the planar embedding, starting from γ_0, consecutively number these boundary components γ_0, $\gamma_1, \ldots, \gamma_{n-1}$. Orient each boundary component γ_i accordingly, i.e. choose a parametrization of γ_i so that the vertex at which γ_i enters $\hat{\Sigma}_P$ appears, in the clockwise ordering, before the vertex at which γ_i leaves $\hat{\Sigma}_P$. Let b_i be the last vertex of γ_i in $\hat{\Sigma}_P$. Every vertex of Σ_P lies on some boundary component. Thus, the consecutive ordering of the boundary components γ_i and the choice of the vertices b_i yield $B := 2 \geqslant d(b_i, \gamma_{i+1})$. See Figure 24 and Figure 25.

Next we argue that all the b_i with $i \notin \{0, n-1\}$ are at distance r from γ_0. More specifically we show that the first vertex of $\gamma_1 \cap \hat{\Sigma}_P$ is 1-close to γ_0, the last vertex b_{n-1} of $\gamma_{n-1} \cap \hat{\Sigma}_P$ is 1-close to γ_0 and all the other endpoints of $\gamma_i \cap \hat{\Sigma}_P$ with $i \neq 0$ are at distance r from γ_0.

The region $\hat{\Sigma}_P$ is bounded, so each geodesic γ_i intersects it in a bounded subinterval (possibly a single vertex). By construction, the extreme vertices of $\hat{\Sigma}_P$ are those that project to either $\gamma_0(0)$ or $\gamma_0(r)$ and those that are at distance r from γ_0. In particular, each endpoint of each $\gamma_i \cap \hat{\Sigma}_P$ either is at distance r from γ_0 or projects to one of the endpoints of $\gamma_0 \cap \hat{\Sigma}_P$. By construction, if $\gamma_0(0) \in \pi_{\gamma_0}(\gamma_i)$ then $\pi_{\gamma_0}(\gamma_i)$ contains the edge $\gamma_0(0,1)$. This edge is either not a face of a square of Σ_P, in which case it is also contained in γ_1, or it is a face of a unique square whose opposite face is a boundary edge, so is contained in γ_1. Thus, γ_1 is the only γ_i for $i \neq 0$ with $\gamma_0(0) \in \pi_{\gamma_0}(\gamma_i)$. Similarly, γ_{n-1} is the only γ_i for $i \neq 0$ such that $\gamma_0(r) \in \pi_{\gamma_0}(\gamma_i)$. Accordingly, the first vertex of $\gamma_1 \cap \hat{\Sigma}_P$ and b_{n-1} are 1-close to γ_0 and any other endpoint of any $\gamma_i \cap \hat{\Sigma}_P$ with $i \neq 0$ leaves the bounded region $\hat{\Sigma}_P$ through a vertex at distance r from γ_0.

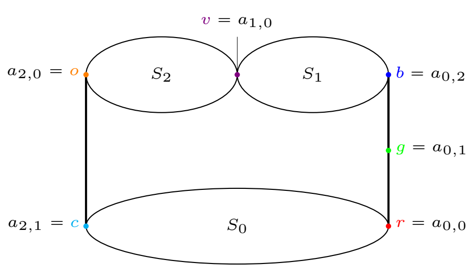

FIGURE 24. Example cycle S_0, P_0, S_1, P_1, S_2, P_2.

The inclusion of Σ_P into Σ_Γ is an isometric embedding, by Hypothesis (1). Hypothesis (2) implies that for each γ_i there is a unique translate X_i of one of the

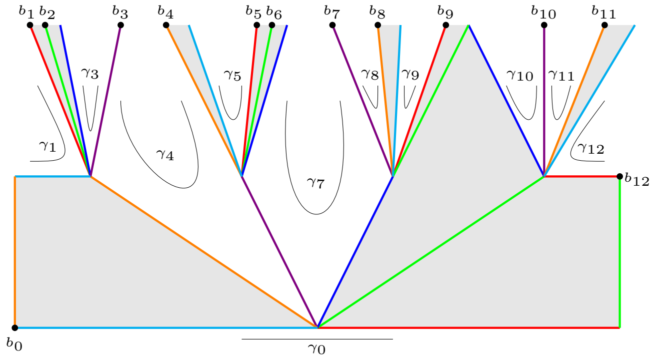

FIGURE 25. $\hat{\Sigma}_P$ for $r = 2$ and P as in Figure 24.

Σ_{S_j} that contains γ_i, since for each j only S_j contains both $a_{j,0}$ and $a_{j-1,\ell(j-1)-1}$. This also implies that $X_i \cap \Sigma_P = \gamma_i$. Furthermore, $b_i \in X_i$ and $B \geqslant d(b_i, \gamma_{i+1}) \geqslant d(b_i, X_{i+1})$ and $d(b_0, b_{n-1})$ is either r or $r+1$. We have also arranged the b_i for $i \notin \{0, n-1\}$ to be far from X_0, since $X_0 \cap \Sigma_P = \gamma_0$ and $d_{\Sigma_P}(b_i, \gamma_0) = r$ imply $d_{\Sigma_\Gamma}(b_i, X_0) = r$. This follows because $X_0 \cap \Sigma_P \neq \varnothing$, so the projection of the convex subcomplex Σ_P to the convex subcomplex X_0 is just their intersection, which is γ_0, so the closest point of X_0 to b_i is a point of γ_0.

To see that the hypotheses of Theorem 7.3 are satisfied, it remains to show that X_0 has bounded coarse intersection with every other X_i.

By Hypothesis (2), if $\mathcal{H}$ is a wall dual to an edge colored by some $a_{j,k}$ that is not on γ_i then $\mathcal{H}$ is not dual to any edge in X_i.

First, suppose γ_0 and γ_i are disjoint, so there is a nontrivial shortest Σ_P–path connecting them. Let $\mathcal{H}$ be the wall dual to the first edge of that path, let $\mathcal{H}^+$ be the halfspace of $\mathcal{H}$ containing 1, let $\mathcal{H}^-$ be the complementary halfspace. Then $X_0 \subset \mathcal{H}^+$. By Proposition 2.31, $\mathcal{H}$ separates X_0 and X_i, so $X_i \subset \mathcal{H}^-$. By Hypothesis (5), the label of $\mathcal{H}$ commutes with at most one edge label of X_0. It then follows from Corollary 2.32 and Lemma 2.33, that $X_0 \overset{c}{\cap} \mathcal{H}^-$ is bounded, so $X_0 \overset{c}{\cap} X_i$ is bounded.

If γ_0 and γ_i are distinct but not disjoint, then X_i is a translate of Σ_{S_j} for some $j \neq 0$, in which case $X_0 \overset{c}{\cap} X_i$ is bounded as a consequence of Hypothesis (4).

We have shown the hypotheses of Theorem 7.3 are satisfied, so if Γ is RAAGedy then there is a quasiisometry $\psi \colon \Sigma_\Gamma \to \Sigma_\Gamma$ taking X_0 to within bounded Hausdorff distance of a compliant subcomplex X' containing a quasigeodesic edge path $\bar{\gamma}'$ that is contained in a bounded neighborhood of X_0 and also comes close to b_{i_0} for some $1 \leqslant i_0 \leqslant n-2$. We will show that either of Hypotheses (a) or (b) leads to a contradiction, so Γ cannot have been RAAGedy.

Suppose $X' = w\Sigma_{S'}$ for some $w \in W_\Gamma$ and $S' \in \mathcal{C}_\Gamma$, and let Γ' be the subgraph of Γ induced by S'. Since $\psi \colon \Sigma_\Gamma \to \Sigma_\Gamma$ takes X_0 to within bounded Hausdorff distance of X' and they are both convex, hence undistorted, X_0 and X' are quasiisometric. Adjusting by the group action, we have a quasiisometry $(w^{-1}\cdot) \circ \psi : \Sigma_\Gamma \to \Sigma_\Gamma$ taking Σ_{Γ_0} Hausdorff close to $\Sigma_{\Gamma'}$, so Γ' satisfies (IV).

The fact that X' and X_0 have unbounded coarse intersection, since their coarse intersection contains $\bar{\gamma}'$, means $\Gamma_0 \cap \Gamma'$ is incomplete. So Γ' satisfies (I).

Theorem 7.3 gives us that $d(\bar{\gamma}', b_{i_0})$ is bounded, independent of r, so by taking r large we can make $r - d(\bar{\gamma}', b_{i_0})$ as large as we like. Assume $r - d(\bar{\gamma}', b_{i_0}) \geqslant 3$. Consider a shortest path ζ in Σ_P from γ_0 to $b_{i_0} \in \gamma_{i_0}$. Let $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{H}_3$ be the walls dual to the first three edges of ζ. For $k \in \{1, 2, 3\}$, let z_k be the generator labelling $\mathcal{H}_k$, and let $\mathcal{H}_k^-$ be the halfspace of $\mathcal{H}_k$ containing b_{i_0}.

By minimality of ζ, $z_1 \notin S_0$, so $X_0 \subset \mathcal{H}_1^+$. If $\mathcal{H}_1$ and $\mathcal{H}_2$ cross then z_1 and z_2 commute and the first two edges of ζ travel along the boundary of a square with opposite side pairs labelled by z_1 and z_2. By replacing the first two edges of ζ by the other two edges of this square we get a path ζ' with the same endpoints and length as ζ that crosses $\mathcal{H}_2$ first and then $\mathcal{H}_1$. If $\mathcal{H}_1$ crosses both $\mathcal{H}_2$ and $\mathcal{H}_3$ then z_1 commutes with z_2 and z_3 and since Σ_P is 2-dimensional, z_2 cannot commute with z_3. Thus, up to exchanging $\mathcal{H}_1$ and $\mathcal{H}_2$ we may assume that either $\mathcal{H}_1$ does not cross $\mathcal{H}_2$ or $\mathcal{H}_1$ crosses $\mathcal{H}_2$ but not $\mathcal{H}_3$.

In the first case $X_0 \subset \mathcal{H}_1^+ \subset \mathcal{H}_2^+$. As argued previously, X_0 has bounded coarse intersection with $\mathcal{H}_1^-$. Likewise $X_0 \overset{c}{\cap} \mathcal{H}_2^-$ is bounded, since $\mathcal{H}_2^- \subset \mathcal{H}_1^-$. In particular, both ends of $\bar{\gamma}'$ are contained in $\mathcal{H}_1^+$, as otherwise we would have an unbounded subset of $\mathcal{H}_1^-$ contained in a bounded neighborhood of X_0, contradicting that their coarse intersection is bounded. We conclude that $\bar{\gamma}'$ enters $\mathcal{H}_1^+$. However, as $\bar{\gamma}'$ comes close to b_{i_0}, it enters $\mathcal{H}_2^-$ as well. Accordingly, $\bar{\gamma}'$ crosses walls $\mathcal{H}_1$ and $\mathcal{H}_2$. Since $\bar{\gamma}' \subset X'$, the set S' contains z_1 and z_2, with $z_1 \in P - S_0$ and $z_2 \in P$ not adjacent to z_1.

In the second case we reach the same conclusions for z_1 and z_3. In either case, S' contains z_1 and a non-adjacent vertex (z_2 or z_3), so $\Gamma' \cap P$ is incomplete. Thus Γ' satisfies (II).

If Hypothesis (b) is true then since Γ' satisfies (I), (II), and (IV), it does not satisfy (III), so $\{a_{0,0}, a_{q-1,\ell(q-1)-1}\} \subset \Gamma'$. Since $S' \in \mathcal{C}_\Gamma$, the minimality condition on S_0 demands $S_0 \cap S' = S_0$. So $S_0 \subseteq S'$ is contained in a join of two anticliques $\Theta_0 * \Theta_1$. Suppose $z_1 \in \Theta_1$. Since $z_1 \in (S' \cap P) - S_0$, Hypothesis (5) implies $S_0 \cap \mathrm{lk}(z_1)$ is at most one vertex. But $S_0 \cap \mathrm{lk}(z_1) = S_0 \cap \Theta_0$, so the anticlique Θ_1 contains all but at most one vertex of S_0. This shows Hypothesis (a) is true.

If Hypothesis (a) is true then X_0 and X' are hyperbolic.

There are only finitely many isometry types of standard subcomplex in Σ_Γ, so there is a uniform bound on hyperbolicity constants that occur, independent of r. The quasigeodesic constants of $\bar{\gamma}'$ are also independent of r. Thus, there is a stability constant, independent of r, bounding the distance between $\bar{\gamma}'$ and a geodesic $\gamma'' \subset X'$ asymptotic to it. Taking $r - d(\bar{\gamma}', b_{i_0})$ larger than this stability constant forces the geodesic γ'' to enter $\mathcal{H}_1^-$, but it still has both ends in $\mathcal{H}_1^+$, since $\bar{\gamma}'$ does. This is a contradiction: walls are convex, so γ'' cannot cross from $\mathcal{H}_1^+$ to $\mathcal{H}_1^-$ and back to $\mathcal{H}_1^+$. □

Observe that Theorem 5.16 is a consequence of Theorem 7.5 and Proposition 7.2, with the poles of the cuts being the compliant sets and their intersections the (singleton) connecting paths.

7.4. **Subcomplexes not in the same quasiisometry orbit.** For W_Γ a RACG and Γ' an induced subgraph of Γ, let $[\![\Gamma']\!]$ denote the quasiisometry type of $\Sigma_{\Gamma'}$. Similarly, if S is a set of vertices of Γ, let $[\![S]\!]$ denote the quasiisometry type of the special subgroup defined by S.

We first state a result that we will use to rule out (IV) of Hypothesis (b) in certain applications of Theorem 7.5. Recall that this condition says there exists a quasiisometry $\psi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ with $\psi(\Sigma_{\Gamma_0}) \overset{c}{=} \Sigma_{\Gamma'}$. In other words, we wish to show that Σ_{Γ_0} and $\Sigma_{\Gamma'}$ are not in the same orbit of the quasiisometry group of Σ_Γ for its action on coarse equivalence classes of subsets of Σ_Γ.

Corollary 7.6 will be a corollary of the subsequent, more general, Lemma 7.7.

Corollary 7.6. *Let Γ be an incomplete triangle-free graph without separating cliques, and let Σ_Γ be the Davis complex of W_Γ. Suppose $\phi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ is a quasiisometry. Suppose S_0 and T_0 are vertex sets of maximal thick joins of Γ such that $\phi(\Sigma_{S_0}) \overset{c}{=} \Sigma_{T_0}$.*

Then $[\![S_0]\!] = [\![T_0]\!]$ *and for every neighbor* S_2 *of* S_0 *in* Ric_Γ *there is a neighbor* T_2 *of* T_0 *in* Ric_Γ *such that* $[\![S_2]\!] = [\![T_2]\!]$ *and* $[\![S_0 \cap S_2]\!] = [\![T_0 \cap T_2]\!]$.

Lemma 7.7. *Let* Γ *be an incomplete triangle-free graph without separating cliques, and let* Σ_Γ *be the Davis complex of* W_Γ. *Suppose* $\phi\colon \Sigma_\Gamma \to \Sigma_\Gamma$ *is a quasiisometry. Suppose* $S_0, S_1, S_2 \in \mathcal{C}_\Gamma$ *such that* $S_0 \subset S_1$, *and each of* S_0 *and* $S_1 \cap S_2$ *contains vertex sets of squares of* Γ. *Suppose* $\phi(\Sigma_{S_0}) \overset{c}{=} \Sigma_{T_0}$. *Then* $T_0 \in \mathcal{C}_\Gamma$ *and there exist* $T_1, T_2 \in \mathcal{C}_\Gamma$ *such that for* $i \in \{0,1,2\}$:

- *Some translate of* Σ_{T_i} *is coarsely equivalent to* $\phi(\Sigma_{S_i})$.
- $[\![S_i]\!] = [\![T_i]\!]$
- $\mathrm{ind}(S_i) = \mathrm{ind}(T_i)$
- $T_0 \cap T_1$ *is either all of* T_0 *or all but a cone vertex of* T_0.
- $[\![S_1 \cap S_2]\!] = [\![T_1 \cap T_2]\!]$

Proof. For any sets A and B of vertices of Γ, if $\phi(\Sigma_A) \overset{c}{=} w\Sigma_B$ then $\pi_{\Sigma_B} \circ (w^{-1}\cdot) \circ \phi|_{\Sigma_A}\colon \Sigma_A \to \Sigma_B$ is a quasiisometry, so $[\![A]\!] = [\![B]\!]$. By Proposition 7.2, for all i there exist $w_i \in W_\Gamma$ (and we take $w_0 = 1$) and $T_i \in \mathcal{C}_\Gamma$ with $\mathrm{ind}(S_i) = \mathrm{ind}(T_i)$ and $\phi(\Sigma_{S_i}) \overset{c}{=} w_i\Sigma_{T_i}$, which implies $[\![S_i]\!] = [\![T_i]\!]$.

We may assume that all of the S_i and T_i have no cone vertex, since removing a cone vertex does not change the coarse equivalence class of the special subgroup, the index in $\mathcal{C}_\Gamma$, or the fact that S_0 and $S_1 \cap S_2$ contain vertex sets of squares.

As in the proof of Proposition 7.2, for V_{ij} defined as the set of generators occurring in minimal length elements of $W_{S_i}(w_i^{-1}w_j)W_{S_j}$ we have:

$$\phi(\Sigma_{S_i \cap S_j}) \overset{c}{=} \pi_{w_i\Sigma_{T_i}}(w_j\Sigma_{T_j}) = w_i\Sigma_{T_i \cap T_j \cap \bigcap_{v \in V_{ij}} \mathrm{lk}(v)}$$

Thus, $[\![S_i \cap S_j]\!] = [\![T_i \cap T_j \cap \bigcap_{v \in V_{ij}} \mathrm{lk}(v)]\!]$. Since Γ is triangle-free, links of vertices are anticliques, so if $V_{ij} \neq \varnothing$ then $T_i \cap T_j \cap \bigcap_{v \in V_{ij}} \mathrm{lk}(v)$ is an anticlique and $[\![T_i \cap T_j \cap \bigcap_{v \in V_{ij}} \mathrm{lk}(v)]\!]$ is one of 'point', 'line', or 'bushy tree'. But $\Sigma_{S_0} = \Sigma_{S_0 \cap S_1}$ and $\Sigma_{S_1 \cap S_2}$ contain 2–flats, so $[\![S_i \cap S_j]\!]$ is not one of these, so V_{01} and V_{12} are empty, which gives $[\![S_0]\!] = [\![S_0 \cap S_1]\!] = [\![T_0 \cap T_1]\!] = [\![T_0]\!]$ and $[\![S_1 \cap S_2]\!] = [\![T_1 \cap T_2]\!]$.

Finally, $S_0 \subset S_1$ implies $\Sigma_{T_0} \overset{c}{=} \pi_{\Sigma_{T_0}}(w_1\Sigma_{T_1}) = \Sigma_{T_0 \cap T_1 \cap \bigcap_{v \in V_{01}} \mathrm{lk}(v)} = \Sigma_{T_0 \cap T_1}$. By Lemma 2.37 and the assumption that T_0 has no cone vertex, $T_0 = T_0 \cap T_1$. ☐

Proof of Corollary 7.6. Being a maximal thick join is the same as being a non-clique of $\mathcal{C}^0_\Gamma$. Suppose there is a T_0 that is in the quasiisometry orbit of S_0. Take $S_1 := S_0$ and any neighbor S_2 of S_0 in Ric_Γ, which shares a square with S_0 by definition of Ric_Γ. Apply Lemma 7.7 to get T_1 and T_2, which are maximal thick joins since S_1 and S_2 were. Since T_0 has no cone vertex, $T_0 \subset T_1$, but T_0 is maximal, so $T_0 = T_1$ and:

$$[\![S_0 \cap S_2]\!] = [\![S_1 \cap S_2]\!] = [\![T_1 \cap T_2]\!] = [\![T_0 \cap T_2]\!]$$ ☐

There are some other ways to rule out standard subcomplexes being in the same quasiisometry orbit. Instead of only considering single neighbor intersection quasiisometry types as in Corollary 7.6, one can consider the pattern of intersection of Σ_{S_0} with all of its neighbors in Π_Γ. Something like this is done in [15]. In another direction, we can use automorphism orbits of Π_Γ to distinguish maximal product regions. For instance, if there are two maximal standard product regions and one gives a cut vertex of Π_Γ and the other does not then a quasiisometry of Σ_Γ cannot take one to the other. Recall that we characterized cut vertices in Lemma 6.8.

7.5. **Examples.** We give some example applications of Theorem 7.5. Example 7.8 gives some of the smallest triangle-free strongly CFS graphs with compliant cycles. It turns out that all three are already known to be non RAAGedy for other reasons. After that we will give examples highlighting different aspects of the theorem.

Example 7.8. Figure 26 shows the smallest triangle-free CFS graphs with compliant cycles, which are therefore not RAAGedy, by Theorem 7.5. In each of these examples, S_0, S_1, and S_2 are the three possible pairs of red vertices and the P_i are single red vertices that are the intersections of consecutive S_i, so $P := \sqcup_i P_i$ is an anticlique. In the first two the S_i are cut pairs. In the third they are poles of maximal suspensions.

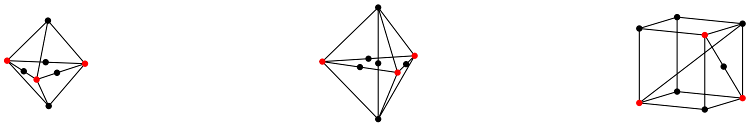

(a) Planar with cycle of cuts.

(b) Nonplanar with cycle of cuts.

(c) Nonplanar, with cycle of suspension poles.

FIGURE 26. Some CFS graphs with compliant cycles where the S_i are pairs and P is an anticlique.

The graph in Figure 26a is planar, so the fact that it is not RAAGedy could have also been deduced by applying a theorem of Nguyen and Tran [73]. The graph in Figure 26b is nonplanar, so Nguyen-Tran does not apply, but it can be shown to be non-RAAGedy by Theorem 5.16. The graph in Figure 26c has a JSJ graph of cylinders with a single cylinder and single non-virtually $\mathbb{Z}^2$ rigid subgroup connected by a virtually $\mathbb{Z}^2$ edge group, so is not RAAGedy by Lemma 5.13. ◇

Example 7.9. Consider the graph Γ of Figure 27. The hypotheses of Theorem 7.5 are satisfied for the path $P = P_0 := (0, 4, 5, 1)$ and the compliant set:

$$S_0 := \{2\} * \{0, 1\} = \{2, 6\} * \{0, 1, 7\} \cap \{2, 8\} * \{0, 1, 3\} \in \mathcal{C}^1_\Gamma \qquad \diamond$$

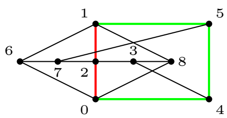

FIGURE 27. An example Γ such that Theorem 7.5 can be satisfied by a single compliant subset S_0 and a single path $P = P_0$.

Example 7.10. Consider the graph Γ of Figure 28. By computer search, there is no single compliant set S_0 and connected P satisfying the hypotheses of Theorem 7.5. For example, consider the maximal thick join $S_0 := \{0, 3\} * \{1, 7, 8, 9\}$. We would need a path P whose endpoints are nonadjacent vertices of S_0 and whose interior vertices have link intersecting S_0 in a clique. The possible interior vertices are 4, 5, and X, which would mean the endpoints of P would be adjacent vertices 0 and 9.

Instead, consider the maximal thick join $S_1 := \{0, 2, 6\} * \{1, 4, 5\}$, which intersects S_0 in a clique, and paths $P_0 := (9, \mathrm{X}, 4)$ and $P_1 := (1)$.

In this example, no other method discussed in this paper works in order to show that it is not RAAGedy. ◇

In the next two examples we use the considerations of Section 7.4 rule out condition (IV) of Theorem 7.5.

Example 7.11. Consider the graph Γ in Figure 29. Take $S_0 := \{0, 5\} * \{2, 6, 7\}$, which is a maximal thick join that is a non-square suspension. Consider $T :=$

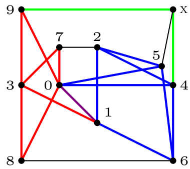

FIGURE 28. A graph with no compliant cycle with connected P, but having a compliant cycle consisting of two compliant sets and two connecting paths.

$\{2,6\} * \{0,4,5\}$ and $T' := \{0,4\} * \{1,2,3,6\}$, which are the two maximal thick joins that intersect S_0 in a non-clique. Both are non-square suspensions. The set of distinct pairs $([\![\Upsilon]\!], [\![\Upsilon \cap S_0]\!])$, where Υ ranges over the neighbors of S_0 in Ric_Γ consists of a single pair $(\text{tree} \times \text{line}, \mathbb{E}^2)$. For T we get the same result, but for T' the answer is different: T' has a neighbor in Ric_Γ such that $([\![\Upsilon]\!], [\![\Upsilon \cap T']\!]) = (\text{tree} \times \text{tree}, \text{tree} \times \text{line})$, coming from $\Upsilon = \{1,2,3\} * \{0,4,9\}$. Thus, we can see by Corollary 7.6 that Σ_{S_0} and $\Sigma_{T'}$ are not in the same orbit under the quasiisometry group of Σ_Γ. Thus, we look for a compliant cycle based on S_0 such that P avoids the vertex $\{4\} = T - S_0$, but we are not forced to avoid $\{1,3\} = T' - S_0$. $P = P_0 = (2,9,1,8,7)$ works.

In this example, no other method discussed in this paper works in order to show that it is not RAAGedy. ◇

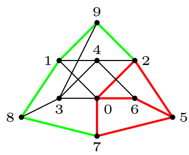

FIGURE 29. A graph with a compliant cycle where we need to recognize different quasiisometry orbits to verify the hypotheses of Theorem 7.5 are satisfied.

Example 7.12. Consider the graph Γ of Figure 30. Let $S_0 := \{2,3\} * \{0,1,4,5\}$ and $S' := \{6,7\} * \{4,5,8,9,14,15\}$, which are maximal thick joins, and let $P := (0,14,12,\mathrm{E},13,15,1)$ be a path. For this S_0 and P, conditions (1)–(5) of Theorem 7.5 are satisfied. Condition (a) is not true; S_0 contains squares. To conclude that Γ is not RAAGedy we show that condition (b) is satisfied. The only potential problem is S', which fulfills (I)-(III) of (b), so we need to show that (IV) is false by showing that Σ_{S_0} and $\Sigma_{S'}$ are not in the same orbit of the quasiisometry group of Σ_Γ.

In this example Lemma 7.7 does not help, but Lemma 6.8 does.

Take R_0 and R_1 to be subgraphs of Ric_Γ induced by vertex sets $\{v_{23}, v_{45}, v_{67}\}$ and $\{v_{67}, v_{89}, v_{\mathrm{XE}}\}$, respectively. As in Lemma 6.8 for $i \in \{0,1\}$ let Γ_i be the subgraph of Γ induced by $\bigcup_{v \in R_i} J_v$, so that Γ_0 is the red/violet subgraph and Γ_1 is the violet/blue subgraph, where the violet subgraph $J_{v_{67}}$ is their intersection. According to Lemma 6.8, the only cut vertices of Π_Γ are those in the W_Γ–orbit of v_{67}. Quasiisometries of Σ_Γ induce automorphisms of Π_Γ, which preserve cut vertices, so the W_Γ–orbit of v_{67} coincides with its orbit under the action of the quasiisometry group of Σ_Γ. Since S_0 corresponds to v_{23} and S' corresponds to v_{67} we conclude that Σ_{S_0} and $\Sigma_{S'}$ are not in the same orbit of the quasiisometry group of Σ_Γ. Thus, (IV) fails. ◇

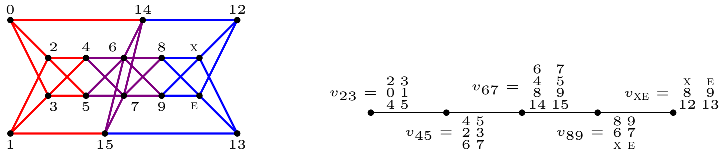

FIGURE 30. A graph Γ and Ric_Γ such that Ric_Γ contains a cut vertex of Π_Γ. The corresponding splitting of Γ is illustrated by the red/violet and blue/violet subgraphs. The cut vertex property is used to distinguish quasiisometry orbits of maximal product regions in the construction of a compliant cycle.

Sometimes passing to a link double is necessary to satisfy Theorem 7.5:

Example 7.13. For the graph Γ on the left of Figure 31, a computer search finds no configurations satisfying Theorem 7.5. After passing to $\mathfrak{D}^\circ_7(\Gamma)$ there is a good choice: $S_0 := \{3_0, 4_0\} * \{0_0, 2_0, 5_0, 8_0\}$ and $P = P_0 := (0_0, \text{E}_0, \text{X}_1, 6_1, 9_1, 2_0)$. ◇

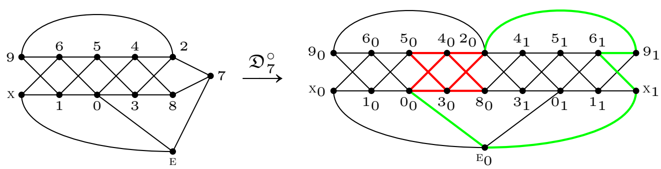

FIGURE 31. An example where taking a link double is necessary to satisfy the conditions of Theorem 7.5.

Example 7.14 (Compliant cycle vs Morse boundary.)**.** The graph of Figure 26a has a compliant cycle, but its Morse boundary is totally disconnected, by Corollary 5.11.

The graph of Figure 32 displays a stable cycle (red), so its Morse boundary contains circles, but it has no compliant cycle (even after passing to a link double). ◇

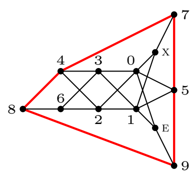

FIGURE 32. Graph with a stable cycle but no compliant cycle.

8. Further questions

Question 8.1. *Quasiisometry versus commensurability:*

- *We have only shown that cloning and unfolding produce groups quasiisometric to the one we started with. Are they actually commensurable?*
- *Does there exist a RACG that is not commensurable to any RAAG, but is quasiisometric to some RAAG?*

Question 8.2. *There are many other finite-index subgroups of RACGs than the ones we have considered. Would using these give any new results about commenurability between RACGs or between RACGs and RAAGs?*

Question 8.3. *Generalize the dimension restrictions:*

- *Is Oh's theorem true when there are no 3–quasiflats (which is the case that* Γ *is icosahedron-free)?*
- *If yes, do all of our arguments generalize to this case?*
- *Oh's theorem is not true in full generality in higher dimensions. What additional hypotheses on* Γ *or* W_Γ *would we need to make to make it be true? Compare Huang assuming* $\mathrm{Out}(A_\Delta)$ *is finite.*

Question 8.4. *So far, we can answer RAAGedy/non-RAAGedy for all of the 533 triangle-free CFS graphs with at most 10 vertices, and all but the following 8 of the 3405 with 11 vertices. Which of them are RAAGedy?*

$\Gamma_{1435} =$ 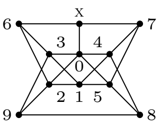

$\Gamma_{1441} =$ 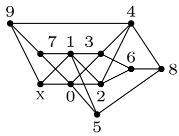

$\Gamma_{1454} =$ 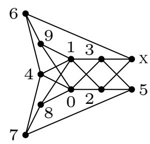

$\Gamma_{2708} =$ 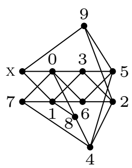

$\Gamma_{2722} =$ 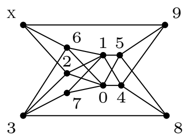

$\Gamma_{3002} =$ 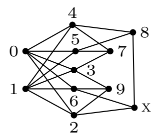

$\Gamma_{3632} =$ 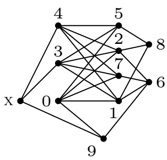

$\Gamma_{3794} =$ 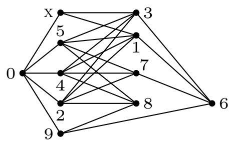

References

[1] C. Abbott, J. Behrstock, and M. G. Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.

[2] J. 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.

[3] J. Behrstock and R. Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356.

[4] J. Behrstock, R. A. Ciceksiz, and V. Falgas-Ravry, *Connectivity for square percolation and coarse cubical rigidity in random right-angled Coxeter groups*, preprint (2025), `arXiv:2502.18165`.

[5] J. Behrstock, V. Falgas-Ravry, M. F. Hagen, and T. Susse, *Global structural properties of random graphs*, Int. Math. Res. Not. IMRN (2018), no. 5, 1411–1441.

[6] J. Behrstock, M. Hagen, and A. Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.

[7] J. Behrstock, M. F. Hagen, and A. Sisto, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 5, 890–926.

[8] J. Behrstock, M. F. Hagen, and A. Sisto, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.

[9] J. Behrstock, M. F. Hagen, and A. Sisto, *Thickness, relative hyperbolicity, and randomness in Coxeter groups*, Algebr. Geom. Topol. **17** (2017), no. 2, 705–740, With an appendix written jointly with Pierre-Emmanuel Caprace.

[10] J. Behrstock, M. F. Hagen, and A. Sisto, *Quasiflats in hierarchically hyperbolic spaces*, Duke Math. J. **170** (2021), no. 5, 909–996.

[11] J. A. Behrstock, T. Januszkiewicz, and W. D. Neumann, *Quasi-isometric classification of some high dimensional right-angled Artin groups*, Groups Geom. Dyn. **4** (2010), no. 4, 681–692.

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

[13] M. Bestvina, B. Kleiner, and M. Sageev, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol. **12** (2008), no. 3, 1653–1699.

[14] J. Bounds and X. Xie, *Quasi-isometric rigidity of a class of right-angled Coxeter groups*, Proc. Amer. Math. Soc. **148** (2020), no. 2, 553–568.

[15] C. H. Cashen, *RAAGedy right-angled Coxeter groups II: in the quasiisometry class of the tree RAAGs*, Proc. Amer. Math. Soc. (in press).

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

[17] 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.

[18] R. Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158.

[19] R. Charney, M. Cordes, and A. Sisto, *Complete topological descriptions of certain Morse boundaries*, Groups Geom. Dyn. **17** (2023), no. 1, 157–184.

[20] I. Chatterji, T. Fernós, and A. Iozzi, *The median class and superrigidity of actions on* CAT(0) *cube complexes*, J. Topol. **9** (2016), no. 2, 349–400, With an appendix by P.-E. Caprace.

[21] V. Chepoi, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. **24** (2000), no. 2, 125–179.

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

[23] M. Cordes, *Morse boundaries of proper geodesic metric spaces*, Groups Geom. Dyn. **11** (2017), no. 4, 1281–1306.

[24] M. Cordes and D. Hume, *Stability and the Morse boundary*, J. Lond. Math. Soc. **95** (2017), no. 3, 963–988.

[25] J. Crisp and L. Paoluzzi, *Commensurability classification of a family of right-angled Coxeter groups*, Proc. Amer. Math. Soc. **136** (2008), no. 7, 2343–2349.

[26] P. Dani, *The large-scale geometry of right-angled Coxeter groups*, Handbook of group actions. V, Adv. Lect. Math. (ALM), vol. 48, Int. Press, Somerville, MA, 2020, pp. 107–141.

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

[28] P. Dani, E. Stark, and A. Thomas, *Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups*, Groups Geom. Dyn. **12** (2018), no. 4, 1273–1341.

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

[30] 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.

[31] M. W. Davis and G. Moussong, *Notes on nonpositively curved polyhedra*, Low dimensional topology (Eger, 1996/Budapest, 1998), Bolyai Soc. Math. Stud., vol. 8, János Bolyai Math. Soc., Budapest, 1999, pp. 11–94.

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

[33] M. W. Davis, *Infinite group actions on polyhedra*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics, vol. 77, Springer, Cham, 2024.

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

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

[36] C. Druţu and M. Kapovich, *Geometric group theory*, American Mathematical Society Colloquium Publications, vol. 63, American Mathematical Society, Providence, RI, 2018, With an appendix by Bogdan Nica.

[37] M. J. Dunwoody and M. E. Sageev, *JSJ-splittings for finitely presented groups over slender groups*, Invent. Math. **135** (1999), no. 1, 25–44.

[38] M. G. Durham and S. J. Taylor, *Convex cocompactness and stability in mapping class groups*, Algebr. Geom. Topol. **15** (2015), no. 5, 2839–2859.

[39] M. Dyer, *Reflection subgroups of Coxeter systems*, J. Algebra **135** (1990), no. 1, 57–73.

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

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

[42] E. Fioravanti and A. Karrer, *Connected components of Morse boundaries of graphs of groups*, Pacific J. Math. **317** (2022), 339–361.

[43] K. Fujiwara and P. Papasoglu, *JSJ-decompositions of finitely presented groups and complexes of groups*, Geom. Funct. Anal. **16** (2006), 70–125.

[44] A. Genevois, *Hyperbolicities in* CAT(0) *cube complexes*, Enseign. Math. **65** (2019), no. 1-2, 33–100.

[45] A. Genevois, *Quasi-isometrically rigid subgroups in right-angled Coxeter groups*, Algebr. Geom. Topol. **22** (2022), no. 2, 657–708.

[46] A. Genevois, *Algebraic properties of groups acting on median graphs*, draft version 3 (2023), `http://drive.google.com/file/d/1skDmnCz9EHHLgMaz1XwQZyRb8CI1QVaF/view`.

[47] V. N. Gerasimov, *Semi-splittings of groups and actions on cubings*, Algebra, geometry, analysis and mathematical physics (Russian) (Novosibirsk, 1996), Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk, 1997, pp. 91–109, 190.

[48] M. Graeber, A. Karrer, N. Lazarovich, and E. Stark, *Surprising circles in Morse boundaries of right-angled Coxeter groups*, Topology Appl. **294** (2021), Paper No. 107645, 3.

[49] D. Groves and M. Hull, *Abelian splittings of right-angled Artin groups*, Hyperbolic geometry and geometric group theory, Adv. Stud. Pure Math., vol. 73, Math. Soc. Japan, Tokyo, 2017, pp. 159–165.

[50] V. Guirardel and G. Levitt, *JSJ decompositions of groups*, Astérisque (2017), no. 395, vii+165.

[51] M. Hagen, *CAT(0) cube complexes, median graph, and cubulating groups*, Lecture notes from a minicourse at *Into the forest* (2019), `https://www.wescac.net/into_the_forest.pdf`.

[52] F. Haglund and D. T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620.

[53] U. Hamenstaedt, *Word hyperbolic extensions of surface groups*, unpublished (2005), `arXiv:math/0505244`.

[54] G. C. Hruska, E. Stark, and H. C. Tran, *Surface group amalgams that (don't) act on 3-manifolds*, Amer. J. Math. **142** (2020), no. 3, 885–921.

[55] J. Huang, *Quasi-isometry classification of right-angled Artin groups II: several infinite out cases*, preprint (2016), `arXiv:1603.02372`.

[56] J. Huang, *Quasi-isometric classification of right-angled Artin groups I: the finite out case*, Geom. Topol. **21** (2017), no. 6, 3467–3537.

[57] J. Huang, *Top-dimensional quasiflats in* CAT(0) *cube complexes*, Geom. Topol. **21** (2017), no. 4, 2281–2352.

[58] J. Huang, *Commensurability of groups quasi-isometric to RAAGs*, Invent. Math. **213** (2018), no. 3, 1179–1247.

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

[60] J. R. Isbell, *Median algebra*, Trans. Amer. Math. Soc. **260** (1980), no. 2, 319–362.

[61] A. Karrer, *Contracting boundaries of amalgamated free products of cat(0) groups with applications for right-angled coxeter groups*, Ph.D. thesis, Karlsruher Institut für Technologie (KIT), 2021.

[62] A. Karrer, *Right-angled Coxeter groups with totally disconnected Morse boundaries*, Geom. Dedicata **217** (2023), no. 4, Paper No. 71, 40.

[63] R. P. Kent, IV and C. J. Leininger, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. **18** (2008), no. 4, 1270–1325.

[64] I. Levcovitz, *Divergence of* CAT(0) *cube complexes and Coxeter groups*, Algebr. Geom. Topol. **18** (2018), no. 3, 1633–1673.

[65] I. Levcovitz, *A quasi-isometry invariant and thickness bounds for right-angled Coxeter groups*, Groups Geom. Dyn. **13** (2019), no. 1, 349–378.

[66] I. Levcovitz, *Characterizing divergence and thickness in right-angled Coxeter groups*, J. Topol. **15** (2022), no. 4, 2143–2173.

[67] J. F. Manning, *Geometry of pseudocharacters*, Geom. Topol. **9** (2005), 1147–1185.

[68] J. F. Manning, *Quasi-actions on trees and property (QFA)*, J. Lond. Math. Soc. **73** (2006), 84–108.

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

[70] MathOverflow, *Right-angled Artin groups that split as direct products*, (version: 2023-02-05), `https://mathoverflow.net/q/439976`.

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

[72] L. Mosher, M. Sageev, and K. Whyte, *Quasi-actions on trees II: Finite depth Bass-Serre trees*, Mem. Amer. Math. Soc. **214** (2011), no. 1008, vi+105.

[73] 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.

[74] S. Oh, *Quasi-isometry invariants of weakly special square complexes*, Topology and its Applications **307** (2022), 107945.

[75] P. Papasoglu, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2) **161** (2005), no. 2, 759–830.
[76] P. Papasoglu and K. Whyte, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv. **77** (2002), no. 1, 133–144.
[77] L. Paris, *Irreducible Coxeter groups*, Internat. J. Algebra Comput. **17** (2007), no. 3, 427–447.
[78] A. Reiter Ahlin, *The large scale geometry of products of trees*, Geom. Dedicata **92** (2002), 179–184.
[79] M. A. Roller, *Poc-sets, median algebras and group actions: An extended study of Dunwoody's construction and Sageev's theorem*, Habilitationsschrift, University of Regensburg, 1998, `arXiv:1607.07747`.
[80] J. Russell, D. Spriano, and H. C. Tran, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. **23** (2023), no. 3, 1167–1248.
[81] H. Tran, *On strongly quasiconvex subgroups*, Geom. Topol. **23** (2019), no. 3, 1173–1235.
[82] L. VanWyk, *Graph groups are biautomatic*, J. Pure Appl. Algebra **94** (1994), no. 3, 341–352.

Faculty of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, 0000-0002-6340-469X
Email address: `christopher.cashen@univie.ac.at`

Department of Mathematics, Louisiana State University, Baton Rouge, LA 70803–4918, USA, 0000-0001-7653-2618
Email address: `pdani@math.lsu.edu`

Faculty of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria, 0000-0002-6584-5149
Email address: `alexandra.edletzberger@gmail.com`

Department of Mathematics, The Ohio State University, 231 W. 18th Ave., Columbus, OH 43210, USA, 0000-0003-3761-5861
Email address: `karrer.14@osu.edu`