DISSERTATION / DOCTORAL THESIS

Titel der Dissertation / Title of the Doctoral Thesis

"Quasi-Isometries for two-dimensional Right-Angled Coxeter Groups"

verfasst von / submitted by

Alexandra Edletzberger, BSc MSc

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Doktorin der Naturwissenschaft (Dr. rer. nat.)

Wien, 2024 / Vienna, 2024

Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet:	A 796 605 405
Studienrichtung lt. Studienblatt / degree programme as it appears on the student record sheet:	Mathematik
Betreut von / Supervisor:	Univ.-Prof. Goulnara Arzhantseva, PhD Christopher Cashen, Privatdoz. PhD

Abstract

The Quasi-Isometry Problem is a fundamental problem in the field of geometric group theory. It asks whether or not two given groups share the same large-scale geometry and it has been investigated for many classes of groups. Due to its geometric origin, the class of Right-Angled Coxeter groups (RACGs), introduced by Coxeter in [Cox34], has received a lot of attention. However, their Quasi-Isometry Problem has only been investigated under additional strong assumptions like hyperbolicity or planarity of the defining graph. In the present thesis, we advance the Quasi-Isometry Problem for a large class of two-dimensional RACGs. In particular, we focus on two specifications of the problem: Finding quasi-isometries within the class of RACGs and between RACGs and the closely related Right-Angled Artin groups (RAAGs). In Section 1, we give an overview of the status quo of the problem.

Our tools of choice to address this problem are the JSJ tree of cylinders and the maximal product region graph. These two decompositions of groups are introduced in Section 2.

Section 3, which is taken from [Edl21], provides the visual construction of the JSJ tree of cylinders of RACGs and establishes it as quasi-isometry-invariant by the use of the structure invariant from [CM17a]. In addition, we show that under a certain additional assumption, the quasi-isometry-invariant is a complete quasi-isometry-invariant for a certain class of RACGs. It is used to provide new examples of non-hyperbolic RACGs that are quasi-isometric but not commensurable.

In Section 4, the difference between RACGs and RAAGs up to quasi-isometry is investigated. The Dani-Levcovitz construction [DL20] for finite index visual RAAG subgroups of RACGs is introduced and their algorithm is improved. Then, by use of the structure invariant as well as the maximal product region graph, new techniques are developed to find RACGs that are not quasi-isometric to any RAAG.

Zusammenfassung

Das Quasi-Isometrie-Problem ist ein fundamentales Problem im Feld der geometrischen Gruppentheorie. Dieses Problem fragt, ob zwei gegebene Gruppen die gleiche großmaßstäbliche Geometrie haben oder nicht und es wurde bereits für viele Klassen von Gruppen untersucht. Aufgrund ihres geometrischen Ursprungs wurde der Klasse der rechtwinkligen Coxeter-Gruppen, die von Coxeter in [Cox34] eingeführt wurden, viel Aufmerksamkeit gewidmet. Ihr Quasi-Isometrie-Problem wurde allerdings nur unter strengen Annahmen wie Hyperbolizität oder Planarität des definierenden Graphs untersucht. In der vorliegenden Dissertation wird das Quasi-Isometrie-Problem für eine große Klasse zwei-dimensionaler rechtwinkliger Coxeter-Gruppen verbessert. Insbesondere fokussieren wir uns auf zwei Spezialisierungen des Problems: Quasi-Isometrien werden zum einen innerhalb der Klasse der rechtwinkligen Coxeter-Gruppen gesucht und zum anderen zwischen rechtwinkligen Coxeter-Gruppen und den eng verwandten rechtwinkligen Artin-Gruppen. In Kapitel 1 geben wir einen Überblick über den Status Quo des Problems.

Die von uns gewählten Methoden, um das Problem zu adressieren, sind der JSJ-Zylinder-Graph und der Graph der maximalen Produkte. Diese beiden Gruppen-Zerlegungen werden in Kapitel 2 eingeführt.

In Kapitel 3, das aus [Edl21] stammt, wird eine visuelle Konstruktion des JSJ-Zylinder-Graphs von rechtwinkligen Coxeter-Gruppen entwickelt und durch die Nutzung der Struktur-Invariante aus [CM17a] als Quasi-Isometrie-Invariante etabliert. Zusätzlich zeigen wir, dass unter einer weiteren Annahme die Quasi-Isometrie-Invariante für eine gewisse Klasse von rechtwinkligen Coxeter-Gruppen vollständig ist. Sie wird genutzt, um neue Beispiele von nicht-hyperbolischen rechtwinkligen Coxeter-Gruppen zu finden, die quasi-isometrisch, aber nicht kommensurabel sind.

In Kapitel 4 wird der Unterschied zwischen rechtwinkligen Coxeter-Gruppen und rechtwinkligen Artin-Gruppen in Bezug auf Quasi-Isometrien untersucht. Die Dani-Levcovitz-Konstruktion [DL20] für visuelle rechtwinklige Artin-Untergruppen von rechtwinkligen Coxeter-Gruppen mit endlichem Index wird eingeführt und deren Algorithmus verbessert. Dann werden mit Hilfe der Struktur-Invariante und des Graphs der maximalen Produkte neue Techniken entwickelt, um rechtwinklige Coxeter-Gruppen zu finden, die zu keiner rechtwinkligen Artin-Gruppe quasi-isometrisch ist.

Acknowledgements

I would like to thank my supervisor Christopher Cashen for his support on my studies as well as for hiring me on his project. He guided me towards equally challenging as exciting mathematical questions, while encouraging my individual development and curiosity. It was a great pleasure to work with and learn from him and I can only recommend him as a supervisor and collaborator. I am also very grateful for the many opportunities to travel to international conferences and that he supported any activity I engaged in beyond the research context.

I am very thankful to my supervisor Goulnara Arzhantseva for inviting me to start working in geometric group theory in the first place. I could always count on her support, especially during the difficult times of the lockdown seasons and she provided a fruitful learning and working environment within her research group. Also, I learned a lot in the weekly research seminar.

It was a great pleasure to spend time with the current and former members of the research group Georg Ehling, Stefan Frank, Charlotte Hoffmann, Merlin Incerti-Medici, Yash Lodha and Markus Steenbock, who always offered advice and expertise as well as a nice chat over lunch.

Throughout my doctoral studies, I got the chance to meet many interesting, knowledgeable and extremely friendly people, who helped me grow mathematically and personally.

I would like to thank Annette Karrer for all the exciting discussions about our favorite groups and for becoming a friend along the way. I am grateful to Davide Spriano for inviting me to my very first international seminar talk and for being one of the most approachable persons I have met in the community. I am thankful to Xabier Legaspi and Stefanie Zbinden for many honest and sympathetic conversations about the academic life.

My appreciation goes to Pallavi Dani for introducing me to new exciting examples and questions to work on as well as for kindly explaining many details. I extend my thanks to Emily Stark for reading and improving my first paper. I am thankful to Jingyin Huang and Sangrok Oh for opening the door to the results in Section 4.3, by introducing me to the maximal product region graph, and to Elia Fioravanti and Jacob Russell for insightful discussions on the extension graph.

A big impact on my PhD journey had the Vienna School of Mathematics (VSM), who provided support, social activities as well as the opportunity to take action as a member of the student board. I am glad that through the VSM I got to spend time with great people, in particular with Mina Monadjem, and I enjoyed working with her, Bettina Hiebl and Stefan Portisch a lot.

Also, I am happy that I have met Anna Breger, with whom fighting for feminism and against the problematic structures of academia has become not only easier but also way more fun. It was a pleasure to work with her and I am looking forward to more projects to come.

Moreover, I would like to thank Clara for all the conversations about university life and beyond, meeting her was one of the highlights of my studies.

Ein riesiges Dankeschön geht außerdem an meine Familie für Eure Unterstützung und dafür, dass Ihr niemals einen Funken Zweifel an mir habt. Außerdem bin ich unglaublich dankbar für den Safe Space meiner Freund*innen und insbesondere dem Herz meines Supportsystems, Nora. Dein Verständnis und Deine Unterstützung für alle Teile meines Lebens, auch den mathematischen, bedeuten mir viel. Und zu guter Letzt: Danke Christoph, dass Du mein Anker bist.

Contents

1 Introduction **1**
1.1 Quasi-Isometry-Problem . . . 1
1.1.1 JSJ graph of cylinders . . . 2
1.1.2 Maximal product region graph . . . 4
1.1.3 QI-Problem of RACGs . . . 4
1.2 RACGs . . . 5
1.2.1 Properties of RACGs . . . 7
1.3 Status Quo of the QI-Problem within the class of RACGs . . . 9
1.3.1 Commensurability classification of RACGs . . . 12
1.3.2 Complete QI-classification of certain non-hyperbolic RACGs . . . 14
1.4 Status Quo of the QI-Problem between RACGs and RAAGs . . . 16
1.4.1 RAAGs . . . 16
1.4.2 Properties and QI-invariants of RAAGs . . . 17
1.4.2.1 Divergence . . . 18
1.4.2.2 Morse boundary . . . 18
1.4.2.3 Morse subgroups . . . 18
1.4.3 Existence of QIs . . . 20
1.4.4 Obstruction for QIs . . . 22

2 Tools **23**
2.1 JSJ graph of cylinders . . . 24
2.1.1 The structure invariant . . . 27
2.2 Maximal Product Region Graph . . . 29
2.2.1 Properties of the MPRG of RAAGs . . . 30
2.2.2 Relation between MPRG and other graphs . . . 32
2.2.2.1 Extension Graph . . . 33
2.2.2.2 Hyperplane Graphs . . . 35

3 QIs within the class of RACGs **38**
3.1 JSJ graph of cylinders of RACGs . . . 38
3.1.1 Hyperbolic case . . . 38
3.1.2 Non-hyperbolic case . . . 41
3.1.2.1 Cylinder vertices . . . 43
3.1.2.2 Non-cylinder vertices . . . 46
3.1.2.3 Edges . . . 47
3.1.2.3.1 Two-ended edge groups . . . 48
3.2 QI-invariance . . . 52
3.2.1 The structure invariant for RACGs . . . 53
3.2.2 Promoting to a QI . . . 54
3.2.2.1 Two-ended cylinder vertices . . . 55
3.2.2.2 VFD cylinder vertices . . . 56
3.2.2.3 VA cylinder vertices . . . 60
3.2.3 Refinement of the structure invariant . . . 65

3.2.4 Complete QI-Invariant 67

4 QIs between RACGs and RAAGs **74**

4.1 DL-Algorithm 74

4.1.1 Construction of a FIDL-Λ 77

4.1.2 Splittings and the DL-Algorithm 83

4.1.3 Beyond the DL-Algorithm 85

4.2 Structure Invariant 86

4.2.1 Comparison 88

4.2.2 New Examples of RACGs not QI to a RAAG 90

4.3 MPRG 92

4.3.1 New Examples of RACGs not QI to a RAAG 95

4.3.2 Is the MPRG a quasi-tree? 99

1 Introduction

From the 1970s to the 1990s, the study of groups experienced a shift from combinatorial methods to a geometric approach. In retrospect, particularly Gromov's contribution in [Gro84] is acknowledged as a landmark in the development of *Geometric Group Theory*. For the geometric approach, a group is interpreted as a geometric object, for instance by considering the *Cayley graph* of its presentation by generators and relators. However, a Cayley graph depends on the choice of the generating set:

Definition 1.1. Let $G = \langle S \mid R \rangle$ be a group with a generating set S and a set of relators R. The *Cayley graph* $Cay(G, S)$ *of* G *with respect to* S is a graph with the following vertex set and edge set:

$$V(Cay(G,S)) = \{g \in G\},$$
$$E(Cay(G,S)) = \{(g, gs) \mid s \in S\}.$$

The graph $Cay(G, S)$ is equipped with the *edge metric*: The distance $d(g, h)$ between two vertices g and h is the number of edges of a shortest path between g and h.

We consider the geometric properties a Cayley graph of G has, independent of the choice of generating set, as *the geometry* of a group G. Thus, an equivalence relation on Cayley graphs is needed: We say that two metric spaces, in particular two graphs, are *quasi-isometric* if there is a map between them that distorts distance at most linearly and that is almost surjective in the sense that in a uniform neighborhood of every point in the target space, there is an image point of the map. Since a group is acting *geometrically*, that is properly discontinuously and cocompactly by isometries, on all of its Cayley graphs, the fundamental *Švarc-Milnor-Lemma* implies that all Cayley graphs of a given group are quasi-isometric to each other:

Theorem 1.2 (Švarc-Milnor-Lemma)**.** [Efr53, Šva55, Mil68] *Let* G *be a group acting geometrically on a geodesic metric space* (X, d)*. Then* G *has a finite generating set* S *and the Cayley graph* $Cay(G, S)$ *of* G *with respect to* S *is quasi-isometric to* X*.*

This leads to one of the fundamental problems in the field:

> **Quasi-Isometry-Problem:** Given two groups G and G', determine whether or not they have quasi-isometric Cayley graphs.

It is the main focus of the present thesis to advance the Quasi-Isometry-Problem in the class of two-dimensional Right-Angled Coxeter Groups, see Section 1.1.3.

1.1 Quasi-Isometry-Problem

Formally, we understand the following map as geometry-preserving:

Definition 1.3. Let (X, d_X) and (Y, d_Y) be two metric spaces. A map $\phi\colon X \to Y$ is a *quasi-isometric embedding* if there are constants $C \geq 1$ and $D \geq 0$ such that for every $x_1, x_2 \in X$ the following holds:

$$\tfrac{1}{C}\, d_X(x_1, x_2) - D \leq d_Y(\phi(x_1), \phi(x_2)) \leq C\, d_X(x_1, x_2) + D\,.$$

A map $\phi\colon X \to Y$ is *quasi-surjective* if there is a constant $C' \geq 1$ such that for every $y \in Y$, there is an $x \in X$ with $d_Y(\phi(x), y) \leq C'$.

A map $\phi\colon X \to Y$ is a *quasi-isometry (QI)* if it is a quasi-surjective quasi-isometric embedding. The metric spaces X and Y are *quasi-isometric (QI)* if there is a quasi-isometry between them.

Being QI is an equivalence relation on geodesic metric spaces. We can think of a QI as a map preserving the geometry at a large scale: Consider for instance the integers $\mathbb{Z}$ as equidistant points on a line. If we zoom out of the image, the points appear to move closer together. Eventually, they look like a line, thus like the reals $\mathbb{R}$. The integers and the reals are QI to each other.

In some cases, there are algebraic reasons for the existence of a QI. For instance:

Definition 1.4. Two groups are *commensurable* if they have isomorphic finite index subgroups.

It is easy to prove that a group is QI to all of its finite index subgroups, which implies:

Lemma 1.5. *Commensurable groups are QI to each other.*

For a negative answer to the QI-Problem, one can use the geometry of the groups: By the Švarc-Milnor-Lemma 1.2, a group is QI to its *model spaces*, these are the spaces it is acting on geometrically. The *QI-invariants* of a group G are the properties that all model spaces of G have in common. Classical examples of such QI-invariants are finite generation, finite presentability, hyperbolicity and divergence. We can distinguish groups up to QI by finding some QI-invariant that differs. If the converse is true, that is, if the fact that two groups exhibit a certain QI-invariant implies that the groups are in fact QI, we say that the QI-invariant is *complete*. A common approach is to start with some QI-invariant and continue to refine it until we can show that it is complete.

A general method to obtain a QI-invariant is to find some features of the group that are invariant under QI and encode their combinatorics in a QI-invariant graph the group acts on. Then we use the structure and the stabilizers of the graph to distinguish groups up to QI. In this thesis, we use the *JSJ graph of cylinders* and the *maximal product region graph (MPRG)* as tools in this framework.

Outline 1.6. The main results of this thesis are technical, but the big picture is to develop the theory of the JSJ graph of cylinders and the MPRG for a certain class of groups, namely for *two-dimensional Right-Angled Coxeter groups*. We establish new, finer QI-invariants, and for some subclasses, we show they are complete invariants. We give new examples of non-hyperbolic Right-Angled Coxeter groups that are QI but not commensurable. By analyzing the structure of the MPRG, we improve a necessary criterion for Right-Angled Coxeter groups to be QI to a *Right-Angled Artin group* from having $\mathcal{CFS}$ defining graph to having strongly $\mathcal{CFS}$ defining graph. Furthermore, we show that even assuming a strongly $\mathcal{CFS}$ defining graph, the geometry of the MPRG distinguishes many Right-Angled Coxeter groups from Right-Angled Artin groups up to QI, even in cases where none of the previously known techniques could do so.

1.1.1 JSJ graph of cylinders

One strategy to learn more about the existence of QIs is to decompose groups into subgroups, whose QI-classification is understood. The interplay of the single pieces is captured by a *graph of groups decomposition*, see [Ser80] for an introduction to *Bass-Serre Theory*:

Definition 1.7. A *graph of groups decomposition* or *splitting* of a group G is a connected, directed graph $\mathcal{G}$, where each vertex $v \in V(\mathcal{G})$ is equipped with a *vertex group* G_v and each edge $e \in E(\mathcal{G})$ is equipped with an *edge group* G_e. In addition, for an edge $e \in E(\mathcal{G})$ with initial vertex $o(e)$ and terminal vertex $t(e)$, the edge group G_e is a subgroup of the vertex group $G_{o(e)}$ and it embeds into the vertex group $G_{t(e)}$ via an *attaching map*.

An *HNN extension* is a graph of groups decomposition with one vertex and one edge. An *amalgamated (free) product* is a graph of groups decomposition with two vertices and one edge.

By the *Structure Theorem* of Bass-Serre theory [Ser80, Theorem 13], from any action of a group G on a simplicial tree without edge inversions, we obtain a graph of groups splitting $\mathcal{G}(G)$ of G, where the vertex and edge stabilizers provide the vertex and edge groups. Conversely, from any graph of groups splitting $\mathcal{G}(G)$ of G, we can define its *Bass-Serre tree*, on which the group acts without edge inversions.

A seminal theorem of Stallings in [Sta71] says that a finitely generated group admits an HNN extension or an amalgamated product over a finite edge group if and only if it has more than one end. Since the number of ends of a group is a QI-invariant (see Section 8 of [Löh17] for an introduction to *ends*), so is the existence of such a decomposition.

By Dunwoody's accessibility [Dun85], any finitely presented group admits a (unique) maximal decomposition over finite edge groups. Then a result of Papasoglu and Whyte [PW02, Theorem 0.4] shows that for finitely presented groups with infinitely many ends, the collection of occurring QI-types of one-ended vertex groups in such a maximal splitting is a QI-invariant. This reduces the QI-Problem to one-ended groups.

In a first step, we aim to restrict ourselves to one-ended groups that split over the most elementary subgroups. By Stallings' theorem, they only split over infinite groups. Hence, we consider two-ended edge groups. These contain the integers $\mathbb{Z}$ as a finite index subgroup, thus are *virtually* $\mathbb{Z}$. Among groups that are not commensurable to surface groups, being one-ended and splitting over a two-ended subgroup is a property which is a QI-invariant by [Pap05].

We want to decompose these groups even further, in a non-trivial and maximal way. A way to do this for one-ended groups splitting over two-ended subgroups is to consider *JSJ decompositions*, produced from *JSJ trees*. These splittings are maximal in some sense and their vertices come in three types: *two-ended*, *hanging* and *rigid*. The terminology of *JSJ theory* has its origin at the decomposition of 3-manifolds, see [GL17] for a survey on the evolution of the theory. For hyperbolic groups, the JSJ tree corresponds to *Bowditch's JSJ tree* from [Bow98].

Like for maximal splittings over finite edge groups, some features of JSJ decompositions are stable under QIs: The QI-equivalence class of non-elementary vertex groups are preserved. In addition, also the patterns coming from the incident edge groups are maintained, see [CM17a, Section 2.3.2]. Thus, we can use this information to distinguish groups up to QI.

However, JSJ decompositions are not unique; usually a group has a whole collection of JSJ decompositions. We desire a canonical representative object for the collection with two key features:

1. It encodes the deformation space of all JSJ decompositions of a group.
2. Quasi-isometric groups have isomorphic representatives.

This canonical representative is the *JSJ tree of cylinders* and its corresponding splitting, the *JSJ graph of cylinders*, can be built from any JSJ decomposition. From the JSJ decomposition, it inherits a categorization of the vertices as *cylinder*, *hanging* or *rigid*.

Feature 1 can be of interest on its own, see [GL17, Part IV] and Section 2.1. We aim to use it as a tool to classify classes of groups up to QI via an application of Feature 2: By [cf. GL11, CM17a], a QI between two groups induces an isomorphism between their JSJ trees of cylinders. In fact, this isomorphism preserves additional information about the vertex groups, like for JSJ decompositions: For instance, the vertex type, the QI-equivalence class of vertex groups and the pattern coming from incident edge groups are preserved. Cashen and Martin use this fact in [CM17a] to introduce the *structure invariant*, see Section 2.1.1.

1.1.2 Maximal product region graph

Another important class of groups is the one of all groups which act on a *CAT(0) cube complex* X. If we restrict to such groups, we can split X into smaller product subcomplexes, for instance flats, each stabilized by a subgroup. The splitting is encoded by a graph, where each product subcomplex corresponds to a vertex and two vertices are connected by an edge if the corresponding subcomplexes share a flat. Ideally, this graph is preserved under QI.

This idea of studying the large-scale geometry of an object by encoding its flats dates back to the Mostow-Prasad rigidity theorem for locally symmetric manifolds of non-positive curvature [Mos73, Pra73]. It is used for the study of *Mapping Class Groups* of surfaces, where coarse product regions occur for short curves and the *curve graph*, introduced in [Har81], captures the relationship between them. For certain CAT(0) cube complexes, the main object in this framework are top-dimensional flats, whose importance was emphasized by work of Huang in [Hua17b] and product subcomplexes, whose significance was highlighted by work of Oh in [Oh22].

As a starting point, we limit ourselves to two-dimensional CAT(0) cube complexes and consider the class of *square complexes*. For their decomposition, we use subcomplexes exhibiting a product structure $P_1 \times P_2$, where P_1 and P_2 are infinite, connected subgraphs of X without a vertex of valence 1, the *maximal product regions*. Then the subgroup stabilizing $P_1 \times P_2$ can be written as the direct product of the stabilizers of the factors. The corresponding graph that describes the decomposition into maximal product regions intersecting in flats is called *maximal product region graph (MPRG)*, see Section 2.2.

In [Oh22, Theorem 3.7], Oh shows that the MPRG provides a QI-invariant: If two square complexes are QI, their MPRGs are isomorphic (cf. Theorem 2.31). However, this QI-invariant is not complete. We aim to fix a class $\mathcal{A}$ of groups and determine a property $\mathcal{P}$ that the MPRG of a group in $\mathcal{A}$ always has. For a given group $G \notin \mathcal{A}$, if we can show that the MPRG of G does not have property $\mathcal{P}$, then G is not QI to any group in $\mathcal{A}$.

1.1.3 QI-Problem of RACGs

This thesis is driven by the following specification of the QI-Problem to Right-Angled Coxeter Groups (RACGs), introduced in Section 1.2:

> **Quasi-Isometry-Problem of RACGs:** Given a RACG, which groups is it QI to?
> *In particular:*
> 1. Given two two-dimensional, one-ended RACGs splitting over two-ended subgroups, are they QI to each other?
> 2. Given a two-dimensional, one-ended RACG splitting over a two-ended subgroup, is it QI to any Right-Angled Artin Group (RAAG)?

To advance Question 1, the JSJ graph of cylinders (see Sections 1.1.1 and 2.1) is used in Section 3, see Section 1.3 for an overview. Question 2 is investigated using the JSJ graph of cylinders as well as the MPRG (see Sections 1.1.2 and 2.2) in Section 4. see Section 1.4 for an overview.

Since RACGs are groups determined by a *defining graph*, see [Rad01], in order to describe their properties in Section 1.2, we need to introduce some graph theoretical terminology first. Unless stated otherwise, every graph is simple and its edges are undirected.

Definition 1.8. Let Ω_1 and Ω_2 be two graphs. The *join* $\Omega = \Omega_1 \circ \Omega_2$ of Ω_1 and Ω_2 is the graph on the following vertex set and edge set:

$$V(\Omega) = V(\Omega_1) \cup V(\Omega_2),$$
$$E(\Omega) = E(\Omega_1) \cup E(\Omega_2) \cup \{(v_1, v_2) \mid v_1 \in V(\Omega_1), v_2 \in V(\Omega_2)\}.$$

We use the following non-standard definition of a *link* and a *star* in a graph:

Definition 1.9. A vertex $s \in V(\Omega)$ in the graph Ω is a *cone* if there is an induced subgraph $\Omega' \leq \Omega$ such that $\Omega = \{s\} \circ \Omega'$.

Definition 1.10. Let Ω be a graph and let $v \in V(\Omega)$. The set $nbs_\Omega(v) = \{v' \in V(\Omega) \mid (v, v') \in E(\Omega)\}$ are the *neighbors of v in* Ω. The *link of v in* Ω is the induced subgraph of Ω on the vertex set $nbs_\Omega(v)$. The *star of v in* Ω is the induced subgraph of Ω on the vertex set $\{v\} \cup nbs_\Omega(v)$.

In addition, we assume familiarity with basic concepts of group theory, see for instance [Bog08] for an introduction. In particular, we require prior knowledge in the field of geometric group theory, introductory material can be found in [CM17b] and [Löh17], for example.

1.2 RACGs

The following section contains parts of [Edl21, Section 2.1] and follows it closely.

Definition 1.11. For a finite, simplicial graph Γ with vertex set S, the *Right-Angled Coxeter Group (RACG)* W_Γ is defined as the group given by the following presentation

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

The graph Γ is called the *defining graph* or *presentation graph*.

Remark 1.12. Note that often in the literature instead of the defining graph, the *Coxeter graph* is used, in particular for general *Coxeter groups*. For RACGs, it is the complement graph of the defining graph.

Throughout this thesis, the notation for a RACG may vary:

Convention. Depending on whether we want to emphasize the defining graph Γ or the generating set $S = V(\Gamma)$ of the RACG we denote it as W_Γ or W_S, respectively.

Example 1.13. We obtain the following 'extrema' as standard examples of RACGs:

- If Γ is a complete graph with $V(\Gamma) = S$, then $W_\Gamma = \mathbb{Z}_2^{|S|}$. Moreover, W_Γ is finite if and only if Γ is complete.
- If Γ does not have any edges and $V(\Gamma) = S$, then $W_\Gamma = *_{|S|}\,\mathbb{Z}_2$. In particular, the infinite dihedral group $D_\infty = \mathbb{Z}_2 * \mathbb{Z}_2$ is a RACG.

Example 1.14. The following are fundamental examples:

- For a graph $\Gamma = \Gamma' \circ \Gamma''$ that is the join of Γ' and Γ'', W_Γ is the direct product $W_{\Gamma'} \times W_{\Gamma''}$.
- For a graph $\Gamma = \Gamma' \sqcup \Gamma''$ that is the disjoint union of two graphs Γ' and Γ'', W_Γ is the free product $W_{\Gamma'} * W_{\Gamma''}$.
- For the graph Γ_1 in Figure 1.2.1, W_{Γ_1} is the direct product $\mathbb{Z}_2 \times \mathbb{Z}_2$.

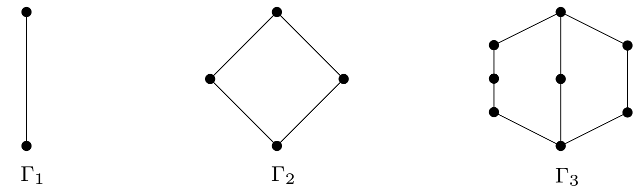

Figure 1.2.1

- For the graph Γ_2 in Figure 1.2.1, opposite vertices do not commute. Thus, they generate a subgroup isomorphic to D_∞. But since Γ_2 is the join of the two pairs of opposite vertices, W_{Γ_2} is the direct product $D_\infty \times D_\infty$.
- The graph Γ_3 in Figure 1.2.1 is a generalized Θ-graph (see Definition 1.28). It is harder to describe its RACG W_{Γ_3} in terms as elementary as the ones for the groups W_{Γ_1} and W_{Γ_2}.

Remark 1.15. By Example 1.14, the class of RACGs is closed under taking direct products by taking the join of defining graphs and under taking free products by taking the disjoint union of defining graphs.

Certain subgroups can be "read off" the defining graph:

Definition 1.16. Given a RACG W_S on $S = V(\Gamma)$, the subgroup W_T generated by $T \subseteq S$ is called a *special subgroup* of W_S.

By Theorem 4.1.6 of [Dav08], W_T is itself a (right-angled) Coxeter group on the defining graph Γ_T which is the induced subgraph of Γ on the vertices labelled by T. Moreover, the intersection of two special subgroups $W_T \cap W_{T'}$ is the special subgroup generated by the intersection $T \cap T'$.

Example 1.17.

- The RACG W_{Γ_2} in Example 1.14 on the defining graph Γ_2 illustrated in Figure 1.2.1 contains for instance the special subgroups $\mathbb{Z}_2 \times \mathbb{Z}_2$ and D_∞.
- In the RACG W_{Γ_3} of Example 1.14 on the defining graph Γ_3 shown in Figure 1.2.1 the two vertices of degree 3 generate a special D_∞ subgroup.

Theorem 1.18. [Kra09, Theorem 6.8.2] *A RACG contains a subgroup isomorphic to $\mathbb{Z}^2$ if and only if it contains a 4-cycle.*

Corollary 1.19. *Let W_Γ be a RACG on a triangle-free defining graph Γ. The intersection of two special subgroups $W_{\Gamma_1}, W_{\Gamma_2} \leq W_\Gamma$ contains a subgroup isomorphic to $\mathbb{Z}^2$ if and only if the intersection of their induced defining graphs Γ_1 and Γ_2 contains a 4-cycle.*

The geometry of a Coxeter group W_S is encoded in a complex, the so-called *Davis complex*. Its construction and properties can be found in [Dav08] and [DT17, Section 2.1].

We outline the following facts relevant for this thesis: The Davis complex of a special subgroup $W_T \subseteq W_S$ embeds isometrically as a convex subcomplex of the Davis complex of W_S. For RACGs, the Davis complex is a CAT(0) cube complex. Its 1-skeleton is precisely the Cayley graph $Cay(W_S, S)$ of W_S with respect to the generating set $S = V(\Gamma)$. Note that in case W_S is infinite, it contains $D_\infty = W_{\{a,b\}}$ as a subgroup, where a and b are non-adjacent vertices in S. Then we find a bi-infinite geodesic in the Cayley graph of W_S that is labelled alternately by a and b. We call such a geodesic *bi-labelled*.

1.2.1 Properties of RACGs

In order to use the JSJ graph of cylinders for the QI-classification of RACGs, we need to restrict ourselves to the subclass of RACGs that are one-ended and split over two-ended subgroups. These conditions are *visual* in the sense that they can be read of the defining graph if we make use of the auxiliary assumption that the RACGs are two-dimensional. While we expect that this additional assumption can be dropped, the generalization is not immediate. This issue is also addressed in [DT17, Section 1] and [Dan20, Question 5.17].

Definition 1.20. A vertex v of Γ is *essential* if it has valence at least 3. We denote the set of all essential vertices in Γ by $EV(\Gamma)$. An embedded path between essential vertices, which does not contain any essential vertices in its interior, is a *branch.*

A vertex a of Γ is a *cut vertex* if $\Gamma \setminus \{a\}$ has at least two connected components.

A pair $\{a, b\}$ of vertices of Γ is a *cut pair* if it separates Γ, that is $\Gamma \setminus \{a, b\}$ has at least two connected components. If both vertices are essential, we call it an *essential cut pair*.

A set $\{a, b, c\}$ of vertices of Γ is called a *cut triple* if a and b are not a cut pair, c is a common adjacent vertex of a and b and the subgraph induced by $\{a, b, c\}$ separates Γ.

Convention. We use the term *cut collection* when referring to both cut pair and cut triple at once and use the notation $\{a - b\}$. The $-$ represents the possibly existing common adjacent vertex c of a and b contributing to the triple.

Example 1.21. In the left graph Γ_1 of Figure 1.2.2, the set $T_1 = \{a, b\}$ is a cut pair. Since a and b are not connected by an edge in Γ_1, the T_1-induced subgraph contains only two disconnected vertices, and thus, the special subgroup generated by T_1 is $W_{\{a,b\}} = D_\infty$. The graph Γ_2 on the right contains two cut triples, one of which is $T_2 = \{a, b, c\}$. The special subgroup on the T_2-induced subgraph is $W_{\{a,b,c\}} = D_\infty \times \mathbb{Z}_2$.

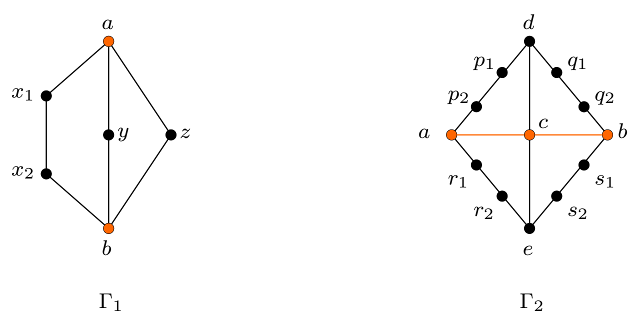

Figure 1.2.2: The orange vertices form a cut pair and a cut triple, respectively.

We aim for the following conditions on W_Γ:

- **The Davis complex of W_Γ is two-dimensional** to simplify the geometry encoded by the group. This is the case if Γ is triangle-free.
- **W_Γ is one-ended**: By [Dav08, Theorem 8.7.2], this is true if Γ is connected and has neither a separating vertex nor a separating edge, under the assumption that Γ has no triangles.
- **W_Γ has a splitting over a two-ended subgroup**: By Theorem 3.5, recalling [MT09, Theorem 1] in our setting, under the assumption that W_Γ is two-dimensional and one-ended, the existence of a splitting over a two-ended subgroup is ensured if Γ has a cut collection

$\{a - b\}$. Indeed, if there is a cut collection $\{a - b\}$ all k components of $\Gamma \setminus \{a - b\}$ attach along the two-ended special subgroup $W_{\{a,b\}} = D_\infty$ or $W_{\{a,b,c\}} = D_\infty \times \mathbb{Z}_2$ as a k-fold amalgamated product.

- W_Γ **is not *cocompact Fuchsian***: That means that W_Γ does not act geometrically on the hyperbolic plane. In the two-dimensional case, this is equivalent to Γ not being a cycle of length ≥ 5 by [DT17, Theorem 4.2]. We can exclude cocompact Fuchsian groups, because the *Švarc-Milnor-Lemma* 1.2 implies that they are QI to each other, thus their QI-Problem is understood.

Thus, to ensure that W_Γ is two-dimensional, one-ended and splitting over a two-ended subgroup, we fix the following:

Standing Assumption 1. Γ is the defining graph of a RACG W_Γ which satisfies:

(1) Γ is triangle-free.
(2) Γ is connected and has neither a separating vertex nor a separating edge.
(3) Γ has a cut collection $\{a - b\}$.
(4) Γ is not a cycle of length ≥ 5.

Convention. Unless stated otherwise, throughout this thesis, every graph Γ satisfies the Standing Assumption 1 and every RACG is defined on a graph satisfying the Standing Assumption 1.

Remark 1.22. Observe the following:

- Under Standing Assumption 1, a cut pair $\{a, b\}$ always consists of non-adjacent vertices and a cut triple $\{a, b, c\}$ forms a segment connecting a and b, where a and b are both adjacent to c and not adjacent to each other. Thus, the special subgroup generated by both a cut pair and a cut triple is two-ended and the elements a and b generate a copy of D_∞.
- For a cut triple $\{a - b\}$, the common adjacent vertex of a and b might not be unique: See for instance Figure 1.2.3, where $\{x, y, b\}$, $\{x, y, c\}$ and $\{x, y, d\}$ are cut triples. We say that the cut triples *overlap*. However, when there are overlapping cut triples the graph necessarily has an induced square, so this configuration does not arise in the hyperbolic case (by [Dav08, Corollary 12.6.3]), but we have to deal with it in our more general setting. In Section 3.1.2.3.1 we make additional assumptions (to guarantee that the graph of cylinders has two-ended edge groups, see Remark 2.14) which exclude overlapping cut triples, see Remark 3.31.

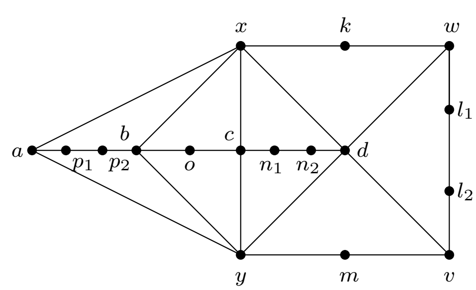

Γ

Figure 1.2.3

1.3 Status Quo of the QI-Problem within the class of RACGs

The class of RACGs is very diverse, and there is no apparent unified way to tackle the QI-Problem within the class of RACG. Rather, the common approach is to pick a QI-invariant, identify graph theoretical properties of the defining graph Γ that imply that the RACG W_Γ has that invariant and then refine the QI-Problem within that class of groups. In this section we highlight and assemble the different strategies and survey what is known about the QI-classification, including the results developed in this thesis in Section 3, see Subsection 1.3.2.

In a hyperbolic group, the JSJ tree of cylinders is equivalent to *Bowditch's JSJ tree* [Bow98], see Section 1.1 and Section 3.1.1, which is defined via the *Gromov boundary* of the group. So, in a natural first step of the QI-classification, Dani-Thomas assume in [DT17] hyperbolicity, in addition to Standing Assumption 1. For two-dimensional RACGs, hyperbolicity is a visual property:

Lemma 1.23. [Mou88, Theorem 17.1] *A RACG W_Γ is hyperbolic if and only if Γ does not contain any cycles of length four.*

In order to give a QI-classification for hyperbolic RACGs, Dani-Thomas describe Bowditch's JSJ tree T_c visually:

Theorem 1.24 (cf. Theorem 3.1)**.** [DST18, cf. Theorem 1.2] *For a hyperbolic RACG W_Γ satisfying Standing Assumption 1, the defining graph visually determines Bowditch's JSJ tree T_c: Subsets of vertices of the defining graph satisfying certain graph theoretical conditions are in bijection with W_Γ-orbits of vertices of T_c and they generate the representatives of the conjugacy classes of the vertex stabilizers.*

This leads to the following QI-classification by the use of the structure invariant (see Section 2.1.1 and Theorem 2.15):

Theorem 1.25. [DT17, Theorem 1.4] *Let W_Γ and $W_{\Gamma'}$ be two hyperbolic RACGs whose defining graphs Γ and Γ' satisfy Standing Assumption 1 and have no induced subgraph which is a subdivided K_4. Then W_Γ and $W_{\Gamma'}$ are QI if and only if they have identical structure invariants.*

Remark 1.26. The assumption that Γ has no induced subgraph that is a subdivided K_4 implies that Bowditch's JSJ tree has no rigid vertices.

Example 1.27. The hyperbolic RACGs W_{Γ_1} and W_{Γ_2} on the graphs Γ_1 and Γ_2, illustrated in Figure 1.3.4, are QI to each other by Theorem 1.25: We can construct their Bowditch's JSJ tree and the corresponding splittings Σ_{c,Γ_1} and Σ_{c,Γ_2}, respectively, shown in Figure 1.3.4, by Theorem 1.24 (see Theorem 3.1 for details). The splittings are isomorphic and the isomorphism sends the cylinder vertices c_1 and c_2 to the cylinder vertices c_1' and c_2' and the hanging vertices h_1, h_2, h_3 and h_4 to the hanging vertices h_1', h_2', h_3' and h_4', respectively. The vertex type is preserved, and the identified vertices have identical relative QI-type. Thus, the two RACGs W_{Γ_1} and W_{Γ_2} have identical structure invariants and hence are QI.

Instead of restricting to a class of graphs by fixing a certain group property, we can also choose a graph property directly. An elementary class of graphs is the one consisting of the following:

Definition 1.28. A graph Γ is a *generalized Θ-graph* if it has two vertices of degree $k \geq 3$ connected with k disjoint paths and the i^{th} path p_i has n_i vertices of valence 2 and $n_i + 1$ edges. The *linear degree l of* Γ is the cardinality of the set $\{n_i \mid n_i = 1\}$. The *hyperbolic degree h of* Γ is $h = k - l$.

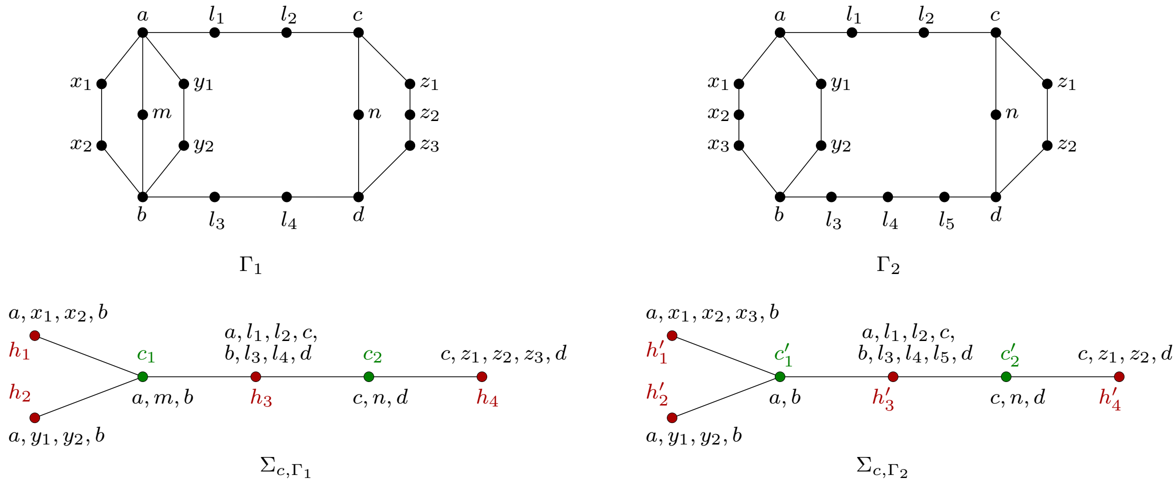

Figure 1.3.4

Example 1.29. The graph Γ_3 on the right of Figure 1.2.1, introduced in Example 1.14, is a generalized Θ-graph with linear degree 1 and hyperbolic degree 2. The graph Γ_1 on the left of Figure 1.2.2, introduced in Example 1.21, is a generalized Θ-graph with linear degree 2 and hyperbolic degree 1.

Remark 1.30. A generalized Θ-graph satisfies Standing Assumption 1. Note that it is hyperbolic if and only if it has linear degree $l \in \{0, 1\}$ by Lemma 1.23.

The complete QI-classification for *hyperbolic* RACGs on generalized Θ-graphs (that do not contain a 4-cycle) is covered by Theorem 1.25, while the following is the complete QI-classification of non-hyperbolic RACGs on generalized Θ-graphs.

Theorem 1.31. [HST20, Theorem A.9] *Let W_Γ and $W_{\Gamma'}$ be two RACGs on generalized Θ-graphs Γ and Γ' with linear degrees $l \geq 2$ and $l' \geq 2$ and hyperbolic degrees $h \geq 0$ and $h' \geq 0$, respectively. Then W_Γ and $W_{\Gamma'}$ are QI if and only if one of the three following conditions hold:*

(1) *$l = l'$ and $h, h' \geq 1$.*
(2) *$l, l' \geq 3$ and $h, h' \geq 1$.*
(3) *$l, l' \geq 3$ and $h = h' = 0$.*

Example 1.32. In Figure 1.3.5, the graph Γ_1 on the left has linear degree $l_1 = 4$ and hyperbolic degree $h_1 = 1$, the graph Γ_2 in the middle has linear degree $l_2 = 3$ and hyperbolic degree $h_2 = 2$ and the graph Γ_3 on the right has linear degree $l_3 = 2$ and hyperbolic degree $h_3 = 3$. Thus, by Theorem 1.31, the RACG W_{Γ_1} is QI to W_{Γ_2} and they are not QI to W_{Γ_3}.

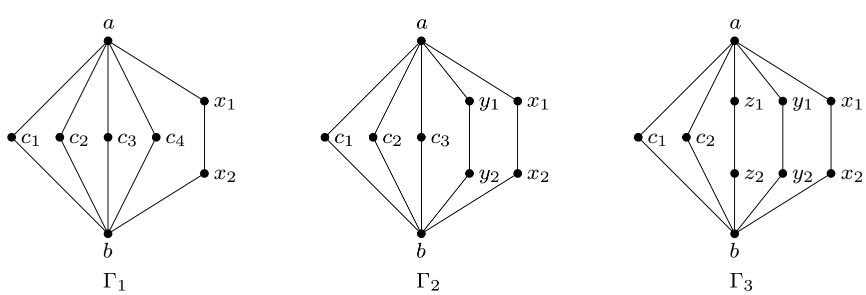

Figure 1.3.5

Another relevant group property is the divergence of a group. It is a function that measures the lengths of paths between two vertices in the Cayley graph that avoid a ball around the identity in terms of their distance, see [Ger94]. Up to equivalence of functions, the divergence is a QI-invariant. Hyperbolic groups all have exponential divergence. But the class of RACGs is much richer:

Theorem 1.33. [DT15, Theorem 1.2] *For all $d \geq 1$, there is a RACG W_{Γ_d} with polynomial divergence of degree d.*

In case the divergence of a RACG W_Γ is at most *quadratic*, it is a visual property that can be described by a graph theoretical condition on Γ.

Definition 1.34. [BFRHS18, Definition 1.3] For a graph Γ, the *4-Cycle Graph* $\square(\Gamma)$ is a graph whose vertices are in one-to-one correspondence with the induced 4-cycles, that is, *squares* in Γ. There is an edge between two vertices in $\square(\Gamma)$ if their corresponding squares share two non-adjacent vertices in Γ. For a vertex $v \in V(\square(\Gamma))$, the *support of* v, denoted by $supp(v)$, are the vertices in Γ contained in the square corresponding to v.

The graph Γ is $\mathcal{CFS}$ (*constructed from squares*) if $\square(\Gamma)$ has a connected component C such that

$$\bigcup_{v \in V(C)} supp(v) = V(\Gamma)\,.$$

The graph Γ is *strongly* $\mathcal{CFS}$ if it is $\mathcal{CFS}$ and $\square(\Gamma)$ is connected. The graph Γ is *min*$\mathcal{CFS}$ if it is $\mathcal{CFS}$ and $\Gamma \setminus \{e\}$ is not $\mathcal{CFS}$ for every edge $e \in E(\Gamma)$.

Theorem 1.35. [DT15, Theorem 1.1] *A RACG W_Γ satisfying Standing Assumption 1 has at most quadratic divergence if and only if Γ is $\mathcal{CFS}$, in particular:*

1. *W_Γ has linear divergence if and only if Γ is a join.*
2. *W_Γ has quadratic divergence if and only if Γ is $\mathcal{CFS}$ and not a join.*

In [NT19], Nguyen and Tran give the complete QI-classification of RACGs on $\mathcal{CFS}$ defining graphs, which are in addition planar. To state it, we need to define additional terminology.

Definition 1.36. A graph Γ is a *suspension* if there are non-adjacent vertices $a, b \in V(\Gamma)$ such that every vertex in $\Gamma \setminus \{a, b\}$ is adjacent to both a and b, that is Γ decomposes as the join $\Gamma = \{a, b\} \circ \Gamma'$. We call a and b the *suspension vertices*.

A *maximal suspension graph* T_Γ of Γ is a graph whose vertices are in one-to-one correspondence with the induced maximal suspension subgraphs of Γ. There is an edge between two vertices if their corresponding maximal suspension subgraphs share a square. Assigned to each vertex is the special subgroup generated by its corresponding maximal suspension subgraph.

A vertex in a maximal suspension graph is *spacious* if in its corresponding induced maximal suspension subgraph of Γ, there are two non-adjacent vertices $c, c' \in V(\Gamma)$ such that c and c' are not a pair of suspension vertices and both c and c' are contained in at most one pair of vertices of T_Γ corresponding to an edge. Otherwise, the vertex is *full*.

Two maximal suspension graphs T_Γ and $T_{\Gamma'}$ are *bisimilar* if there are a graph T whose vertices are labelled *spacious* and *full* and two maps $f\colon T_\Gamma \to T$ and $f'\colon T_{\Gamma'} \to T$ such that the labelling is preserved and for every $v \in V(T_\Gamma)$ and every edge $t \in E(T)$ at $f(v)$, there is an edge $e \in E(T_\Gamma)$ at v with $f(e) = t$, and such that the same property holds for f'.

Theorem 1.37. [NT19, Theorem 1.1] *Let W_Γ and $W_{\Gamma'}$ be two RACGs on planar, $\mathcal{CFS}$ defining graphs Γ and Γ' satisfying Standing Assumption 1, respectively. Then W_Γ and $W_{\Gamma'}$ are QI if and only if their maximal suspension graphs T_Γ and $T_{\Gamma'}$ are bisimilar.*

Idea of the Proof. The assumption that Γ is planar and $\mathcal{CFS}$ implies that every vertex and edge of Γ occurs in some defining graph of a vertex group of the maximal suspension graph of Γ, thus, it is a decomposition of the whole group. Nguyen and Tran show that every vertex group in the maximal suspension graph of Γ acts geometrically on a 3-manifold, hence such a RACG is a *3-manifold group*. They conclude that if the corresponding vertex is spacious, the 3-manifold has a boundary, if it is full, it does not. Then mirroring the proof of Theorem 3.2 in [BN08] about the QI-classification of such 3-manifold groups finishes the proof. □

Remark 1.38. In fact, the suspension decomposition introduced in [NT19] to prove Theorem 1.37 and also Theorem 1.77 in Section 1.4.3, is in correspondence with the JSJ graph of cylinders. The classification they use in terms of the boundary of the manifold is actually a use of the structure invariant with the relative QI-type of the vertex groups as choice of decoration.

Example 1.39. While the RACGs W_{Γ_1} and W_{Γ_3} on the defining graphs Γ_1 and Γ_3 in Figure 1.3.6 have a full (white) vertex in their maximal suspension graph T_{Γ_1} and T_{Γ_3}, respectively, the RACG W_{Γ_2} on the defining graph Γ_2 has only spacious (black) vertices in its maximal suspension graph T_{Γ_2} (cf. Example 4.2 of [NT19]). Thus, by Theorem 1.37, W_{Γ_2} is not QI to W_{Γ_1} and W_{Γ_3}. Since T_{Γ_1} and T_{Γ_3} are bisimilar, W_{Γ_1} and W_{Γ_3} are QI to each other by Theorem 1.37.

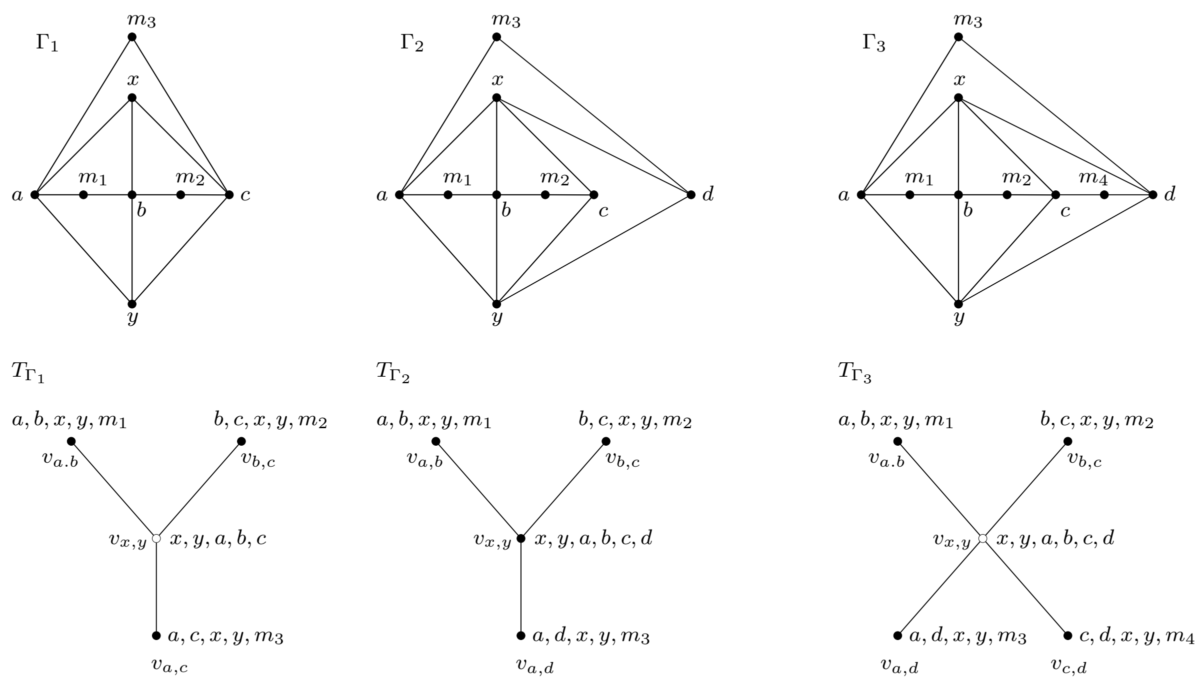

Figure 1.3.6

1.3.1 Commensurability classification of RACGs

Recall from Section 1.1 that the commensurability problem is closely related to the QI-Problem, since commensurability implies QI. For RACGs on generalized Θ-graphs, a (partial) commensurability classification is due to Dani, Stark and Thomas [DST18] and Hruska, Stark and Tran [HST20].

Definition 1.40. Let Γ be a generalized Θ-graph with k paths $\{p_1, \ldots, p_k\}$. The *Euler characteristic vector v of W_Γ* is the vector $v = (\chi(W_{p_1}), \chi(W_{p_2}), \ldots, \chi(W_{p_k})) \in \mathbb{Q}^k$, where each $\chi(W_{p_i})$ is given by

$$\chi(W_{p_i}) = 1 - \frac{n_i}{2} + \frac{n_i + 1}{4} \in \mathbb{Q} \quad \text{for } i \in \{1, \ldots, k\}\,.$$

Two Euler characteristic vectors $v \in \mathbb{Q}^k$ and $v' \in \mathbb{Q}^{k'}$ are *commensurable* if $k = k'$ and there are integers $\lambda, \kappa \in \mathbb{Z}$ such that $\lambda v = \kappa v'$.

This characterizes the commensurability classification of RACGs on certain generalized Θ-graphs:

Theorem 1.41. [DST18, Theorem 1.8; HST20, Corollary 4.10; Dan20, cf. Theorem 6.1] *Let W_Γ and $W_{\Gamma'}$ be two RACGs with Γ and Γ' generalized Θ-graphs and Euler characteristic vectors v and v', respectively. If v and v' are commensurable, then W_Γ and $W_{\Gamma'}$ are commensurable. Moreover, if W_Γ and $W_{\Gamma'}$ both have linear degree 0, then W_Γ and $W_{\Gamma'}$ are commensurable if and only if v and v' are commensurable.*

In fact, Dani-Stark-Thomas also give a commensurability classification for RACGs on graphs that are a *cycle of generalized Θ-graphs*, see [DST18, Definition 1.10] for the definition. There is an analogue of Euler characteristic vectors for such graphs and they give a classification in terms of these [DST18, Theorem 1.12]. Since the result is quite technical to state, we refer readers for the details to their paper [DST18].

In addition, for a given RACG W_Γ, there is a procedure to find infinitely many finite index RACG subgroups, thus, it provides an infinite collection of RACGs that W_Γ is commensurable to. The following algorithm was introduced in Section 2.2 of the preliminary version [DT14] of [DT17] and brought to the author's attention by Annette Karrer.

Definition 1.42. Let Γ be a graph and let $v \in V(\Gamma)$ be some vertex. The *double of Γ along v* is the graph Γ_v that has the following vertex set and edge set:

$$V(\Gamma_v) = V(\Gamma) \setminus \{v\} \cup \{s' \mid s \in V(\Gamma) \setminus st(v)\},$$
$$E(\Gamma_v) = E(\Gamma \setminus \{v\}) \cup \{(s', t') \mid (s, t) \in E(\Gamma)\} \cup \{(s', l) \mid l \in lk(v), (s, l) \in E(\Gamma)\},$$

where for every vertex $s \in V(\Gamma) \setminus st(v)$, we define a new vertex $s' \in V(\Gamma_v)$ as the double of v.

Lemma 1.43. [DT14, Lemma 2.3] *Let W_Γ be a RACG and Γ_v be the double of Γ along some $v \in V(\Gamma)$. Then W_{Γ_v} is an index 2 subgroup of W_Γ.*

Sketch of the Proof. Define the following map:

$$\begin{array}{rccl} \phi\colon & W_\Gamma & \to & \mathbb{Z}_2 \\ & v & \mapsto & 1 \\ & s & \mapsto & 0 \quad \text{for } \; s \in V(\Gamma) \setminus \{v\}\,. \end{array}$$

The kernel of the map ϕ is generated by $\langle \{s \mid s \in V(\Gamma) \setminus \{v\}\} \cup \{vsv \mid s \in V(\Gamma) \setminus \{v\}\} \rangle$ and is isomorphic to the RACG on the double Γ_v of Γ along v, where each new vertex s' corresponds to the generator vsv. □

Example 1.44. For the RACG W_Γ with Γ on the left of Figure 1.3.7, the RACG W_{Γ_c} on the double Γ_c of Γ along c illustrated on the right of Figure 1.3.7 is a finite index subgroup by Lemma 1.43.

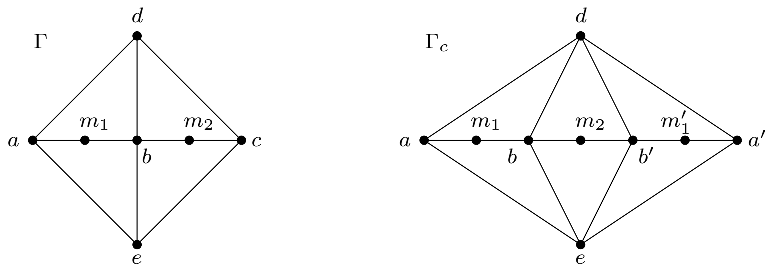

Figure 1.3.7

While with Lemma 1.43, it is easy to construct a pair of commensurable RACGs, the other direction is significantly harder. If we consider a pair of RACGs which are QI to each other and whose defining graphs are not (cycles of) generalized Θ-graphs, thus lying outside the scope of Theorem 1.41, there is no general way of determining whether or not these groups are commensurable. We give new such examples of non-hyperbolic RACGs that are QI but not commensurable: In Lemma 3.52 the JSJ graph of cylinders is used to show that the pairs of RACGs in Examples 3.49 and 3.50 on the defining graphs in Figure 3.2.13 and 3.2.14 which are not (cycles of) generalized Θ-graphs, are QI but not commensurable. The proof of Lemma 3.52 is applicable to produce more examples of that kind, see Remark 3.53.

1.3.2 Complete QI-classification of certain non-hyperbolic RACGs

The results in this Section were published in [Edl21], and thus, the following summary closely follows the introductory Section 1 of the author's final version of [Edl21]. The detailed statements and all proofs of the mentioned results can be found in Section 3.

The JSJ graph of cylinders, introduced in Section 2.1, is the main tool for the QI-classification of a wide family of RACGs. Its construction, as proven in Section 3.1, is visual:

Theorem 1.45 (cf. Theorem 3.33)**.** *For a RACG W_Γ satisfying Standing Assumption 1, the defining graph visually determines the JSJ tree of cylinders T_c: Subsets of vertices of the defining graph satisfying certain graph theoretical conditions are in bijection with W-orbits of vertices of T_c and they generate the representatives of the conjugacy classes of the vertex stabilizers.*

This construction generalizes the one by Dani-Thomas [DT17] in Theorem 1.24 and Theorem 3.1 to a class of non-hyperbolic RACGs. As illustrated with examples throughout Section 3, it is particularly convenient that all vertex and edge groups can be "read off" the defining graph. The cylinder vertices are produced by a simple process, see Section 3.1.2.1: Each comes from an *uncrossed* cut collection and its common adjacent vertices. This implies that cylinder vertices occur only in three types, see Lemma 3.15: Two-ended, virtually $\mathbb{Z}^2$ or the direct product of a virtually non-abelian free group and an infinite dihedral group.

It is highlighted in Remark 2.14 that the edge groups in a graph of cylinders are not necessarily two-ended. However, in the case of RACGs, we characterize the edge stabilizers of the JSJ tree of cylinders visually, shown by combining Lemma 3.25 and Theorem 3.28 in Remark 3.30:

Theorem 1.46. *All the edge stabilizers of the JSJ tree of cylinders of a RACG W_Γ satisfying Standing Assumption 1 are two-ended if and only if in the defining graph no uncrossed cut collection contains opposite corners of a square whose other two corners are connected by a subdivided diagonal.*

With the JSJ graph of cylinders of a RACG given, the structure invariant defined in [CM17a] and introduced in Section 2.1.1 comes into play: One glance suffices to conclude that RACGs with rather basic defining graphs from Figure 3.2.6 of Example 3.36 are not QI:

The graph on the left has two uncrossed cut pairs, coloured in blue and red, which both have three common adjacent vertices. This implies that the corresponding cylinder vertices both have vertex groups that are the direct product of a virtually non-abelian free group and an infinite dihedral group. The red cut pair of the right graph, however, has only two common adjacent vertices. Thus, the corresponding cylinder vertex group is virtually $\mathbb{Z}^2$. This is an obstruction for the existence of a QI between the corresponding RACGs.

It is important to keep in mind that identical structure invariants in general do not imply that two groups are QI to each other. However, we can adjust it in the setting of RACGs by refining it to the *modified* structure invariant to make it a complete QI-invariant.

Theorem 1.47 (cf. Theorem 3.45)**.** *Let W and W' be two finitely presented, one-ended RACGs with non-trivial JSJ decompositions over two-ended subgroups, both without rigid vertices. Let T and T' be the JSJ trees of cylinders of W and W', respectively. Then, W and W' are QI if and only if T and T' have the same structure invariant up to reordering and QI-equivalence of vertex groups.*

Remark 1.48. The class of RACGs classified by Theorem 1.47 includes the hyperbolic RACGs covered by Theorem 1.25 and the RACGs on generalized Θ-graphs covered by Theorem 1.31.

With this Theorem 1.47 at hand, we can now immediately see that RACGs corresponding to defining graphs such as the following from Figure 3.2.12 of Example 3.46 are indeed QI:

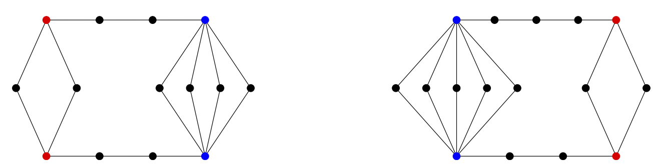

Both graphs have one red uncrossed cut pair with two common adjacent vertices producing a virtually $\mathbb{Z}^2$ cylinder vertex group and a blue uncrossed cut pair with more than two common adjacent vertices producing a cylinder vertex group that is the direct product of a virtually non-abelian free group and an infinite dihedral group. So, the two defining graphs produce the same (modified) structure invariant.

Additionally, Theorem 1.47 and its proof in Section 3.2 can be exploited to obtain various examples of RACGs that are QI, see Examples 3.49 and 3.50: Starting from a defining graph, we perform reflections and duplications of subgraphs to produce new graphs whose corresponding RACGs are QI to the original one. This method is even applicable to groups with rigid vertices, as long as these remain unaltered or have additional properties (see Remarks 3.37 and 3.47).

1.4 Status Quo of the QI-Problem between RACGs and RAAGs

Since each RACG is the quotient of a Right-Angled Artin group (RAAG), these classes of groups are strongly related. However, within the class of RAAGs, in comparison to the class of RACGs, the QI-Problem is significantly more advanced due to work of Huang [Hua17a, Hua16] and Margolis [Mar20]. On the other hand, the QI-Problem between RACGs and RAAGs has not received as much attention and was focused on finding examples with a positive answer. As for the QI-Problem within the class of RACGs, most results require an additional property like, for instance, the planarity of the defining graph.

We aim to give an overview of what is known in this section, including the results and in particular the answers in the negative produced in Section 4 of this thesis, see Subsection 1.4.4. We start with an introduction of RAAGs.

1.4.1 RAAGs

We introduce RAAGs along the lines of the introduction of RACGs in Section 1.2, see [Cha07] for a detailed survey of RAAGs.

Definition 1.49. For a finite, simplicial graph Δ with vertex set M, the *Right-Angled Artin Group (RAAG)* A_Δ is defined as the group given by the following presentation

$$A_\Delta = \langle m \in M \mid mn = nm \text{ if } (m,n) \in E(\Delta) \rangle .$$

The graph Δ is called the *defining graph* or *presentation graph*.

Remark 1.50. Like for RACGs, instead of the defining graph, the *Coxeter graph* is often used for general *Artin groups*. For RAAGs, it is the complement graph of the defining graph.

Throughout this thesis, we use the following notation for RAAGs.

Convention. Depending on whether we want to emphasize the defining graph Δ or the generating set $M = V(\Delta)$ of the RAAG we denote it as A_Δ or A_M, respectively.

Example 1.51. We obtain the following 'extrema' as standard examples of RAAGs:

- If Δ is a complete graph with $V(\Delta) = M$, then A_Δ is the free abelian group on $|M|$ generators, $A_\Delta = \mathbb{Z}^{|M|}$. Moreover, A_Δ is abelian if and only if Δ is complete.
- If Δ does not have any edges and $V(\Delta) = M$, then A_Δ is a free group on $|M|$ generators, $A_\Delta = *_{|M|}\,\mathbb{Z} = F_{|M|}$.

Example 1.52. The following are fundamental examples:

- For a graph $\Delta = \Delta' \circ \Delta''$ that is the join of Δ' and Δ'', A_Δ is the direct product $A_{\Delta'} \times A_{\Delta''}$.
- For a graph $\Delta = \Delta' \sqcup \Delta$ that is the disjoint union of two graphs Δ' and Δ'', A_Δ is the free product $A_{\Delta'} * A_{\Delta''}$.
- For the graph Δ_1 in Figure 1.4.8, A_{Δ_1} is the direct product $\mathbb{Z} \times \mathbb{Z} = \mathbb{Z}^2$.
- For the graph Δ_2 in Figure 1.4.8, opposite vertices do not commute, thus, they generate a non-abelian free group F_2 of rank 2. But since Δ_2 is the join of the two pairs of opposite vertices, A_{Δ_2} is the direct product $F_2 \times F_2$.

Remark 1.53. We deduce from Example 1.52 that the class of RAAGs is closed under taking direct products by taking the join and under taking free products by taking the disjoint union of defining graphs.

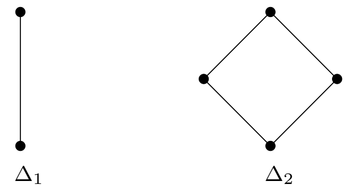

Figure 1.4.8

Like for RACGs, certain subgroups can be “read off” the defining graph:

Definition 1.54. Given a RAAG A_M on $M = V(\Delta)$, the subgroup A_N generated by $N \subseteq M$ is called a *special subgroup* of A_M.

A_N is itself a (right-angled) Artin group on the defining graph Δ_N, which is the induced subgraph of Δ on the vertices labelled by N, see [Cha07, Section 2.2]. Hence, the intersection of two special subgroups $A_N \cap A_{N'}$ is the special subgroup generated by the intersection $N \cap N'$.

Example 1.55. The RAAG A_{Δ_2} in Example 1.52 on the defining graph Δ_2 from Figure 1.4.8 contains for instance the special subgroups $\mathbb{Z} \times \mathbb{Z} = \mathbb{Z}^2$ and F_2.

Salvetti introduced a cube complex that encodes the presentation of the RAAG A_Δ, and its universal cover that encodes the geometry of the RAAG A_Δ. We will always be interested in the universal cover, so we will refer to that one as *the* Salvetti complex S_Δ. Its construction and properties can be found in [Cha07, Section 2.6].

We outline the following facts about the Salvetti complex of a RAAG relevant for this thesis: The Salvetti complex of a special subgroup $A_N \subseteq A_M$ embeds locally isometrically into the Salvetti complex of A_M. For RAAGs, the Salvetti complex is a CAT(0) cube complex. Its 1-skeleton is the Cayley graph $Cay(A_M, M)$ of A_M with respect to the generating set $M = V(\Delta)$. There is a geodesic of arbitrary length in the Cayley graph of A_M labelled by a single generator $m \in M$. We call such a geodesic a *standard geodesic*. Analogously, the subcomplex corresponding to an induced subgraph of Δ is called *standard subcomplex*.

1.4.2 Properties and QI-invariants of RAAGs

In this section we aim to give an overview of the properties of RAAGs that are helpful for the QI-problem between RACGs and RAAGs.

Since we aim to use the JSJ graph of cylinders to compare RACGs and RAAGs up to QI, we first need the analogous properties for RAAGs as for RACGs in Standing Assumption 1:

- **The Salvetti complex of A_Γ is two-dimensional** to simplify the geometry encoded by the group. This is the case if Δ is triangle-free.
- A_Δ **is one-ended**: In [BM01] Brady-Meier show that this is true if and only if Δ is connected and not a single vertex.
- A_Δ **has a splitting over a two-ended subgroup**: By Theorem A of [Cla14], the existence of a splitting over a two-ended subgroup is ensured if Δ has a cut vertex v. Indeed, if there is a cut vertex v, all k parts of $\Delta \setminus \{v\}$ attach along the two-ended special subgroup $A_{\{v\}} = \mathbb{Z}$ as a k-fold amalgamated product.

To ensure that A_Δ is two-dimensional, one-ended and splitting over two-ended subgroups, we fix:

Standing Assumption 2. Δ denotes the defining graph of a RAAG A_Δ which satisfies:

(1) Δ is triangle-free.
(2) Δ is connected and not a single vertex.
(3) Δ has a cut vertex.

Convention. Unless stated otherwise, throughout this thesis, every graph Δ satisfies the Standing Assumption 2 and every RAAG is defined on a graph satisfying the Standing Assumption 2.

Remark 1.56. If Δ is connected and not a single vertex, A_Δ is not hyperbolic: Since Δ contains at least one edge $(m_1, m_2) \in E(\Delta)$, there is a special subgroup $\mathbb{Z}^2 \cong A_{\{m_1,m_2\}} \leq A_\Delta$. However, it is a well-known fact that a hyperbolic group does not have a subgroup isomorphic to $\mathbb{Z}^2$.

For RAAGs satisfying Standing Assumption 2, several properties are established that are known to be a QI-invariant. We aim to use them to restrict the class of RACGs we need to consider.

1.4.2.1 Divergence The asymptotic type of the divergence of one-ended RAAGs is known:

Theorem 1.57. [BC12, Corollary 4.8] *A one-ended RAAG has at most quadratic divergence.*

Since by [Ger94], the divergence provides a QI-invariant, this implies:

Corollary 1.58. *If a group G is QI to a one-ended RAAG, then G has at most quadratic divergence.*

By Theorem 1.35, W_Γ is of at most quadratic divergence if and only if Γ is $\mathcal{CFS}$. Hence, we get:

Corollary 1.59. *If a RACG W_Γ is QI to a one-ended RAAG, then Γ is $\mathcal{CFS}$.*

1.4.2.2 Morse boundary Another way to distinguish groups up to QI is by comparing their *Morse boundaries*.

Theorem 1.60. [CCS23, Theorem 1.1] *The Morse boundary of a RAAG is totally disconnected.*

By [Cor17, Proposition 3.7], QI groups have isomorphic Morse boundaries, implying the following:

Corollary 1.61. *If a group G is QI to a RAAG, then its Morse boundary is totally disconnected.*

There are several criteria to check whether the Morse boundary of a RACG W_Γ is totally disconnected or contains a circle. For instance, by [CS11], the Morse boundary of W_Γ is empty if and only if its defining graph Γ is a clique or a non-trivial join. In [Kar23, Corollary 1.8], Karrer shows that the class of *clique-square-decomposable* groups have totally disconnected Morse boundary. On the other hand, Behrstock gave in [Beh19] the first class of RACGs on $\mathcal{CFS}$ defining graphs that contain an embedded circle in their Morse boundary formed by a certain induced n-cycle in Γ with $n \geq 5$. Another class of groups with an embedded circle in their Morse boundaries was given in [GKLS21] via a certain 3-path condition.

1.4.2.3 Morse subgroups It follows immediately form the definition that subsets with the following property are preserved under QI:

Definition 1.62. Let X be a quasi-geodesic metric space and $Y \subseteq X$ a subset. Then Y is *Morse* if there is a function $Q\colon [1,\infty) \times [0,\infty) \to [0,\infty)$ such that for every (C, D)-quasi-geodesic γ with endpoints in Y, γ is contained in the neighborhood $N_{Q(C,D)}(Y)$.

Remark 1.63. In [Tra19], Tran uses the term *strongly quasi-convex* for Morse subsets.

In RAAGs, Morse subsets satisfy a dichotomy:

Theorem 1.64. [RST23, Corollary 7.4; CH17, Theorem F] *In a one-ended RAAG every Morse subset is either a quasi-tree or coarsely covers the whole space. In particular, every Morse subgroup is either free or of finite index.*

Since the property that a subset is hyperbolic or coarsely covering the whole space is invariant under QI, any Morse subset of a group QI to a RAAG exhibits the same dichotomy:

Corollary 1.65. *Let G be group QI to a one-ended RAAG. Then every Morse subgroup of G is either hyperbolic or of finite index.*

The Morse property mimics quasi-convexity in hyperbolic spaces, so we aim to distinguish between Morse subsets that are hyperbolic and those that are not:

Definition 1.66. Let X be a quasi-geodesic metric space. The subset $Y \subseteq X$ is *stable* if it is Morse and hyperbolic and *eccentric* if it is minimally Morse unstable.

We can detect all Morse and all stable special subgroups of a RACG W_Γ in the defining graph Γ:

Definition 1.67. Let Γ be a graph and $\Gamma' \leq \Gamma$ an induced subgraph. Then Γ' is *square-complete* if it has the following property: If Γ' contains two non-adjacent vertices of an induced square σ, then Γ' contains all vertices of σ.

If Γ' contains at least one square, is square-complete and does not contain any proper induced square-complete subgraph still containing a square, Γ' is *minsquare*.

Theorem 1.68. [Tra19, Theorem 1.11] *A subgroup $W_{\Gamma'} \leq W_\Gamma$ is Morse if and only if Γ' is square-complete.*

By Lemma 1.23, a RACG is hyperbolic if and only if its defining graph does not have any square, implying that stable subgroups are visible in the defining graph:

Corollary 1.69. *The special subgroup $W_{\Gamma'} \leq W_\Gamma$ is stable if and only if Γ' is square-complete and does not contain any squares.*

The example of Behrstock in [Beh19] is also the first example of a stable special subgroup in a RACG whose defining graph is $\mathcal{CFS}$.

Corollary 1.70. *If W_Γ satisfies Standing Assumption 1 and is QI to a RAAG, then Γ is minsquare.*

Proof. Suppose Γ is not minsquare. Let $\Gamma' \lneq \Gamma$ be a proper induced minsquare subgraph. Such a Γ' exists, because as W_Γ is QI to a RAAG, Γ is $\mathcal{CFS}$ by Corollary 1.59. By Theorem 1.68, $W_{\Gamma'}$ is Morse. Thus, by Corollary 1.65, $W_{\Gamma'}$ is either hyperbolic or of finite index. However, Γ' contains a square, so, by Lemma 1.23, $W_{\Gamma'}$ is not hyperbolic.

Minimality in the minsquare condition implies that every vertex of Γ' is contained in a square in Γ'. Since Γ is connected and Γ' is proper, there is some vertex $s \in \Gamma \setminus \Gamma'$ adjacent to some vertex $v \in \Gamma'$. Let $(v, w_1, w_2, w_3) \subseteq \Gamma'$ be an induced square in Γ' containing v. By Standing Assumption 1, Γ is triangle-free, so s is not adjacent to neither w_1 nor w_3. Thus, elements of the form $p \cdot s$ with $p \in W_{\{w_1,w_3\}}$ give infinitely many cosets of $W_{\Gamma'}$. Thus, $W_{\Gamma'}$ is also not of finite index. Hence, such a subgraph Γ' does not exist and Γ is minsquare. ☐

Also eccentric subgroups can be recognized in the defining graph:

Theorem 1.71. [Gen22, Theorem 1.8] *Let W_Γ be a RACG with defining graph Γ. A subspace $Y \subseteq Cay(W_\Gamma, V(\Gamma))$ is eccentric if and only if there is an induced minsquare subgraph $\Gamma' \leq \Gamma$ such that Y is finite Hausdorff distance from a coset of $W_{\Gamma'}$.*

The following example brought to the author's attention by Pallavi Dani has an eccentric subgroup of infinite index and thus is not QI to any RAAG by Corollary 1.70:

Example 1.72. The square (c, f, i, l) in the defining graph Γ in Figure 1.4.9 is minsquare, thus Γ itself is not. By Corollary 1.70, W_Γ is not QI to a RAAG. Note that the special subgroup $W_{\{c,f,i,l\}}$ is not hyperbolic by Lemma 1.23, eccentric by Theorem 1.71 and of infinite index in W_Γ.

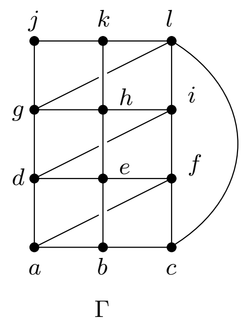

Figure 1.4.9

Outline 1.73. Suppose a RACG W_Γ is QI to a RAAG satisfying Standing Assumption 2. Then we can assume that Γ has the following properties:

- By Corollary 1.59, Γ is $\mathcal{CFS}$.
- By Corollary 1.61, W_Γ has totally disconnected Morse boundary.
- By Corollary 1.70, Γ is minsquare.

1.4.3 Existence of QIs

One approach to the QI-problem between RACGs and RAAGs is to search for pairs consisting of a RACG and a RAAG that are in fact QI to each other. When starting with a RAAG, we can always provide such a pair by a fundamental result of Davis-Januszkiewicz:

Theorem 1.74. [DJ00] *Every RAAG is commensurable to a RACG.*

This is why it is sufficient for the QI-classification to ask for the converse of when a RACG W_Γ is QI to a RAAG A_Δ. The explicit construction in [DJ00] of a RACG QI to a given RAAG can be beneficial for investigating the converse problem:

- Given a RACG, identifying it as a result of the Davis-Januszkiewicz construction ensures that it is QI to a RAAG, see for instance Example 4.61.
- In trying to find new QI-invariants, QI groups A_Δ and W_Γ determined by the Davis-Januszkiewicz construction provide excellent test examples.

When starting with a RACG W_Γ, a way to find a RAAG commensurable to W_Γ from looking at Γ was given by Dani-Levcovitz (see Section 4.1 for details):

Theorem 1.75 (cf. Theorem 4.13)**.** [DL20] *There is a graph theoretical algorithm on the graph Γ satisfying Standing Assumption 1 that, if it succeeds, finds a visual RAAG subgroup of finite index of the RACG W_Γ.*

Theorem 1.75 can be used to classify RACGs on planar defining graphs QI to RAAGs:

Theorem 1.76 (cf. Theorem 4.17)**.** [DL20] *A RACG W_Γ on a planar defining graph satisfying Standing Assumption 1 is QI to a RAAG if and only if it has a finite index visual RAAG subgroup found by the graph theoretical algorithm.*

Their proof of Theorem 1.76 (see Theorem 4.17 for the full statement and a sketch of the proof) uses the following QI-classification for RAAGs, whose defining graphs are trees:

Theorem 1.77. [NT19, Theorem 1.2] *Let W_Γ be a RACG on a connected, $\mathcal{CFS}$, non-join, triangle-free, planar defining graph with at least 5 vertices and no separating vertex or edge. Then the following statements are equivalent:*

(1) *W_Γ is QI to a RAAG A_Δ.*
(2) *W_Γ is QI to a RAAG A_Δ on a tree Δ of diameter at least 3.*
(3) *All vertices in the maximal suspension graph of W_Γ are spacious.*

Idea of the Proof. By the proof of Theorem 1.37, we know that a RACG W_Γ satisfying the given assumptions is a 3-manifold group with a 3-manifold decomposition given by the maximal suspension graph T_Γ. Suppose that W_Γ is QI to a RAAG A_Δ. By [BN08], a RAAG A_Δ is QI to a 3-manifold group if and only if it is a 3-manifold group itself. In this case, [Gor04] implies that the defining graph Δ of the RAAG A_Δ is a tree of diameter at least 3.

Recall that if a vertex in T_Γ is spacious, its corresponding 3-manifold has boundary, if it is full it does not. However, by [BN08], all vertex groups in the 3-manifold decomposition of a RAAG on a tree correspond to a 3-manifold with boundary. Since the boundary is preserved under QI, this implies that if W_Γ is QI to such a RAAG, it can only have spacious vertices. Conversely, if W_Γ only has spacious vertices, its 3-manifold decomposition also provides a RAAG QI to W_Γ by [BN08]. □

In fact, the *Dani-Levcovitz-Algorithm* (see Definition 4.1) of Theorem 1.75 also provides insight in the non-planar case:

Theorem 1.78 (cf. Theorem 4.23)**.** *If Γ satisfies Standing Assumption 1 and can be constructed by the Coning Algorithm 4.21.1, then the RACG W_Γ has a finite index visual RAAG subgroup provided by the Dani-Levcovitz-Algorithm.*

Moreover, via this method of *coning*, we can also construct from Γ, whose RACG W_Γ has a visual RAAG subgroup, a new graph Γ', whose RACG $W_{\Gamma'}$ has a visual RAAG subgroup as well, see Corollary 4.24.

Also, we can simplify the application of the Dani-Levcovitz-Algorithm:

Proposition 1.79 (see Proposition 4.30 and Corollary 4.31 for the precise statement)**.** *The Dani-Levcovitz-Algorithm provides a visual RAAG subgroup of W_Γ, where Γ satisfies Standing Assumption 1, if it provides a visual RAAG subgroup for every vertex group in the JSJ graph of cylinders of W_Γ and the defining graphs of these RAAGs match up in a certain way.*

Given a RACG W_Γ with a visual RAAG subgroup A_Δ of finite index, any finite index subgroup of A_Δ is a finite index RAAG subgroup of W_Γ as well. It was brought to the author's attention by Pallavi Dani and Annette Karrer that, like for RACGs (see Definition 1.42 and Lemma 1.43), we can produce finite index RAAG subgroups of a RAAG A_Δ via doubling:

Definition 1.80. Let Δ be a graph and let $v \in V(\Delta)$ be a vertex. The *double of* Δ *over* v is the graph Δ_v that has the following vertex set and edge set:

$$V(\Delta_v) = V(\Delta) \cup \{m' \mid m \in V(\Delta \setminus st(v))\},$$
$$E(\Delta_v) = E(\Delta) \cup \{(m', n') \mid (m, n) \in E(\Delta)\} \cup \{(m', l) \mid l \in lk(v), (m, l) \in E(\Delta)\},$$

where for every vertex $m \in V(\Delta \setminus st(v))$, we define a new vertex $m' \in V(\Delta_v)$ as the double of v.

Lemma 1.81. *Let A_Δ be a RAAG and Δ_v be the double of Δ over some $v \in V(\Delta)$. Then A_{Δ_v} is an index 2 subgroup of A_Δ.*

Sketch of the Proof. Define the following map:

$$\begin{array}{rccll} \phi\colon & A_\Delta & \to & \mathbb{Z}_2 & \\ & v & \mapsto & 1 & \\ & m & \mapsto & 0 & \text{for } \ m \in V(\Delta) \setminus \{v\} \end{array} .$$

The kernel of the map ϕ is generated by $\langle \{m \mid m \in V(\Delta) \setminus \{v\}\} \cup \{v^2\} \cup \{vmv \mid m \in V(\Delta) \setminus \{v\}\}\rangle$ and is isomorphic to the RAAG on the double Δ_v of Δ over v, where the vertex v corresponds to the generator v^2 and each new vertex m' corresponds to the generator vmv. □

Example 1.82. For the RAAG A_Δ with Δ on the left of Figure 1.4.10, the RAAG A_{Δ_c} on the double Δ_c of Δ over c illustrated on the right of Figure 1.4.10 is a subgroup of index 2.

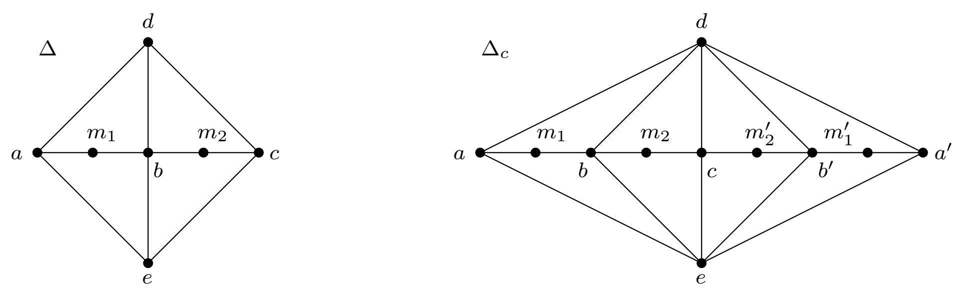

Figure 1.4.10

1.4.4 Obstruction for QIs

In Section 1.4.2 we have seen that not all RACGs are QI to a RAAG, since there are for instance RACGs with polynomial divergence of degree at least 3 (cf. Theorem 1.33), connected Morse boundary or eccentric subgroups of infinite index (see Outline 1.73). Thus, it is natural to search further for new examples and classes of RACGs, for which there are obstruction for QIs to RAAGs. There are two new approaches in this direction: Using the structure invariant and the MPRG.

In [Mar20], Margolis gives a description of the JSJ graph of cylinders of a RAAG satisfying Standing Assumption 2, see Theorem 4.34. Thus, we can use the structure invariant to compare the JSJ graphs of cylinders of RAAGs and RACGs with each other, see Section 4.2.1.

As highlighted in Outline 4.46, if W_Γ is QI to a RAAG, we can use the structure invariant to draw the following conclusions about W_Γ and its JSJ graph of cylinders $\Sigma_{c,\Gamma}$:

- By Proposition 4.40, every cylinder vertex group of $\Sigma_{c,\Gamma}$ is the direct product of a virtually non-abelian free group and an infinite dihedral group.
- By Proposition 4.41, the defining graph of every rigid vertex group of $\Sigma_{c,\Gamma}$ is $\mathcal{CFS}$.
- By Proposition 4.42, every rigid vertex in $\Sigma_{c,\Gamma}$ has relative QI type $[\![(G_r, \mathcal{P}_r)]\!]$ with either G_r and all edge groups in $\mathcal{P}_r$ virtually $\mathbb{Z}^2$ or neither G_r nor any edge group in $\mathcal{P}_r$ virtually $\mathbb{Z}^2$.
- By Proposition 4.43, no rigid vertex group of $\Sigma_{c,\Gamma}$ splits over a two-ended subgroup.

By using these conclusions, we give several new examples of RACGs not QI to any RAAG in Section 4.2.2.

Besides the structure invariant, we can use the MPRG to distinguish RACGs and RAAGs up to QI as suggested in Section 1.1.2: We focus on the case, where $\mathcal{A}$ is the class of two-dimensional RAAGs. By an application of [Oh22, Lemma 4.11], see Theorem 2.36, the MPRG of a two-dimensional RAAG has the property $\mathcal{P}$ that removing the star of certain vertices disconnects it. Thus, any two-dimensional RACGs, whose MPRG does not have this property, is not QI to any RAAG.

This fact is exploited in Section 4.3: We focus on a graph Γ with an edge that is not contained in any square. In this case, the MPRG of the RACG W_Γ exhibits a certain structure:

Lemma 1.83 (cf. Lemma 4.53)**.** *Let Γ be a triangle-free $\mathcal{CFS}$ graph with an edge $(s,t) \in E(\Gamma)$ not contained in a square. Then this edge creates a (subdivided) (s,t)-*square *$S_{(s,t)}$ in the MPRG Γ^p. In particular, if the edge (s,t) is removed from Γ, so is the (s,t)-square $S_{(s,t)}$ from Γ^p.*

In certain cases, such a square then creates an infinitely long *ladder* in the MPRG that contradicts the property $\mathcal{P}$ of RAAGs that every sufficiently far apart pair of points is separated by the star of some vertex.

Theorem 1.84 (cf. Theorem 4.55 and Remark 4.56)**.** *Let Γ be a triangle-free $\mathcal{CFS}$ graph with an edge $(s,t) \in E(\Gamma)$ not contained in a square. Under certain conditions, conjugates of the corresponding (possibly subdivided) (s,t)-square $S_{(s,t)}$ create a subgraph of the MPRG Γ^p in the shape of an infinitely long* ladder. *Assuming that the* rungs *of the ladder are wide, there is no vertex in Γ^p whose star disconnects the ladder, and thus, the RACG W_Γ is not QI to any RAAG.*

We give new examples of graphs containing an edge that is not contained in a square in Section 4.3.1, where an application of Theorem 1.84 leads to the conclusion that the RACG is not QI to a RAAG. In Example 4.61, however, we point out that an edge not contained in a square does not always cause an obstruction for the existence of a QI: The square created by Lemma 1.83 has only side length 2, which prevents Theorem 1.84 from being applicable. In fact, by doubling (cf. Lemma 1.43) and a comparison with the Davis-Januszkiewicz construction [DJ00] (cf. Theorem 1.74), we can show that the RACG in Example 4.61 is in fact commensurable and thus QI to a RAAG. Hence, we are curious to establish when such an edge is an indication for the existence or non-existence of a QI, see Outline 4.62 and Question 9.

2 Tools

As described in Section 1.1, the JSJ graph of cylinders (see Section 1.1.1) and the maximal product region graph (see Section 1.1.2) are valuable tools to advance the QI-Problem. This section is dedicated to a detailed introduction of these two decompositions.

2.1 JSJ graph of cylinders

The following subsection about the *JSJ graph of cylinders* is from Section 2.2 of the author's final version of [Edl21].

Throughout this subsection, let T be a simplicial tree and G a finitely generated group acting on T by isometries and without edge inversions. The stabilizer of any element t in T is denoted as G_t, geodesic paths in T starting at vertex a and ending at vertex b are denoted as $[a, b]$. Let $\mathcal{A}$ be a class of infinite subgroups of G that is stable under conjugation. T is an *$\mathcal{A}$-tree* if all the edge stabilizers G_e of T are contained in $\mathcal{A}$.

Example 2.1. Since we split RACGs over two-ended subgroups, the class of subgroups we have in mind as $\mathcal{A}$ is the class $\mathcal{VC}$ of virtually infinite cyclic (or equivalently two-ended) subgroups. Note that $\mathcal{VC}$ is invariant under conjugation, but not under taking subgroups.

Our main tool is a universal tree on which G acts with vertex stabilizers as small as possible:

Definition 2.2.

1. A subgroup H of G is *elliptic* in T if it fixes a point in T. It is a *universally elliptic* subgroup if it fixes a point in any $\mathcal{A}$-tree. An $\mathcal{A}$-tree is *universally elliptic* if all its edge stabilizers are universally elliptic subgroups of G.
2. An $\mathcal{A}$-tree T *dominates* another $\mathcal{A}$-tree T' if every vertex stabilizer of T is elliptic in T'.
3. A *JSJ tree* of G is an $\mathcal{A}$-tree T that is universally elliptic and that dominates any other universally elliptic $\mathcal{A}$-tree T'. The quotient graph $\Sigma = T/G$ is called a *JSJ decomposition* or *JSJ splitting* of G.

JSJ trees are extensively surveyed in [GL17]. Unfortunately, the JSJ tree is not as universal as we would like it to be. It does not even always exist, nor is it unique if it does. It rather happens that we find a collection of universally elliptic trees, which are pairwise dominating each other. This collection then is called the *JSJ deformation space* [GL17, Section 2.3].

Remark 2.3. If in a graph of groups $\Sigma = T/G$ of G, whose edge groups are all universally elliptic, also up to conjugation all universally elliptic subgroups of G occur as edge groups, Σ is a JSJ decomposition of G. Indeed, if all edge groups in Σ are universally elliptic, so are the edge stabilizers of T, thus T is universally elliptic. Furthermore, given any other universally elliptic tree T', we can refine it to T, and T' is therefore dominated by T [GL17, Lemma 2.15]. Thus, T is a JSJ tree.

We aim to obtain a more accessible equivalent definition, when restricting to one-ended groups splitting over two-ended subgroups. For that, we introduce the following terminology:

Definition 2.4. [cf. GL17, Definition 5.13] A vertex v of a graph of groups Σ over two-ended edge groups and its vertex group G_v are called *hanging* if G_v maps onto the fundamental group $\pi_1(X_v)$ of a hyperbolic, compact, two-dimensional orbifold X_v and the image of every edge group incident to G_v in $\pi_1(X_v)$ is either finite or contained in a boundary subgroup of $\pi_1(X_v)$. We call v and G_v *maximal hanging* if there is no other hanging vertex group G_w such that the corresponding orbifold X_w can be glued to X_v along identical boundary components to obtain a new splitting of the group.

Remark 2.5. While the interpretation of a *hanging* subgroup is not universal, in the setting of RACGs all existing versions are equivalent: For instance, suppose, following [Bow98], we see a vertex group G_v which is non-elementary, finitely generated and which acts properly discontinuously on the hyperbolic plane $\mathbb{H}^2$. This is equivalent to saying that G_v surjects with finite kernel onto the fundamental group of a hyperbolic, compact, two-dimensional orbifold X_v [cf. Bar18, Defnition 3.2.]. If additionally all the incident edge groups of G_v map onto the fundamental groups of the boundary components of X_v, Bowditch calls G_v *hanging Fuchsian*. However, then G_v meets the Definition 2.4 of a hanging vertex group as well.

Also, it is worth noting that in their Definition 5.13 in [GL17], Guirardel and Levitt define the vertex and vertex group we call hanging as *quadratically hanging (QH)*, to extend the definition of quadratically hanging subgroups given by Rips and Sela in [RS97]. Moreover, various authors call vertex groups meeting the properties of Definition 2.4 along similar lines as the hanging Fuchsian groups, *hanging surface groups* for instance.

Definition 2.6. A vertex v of a graph of groups Σ over two-ended edge groups and its vertex group G_v are called *rigid* if G_v is not two-ended, not hanging and does not split over a two-ended subgroup relative to its incident edge groups.

By piecing together Theorem 6.5, Corollary 6.3, Section 2.6 and Proposition 5 of [GL17], which rely on work of Fujiwara and Papasoglu [FP06], and the results of [Pap05], we can describe certain JSJ decompositions neatly in terms of graphs of groups:

Lemma 2.7. *If G is a finitely presented, one-ended group not commensurable to a surface group, a graph of groups decomposition with two-ended edge groups is a JSJ decomposition if and only if the following conditions hold:*

- *Each vertex group is either two-ended, hanging or rigid.*
- *Any valence one vertex v with two-ended vertex group does not have an incident edge group surjecting onto G_v.*
- *All hanging vertex groups are maximal.*

Even though JSJ decompositions are not unique, under certain conditions, we can produce a canonical representative of the JSJ deformation space, the so-called *tree of cylinders* T_c. The rest of this subsection gives a short overview of its construction. For all details, see [GL11].

Definition 2.8. An equivalence relation $\sim$ on $\mathcal{A}$ is called *admissible* if for all $A, B \in \mathcal{A}$ the following axioms hold:

1. If $A \sim B$ and $g \in G$ then $gAg^{-1} \sim gBg^{-1}$.
2. If $A \subseteq B$, then $A \sim B$.
3. Given an $\mathcal{A}$-tree T and $a, b \in V(T)$ that are fixed by $A, B \in \mathcal{A}$, respectively, then for every edge $e \subseteq [a, b]$ we have $A \sim G_e \sim B$.

Definition 2.9. Two subgroups H and K of a group G are called *commensurable* if their intersection $H \cap K$ has finite index in both H and K. The *commensurator* of a subgroup H in G is the set

$$Comm_G(H) = \{g \in G \mid gHg^{-1} \text{ and } H \text{ are commensurable}\}.$$

Commensurability is an equivalence relation on subgroups. We denote the equivalence class of $A \in \mathcal{A}$ by $[A]$. The stabilizer of $[A]$ under the action of G on $\mathcal{A}/\sim$ by conjugation is denoted as $G_{[A]}$.

Example 2.10. On the class $\mathcal{VC}$ of two-ended subgroups of G, commensurability is an admissible equivalence relation. For $A \in \mathcal{VC}$, we obtain $G_{[A]} = Comm_G(A)$.

Construction 2.11. Given an $\mathcal{A}$-tree T, we construct the object of interest, the *cylinder*, in the following way:

- Start with an admissible equivalence relation $\sim$ on $\mathcal{A}$.
- Define two edges $e, f \in E(T)$ to be equivalent if their edge stabilizers G_e and G_f are equivalent, that is, $e \sim f$ if $G_e \sim G_f$.
- If G_e fixes the edge $e \in E(T)$, in particular, it fixes its endpoints $o(e), t(e) \in V(T)$. Thus, by axiom (3) for an admissible relation, all the edges on a path between two equivalent edges are in the same equivalence class as well. Thus, this equivalence class forms a subtree Y of T, called a *cylinder* of T.
- By construction, two distinct cylinders can intersect at most in one common vertex.
- We refer to the equivalence class in $\mathcal{A}/{\sim}$ containing all edge stabilizers of edges in Y as $[Y]$.

Definition 2.12. Given an admissible equivalence relation on $\mathcal{A}$ and an $\mathcal{A}$-tree T, its *tree of cylinders* T_c is the following bipartite tree with vertex set $V_1 \sqcup V_2$: The vertex set V_1 contains one vertex v_Y per cylinder Y, the *cylinder vertices*. The vertex set V_2 contains all the vertices of T that belong to at least two cylinders. There is an edge $(v_Y, v) \in E(T_c)$ between v_Y and every vertex v contained in Y. The graph of groups decomposition of G coming from the quotient of the action of G on T_c is the *graph of cylinders* Σ_c.

The stabilizer G_Y of a cylinder vertex v_Y in V_1 is $G_{[Y]}$. The stabilizer G_v of a vertex v in V_2 is the stabilizer G_v of v viewed as a vertex of T. An edge (v_Y, v) in $E(T_c)$ is stabilized by the intersection of $G_{[Y]}$ and G_v.

Example 2.13. Consider the Baumslag-Solitar group $BS(m,n) = \langle a, b \mid b^{-1}a^m b = a^n \rangle$ defined for the integers $m, n \in \mathbb{Z} \setminus \{0\}$. We view it as an HNN-extension with stable letter b and consider its action on the corresponding Bass-Serre tree T. All the edge stabilizers are of the form $g\langle a^m \rangle g^{-1}$ for $g \in BS(m,n)$, thus they are contained in $\mathcal{VC}$. By use of the inductive consequence

$$b^{-k} a^{m^k y} b^k = a^{n^k y} \quad \text{for any } k, y \in \mathbb{N}$$

of the relation, one shows that $\langle a^m \rangle$ is commensurable to $g\langle a^m \rangle g^{-1}$ for any $g \in BS(m,n)$. Hence, all edges are part of the same commensurability-cylinder and T_c consists of only one vertex.

Remark 2.14. T_c is not necessarily an $\mathcal{A}$-tree. This problem is resolved by collapsing all edges that have edge stabilizers not in $\mathcal{A}$ to obtain the *collapsed tree of cylinders* T_c^*. However, in our application of the construction, we aim to bypass this complication.

Convention. Henceforth, when the set $\mathcal{A}$ and the admissible equivalence relation on it are not specified, it is fixed to be $\mathcal{VC}$ with the commensurability relation, as in Example 2.10.

The question left to answer is how the construction of the tree of cylinders gives a canonical object encoding the structure of the group. Starting from a finitely presented, one-ended group G, we pick some JSJ tree T of the JSJ deformation space, which exists by [GL17, Theorem 1]. For T, we construct the tree of cylinders T_c, which by [GL11, Theorem 1] does not depend on the choice of T but only on the deformation space itself. Thus, it makes sense to call it *the JSJ tree of cylinders* and the corresponding graph of cylinders Σ_c *the JSJ graph of cylinders*. While for instance for hyperbolic groups, Σ_c is itself a JSJ decomposition [GL17, Theorem 9.18], this is not true in general. However, by construction its Bass-Serre tree is G-equivariantly isomorphic to the tree of cylinders of any JSJ tree. Hence, from the JSJ graph of cylinders Σ_c, we can essentially recover the deformation space of JSJ splittings.

Moreover the JSJ tree of cylinders produces a QI-invariant for groups, by a result of Cashen and Martin based on work of Papasoglu [Pap05, Theorem 7.1] with a correction made by Shepherd and Woodhouse in [SW22]:

Theorem 2.15. [CM17a, Theorem 2.9; SW22, Theorem 2.8] *Given two finitely presented, one-ended groups G and G' splitting over two-ended subgroups which are quasi-isometric via $\phi : G \to G'$, then ϕ induces a tree isomorphism $\phi_* : T_c \to T'_c$. Moreover, ϕ_* is vertex-type preserving and for every vertex $v \in V(T)$ with vertex group G_v, there is a real constant $C_v \geq 0$ such that ϕ maps G_v within distance C_v of $G'_{\phi_*(v)}$.*

Thus, ideally, we construct the JSJ graphs of cylinders directly from the groups, in our case from the defining graphs of the RACGs. Deducing from them that the corresponding JSJ trees of cylinders are not isomorphic then implies that the groups we started with are not QI. On the other hand, if there is an isomorphism between the JSJ trees of cylinders, we try to promote it to a QI of the groups.

Outline 2.16. To summarize, the framework we focus on is the following: The group G is finitely presented, one-ended and splits over the set of two-ended subgroups $\mathcal{VC}$. We obtain a JSJ splitting Σ, in which all vertex groups are either two-ended, hanging or rigid by Lemma 2.7. By considering the commensurability relation on the corresponding JSJ tree, we produce the JSJ graph of cylinders Σ_c, whose cylinder vertex groups are the commensurators of the two-ended groups of Σ and whose non-cylinder vertex groups are precisely the hanging and rigid vertices of Σ.

2.1.1 The structure invariant

To see whether such a tree isomorphism as in Theorem 2.15 between the JSJ trees of cylinders of two groups can exist, Cashen and Martin introduced in [CM17a] the *structure invariant*. The following subsection recalls their construction and was taken from Section 3.2.1 of the author's final version of [Edl21].

We fix T to be a simplicial tree of countable valence and G to be a group acting on T cocompactly and without edge inversions. We introduce some terminology following [CM17a, Section 3].

Definition 2.17. Given an arbitrary set $\mathcal{O}$ of *ornaments*, a G-invariant map $\delta \colon V(T) \to \mathcal{O}$ is called a *decoration*. The tree T is said to be *decorated*.

Example 2.18. A standard set of ornaments for a JSJ tree of cylinders T_c is the *vertex type*, that is $\mathcal{O} = \{\text{'cylinder'}, \text{'hanging'}, \text{'rigid'}\}$. A possibly finer decoration is obtained by equipping each vertex v with the ornament consisting of the vertex type and the so-called *relative QI-type* of the corresponding vertex group G_v. This relative QI-type is determined as follows: Given the vertex group G_v, we consider the set $\mathcal{P}_v$ of distinct Hausdorff equivalence classes in G_v of G_v-conjugates of images of the edge injections $\alpha_e \colon G_e \hookrightarrow G_v$, where $e \in E(T_c)$ is an edge incident to v. $\mathcal{P}_v$ is often referred to as the *peripheral structure of G_v coming from incident edge groups* or just as the *peripheral structure of G_v*. Then the relative QI-type $[\![(G_v, \mathcal{P}_v)]\!]$ of G_v is the set of all pairs (Y, P), where Y is a geodesic metric space and P is a collection of Hausdorff equivalence classes of subsets of Y such that there is a QI from G_v to Y inducing a bijection from $\mathcal{P}_v$ to P. Thus, the relative QI-type captures the structure of the vertex group with respect to its incident edge groups up to QI.

Definition 2.19. A decoration $\delta' \colon V(T) \to \mathcal{O}'$ is called a *(strict) refinement* of the decoration $\delta \colon V(T) \to \mathcal{O}$ if the δ'-partition $\bigsqcup_{o' \in \mathcal{O}'} (\delta')^{-1}(o')$ of $V(T)$ (strictly) refines the δ-partition $\bigsqcup_{o \in \mathcal{O}} \delta^{-1}(o)$. A non-strict refinement is called *trivial*.

The refinement process used to obtain the structure invariant is the *neighbor refinement*, which is an idea generalizing the degree refinement algorithm known from graph theory. It works as follows:

Construction 2.20. Let $\overline{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$ and call $\mathcal{O}_0 = \mathcal{O}$ the *initial set of ornaments* and $\delta_0 = \delta$ the *initial decoration*. Starting from $i = 0$, we define for each $i \in \mathbb{N}$ and each $v \in V(T)$ the map

$$\begin{array}{rccl} f_{v,i}\colon & \mathcal{O}_i & \to & \overline{\mathbb{N}} \\ & o & \mapsto & \left|\{w \in \delta_i^{-1}(o) \mid (w,v) \in E(T)\}\right| . \end{array}$$

Define $\mathcal{O}_{i+1}$ as $\mathcal{O}_0 \times \overline{\mathbb{N}}^{\mathcal{O}_i}$ and δ_{i+1} as the pair $(\delta_0(v), f_{v,i}) \in \mathcal{O}_0 \times \overline{\mathbb{N}}^{\mathcal{O}_i}$.

Cashen and Martin prove the following facts about the maps defined in Construction 2.20:

Lemma 2.21. [CM17a, Lemma 3.2, Proposition 3.3] *The map $\delta_{i+1}\colon V(T) \to \mathcal{O}_{i+1}$ is a decoration refining $\delta_i\colon V(T) \to \mathcal{O}_i$ for all $i \in \mathbb{N}$. Furthermore, this refinement process stabilizes. That is, there is an $s \in \mathbb{N}$ such that for any $i + 1 \leq s$, the decoration δ_{i+1} is a strict refinement of δ_i, but for any $i \geq s$, the refinement δ_{i+1} is trivial.*

Definition 2.22. The decoration $\delta_s\colon V(T) \to \mathcal{O}_s$ at which the neighbor refinement process stabilizes, is called the *neighbor refinement* of δ.

To capture the information contained in the neighbor refinement, we define $\pi_0\colon \mathcal{O}_s \to \mathcal{O}$ to be the projection to the first coordinate. After choosing an ordering on the image $\delta(V(T))$, we denote the j-th element as $\mathcal{O}[j]$. Then we can choose an ordering of $\pi_0^{-1}(\mathcal{O}[j]) \cap \delta_s(V(T))$. We order $\delta_s(V(T))$ lexicographically and denote the i-th element as $\mathcal{O}_s[i]$.

Definition 2.23. A *structure invariant* $S = S(T, \delta, \mathcal{O})$ is the $|\delta_s(V(T))|^2$-matrix, where

$$S_{j,k} = (n_{j,k}, \pi_0(\mathcal{O}_s[j]), \pi_0(\mathcal{O}_s[k]))\,,$$

with $n_{j,k}$ the number of vertices in $\delta_s^{-1}(\mathcal{O}_s[j])$ adjacent to $\delta_s^{-1}(\mathcal{O}_s[k])$. The second entry of the tuple $S_{j,k}$ is the *row* and the third entry the *column ornament*.

We can view $S(T, \delta, \mathcal{O})$ as a block matrix, which is well-defined up to block permutations and the choice of ordering on $\delta(V(T))$ and $\pi_0^{-1}(\mathcal{O}[j])$. We denote a structure invariant in a table with entries $n_{j,k}$, whose rows and columns are labelled by the initial decoration $\delta(V(T))$, as illustrated in Example 3.36 or labelled by the vertex orbit representatives carrying the same ornaments, as illustrated in Example 3.44.

As indicated in the definition, a structure invariant depends on the initial choice of ornaments and decoration. When we refer to *the* structure invariant, the initial decoration is the one introduced in Example 2.18: the ornaments consist of vertex type and relative QI-type. We call two vertices in the JSJ graph of cylinders *indistinguishable* if they have the same image under δ_s.

By construction, the structure invariant relates to the existence of a tree isomorphim between the JSJ tree of cylinders:

Proposition 2.24. [cf. CM17a, Proposition 3.7] *Given two groups G and G' with JSJ trees of cylinders T_c and T_c', and G- and G'-invariant decorations $\delta\colon V(T_c) \to \mathcal{O}$ and $\delta'\colon V(T_c') \to \mathcal{O}$, respectively, there is a decoration-preserving isomorphism $\phi\colon T_c \to T_c'$ if and only if up to permuting rows and columns within $\mathcal{O}$-blocks, $S(T_c, \delta, \mathcal{O}) = S(T_c', \delta', \mathcal{O})$.*

We use Proposition 2.24 to distinguish RACGs up to QI in Section 3.2 and to distinguish RACGs and RAAGs up to QI in Section 4.2.

2.2 Maximal Product Region Graph

Since we are interested in RACGs and RAAGs, which act geometrically on their associated CAT(0) cube complex, another way to investigate their geometry is via this complex. As highlighted in the introductory Section 1.1.2, we can translate the algebraic approach described in Section 2.1 of decomposing a group into a collection of vertex subgroups that is invariant under QI to the setting of the CAT(0) cube complexes: We determine a collection of special subcomplexes that is preserved under QI. For instance, in [Hua17b, Theorem 1.3] (see Theorem 2.32), Huang shows that the collection of top-dimensional flats is preserved under QI. Oh applies this to show in [Oh22, Theorem C] that all maximal standard subcomplexes in a square complex exhibiting a certain product structure form such a collection (see Theorem 2.31). In this section we introduce these results and the necessary terminology in the setting of RACGs and RAAGs.

To discuss both RACGs and RAAGs at once, we fix Ω to be the defining graph of either a RACG or a RAAG and refer to the corresponding group as G_Ω.

Definition 2.25. Let Ω be a graph.

- A *product subgraph* of Ω is an induced subgraph I such that I decomposes as the join of two graphs I_1 and I_2.
- A product subgraph $I \leq \Omega$ is called *maximal* if I is not properly contained in any other product subgraph of Ω.
- A maximal product subgraph is *essential* if the special subgroups G_{I_1} and G_{I_2} generated by I_1 and by I_2, respectively, are both infinite.
- We denote the collection of all essential maximal product subgraphs of Ω by $\mathcal{M}(\Omega)$. The collection of all special subgroups G_M generated by some $M \in \mathcal{M}(\Omega)$ is called *special maximal product subgroups* and is denoted as $\mathcal{M}(G_\Omega)$.

Remark 2.26. For a product subgraph $I = I_1 \circ I_2 \leq \Omega$, the corresponding special subgroup G_I decomposes as the direct product of G_{I_1} and G_{I_2}. For a RAAG, every maximal product subgraph is essential, since every special subgroup is infinite.

Definition 2.27. The *maximal product region graph (MPRG)* Ω^p of Ω is the graph with the following vertex set and edge set:

$$V(\Omega^p) = \{pG_Mp^{-1} \mid G_M \in \mathcal{M}(G_\Omega),\, p \in G_\Omega\},$$
$$E(\Omega^p) = \{(p_1G_{M_1}p_1^{-1}, p_2G_{M_2}p_2^{-1}) \mid \exists\, H \leq p_1G_{M_1}p_1^{-1} \cap p_2G_{M_2}p_2^{-1} : H \cong \mathbb{Z}^2\}.$$

The induced subgraph $R_\Omega \leq \Omega^p$ on the vertex set corresponding to the trivial conjugates of the special maximal product subgroups in $\mathcal{M}(G_\Omega)$ is called the *fundamental domain* of Ω^p.

Remark 2.28. G_Ω acts on Ω^p by conjugation with finite fundamental domain R_Ω.

Convention. We denote the action by conjugation by powers of elements: Let $v \in V(R_\Omega)$ be a vertex in the fundamental domain with corresponding maximal product subgroup $G_M \in \mathcal{M}(G_\Omega)$. Then the vertex in Ω^p corresponding to the conjugate pG_Mp^{-1} with $p \in G_\Omega$, is denoted by ${}^pv$.

Remark 2.29. To produce the fundamental domain R_Ω, we only need to consider intersections of the essential maximal product subgraphs in Ω: In a RACG W_Γ, the intersection $W_{M_1\cap M_2}$ of W_{M_1} and W_{M_2} contains a subgroup isomorphic to $\mathbb{Z}^2$ if and only if M_1 and M_2 share a square in Γ, by Corollary 1.19. In a RAAG A_Δ, the intersection $A_{M_1\cap M_2}$ of A_{M_1} and A_{M_2} contains a subgroup isomorphic to $\mathbb{Z}^2$ if and only if M_1 and M_2 share an edge in Δ.

Remark 2.30. In [Oh22, Definition 3.5], Oh gives a geometric interpretation of the MPRG by defining the *intersection complex* of a special square complex as the graph, where every maximal standard product subcomplex corresponds to a vertex and every intersection in a 2-flat corresponds to an edge. The MPRG of a RACG or a RAAG G_Ω in Definition 2.27 is in correspondence with this intersection complex of the Davis or Salvetti complex of G_Ω, respectively.

With this correspondence, we can reformulate the results of [Oh22] to provide a QI-invariant between RACGs and RAAGs, given by the following theorem:

Theorem 2.31. [Oh22, cf. Theorem 3.7] *Let $\phi : W_\Gamma \to A_\Delta$ be a QI between a RACG W_Γ satisfying Standing Assumption 1 and a RAAG A_Δ satisfying Standing Assumption 2, then there is an induced isomorphism $\phi^* : \Gamma^p \to \Delta^p$ between their MPRGs. Specifically, ϕ^* is an isomorphism of graphs that also preserves the QI-types of the maximal product regions corresponding to each vertex.*

We emphasize that Theorem 2.31 is limited to the 2-dimensional setting, because then all flats are naturally top-dimensional and the following key result about *weakly special* cube complexes, which include Davis and Salvetti complexes, can be used to control all standard flats:

Theorem 2.32. [Hua17b, Theorem 1.3] *If there is a QI ϕ between the universal covers X_1 and X_2 of two compact weakly special cube complexes with the same dimension, then there is a constant $C > 0$ such that for every top-dimensional flat $F_1 \subseteq X_1$, there is a top-dimensional flat $F_2 \subseteq X_2$ such that $d_H(\phi(F_1), F_2) < C$, where d_H denotes the Hausdorff distance.*

It would be very interesting to have a higher dimensional version of Theorem 2.31, but dealing with maximal flats that are not top-dimensional is considerably more difficult.

2.2.1 Properties of the MPRG of RAAGs

We aim to use Theorem 2.31 as a QI-invariant between RACGs and RAAGs by constructing the MPRG of a given RACG and distinguishing it from the MPRG of any RAAG.

A fundamental fact is the connectedness of the MPRG:

Proposition 2.33. [Oh22, Proposition 4.6] *The MPRG Δ^p of a RAAG A_Δ satisfying Standing Assumption 2 is connected.*

In addition, we obtain our key criteria from the following properties of the MPRG:

Proposition 2.34. [cf. Oh22, Lemma 4.11] *Let Δ^p be the MPRG of the RAAG A_Δ satisfying the Standing Assumption 2 and let $v \in V(\Delta^p)$ be a vertex corresponding to a maximal product subgraph of Δ that contains the star of a vertex in Δ. Then the star $st_{\Delta^p}(v)$ of v separates Δ^p. In particular, the complement $\Delta^p \setminus st_{\Delta^p}(v)$ of the star $st_{\Delta^p}(v)$ of v has infinitely many connected components.*

The condition in Proposition 2.34 that a maximal product subgraph of Δ contains the star of a vertex in Δ is not always satisfied. Sangrok Oh provided examples to the author, where the MPRG has a vertex whose star is not separating it. However, Proposition 2.34 implies that the 2-ball around a vertex in Δ^p is always sufficient to separate it:

Corollary 2.35. *Let Δ^p be the MPRG of the RAAG A_Δ satisfying the Standing Assumption 2 and let $v \in V(\Delta^p)$ be a vertex in Δ^p. Then the ball in Δ^p around v of radius 2 separates Δ^p.*

Proof. Let $v_1 \in V(\Delta^p)$ be a vertex in the MPRG graph and let $M_1 \leq \Delta$ be its corresponding maximal product subgraph in Δ. If M_1 contains the star in Δ of some vertex, by Proposition 2.34, the star $st_{\Delta^p}(v_1)$ separates Δ^p, and thus, so does the ball in Δ^p around v_1 of radius 2.

So, suppose that M_1 does not contain the star in Δ of a vertex and let $m \in V(M_1)$ be some vertex in M_1. Then the star $st_\Delta(m)$ of m in Δ is contained in some other maximal product subgraph $M_2 \leq \Delta$. Let $v_2 \in V(\Delta^p)$ be the vertex in Δ^p corresponding to M_2 that is closest to v_1. Since $st_{M_1}(m) \subseteq M_1 \cap M_2$ contains an edge, the vertices v_1 and v_2 share an edge in Δ^p. But by Proposition 2.34, the star $st_{\Delta^p}(v_2)$ is separating Δ^p, thus, so is the ball in Δ^p around v_1 of radius 2. □

However, the star of a lot of vertices in the MPRG are separating it by the following theorem, which is a consequence of the proof of Proposition 2.34 and was obtained in collaboration with Christopher Cashen and Sangrok Oh:

Theorem 2.36. *Let Δ^p be the MPRG of the RAAG A_Δ satisfying the Standing Assumption 2 with fundamental domain R_Δ. Then for any pair of distinct conjugates of R_Δ, there is a vertex $v \in V(\Delta^p)$ such that the star $st_{\Delta^p}(v)$ of v separates them.*

Proof. We fix the notation: The x-conjugate of the fundamental domain R_Δ in the MPRG Δ^p is denoted by ${}^xR_\Delta$, the corresponding subcomplex in the Salvetti complex S_Δ is denoted by the translate xR_Δ.

Let $v_1 \in {}^{x_1}R_\Delta$ and $v_2 \in {}^{x_2}R_\Delta$ be two vertices in Δ^p in two different conjugates in Δ^p of the fundamental domain R_Δ for $x_1, x_2 \in A_\Delta$. Consider a combinatorial path p in the Salvetti complex S_Δ between x_1 and x_2. Let $e \in E(\Delta^p)$ be the first edge in p that does not lie in the translate x_1R_Δ in S_Δ corresponding to the conjugate ${}^{x_1}R_\Delta$. Let h_e be the hyperplane dual to e, let $m \in V(\Delta)$ be the label of the edge e and let l be the standard geodesic labelled by m containing the edge e. Every combinatorial path between x_1 and x_2 in the Salvetti complex S_Δ crosses the hyperplane h_e, and h_e separates S_Δ into two parts X_1 and X_2 with $x_1 \in X_1$ and $x_2 \in X_2$.

Let $v \in V(\Delta^p)$ be the vertex in the MPRG Δ^p corresponding to the maximal product subgraph containing the star $st_\Delta(m)$ and whose corresponding maximal product subcomplex of S_Δ contains the edge e. Let c be a combinatorial path in Δ^p connecting v_1 to v_2. We claim that c passes through the star $st_{\Delta^p}(v)$, implying that $st_{\Delta^p}(v)$ separates the MPRG Δ^p: Indeed, c passes through a sequence of conjugates $({}^{x_1}R_\Delta = {}^{y_1}R_\Delta, \dots, {}^{y_n}R_\Delta = {}^{x_2}R_\Delta)$ with $y_i \in A_\Delta$ and such that $y_iR_\Delta \cap y_{i+1}R_\Delta \neq \emptyset$ for every $i \in \{1, \dots, n\}$. By the choice of the hyperplane h_e, there is some $i_0 \in \{1, \dots, n-1\}$ such that $y_{i_0} \in X_1$ and $y_{i_0+1} \in X_2$. Thus, the maximal product region $K \subseteq S_\Delta$ contained in the intersection $y_{i_0}R_\Delta \cap y_{i_0+1}R_\Delta$ of the subcomplexes $y_{i_0}R_\Delta$ and $y_{i_0+1}R_\Delta$ contains a standard geodesics dual to h_e and thus parallel to l. However, this implies that K intersects the maximal product region corresponding to the vertex $v \in V(\Delta^p)$ in a flat. Hence, K corresponds to a vertex in the MPRG Δ^p contained in the star $st_{\Delta^p}(v)$, and thus, c passes through $st_{\Delta^p}(v)$. □

Corollary 2.37. *Let Δ^p be the MPRG of the RAAG A_Δ satisfying the Standing Assumption 2. Then there is a constant $D \geq 4$ such that for every pair of vertices $u, w \in V(\Delta^p)$ at distance at least D from one another, there is a vertex $v \in V(\Delta^p)$ whose star in Δ^p separates u and w.*

The separating property described in Proposition 2.34 has to be preserved under QI:

Corollary 2.38. *Let W_Γ be a RACG that satisfies Standing Assumption 1 and is QI to a RAAG. Then the star $st_{\Delta^p}(v)$ of a vertex $v \in V(\Gamma^p)$ in the MPRG corresponding to an essential maximal product subgraph that is a suspension separates the MPRG Δ^p. In particular, the complement $\Delta^p \setminus st_{\Delta^p}(v)$ has infinitely many connected components.*

Proof. Let $v \in V(\Gamma^p)$ be a vertex in the MPRG corresponding to an essential maximal product subgraph that is a suspension. Thus, v corresponds to a maximal product region that is virtually the product of a tree and a line. Then, if there is a QI ϕ between W_Γ and a RAAG A_Δ, the induced isomorphism $\phi^* : \Gamma^p \to \Delta^p$ between the MPRGs preserves the QI-type of the maximal product regions by Theorem 2.31. So, as a vertex in Δ^p, v corresponds to a maximal product region in A_Δ which is virtually the product of a tree and a line as well. This implies that there is a vertex $m \in V(\Delta)$ such that the maximal product region of A_Δ corresponding to v is given by the star $st_\Delta(m)$. Hence, the property of Proposition 2.34 is satisfied and the complement of the star of the vertex v in the MPRG Γ^p has infinitely many connected components. □

Another essential property of the MPRG of RAAGs is the following:

Theorem 2.39. [Oh22, Corollary 4.9] *For a RAAG A_Δ satisfying Standing Assumption 2, the MPRG Δ^p is a quasi-tree.*

These properties lead to the following strategies to distinguish RACGs and RAAGs up to QI by the use of the MPRG:

Outline 2.40. Given a RACG W_Γ, we aim to investigate its MPRG Γ^p. If we can show that one of the properties of the MPRG of a RAAG is not satisfied, by Theorem 2.31, we conclude that W_Γ cannot be QI to any RAAG. In particular, we aim to use one of the following four things:

1. Show that the MPRG Γ^p is not connected, then it cannot be isomorphic to the MPRG of a RAAG by Proposition 2.33, and thus, W_Γ is not QI to any RAAG.
2. Find a sequence $(v_i)_{i=0}^{\infty} \subseteq V(\Gamma^p)$ of infinitely many distinct vertices in Γ^p such that for every $i \in \mathbb{N}$, v_0 and v_i are connected in $\Gamma^p \setminus st(v)$ for every $v \in \Gamma^p \setminus \{st(v_i) \cup st(v_0)\}$. If Γ^p is isomorphic to the MPRG of some RAAG A_Δ, there is a finite fundamental domain R_Δ such that by Theorem 2.36, two vertices in different translates of R_Δ are separated by the star of a vertex. However, since the sequence $(v_i)_{i=0}^{\infty}$ is infinite, some v_j for $j \in \mathbb{N}$ is in a different translate of R_Δ than v_0, but v_0 and v_j are not separated by any star. So, Γ^p is not isomorphic to the MPRG of a RAAG, and thus, W_Γ is not QI to any RAAG.
3. Find a vertex $v \in V(\Gamma^p)$, whose corresponding essential maximal product subgraph is a suspension, such that $\Gamma^p \setminus st(v)$ has finitely many connected components. Then, by Corollary 2.38, Γ^p cannot be isomorphic to the MPRG of a RAAG, and thus, W_Γ is not QI to any RAAG.
4. Show that the MPRG Γ^p is not a quasi-tree, then it cannot be isomorphic to the MPRG of a RAAG by Theorem 2.39, and thus, W_Γ is not QI to any RAAG.

2.2.2 Relation between MPRG and other graphs

For RAAGs, there is an important construction of a graph, called *extension graph*, which is related to the construction of the MPRG. The extension graph has equivalent algebraic and geometric definitions, but both depend on the RAAG presentation, and it is not clear how to adapt these definitions for RACGs.

In the next two subsections we recall the main results on extension graphs and some related attempts to define graphs in terms of cube complexes by the use of hyperplanes, like the *contact graph*. While it is yet to be determined how the MPRG fits into the big picture, we conclude that for our purpose the MPRG is the best choice, because it is a QI-invariant for both RAAGs and RACGs by Theorem 2.31.

2.2.2.1 Extension Graph The following well-studied and useful relative of the MPRG graph is given by two equivalent definitions:

Definition 2.41. [KK13, Definition 1] Let A_Δ be a RAAG. The *extension graph* Δ^e is the graph with the following vertex set and edge set:

$$V(\Delta^e) = \{pvp^{-1} \mid v \in V(\Delta),\, p \in A_\Delta\},$$
$$E(\Delta^e) = \{(pvp^{-1}, qwq^{-1}) \mid [pvp^{-1}, qwq^{-1}] = 1\}.$$

In [Hua17a, Lemma 4.2], Huang shows that this Definition 2.41 is equivalent to the following:

Definition 2.42. [Hua17a, Section 4.1.1] Let A_Δ be a RAAG with Salvetti complex S_Δ. The *extension graph* $\mathcal{P}(\Delta)$ is the graph with a vertex for every class of parallel standard geodesics in S_Δ. Two distinct vertices v_1 and v_2 are connected by an edge if there is a representative l_i per corresponding class of standard geodesics for $i \in \{1, 2\}$ such that l_1 and l_2 span a 2-flat.

Convention. Since both Definition 2.41 and Definition 2.42 of the extension graph of a RAAG A_Δ are equivalent, we use the notation Δ^e uniformly in both cases and if necessary state which definition is used explicitly.

Remark 2.43. It is highlighted in [KK14] that the extension graph of the Salvetti complex of a RAAG is the analogue of the *curve graph* of a manifold. A broader generalization of the curve graph are *hierarchically hyperbolic spaces*, thus, the extension graph also resembles their machinery, as emphasized in Section 1.3 of [BHS17].

It was pointed out to the author by Jingyin Huang that in a certain set-up, the MPRG of a RAAG can be constructed directly from its extension graph by removal of its leaves. This lets us exploit the known properties of the extension graph for the study of the MPRG.

Definition 2.44. Let Ω be a graph. Then the *core* $Core(\Omega)$ of Ω is the induced subgraph of Ω on all vertices except the vertices of valence 1.

Proposition 2.45. *If A_Δ is a RAAG with Δ a triangle-free graph such that all maximal product subgraphs are stars, then $\Delta^p = Core(\Delta^e)$.*

Sketch of the Proof. Since by assumption Δ is triangle-free and every maximal product subgraph is a star, every maximal product subgraph is a join J_m of its middle vertex $m \in V(\Delta)$ and its pairwise non-adjacent neighbors. This implies that Δ is square-free. Thus, every vertex $m \in Core(\Delta)$ corresponds to the middle vertex of a maximal product subgraph J_m. Therefore, the conjugates of m in Δ^e are in correspondence with the conjugates of A_{J_m} in Δ^p. Only the vertices in $\Delta \setminus Core(\Delta)$ are accounted for in Δ^e but not in Δ^p. This implies the correspondence between $Core(\Delta^e)$ and Δ^p. □

Remark 2.46. For one-ended RAAGs, the assumption in Proposition 2.45 that every maximal product subgraph is a star is equivalent to assuming that every maximal product subgroup is of the form $F \times \mathbb{Z}$, for F a non-abelian free group.

The extension graph can be constructed from Δ algorithmically:

Lemma 2.47. [KK13, Lemma 22] *Let A_Δ be a RAAG with extension graph Δ^e and let $D \leq \Delta^e$ be an induced subgraph. Then there is an $n > 0$, a sequence of vertices $(v_1, v_2, \ldots, v_n) \subseteq V(\Delta^e)$ and a sequence of finite induced subgraphs $\Delta = \Delta_0 \leq \Delta_1 \leq \cdots \leq \Delta_n$, where Δ_i is obtained by doubling Δ_{i-1} over v_i (as described in Definition 1.80) for every $i \in \{1, \ldots, n\}$ such that $D \leq \Delta_n$.*

From this algorithm, we can deduce that the extension graph has the same separability properties as the MPRG described in Proposition 2.33:

Lemma 2.48. [KK13, Lemma 26.*(6)* and *(7)*] *The extension graph Δ^e of a one-ended RAAG A_Δ is separated by the star of any vertex $v \in V(\Delta)$. In particular, the complement $\Delta^e \setminus st(v)$ of the star $st(v)$ of v has infinitely many connected components. Moreover, the extension graph Δ^e is a quasi-tree.*

Remark 2.49. The algorithm to construct a finite induced subgraph of the extension graph Δ^e described in Lemma 2.47 is the same process of doubling (introduced in Definition 1.80) as performed in Lemma 1.81 to obtain finite index RAAG subgroups.

In fact, the extension graph detects all RAAG subgroups of RAAGs:

Theorem 2.50. [KK13, Theorem 1.11] *Let Δ and E be two finite, triangle-free graphs. Then the RAAG A_E embeds into the RAAG A_Δ if and only if $E \leq \Delta^e$.*

Remark 2.51. Theorem 2.50 resembles the search for visual finite index RAAG subgroups in RACGs in [DL20] by constructing a FIDL-Λ (see Definition 4.14), introduced in Section 4.1. See Remark 4.16 for more details.

The extension graph is an invariant used to show the following QI-classification of RAAGs:

Theorem 2.52. [Hua17a, Theorem 1.1] *Two RAAGs with finite outer automorphism groups are QI if and only if they are isomorphic.*

Theorem 2.53. [Hua17a, Theorem 1.2] *Suppose A_{Δ_1} is a RAAG with finite outer automorphism group and A_{Δ_2} is any RAAG. Then the following statements are equivalent:*

(1) A_{Δ_2} is QI to A_{Δ_1}.
(2) A_{Δ_2} is isomorphic to a finite index subgroup of A_{Δ_1}.
(3) Δ_2^e is isomorphic to Δ_1^e.

In case the outer automorphism group of the RAAG is infinite, additional properties are needed:

Definition 2.54. [Hua16, Definition 1.1 and Definition 1.5] A RAAG A_Δ is

- of *weak type I* if:
  - (i) Δ is connected and does not contain any separating star.
  - (ii) There do not exist vertices $v, w \in V(\Delta)$ such that $d(v, w) = 2$ and $\Delta = st(v) \cup st(w)$.
- of *type II* if Δ is connected and for every pair of distinct vertices $v, w \in V(\Delta)$, $lk(v) \cap lk(w)$ does not separate Δ.

Theorem 2.55. [Hua16, Theorem 1.2] *If the RAAGs A_{Δ_1} and A_{Δ_2} are of weak type I, then they are quasi-isometric if and only if they are isomorphic.*

Theorem 2.56. [Hua16, Theorem 1.3] *Suppose A_{Δ_1} is a RAAG of weak type I and A_{Δ_2} is a RAAG. Then the following statements are equivalent:*

(1) A_{Δ_2} is QI to A_{Δ_1}.
(2) A_{Δ_2} is isomorphic to a finite index subgroup of A_{Δ_1}.
(3) A_{Δ_2} is isomorphic to a special subgroup of A_{Δ_1}.

Theorem 2.57. [Hua16, Theorem 1.2] *If the RAAG A_{Δ_1} is of type II then a RAAG A_{Δ_2} is QI to A_{Δ_1} if and only if A_{Δ_2} is commensurable to A_{Δ_1}. Moreover, there exists a RAAG A_Δ such that A_{Δ_1} and A_{Δ_2} are special subgroups of A_Δ.*

Remark 2.58. Being weak type I or type II is strongly related to the outer automorphism group $\mathrm{Out}(A_\Delta)$ of A_Δ (see [Hua16, Section 1.2]):

- If $\mathrm{Out}(A_\Delta)$ is finite, then A_Δ is of weak type I.
- If A_Δ is of type II, then $\mathrm{Out}(A_\Delta)$ does not contain *non-adjacent transvections*, but only *partial conjugations* and *adjacent transvections* (see [Hua16, Section 2.3] for definitions).

In light of the utility of the extension graph for the QI-classification and for finding RAAG subgroups, as well as of the Remarks 2.51 and 4.16 about finding finite index RAAG subgroups in RACGs, it is natural to ask:

Question 1. Is there an analogous definition of the extension graph of a RACG such that

(i) it is a QI-invariant for RACGs;
(ii) it detects RAAG subgroups of RACGs?

2.2.2.2 Hyperplane Graphs

Another method to define a graph from a CAT(0) cube complex is to consider classes of hyperplanes as vertices and use the edges to describe their interplay. There are several useful ways to do this:

Definition 2.59. [Hag14, Definition 2.16] The *contact graph* $\mathcal{C}(X)$ of a CAT(0) cube complex X is a graph that has a vertex for every hyperplane in X and two vertices $v, w \in V(\mathcal{C}(X))$ are connected by an edge if the carriers of their corresponding hyperplanes h_v and h_w intersect.

Remark 2.60. The contact graph of Definition 2.59 can be viewed as the CAT(0) cube complex analogue of the *curve graph* of a surface. This comparison motivated the introduction of *hierarchically hyperbolic spaces* in [BHS17] and [BHS19], which provide a general framework for these concepts.

Remark 2.61. If Δ does not have an isolated vertex, then the contact graph $\mathcal{C}(S_\Delta)$ of the Salvetti complex S_Δ of the RAAG A_Δ is quasi-isometric to the extension graph via the following map:

$$f\colon \quad \begin{array}{rcl} \Delta^e & \to & \mathcal{C}(S_\Delta) \\ gvg^{-1} & \mapsto & gh_v \end{array} .$$

This map is introduced in Section 7 of [KK14]. It is also one way to see by Lemma 2.48 that the contact graph $\mathcal{C}(S_\Delta)$ of a RAAG A_Δ is a quasi-tree, which was first proven in [Hag14, Theorem 4.1].

Similar to the contact graph we can define the crossing graph:

Definition 2.62. [Hag14, Definition 2.16] The *crossing graph* $\mathcal{X}(X)$ of a CAT(0) cube complex X is a graph that has a vertex for every hyperplane in X and two vertices $v, w \in V(\mathcal{C}(X))$ are connected by an edge if their corresponding hyperplanes h_v and h_w cross.

We can quotient the crossing graph by an equivalence relation for hyperplanes:

Definition 2.63. [Fio22, Definition 1] The *reduced crossing graph* $\mathcal{X}_r(X)$ of a CAT(0) cube complex X is a graph that has a vertex for every maximal collection $\mathcal{V} \subseteq \mathcal{H}(X)$ of hyperplanes with the property that for any two hyperplanes $v_1, v_2 \in \mathcal{V}$, each hyperplane $h \in \mathcal{H}(X)$ of X is crossing v_1 if and only if it is crossing v_2. Vertices corresponding to subsets $\mathcal{V}, \mathcal{W} \subseteq \mathcal{H}(X)$ are joined by an edge if there are hyperplanes $v \in \mathcal{V}$ and $w \in \mathcal{W}$ such that v and w cross.

The following corrected version of a statement in [Fio22] and its proof was brought to the author's attention in private communication with Elia Fioravanti:

Lemma 2.64. [cf. Fio22, Introduction, item (iii)] *Let A_Δ be a RAAG. Then $\mathcal{X}_r(S_\Delta)$ is isomorphic to Δ^e if Δ satisfies the following conditions:*

(i) There are no inclusions between links in Δ.
(ii) There are no two vertices with coinciding stars in Δ.

Proof. We introduce another graph, the *coset graph* $\mathcal{CS}(\Delta)$ (by using the $\mathcal{C}$osets of $\mathcal{S}$tars), as an auxiliary tool. It has has the following vertex set and edge set:

$$\begin{aligned} V(\mathcal{CS}(\Delta)) &= \{gK_v \mid g \in A_\Delta, v \in V(\Delta), K_v = \langle st(v)\rangle\}, \\ E(\mathcal{CS}(\Delta)) &= \{(g_1K_v, g_2K_w) \mid [g_1vg_1^{-1}, g_2wg_2^{-1}] = 1\}. \end{aligned}$$

We define a map ϕ_1 between the coset graph $\mathcal{CS}(\Delta)$ and the extension graph Δ^e:

$$\begin{array}{rccc} \phi_1: & \mathcal{CS}(\Delta) & \to & \Delta^e \\ & gK_v & \mapsto & gvg^{-1} \end{array}$$

We check that ϕ_1 is a graph isomorphism:

- ϕ_1 is a well-defined graph homomorphism: Since the edge conditions for $\mathcal{CS}(\Delta)$ and Δ^e are the same, ϕ_1 clearly maps edges to edges. Given $gK_v = g'K_{v'}$, then $K_v = K_{v'}$ and thus, by (ii), $v = v'$, and $g = g'k$ with $k \in K_v$. Hence, k commutes with v and thus $g'v'g'^{-1} = gkvk^{-1}g^{-1} = gvg^{-1}$.
- ϕ_1 is surjective, since A_Δ is acting equivariantly on $\mathcal{CS}(\Delta)$.
- ϕ_1 is injective: Consider

$$gvg^{-1} = \phi_1(gK_v) = \phi_1(g'K_{v'}) = g'v'g'^{-1}\,.$$

Then, since v and v' are generators, $v = v'$ and $g' = g_1kg_2$, where $k \in \langle st(v)\rangle = K_v$, k commutes with g_2 and $g_1g_2 = g$. Hence,

$$g'K_{v'} = g_1kg_2K_v = g_1g_2kK_v = g_1g_2K_v = gK_v\,.$$

Define now a map ϕ_2 between the coset graph $\mathcal{CS}(\Delta)$ and the reduced crossing graph $\mathcal{X}_r(S_\Delta)$:

$$\begin{array}{rccc} \phi_2: & \mathcal{CS}(\Delta) & \to & \mathcal{X}_r(S_\Delta) \\ & gK_v & \mapsto & [gh_v], \end{array}$$

where $[gh_v]$ denotes the collection of hyperplanes in Definition 2.63 corresponding to a vertex, containing the representative gh_v, that is, the hyperplane crossing the edge labelled by v at the vertex g. Also ϕ_2 is a graph isomorphism:

- ϕ_2 is well-defined: If gvg^{-1} commutes with $g'v'g'^{-1}$, then gh_v crosses $g'h_{v'}$, hence edges get mapped to edge. Given $gK_v = g'K_{v'}$, then $K_v = K_{v'}$ and thus, by (ii), $v = v'$, and $g = g'k$ with $k \in K_v$. Hence, k commutes with v, and the hyperplane gh_v passing through the edge labelled by v at the vertex g also passes through the edge labelled by v at the vertex gk. Hence,

$$\phi_2(gK_v) = [gh_v] = [gkh_v] = [g'h_v] = [g'h_{v'}] = \phi_2(g'K_{v'})\,.$$

- ϕ_2 is surjective, since A_Δ is acting equivariantly on $\mathcal{CS}(\Delta)$.

- ϕ_2 is injective: Assume that ϕ_2 is not injective and let $gK_v \neq g'K_{v'} \in V(\mathcal{CS}(\Delta))$ such that

$$[gh_v] = \phi_2(gK_v) = \phi_2(g'K_{v'}) = [g'h_{v'}]\,.$$

  Without loss of generality we assume that $g' = 1$. Let $g = m_1 \cdots m_n$ with $m_i \in V(\Delta)$ for every $i \in \{1, \ldots, n\}$. We claim that $[gh_v] = [h_{v'}]$ implies that $lk(v) = lk(v')$.
  Suppose this is not true. If $lk(v) \neq lk(v')$, without loss of generality, there is $x \in lk(v') \setminus lk(v)$. But then h_x crosses $h_{v'}$ but does not cross gh_v. Hence, $[gh_v] \neq [h_{v'}]$, in contradiction to our assumption.
  Thus, we can assume that $lk(v) = lk(v')$. If $v \neq v'$, this is in contradiction to (i), so $v = v'$ and $[h_v] = [h_{v'}] = [gh_v]$.
  Now, we show that $m_i \in (lk(v) \cup (\bigcap_{l \in lk(v)} lk(l))$ for every $i \in \{1, \ldots, n\}$. Suppose there is some $i_0 \in \{1, \ldots, n\}$ such that $m_{i_0} \notin (lk(v) \cup (\bigcap_{l \in lk(v)} lk(l))$. Hence, $m_{i_0} \notin lk(v)$ and $m_{i_0} \notin \bigcap_{l \in lk(v)} lk(l)$, which implies that there is $l_0 \in lk(v)$ such that $m_{i_0} \notin st(l_0)$. However, then h_{l_0} crosses h_v but does not cross gh_v, which contradicts the assumption that $[gh_v] = [h_v]$.
  So, as $g = m_1 \cdots m_n$ with $m_i \in (lk(v) \cup (\bigcap_{l \in lk(v)} lk(l))$ for every $i \in \{1, \ldots, n\}$, there are two options: Either for every $i \in \{1, \ldots, n\}$ we have that $m_i \in lk(v)$, but then $g \in \langle st(v) \rangle = K_v$, thus, $gK_v = K_v$ and ϕ_2 is injective. Or we find some m_{i_0} for $i_0 \in \{1, \ldots, n\}$ such that $m_{i_0} \notin lk(v)$ and $m_{i_0} \in \bigcap_{l \in lk(v)} lk(l)$. But then $lk(v) \subseteq lk(m_{i_0})$, which again contradicts assumption (i). So, $[gh_v] = [g'h_{v'}]$ implies that $gK_v = g'K_{v'}$, thus ϕ_2 is injective.

In conclusion, the extension graph Δ^e and the reduced crossing graph $\mathcal{X}_r(\Delta)$ are isomorphic as they are both isomorphic to the coset graph $\mathcal{CS}(\Delta)$. □

Example 2.65. Consider the RAAG on the graph Δ in the first column of Table 2.2.1, consisting of one edge, $A_\Delta = \langle a, b \mid ab = ba \rangle \cong \mathbb{Z}^2$. We obtain the following Table 2.2.1 of associated graphs:

Δ	$\mathcal{C}(S_\Delta)$	$\mathcal{X}(S_\Delta)$	$\mathcal{X}_r(S_\Delta)$	Δ^e	Δ^p

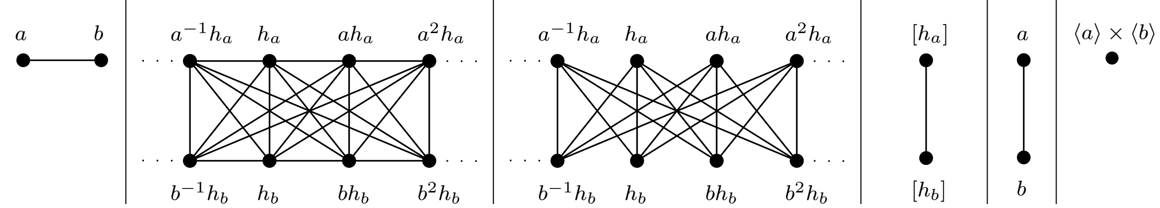

Table 2.2.1: Graphs associated to the RAAG A_Δ.

Naturally, we ask again about RACGs:

Question 2. Can this framework of the reduced crossing graph be also used on RACGs to establish a correspondence with an analogue of the extension graph?

Outline 2.66. As highlighted in Section 2.2.2 and Example 2.65, we start building graphs from hyperplanes of the Salvetti complex S_Δ of a RAAG A_Δ satisfying Standing Assumption 2 and obtain (sometimes) new graphs by coarsening the recorded information to eventually obtain QI-invariants:

1. Contact graph $\mathcal{C}(S_\Delta)$:
   - quasi-tree
2. Crossing graph $\mathcal{X}(S_\Delta)$:
   - remove edges in $\mathcal{C}(S_\Delta)$ corresponding to osculation between hyperplanes

3. Reduced crossing graph $\mathcal{X}_r(S_\Delta)$:
   - identify parallel hyperplanes in $\mathcal{X}(S_\Delta)$ with each other
4. Extension graph Δ^e:
   - equivalent to $\mathcal{X}_r(S_\Delta)$ if conditions (i) and (ii) in Lemma 2.64 on links and stars hold
   - QI-invariant
   - QI to contact graph $\mathcal{C}(S_\Delta)$
   - quasi-tree
5. Maximal product region graph Δ^p:
   - equivalent to $Core(\Delta^e)$ if every maximal product subgraph is a star
   - QI-invariant
   - quasi-tree

As pointed out in Remarks 2.60 and 2.43, these graphs, in particular the contact and the extension graph, are relevant within the study of hierarchically hyperbolic spaces. So, we would like to know:

Question 3. How does the MPRG fit into the framework of hierarchically hyperbolic spaces?

3 QIs within the class of RACGs

The following chapter was taken from Sections 2.3 to 4 with the exception of Section 3.2.1 (which can be found in Section 2.1) from the author's final version of [Edl21]. The only changes are in notations (Λ replaced by Σ and Δ replaced by Ω) and adjustments of references and section titles.

3.1 JSJ graph of cylinders of RACGs

3.1.1 Hyperbolic case

For one-ended, two-dimensional, hyperbolic RACGs whose defining graphs do not have any cut triples, a way to construct the JSJ graph of cylinders directly from the defining graph Γ is given in [DT17]. By [Dav08, Corollary 12.6.3], a RACG W_Γ is hyperbolic if and only if Γ has no squares (see Lemma 1.23). Although Dani and Thomas's construction follows the one for *Bowditch's JSJ tree* described in [Bow98], it turns out that the tree they produce in their (main) Theorem 3.37 corresponds to the JSJ tree of cylinders of W_Γ. More precisely, since W_Γ is hyperbolic, it follows from [GL17, Theorem 9.18] that both trees, and thus, their corresponding decompositions are W_Γ-equivariantly isomorphic.

Dani and Thomas claim in [DT17] that they give a construction of Bowditch's JSJ tree for *all* one-ended, two-dimensional, hyperbolic RACGs splitting over two-ended subgroups. However, they miss the fact that a RACG can not only split over a two-ended D_∞-subgroup coming from a cut pair, but also over a two-ended $D_\infty \times \mathbb{Z}_2$-subgroup coming from a cut triple. The origin of this problem is a miscitation of Theorem 1 of [MT09] as Theorem 2.1 in [DT17] claiming that every splitting over a two-ended subgroup corresponds to a cut pair. Example 1.21 gives a counterexample to this claim.

However, under the mild additional assumption that the defining graph Γ does not have any cut triples, all the results in [DT17] remain valid. We add this assumption whenever referring to results in [DT17]. This additional assumption was also implicitly used in an earlier version of [Edl21], however, the error has been removed as Theorem 3.33 now also includes the construction of the JSJ tree of cylinders of RACGs splitting over two-ended subgroups coming from cut triples in both the hyperbolic case and the non-hyperbolic case. In particular, removing this assumption does not affect the strategy and large-scale geometry results developed in [DT17] and in [Edl21], but only certain descriptions of the subgraphs of Γ corresponding to the large-scale structures of interest.

Most of the proofs in [DT17] depend on the hyperbolicity, in particular on the existence of the Gromov boundary, of the group W_Γ. Before we can produce the broader result, we want to understand the correspondence between the two constructions of the JSJ tree of cylinders. This subsection is dedicated to this task.

In our terminology, the JSJ tree of cylinders of one-ended, two-dimensional, hyperbolic RACGs splitting over D_∞-subgroups is produced by the following theorem:

Theorem 3.1. [cf. DT17, Theorem 3.37] *Let W_Γ be a hyperbolic RACG with Γ satisfying the Standing Assumption 1 and in addition let Γ have no cut triples. Then its JSJ tree of cylinders T_c has vertices and associated vertex groups in the JSJ graph of cylinders Σ_c as follows:*

1. *Type 1 vertex:*
   - *(a) For any cut pair $\{a, b\}$ such that $\Gamma \setminus \{a, b\}$ has $k \geq 3$ connected components, none of which consists of only one single vertex, there is a valence k vertex in T_c. The associated vertex group in Σ_c is the subgroup of W_Γ generated by $\{a, b\}$, unless a and b have a common adjacent vertex c, then it is generated by $\{a, b, c\}$.*
   - *(b) For any cut pair $\{a, b\}$ such that $\Gamma \setminus \{a, b\}$ has $k \geq 3$ connected components, one of which consists of only one vertex c, there is a valence $2 \cdot (k - 1)$ vertex in T_c. The associated vertex group is the subgroup of W_Γ generated by $\{a, b, c\}$.*
   - *(c) For any set $A \subseteq V(\Gamma)$ satisfying the properties (A1), (A2) and (A3) and which generates a two-ended subgroup not occurring in 1.(a) nor in 1.(b), there is a valence 2 vertex in T_c, where the properties (A1), (A2) and (A3) are the following:*
     - *(A1) Elements of A pairwise separate the geometric realization $|\Gamma|$.*
     - *(A2) If any subgraph Γ' of Γ that is a subdivided K_4 contains more than 2 vertices of A, all vertices of A lie on the same branch of the graph Γ'.*
     - *(A3) The set A is maximal among all sets satisfying (A1) and (A2).*

     If either $|A| = 2$ and there is no third vertex c adjacent to both elements in A or $|A| = 3$, the associated vertex group in Σ_c is the subgroup of W_Γ generated by A. If $|A| = 2$ and the two elements in A have a common adjacent vertex c, then the associated vertex group in Σ_c is the subgroup of W_Γ generated by $A \cup \{c\}$.
   - *(d) On any edge between a type 2 and a type 3 vertex, there is a valence 2 vertex added in T_c. The associated vertex group in Σ_c is the intersection of the vertex groups of its neighbors.*
2. *Type 2 vertex: For any set $A \subseteq V(\Gamma)$ satisfying the properties (A1), (A2) and (A3) such that the subgroup generated by A is infinite but not two-ended, there is a vertex in T_c with associated vertex group W_A in Σ_c.*
3. *Type 3 vertex: For any set $B \subseteq EV(\Gamma)$ of essential vertices in Γ satisfying the properties (B1), (B2) and (B3), there is a vertex in T_c whose associated vertex group in Σ_c is the subgroup W_B generated by B, where the properties (B1), (B2) and (B3) are the following:*
   - *(B1) For any pair $C = \{c, d\} \subseteq EV(\Gamma)$ of essential vertices, $B \setminus C$ is contained in one single connected component of $\Gamma \setminus C$.*
   - *(B2) The set B is maximal among all sets satisfying (B1).*
   - *(B3) $|B| \geq 4$.*

Between a vertex v of type 1 and a vertex v' of type 2 or 3 in $V(T_c)$, there is an edge connecting them if and only if their corresponding vertex groups intersect in an infinite subgroup.

Convention. Whenever we illustrate a JSJ graph of cylinders Σ_c of a RACG, see Figure 3.1.1 for instance, for economy of notation we omit the brackets of the vertex and edge groups and just write down the collection of generating vertices. We mark cylinder vertices in green, hanging vertices in red and rigid vertices in blue.

Remark 3.2. Not only type 1.(a) or type 1.(b) vertices correspond to essential cut pairs, but all type 1 vertices in Theorem 3.1 do.

Indeed, any set $A \subseteq V(\Gamma)$ satisfying properties (A1), (A2) and (A3) must contain an essential cut pair as shown in Lemma 3.26. But for a vertex of type 1.(c), we need that W_A is two-ended. By Theorem 3.18, we see that the only two options for a special subgroup of Γ satisfying Standing Assumption 1 to be two-ended is that it is generated either by two non-adjacent vertices of Γ or by two vertices connected via one common adjacent vertex in Γ. So, either $|A| = 2$, then it is precisely an essential cut pair. Or, $|A| = 3$, thus, it contains an essential cut pair and one common adjacent vertex in-between.

By [DT17, Lemma 3.30], the intersection of a set A satisfying properties (A1), (A2) and (A3) and a set B satisfying properties (B1), (B2) and (B3) contains at most two vertices. Thus, A and B can intersect at most in an essential cut pair. But in case their associated vertex groups intersect non-trivially, this intersection cannot be finite, implying that it must contain precisely the essential cut pair. The vertex of type 1.(d) can therefore be detected from an essential cut pair as well.

However, not all essential cut pairs contribute to a type 1 vertex, as illustrated in Example 3.3. The question on how to distinguish the ones contributing from the ones that do not is addressed in Section 3.1 in Proposition 3.9.

Example 3.3. In Figure 3.1.1, we see on the left side a square-free graph Γ satisfying the Standing Assumption 1. On the right side, the JSJ graph of cylinders Σ_c of W_Γ is illustrated. It is obtained by Theorem 3.1 with the following considerations: There is no cut pair of type 1.(a) and the cut pairs $\{u, y\}$ and $\{v, y\}$ give vertices of type 1.(b). From the cut pairs $\{v, w\}$ and $\{w, x\}$ we obtain a vertex of type 1.(c) and $\{v, x\}$, $\{w, z\}$ and $\{y, z\}$ add vertices of type 1.(d). Of type 2, there are the five vertex sets $\{v, n_1, n_2, x\}$, $\{w, r_1, r_2, z\}$, $\{y, s_1, s_2, z\}$, $\{u, p_1, p_2, y\}$ and $\{v, l_1, l_2, u, y\}$. The only vertex of type 3 is given by $\{v, w, x, y, z\}$. Note that for instance the set $\{v, x, w\}$ does not give a vertex of type 2 as property (A2) is violated by the subdivided K_4 with corners w, v, x and z. Furthermore, while the set $\{v, l_1, l_2, u\}$ is a pairwise separating branch, it does not satisfy (A3). Thus, even though $\{u, v\}$ is an essential cut pair and thus gives a two-ended subgroup over which W_Γ splits, it does not give a type 1 vertex. As proved in Proposition 3.9, this is due to the fact that there are other cut pairs, for instance $\{y, l_1\}$, separating u and v (see also Example 3.7). This implies that $W_{\{u,v\}}$ is not universally elliptic and therefore contained within a hanging subgroup.

Also, the type of a vertex determines a key property:

Theorem 3.4. [cf. Bow98, Theorem 5.28] *Let W_Γ be a hyperbolic RACG with Γ satisfying the Standing Assumption 1. Let Σ_c be the JSJ graph of cylinders given by Theorem 3.1, then:*

- *The vertex group associated to a type 1 vertex is two-ended.*
- *The vertex group associated to a type 2 vertex is hanging.*
- *The vertex group associated to a type 3 vertex is rigid.*

Comparing this Theorem 3.4 and Outline 2.16, we can now establish the following correspondence between the JSJ tree of cylinders given by Construction 2.11 and the JSJ tree of cylinders constructed in Theorem 3.1 for hyperbolic RACGs:

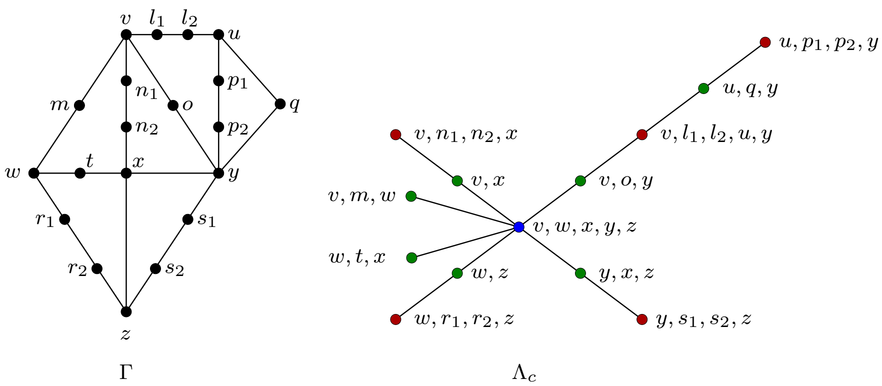

Figure 3.1.1

- *Type 1 vertices correspond to cylinder vertices:* Type 1 vertices in Σ_c lift to vertices of finite valence in T_c with two-ended vertex stabilizer. These properties can only hold for cylinder vertices. Furthermore, by existence of vertices of type 1.(d), the JSJ tree of cylinders constructed in Theorem 3.1 is bipartite with $V = V(1) \sqcup V(2,3)$, where $V(1)$ are all the vertices of type 1 and $V(2,3)$ contains vertices of type 2 and 3. Thus, as the JSJ tree of cylinders is also bipartite, no other than the type 1 vertices can correspond to cylinder vertices.
- *Type 2 vertices correspond to hanging vertices:* By Theorem 3.4, type 2 vertices are hanging, thus they are the hanging non-cylinder vertices.
- *Type 3 vertices correspond to rigid vertices:* Again, by Theorem 3.4, type 3 vertices are rigid, thus they are the rigid non-cylinder vertices.

3.1.2 Non-hyperbolic case

Since in the non-hyperbolic case, there is no universal construction of a JSJ tree and its tree of cylinders like the one given by Bowditch, for arbitrary RACGs we need to start from scratch: We first determine how to find a JSJ decomposition in terms of the defining graph Γ and then produce a construction of the JSJ graph of cylinders from there. In fact, any decomposition of a (right-angled) Coxeter group over two-ended subgroups is visible in the defining graph Γ:

Theorem 3.5. [MT09, Theorem 1] *For a simplicial graph Γ which is triangle-free and which has no separating vertices or edges (that is, Γ satisfies Standing Assumptions 1, (1) and (2)), W_Γ splits over a two-ended subgroup H if and only if Γ has a cut collection $\{a - b\}$.*

Moreover, given some decomposition Σ of W_Γ with two-ended edge groups, there is a visual decomposition Ψ of W_Γ such that:

- *All occurring subgroups in Ψ are special.*
- *Each vertex group of Ψ is a subgroup of a conjugate of a vertex group of Σ.*
- *Each edge group of Ψ is a subgroup of a conjugate of an edge group of Σ.*
- *In particular, for each two-ended edge group H of Σ, there is a unique cut collection $\{a - b\}$ such that some conjugate of H contains $W_{\{a-b\}}$.*

Thus, in order to produce a splitting over two-ended special subgroups, by Theorem 3.5, we need to collect all cut collections of Γ. Then, by Remark 2.3, we are left to eliminate the cut collections that produce a subgroup that is not universally elliptic and thus belongs inside a hanging subgroup. To be able to do that, we need the following terminology:

Definition 3.6. A cut pair $\{a, b\} \in V(\Gamma)$ is said to be *crossed* by another, disjoint cut pair $\{c, d\} \in V(\Gamma) \setminus \{a, b\}$ if a and b lie in different connected components of $\Gamma \setminus \{c, d\}$. We say $\{c, d\}$ is *crossing* $\{a, b\}$. If there is no cut pair crossing $\{a, b\}$, then $\{a, b\}$ is *uncrossed*.

A cut triple $\{a, b, c\}$, where c is the common adjacent vertex of the non-adjacent vertices a and b, is said to be *crossed* by another cut triple $\{d, e, f\}$, where f is the common adjacent vertex of the non-adjacent vertices d and e, if c is equal to f and a and b lie in different connected components of $\Gamma \setminus \{d, e, c\}$. We say $\{d, e, f\}$ is *crossing* $\{a, b, c\}$. If there is no cut triple crossing $\{a, b, c\}$, then $\{a, b, c\}$ is *uncrossed*.

Example 3.7. In Figure 3.1.1 of Example 3.3, while for instance $\{w, z\}$ is an uncrossed cut pair, $\{u, v\}$ is not as it is crossed by $\{l_1, y\}$, for example. In the right graph of Figure 1.2.2 considered in Example 1.21, the cut triple $\{a, b, c\}$ is crossed by the cut triple $\{d, e, c\}$.

Remark 3.8. Any uncrossed cut pair is essential, but not every essential cut pair is uncrossed, see Example 3.7. Moreover, it is not necessary to define a notion of a crossing between a cut pair and a cut triple, because it is obvious that this situation cannot happen.

It turns out that all the two-ended edge groups of a JSJ splitting are detected by the uncrossed cut collections of Γ:

Proposition 3.9. *If Γ is a graph which satisfies Standing Assumption 1 and which has at least one uncrossed cut collection, then:*

(a) For every special subgroup $W_{\{a-b\}}$ generated by an uncrossed cut collection $\{a - b\}$ of Γ, there is a JSJ splitting Σ such that $W_{\{a-b\}}$ is contained in a special, two-ended edge group of Σ.
(b) Given a two-ended edge group of a JSJ splitting Σ of W_Γ that is special and contains $W_{\{a-b\}}$, where $\{a - b\}$ is a cut collection, then $\{a - b\}$ is uncrossed.

Proof. For (a), let $\{a - b\}$ be a cut collection of Γ. Let Σ_1 be a splitting of W_Γ over a two-ended subgroup containing $W_{\{a-b\}}$, which exists by Theorem 3.5. Suppose that Σ_1 is not a JSJ splitting. If the Bass-Serre tree T_1 of Σ_1 is universally elliptic, but not dominating every other universally elliptic tree, then, by [GL11, Lemma 2.15], T_1 can be refined to a JSJ tree T_1' with a two-ended edge stabilizer containing $W_{\{a-b\}}$ and the claim follows. If, on the other hand, the subgroup containing $W_{\{a-b\}}$ is not universally elliptic, there must be another splitting Σ_2 of W_Γ in whose Bass-Serre tree the group $W_{\{a-b\}}$ is not elliptic. Hence,the infinite-order element $ab \in W_{\{a-b\}}$ cannot fix a point in it. Now, we can refine Σ_2 by Theorem 3.5 to a visual splitting Ψ of W_Γ. Because Ψ is visual, we know that we can find a unique cut collection $\{c - d\}$ in Γ such that $W_{\{c-d\}}$ is contained in a two-ended edge group of Ψ. Also, every edge group of Ψ is a subgroup of a conjugate of an edge group of Σ_2 and every vertex group of Ψ is a subgroup of a conjugate of a vertex group of Σ_2. Thus, the element ab does not fix a point in the Bass-Serre tree of Ψ either, implying that the elements a and b must be in different vertex groups of Ψ. This implies that the vertices a and b must be separated in Γ by the cut collection $\{c - d\}$, thus the cut collection $\{a - b\}$ is not uncrossed.

Assume conversely for (b) that we have a cut collection $\{a - b\}$ crossed by another cut collection $\{c - d\}$, respectively, with splittings Σ_1 and Σ_2 over two-ended subgroups containing $W_{\{a-b\}}$ and $W_{\{c-d\}}$, respectively. Since the cut collection $\{a - b\}$ is crossed by $\{c - d\}$, the elements a and b are in different vertex groups of the splitting Σ_2. Thus, the infinite order element $ab \in W_{\{a-b\}}$ cannot fix a point in the Bass-Serre tree of Σ_2, implying that the edge group of Σ_1 containing $W_{\{a-b\}}$ is not universally elliptic and therefore Σ_1 is no JSJ decomposition. □

Remark 3.10. In Proposition 3.9 the assumption that Γ must contain an uncrossed cut collection excludes the case, where Γ is a square. This is important, because for Γ a square, the corresponding RACG $W_\Gamma = D_\infty \times D_\infty$ is virtually $\mathbb{Z}^2$. Such W_Γ is commensurable to the fundamental group of a surface, in this case a torus, which have to be treated separately, cf. [Pap05]. However, this is the only case we need to rule out additionally, since by the Standing Assumption 1, W_Γ is not cocompact Fuchsian and thus never commensurable to a surface group of higher genus.

Also, excluding the case that Γ is not a square is no obstacle for the QI-classification, since the property of being virtually $\mathbb{Z}^2$ determines the QI-type of the group. Thus, we refine the standing assumption by modifying (4) of Standing Assumption 1 to exclude squares:

Standing Assumption 1.2. Γ is the defining graph of a RACG W_Γ which satisfies:

(1) Γ is triangle-free.
(2) Γ is connected and has neither a separating vertex nor a separating edge.
(3) Γ has a cut collection.
(4) Γ is not a cycle.

Now, starting from a visual JSJ decomposition over all uncrossed cut collections, we can determine how to produce the different vertices and the edges of the JSJ graph of cylinders.

3.1.2.1 Cylinder vertices

Given a JSJ decomposition Σ, by Outline 2.16, we know that we can find all cylinder vertices of the JSJ graph of groups Σ_c and their vertex groups G_Y by running through all edge groups visible in Σ.

Thus, in light of Proposition 3.9, we pick up all uncrossed cut collections in Γ and compute the commensurators of the special subgroups they generate. It turns out that the commensurator of a special subgroup is also visible from the defining graph Γ:

Theorem 3.11. [Par97, Theorem 2.1] *Let W be a RACG on the defining graph Γ with finite generating set S and let $T \subseteq S$ be a subset of S. Consider the maximal decomposition $W_T = W_{T_1} \times \cdots \times W_{T_n}$ of W_T as a direct product of subgroups, where $W_{T_1}, \ldots, W_{T_r}$ are finite for some $r \in \{1, \ldots, n\}$ and all the other subgroups are infinite. Then the commensurator of W_T in W is given by*

$$Comm_W(W_T) = W_{T^\infty} \times W_{Y^\infty}$$

$$\text{with} \quad T^\infty = T_{r+1} \cup \cdots \cup T_n \quad \text{and} \quad Y^\infty = \{s \in S \mid e = (t,s) \in E(\Gamma) \text{ for all } t \in T^\infty\}.$$

Convention. To simplify terminology, we refer to the vertices of the defining graph of the commensurator of the special subgroup given by a cut collection as *commensurator of the cut collection.*

Remark 3.12. We encounter the following situations:

- For a cut pair $\{a, b\}$ in Γ, Theorem 3.11 implies that the commensurator is generated by $\{a, b\} \cup \mathcal{C}$, where $\mathcal{C}$ contains all the common adjacent vertices of a and b. That is,

$$Comm_W(W_{\{a,b\}}) = W_{\{a,b\}} \times W_{\mathcal{C}}\,.$$

- For a cut triple $\{a, b, c\}$, where a and b are non-adjacent and c is a common adjacent vertex, the special subgroup $W_{\{a,b,c\}}$ decomposes as $W_{\{a,b\}} \times W_{\{c\}}$. Thus, also in this case, the commensurator is generated by $\{a, b\} \cup \mathcal{C}$, where $\mathcal{C}$ contains all the common adjacent vertices of a and b, in particular, $\mathcal{C}$ contains c.

- In case there are two overlapping cut triples $\{a, b, c\}$ and $\{a, b, c'\}$ sharing the same non-adjacent vertex pair $\{a, b\}$, for both cut triples we obtain the same commensurator. Hence, their corresponding edges in a JSJ decomposition are equivalent under commensurability, thus they lie in the same cylinder. Therefore, such two cut triples only give one cylinder.
- It is immediate from Theorem 3.11 that a cut collection $\{a - b\}$ in a hyperbolic RACG always has a two-ended commensurator. This is because a and b can have at most one common adjacent vertex c, as otherwise two common adjacent vertices and a and b give a square in contradiction to hyperbolicity. But both $W_{\{a,b\}} \simeq D_\infty$ and $W_{\{a,b,c\}} \simeq D_\infty \times \mathbb{Z}_2$ are two-ended (cf. Theorem 3.18), thus hyperbolic RACGs have two-ended cylinder vertices.

Example 3.13. In the non-hyperbolic defining graph Γ in Figure 3.1.2 the orange cut pair $\{v, x\}$ has three purple common adjacent vertices $\mathcal{C} = \{w, m, y\}$, thus

$$Comm_{W_\Gamma}(W_{\{v,x\}}) = W_{\{v,x\}} \times W_{\{w,m,y\}} = W_{\{v,w,m,y,x\}}\,.$$

The other two cut pairs $\{w, z\}$ and $\{y, z\}$ correspond to special subgroups with commensurators $W_{\{w,z,n,x\}}$ and $W_{\{y,z,o,x\}}$, respectively.

The commensurator of the special subgroup corresponding to the cut triple $\{w, x, y\}$, on the other hand, is $W_{\{v,w,x,y\}}$, since v and x are the common adjacent vertices of w and y. This is the only cut triple in Γ: Recall that for instance the vertices $\{w, m, y\}$ are not a cut triple, despite separating v from the the rest of Γ, because $W_{\{w,m,y\}}$ is not two-ended.

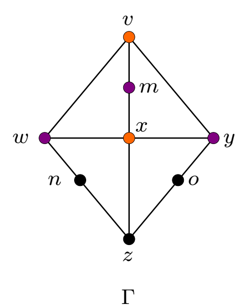

Figure 3.1.2

For the sake of completeness, we summarize this insight as a Proposition:

Proposition 3.14. *Let $\mathcal{S}$ be the following set: For every uncrossed cut collection $\{a - b\}$ of Γ, the set $\{a, b\} \cup \mathcal{C}$, where $\mathcal{C}$ is the set of common adjacent vertices of a and b, is an element in $\mathcal{S}$. Then every set S in $\mathcal{S}$ corresponds to a cylinder vertex in the JSJ graph of cylinders of W_Γ with vertex group the special subgroup generated by S.*

Lemma 3.15. *Every cylinder vertex group of the JSJ graph of cylinders of a RACG W_Γ, where Γ satisfies the Standing Assumption 1.2, in particular is triangle-free, is either*

- *virtually cyclic;*
- *virtually $\mathbb{Z}^2$; or*
- *the direct product of a virtually non-abelian free group and the infinite dihedral group D_∞.*

For the proof, we need the following characterizations:

Theorem 3.16. [Dav08, Theorem 17.2.1] *A RACG W_Γ is virtually abelian if and only if it decomposes as the direct product of finitely many infinite dihedral groups D_∞ and a finite RACG.*

Theorem 3.17. [MT09, Theorem 8.34] *A RACG W_Γ is virtually free if and only if no induced subgraph is a circuit of more than three vertices.*

We can detect the intersection of the above two classes of groups by:

Theorem 3.18. *[Dav08, Theorem 8.7.3] A RACG W_Γ is two-ended if and only if it is the direct product of one infinite dihedral group D_∞ and a finite RACG. In terms of the defining graph this means that Γ is a two-point suspension of a complete graph.*

Proof of Lemma 3.15. Consider a cut collection $\{a-b\}$ of Γ with a and b non-adjacent, then $\Gamma \setminus \{a-b\}$ has $i \geq 2$ connected components $\Gamma_1, \ldots, \Gamma_i$. We distinguish the contribution of some component Γ_j for $j \in \{1, \ldots, i\}$ to the commensurator G_Y of $W_{\{a-b\}}$:

- If Γ_j does not contain any common adjacent vertex of $\{a, b\}$, no vertex contributes to G_Y.
- If Γ_j contains one common adjacent vertex c of $\{a, b\}$, the contribution to G_Y is a direct product with $\mathbb{Z}_2$.
- If Γ_j contains at least two common adjacent vertices c_1 and c_2 of $\{a, b\}$, then they must be connected by a path not passing through a or b. Otherwise, they would not lie in the same connected component of $\Gamma \setminus \{a-b\}$. However, there cannot be an edge between c_1 and c_2, as otherwise $\{a, c_1, c_2\}$ would form a triangle. Thus, there is no relation between c_1 and c_2 in G_Y.

Moreover, if $\{a-b\}$ is a cut triple, the third vertex c of the triple contributes a direct product with $\mathbb{Z}_2$ to G_Y. In conclusion, the commensurator G_Y of the cut collection $\{a-b\}$ is the RACG given by a defining graph Γ_Y consisting of a and b with k common adjacent vertices $\{c_1, \ldots, c_k\} =: \mathcal{C}$, which are all only connected to a and b in G_Y, see Figure 3.1.3. Thus, we can consider the following cases:

- $\mathcal{C} = \emptyset$: $G_Y = W_{\{a,b\}} \simeq D_\infty$, thus virtually cyclic.
- $k = 1$: $G_Y = W_{\{a,b,c_1\}} \simeq D_\infty \times \mathbb{Z}_2$, thus virtually cyclic.
- $k = 2$: $G_Y = W_{\{a,b,c_1,c_2\}} \simeq D_\infty \times D_\infty$, thus virtually abelian, in particular virtually $\mathbb{Z}^2$.
- $k > 2$: $G_Y = W_{\{a,b,c_1,\ldots,c_k\}} \simeq D_\infty \times F$, where F is virtually a non-abelian free group.

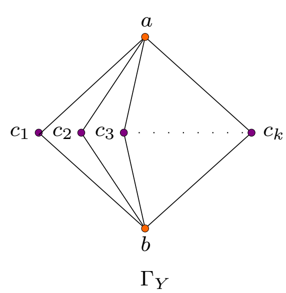

Figure 3.1.3

□

Example 3.19. The commensurator of $W_{\{v,x\}}$ in Figure 3.1.2 is the direct product of the virtually non-abelian free group $W_{\{w,m,y\}}$ and the infinite dihedral group $W_{\{v,x\}}$.

Convention. From now on we refer to the two "new" types of cylinder vertex groups and their corresponding cylinder vertex as

- *VA*, if the cylinder vertex group is virtually $\mathbb{Z}^2$.
- *VFD*, if the cylinder vertex group is the direct product of a virtually non-abelian free group and an infinite dihedral group.

3.1.2.2 Non-cylinder vertices The fact that a certain collection of vertices gives a hanging or rigid vertex group in a graph of groups with respect to incident two-ended edge groups is intrinsic to this collection in the sense that it is independent of the existence of squares in the defining graph Γ. Furthermore, by Outline 2.16, if we see a hanging or rigid vertex in the JSJ decomposition, it transfers over to the JSJ graph of cylinders. So, the results of [DT17] in Theorem 3.1 translate to the general setting:

Proposition 3.20. *Let $A \subseteq V(\Gamma)$ be a set of vertices such that the A-induced subgraph Γ_A is not a complete graph and A satisfies either the conditions:*

(A1) Elements of A pairwise separate the geometric realization $|\Gamma|$.
(A2) If any subgraph Γ' of Γ that is a subdivided K_4 contains more than 2 vertices of A, all vertices of A lie on the same branch of Γ'.
(A3) The set A is maximal among all sets satisfying (A1) and (A2).

Or A satisfies the condition:

(A) The set A is a maximal collection of pairwise crossing cut triples.*

If A is not contained in a vertex set corresponding to a cylinder vertex, then A corresponds to a vertex in the JSJ graph of cylinders Σ_c with hanging vertex group W_A.

Sketch of the Proof. Recall that the assumption that Γ_A is not a complete graph ensures that W_A is infinite. Now, we want to give some motivation on how the graph theoretical conditions (A1), (A2) and (A3) and the graph theoretical condition (A*) produce a hanging subgroup W_A. Intuitively, the object to have in mind as the hanging subgroup is a surface with boundary.

Let us first consider crossing cut pairs. By the proof of Proposition 3.9, they are not universally elliptic and thus belong in a hanging subgroup. They give crossing curves corresponding to different interfering splittings, which are thus not part of a JSJ decomposition. In particular, a collection of pairwise separating vertices as forced by condition (A1) contains all pairwise crossing cut pairs within a branch and at least one uncrossed essential cut pair (cf. Remark 3.27). Such an uncrossed cut pair then generates precisely a universally elliptic subgroup as the boundary component. If we see however a subdivided K_4 in Γ, we could choose three or all four corner vertices as a collection of pairwise $|\Gamma|$-separating vertices. But then this collection cannot contain any non-essential vertex contained in a branch. So, there are no crossing cut pairs in the collection producing crossing curves. Thus, the resulting group is not a hanging, but rather a candidate for a rigid vertex group. Therefore, we need to exclude such a collection by condition (A2). Maximality needs to be ensured by condition (A3), since a JSJ decomposition is maximal (cf. Definition 2.4 and Lemma 2.7).

Consider now crossing cut triples contained in a collection A satisfying condition (A*). Again, by the proof of Proposition 3.9, they are not universally elliptic and thus belong in a hanging subgroup. By the definition of a JSJ decomposition, we again need maximality.

Since all the cut triples in A cross pairwise, all share their "middle" common adjacent vertex, call it c. Thus, the subgraph induced on the collection A is a graph theoretical star based at c. Since by Standing Assumption 1.2, Γ has no triangles and no separating edge, for every pair $\{x, y\} \in A \setminus \{c\}$ of leaves, x and y are not adjacent and there are at least three disjoint paths connecting x and y: One is the segment $\{x, y, c\}$ and two paths do not contain c, call them p_1 and p_2.

We claim that either $\{x, y\}$ is an uncrossed cut pair or $\{x, y, c\}$ is a cut triple: If every path connecting the interior of p_1 (or analogously p_2) with c passes through x or y, removing $\{x, y\}$

separates the interior of p_1 from $\Gamma_A \setminus \{x, y\}$. Hence, $\{x, y\}$ is a cut pair. In fact, $\{x, y\}$ is an uncrossed cut pair, because no matter which other cut pair is removed, x and y will stay connected with each other via either p_1, p_2 or the segment $\{x, y, c\}$. Thus, the cut pair $\{x, y\}$ generates a universally elliptic subgroup representing the boundary component of the surface.

If both p_1 and p_2 contain an interior vertex that is connected to c via a path not passing through x or y, then we need to show that removing $\{x, y, x\}$ separates Γ. Since x is leaf in $A \setminus \{c\}$, there exist $x' \in A \setminus \{c, x\}$ such that $\{x, x', c\}$ is a cut triple separating Γ into two connected components C and C' of $\Gamma \setminus \{x, x', c\}$. Then either the interior of p_1 or p_2 must be contained in one of the connected components of $\Gamma \setminus \{x, x', c\}$. Without loss of generality, assume that the interior of p_1 is contained in C. Thus, there is a vertex l_1 in $V(C) \cap V(p_1)$ which is not connected in $\Gamma \setminus \{x, x', c\}$ to some $l_2 \in V(C')$. However, since l_1 is in the interior of p_1, l_1 will also not be connected to l_2 in $\Gamma \setminus \{x, y, c\}$. Thus, also $\{x, y, c\}$ is a cut triple contained in A. If it is uncrossed, it represents a boundary component. □

Corollary 3.21. *Let $A \subseteq V(\Gamma)$ be a set of vertices satisfying the condition (A*). Then the A-induced subgraph Γ_A of Γ is a star.*

Proposition 3.22. *For any set $B \subseteq EV(\Gamma)$ of essential vertices in Γ satisfying the properties (B1), (B2) and (B3), there is a vertex in the JSJ graph of cylinders Σ_c with rigid vertex group W_B, where the properties (B1), (B2) and (B3) are the following:*

- *(B1) For any set C that is a pair $\{c, d\} \subseteq EV(\Gamma)$ of essential vertices or a path $\{c, d, e\} \subseteq EV(\Gamma)$ of length 2 of essential vertices, $B \setminus C$ is contained in one single connected component of $\Gamma \setminus C$.*
- *(B2) The set B is maximal among all sets satisfying (B1).*
- *(B3) $|B| \geq 4$.*

Sketch of the Proof. Again, we want to give some motivation on how the graph theoretical conditions (B1), (B2) and (B3) produce a rigid subgroup W_B. The key feature of a rigid vertex group is that it cannot be split any further. This is precisely captured by condition (B1): We consider the collection of cut pairs and cut triples which are pairwise not separating the collection. We want a maximal such collection and thus impose condition (B2). Suppose now we find a collection $B = \{x, y, z\}$ with only three essential vertices satisfying conditions (B1) and (B2). Then, since we restrict to special subgroups, the virtually free RACG W_B can have the adjacent edge groups $W_{\{x,y\}}$, $W_{\{y,z\}}$ and $W_{\{x,z\}}$. However, such a group is then virtually a surface with boundary and thus not considered to be rigid. This case is excluded by condition (B3). □

3.1.2.3 Edges It remains to be determined which vertices in the JSJ graph of cylinders are connected by an edge:

Lemma 3.23. *For any pair of vertices in the JSJ graph of cylinders Σ_c, there is an edge connecting them if and only if the pair consists of one cylinder vertex corresponding to the cut collection $\{a - b\}$ and one non-cylinder vertex and their vertex groups intersect in a special subgroup containing $W_{\{a-b\}}$. The edge group is the special subgroup generated by the intersection of the corresponding vertex sets.*

Proof. Since the JSJ graph of cylinders Σ_c is bipartite, edges can only connect cylinder with non-cylinder vertices. Suppose there is an edge connecting a cylinder vertex corresponding to a cut collection $\{a - b\}$ and a non-cylinder vertex. By definition of the fundamental group of a graph of groups, the edge group is the intersection of their vertex groups. If this intersection would be a finite group, the group W cannot be one-ended by Stallings' Theorem [Sta71], in contradiction to the Standing Assumption 1.2. Thus, the intersection is infinite. Furthermore, the vertex groups are special subgroups, thus so is their intersection by [Dav08, Theorem 4.1.6]. Since the structure of Σ_c comes from a JSJ decomposition Σ, by Proposition 3.9, the edge group in Σ must contain $W_{\{a-b\}}$, thus so does the edge group in Σ_c.

Assume conversely that the vertex group of a cylinder vertex corresponding to a cut collection $\{a-b\}$ and a non-cylinder vertex intersect in a special subgroup containing $W_{\{a-b\}}$. Then they are connected by an edge by the definition of the fundamental group of a graph of groups. □

Example 3.24. For the graph Γ shown in Figure 1.2.3, which satisfies Standing Assumption 1.2, we can construct the corresponding JSJ graph of cylinders in Figure 3.1.4 by reading off the following collections of vertices according to Proposition 3.14, Theorem 3.20, Theorem 3.22 and connect them with edges according to Lemma 3.23:

uncrossed cut collection	commensurator	hanging	rigid
x, w	x, w, k, d		
v, w	v, w, d	v, w, l_1, l_2	
v, y	v, y, m, d		
$x, y, b \mid x, y, c \mid x, y, d$	x, y, a, b, c, d	v, w, x, y, d	
c, d	c, d, x, y	c, d, n_1, n_2	c, x, d, y
b, c	b, c, x, o, y		b, x, c, y
a, b	a, b, x, y	a, b, p_1, p_2	a, x, b, y

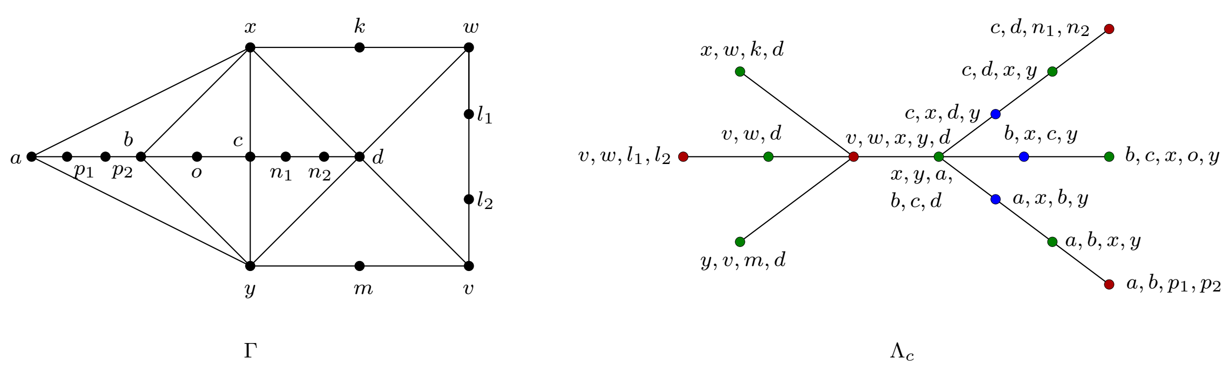

Figure 3.1.4

3.1.2.3.1 Two-ended edge groups As indicated in Remark 2.14, we aim to restrict to trees of cylinders that are $\mathcal{VC}$-trees themselves. This is not always the case, as we can see in Example 3.24, Figure 3.1.4: The vertex set $\{x, y, a, b, c, d\}$ generating the commensurator of the uncrossed cut triples $\{x, y, b\}$, $\{x, y, c\}$ and $\{x, y, d\}$ contains for instance the collection $\{c, x, d, y\}$, which corresponds to an adjacent rigid vertex. Thus, the connecting edge group generated by the intersection by Lemma 3.23 is $W_{\{c,x,d,y\}} = D_\infty \times D_\infty$, which is not two-ended.

Therefore, we need to impose assumptions on the defining graph Γ to ensure that the intersection of vertex groups is two-ended. Recall that by the bipartiteness of the JSJ tree of cylinders, the only intersections we need to consider are between cylinder and non-cylinder vertices.

Lemma 3.25. *If the intersection of a cylinder vertex group G_Y corresponding to a cut collection $\{a - b\}$ and a hanging vertex group W_A contains $W_{\{a-b\}}$ and*

(a) the set A satisfies the conditions (A1), (A2) and (A3), then the intersection is the infinite dihedral group D_∞.
(b) the set A satisfies the condition (A), then the intersection is two-ended.*

The proof of Lemma 3.25 relies on the following properties:

Lemma 3.26. *If $A \subseteq V(\Gamma)$ is a set that satisfies conditions (A1), (A2) and (A3), the A-induced subgraph Γ_A is not a complete graph and W_A is not contained in a cylinder vertex group, then:*

(1) A does not contain a cut triple.
(2) A does not contain two branches which share a common endpoint.

Proof. By definition, the non-adjacent vertices a and b of a cut triple $\{a - b\}$ are not a cut pair and by Standing Assumption 1.2, a and b do not share an edge. Thus, a and b do not separate $|\Gamma|$. Therefore, a set satisfying (A1) cannot contain a cut triple, implying (1).

For (2), suppose that two vertices x and y of degree 2 lie in different branches contained in A meeting at an essential vertex a. Let b_x and b_y be the other endpoint of these branches, respectively. Since a, b_x and b_y are essential, a is connected to both b_x and b_y via a path neither passing through x nor y. Thus, $|\Gamma| \setminus \{x, y\}$ is connected, in contradiction to condition (A1). □

Proof of Lemma 3.25. For the proof of (a), assume that A satisfies the conditions (A1), (A2) and (A3). Let W_A be the corresponding hanging vertex group on the defining graph Γ_A intersecting the cylinder vertex group G_Y corresponding to the cut collection $\{a - b\}$ non-trivially. Recall that by Lemma 3.23 the intersection must contain $W_{\{a-b\}}$. Then, by Lemma 3.26.1., $\{a - b\}$ cannot be a cut triple, so G_Y must correspond to a cut pair $\{a, b\}$.

If G_Y is $W_{\{a,b\}} = D_\infty$, so is the intersection. Therefore, we can assume that G_Y is not D_∞. Thus, the cylinder vertex group is the special subgroup on the defining graph Γ_Y consisting of the pair $\{a, b\}$ with a non-empty common adjacent vertex set $\{c_1, c_2, \dots, c_k\}$ for $k \geq 1$ and the degree of every vertex in $\mathcal{C}$ in Γ_Y is 2. Since by [Dav08, Theorem 4.1.6] the intersection of two special subgroups is the special subgroup defined on the induced graph given by the intersection, we need to consider how $\Gamma_A \cap \Gamma_Y$ can look like. Recall that the intersection $I = V(\Gamma_A \cap \Gamma_Y)$ contains a and b. We distinguish the following cases:

1. $I = \{a, b\}$: The corresponding group $W_{\{a,b\}} \cong D_\infty$ is two-ended.
2. $c_i \in I$ and c_i has degree 2 in Γ for $i \in \{1, \dots, k\}$: Then I contains the whole branch $\{a, c_i, b\}$. No other branch in A can be attached at a or b by Lemma 3.26.2., implying $I = \{a, c_i, b\}$. But A cannot be equal to $I = \{a, c_i, b\}$, since the hanging vertex corresponding to A is not two-ended. However, supposing that there is another vertex $v \in A \setminus \{a, c_i, b\}$ such that $|\Gamma| \setminus \{c_i, v\}$ is separated, contradicts the fact that a and b are uncrossed and c_i has degree 2. This implies that this case cannot occur.
3. $c_i \in I$ for $i \in \{1, \dots, k\}$ and c_i essential in Γ: We argue as in case 2. that there must exist a $v \in A \setminus \{a, c_i, b\}$ such that $|\Gamma| \setminus \{c_i, v\}$ is separated. Since a and b are an uncrossed cut pair, there is a path between v and c_i not passing through a nor b and another path connecting a and b not passing through c_i nor v. This means that we have a subdivided K_4 with corners a, c_i, b and v, in contradiction to (A2). So again, this case cannot occur.

To conclude, in case (a) the special subgroup W_I generated by the intersection I is always D_∞.

Assume now for the proof of (b) that A satisfies the condition (A*), and that the corresponding vertex group W_A on the defining graph Γ_A is infinite. By Corollary 3.21, Γ_A is a graph theoretical star based at the vertex c, where all the cut triples contained in A meet. Suppose W_A intersects the cylinder vertex group G_Y corresponding to a cut collection $\{a-b\}$ in a subgroup containing $W_{\{a-b\}}$. Thus, if G_Y is two-ended, so is the intersection and we can assume that G_Y is not two-ended. That means that the cylinder vertex group is the special subgroup on the defining graph Γ_Y consisting of the two non-adjacent vertices $\{a, b\}$ of the cut collection and their common adjacent vertices $\mathcal{C} = \{c_1, \dots, c_k\}$ with $k \geq 2$, which all have degree 2 in Γ_Y. As above, we need to consider the graph $\Gamma_A \cap \Gamma_Y$. Define $I = V(\Gamma_A) \cap V(\Gamma_Y)$ which contains $\{a-b\}$ by Lemma 3.23 and consider the following cases for I:

1. $I = \{a, b\}$: In this case W_I is D_∞ thus two-ended.
2. $I = \{a, b, c_i\}$ for some $i \in \{1, \dots, k\}$: In this case W_I is $D_\infty \times \mathbb{Z}_2$ and thus two-ended.
3. $\{a, b\} \in I$ and $|I \cap \mathcal{C}| \geq 2$: Then I contains a square, thus so does Γ_A in contradiction to the fact that Γ_A is a triangle-free star. Thus, this case cannot occur.

In conclusion, also in case (b) the special subgroup W_I generated by the intersection I is always two-ended. This finishes the proof. □

Remark 3.27. By the Standing Assumption 1.2, the defining graph Γ is never a cycle. Thus, in case A satisfies the conditions (A1), (A2) and (A3) by Lemma 3.26.2, A cannot contain a cycle. This is also true in case (b), where A satisfies *(A*)*, since the A-induced subgraph Γ_A is a graph theoretical star by Corollary 3.21. Therefore, by Theorem 3.17 any hanging vertex group is virtually free.

Theorem 3.28. *Let G_Y be the vertex group of the cylinder vertex v_Y in Σ_c corresponding to the cut collection $\{a-b\}$ with defining graph $\Gamma_Y \subseteq \Gamma$ on the vertex set $V(\Gamma_Y) = \{a, b\} \cup \mathcal{C}$, where $\mathcal{C}$ is the set of common adjacent vertices of a and b. Then every rigid vertex group W_B adjacent to the cylinder vertex group G_Y intersects G_Y in a two-ended subgroup if and only if for any pair of vertices in $\mathcal{C}$ every path connecting them in Γ passes through a or b.*

Example 3.29. In Figure 3.1.4, the rigid vertex group generated by $\{c, d, x, y\}$ is adjacent to the cylinder vertex group corresponding to the cut triple $\{x-y\}$, which has $\{c, d\}$ as common adjacent vertices. Because there is a path through the vertices $\{c, n_1, n_2, d\}$ connecting c and d without passing through x nor y, they intersect in the non-two-ended edge group generated by $\{c, d, x, y\}$.

Proof. Suppose first that there is a pair $\{c_i, c_j\} \subseteq \mathcal{C}$ of distinct vertices that are connected by a path in Γ not passing through a nor b nor any other common adjacent vertex of a and b. There must be a path between a and b not passing through c_i nor c_j, as otherwise $\{a-b\}$ would be crossed by $\{c_i - c_j\}$. However, this implies that the vertex collection $\{a, b, c_i, c_j\}$ forming a square in Γ_Y satisfies condition (B1). While this set might not be maximal with respect to this condition, it is for sure contained in a maximal collection B satisfying conditions (B1), (B2) and (B3), corresponding to a rigid vertex group W_B. Thus, G_Y is adjacent to the rigid vertex group W_B which it intersects in a subgroup containing $W_{\{a,b,c_i,c_j\}} = D_\infty \times D_\infty$. Hence, the intersection is not two-ended.

Assume conversely that G_Y is adjacent to the rigid vertex group W_B and that no pair of vertices in $\mathcal{C}$ is connected by a path in Γ that is not passing through a nor b. Then each such pair is separated when a and b are removed. Thus, at most one of the $c \in \mathcal{C}$ can be contained in B, as otherwise condition (B1) would be violated. Since the intersection of G_Y and W_B must be infinite, we conclude that $\{a, b\} \subseteq B$. Thus, the intersection is either $W_{\{a,b\}}$ or $W_{\{a,b,c\}}$ and therefore two-ended. □

Remark 3.30. Combining Lemma 3.25 and Theorem 3.28 implies Theorem 1.46, stating that the JSJ tree of cylinders has two-ended edge stabilizers if and only if there is no uncrossed cut collection containing the corners of a square in the defining graph, where the other two corners are connected by a subdivided diagonal. Note that this can be interpreted as a condition about a subdivided K_4.

Remark 3.31. If Γ contains two overlapping cut triples $\{a,b,c\}$ and $\{a,b,c'\}$, then c and c' are connected by a path not passing through a or b. Otherwise, a and b would be a cut pair, in contradiction to the definition of a cut triple. Thus, if we only consider graphs Γ, where the JSJ graph of cylinders has two-ended edge groups, overlapping cut triples do not occur in Γ.

This has an impact on Proposition 3.14: Recall that the set $\mathcal{S}$ contains as elements the sets $\{a,b\} \cup \mathcal{C}$, where $\{a-b\}$ is an uncrossed cut collection and $\mathcal{C}$ is the set of common adjacent vertices of a and b. Excluding overlapping cut triples implies that every uncrossed cut collection contributes a new element to $\mathcal{S}$. Hence, every uncrossed cut collection is in one-to-one correspondence with a cylinder vertex.

To obtain a JSJ graph of cylinders with two-ended edge groups, we need to refine the Standing Assumption 1.2 to:

Standing Assumption 1.3. *Γ is the defining graph of a RACG W_Γ which satisfies:*

(1) Γ is triangle-free.
(2) Γ is connected and has neither a separating vertex nor a separating edge.
(3) Γ has a cut collection.
(4) Γ is not a cycle.
(5) Γ has only uncrossed cut collections $\{a-b\}$ for which for any pair $\{c_1,c_2\} \in \mathcal{C}$ of common adjacent vertices of a and b, every path in Γ connecting c_1 and c_2 passes through a or b.

Remark 3.32. Under Standing Assumption 1.3, the proofs of Lemma 3.25 and Theorem 3.28 imply that the edge stabilizers are either of the form $W_{\{a,b\}}$ or $W_{\{a,b\}} \times W_{\{c\}}$, where $\{a-b\}$ is an uncrossed cut collection and c is a common adjacent vertex of a and b. In particular, the latter case can only happen when a, b and c are the corners of a subdivided K_4.

To conclude, we summarize the construction of the JSJ graph of cylinders Σ_c in the following theorem:

Theorem 3.33. *Let W_Γ be a RACG with Γ satisfying the Standing Assumption 1.3. Then its JSJ graph of cylinders Σ_c consists of the following vertices:*

- *For any uncrossed cut collection $\{a-b\} \subseteq EV(\Gamma)$, there is a cylinder vertex with vertex group $W_{\{a,b\}\cup\mathcal{C}}$, where $\mathcal{C}$ is the collection of common adjacent vertices of a and b in Γ. All the cylinder vertices are either two-ended, VA or VFD.*
- *For any set $A \subseteq V(\Gamma)$ of vertices such that W_A is infinite, A satisfies either conditions (A1), (A2) and (A3) or condition (A*) and A is not contained in a vertex set corresponding to a cylinder vertex group, there is a hanging vertex with vertex group W_A. The vertex group is virtually free.*
- *For any set $B \subseteq EV(\Gamma)$ of essential vertices in Γ satisfying the conditions (B1), (B2) and (B3), there is a rigid vertex with vertex group W_B.*

Furthermore a pair of vertices is connected by an edge if and only if the pair consists of one cylinder vertex corresponding to the cut collection $\{a-b\}$ and one non-cylinder vertex, and their vertex groups intersect in a special subgroup containing $W_{\{a-b\}}$. The edge group is the special subgroup generated by the intersection of the corresponding vertex sets. It is two-ended.

Example 3.34. For the graph Γ shown in Figure 3.1.5, which satisfies Standing Assumption 1.3, we can construct the corresponding JSJ graph of cylinders by reading off the following collections of vertices:

uncrossed cut collection	commensurator	hanging	rigid
a, b	a, m_1, m_2, m_3, b		a, b, c, d
a, c	a, c	a, k_1, k_2, c	
a, d	a, d	a, l_1, l_2, d	
b, c	b, o, c		
b, d	b, p, d		
c, d	c, n_1, n_5, d	c, n_2, n_3, n_4, d	

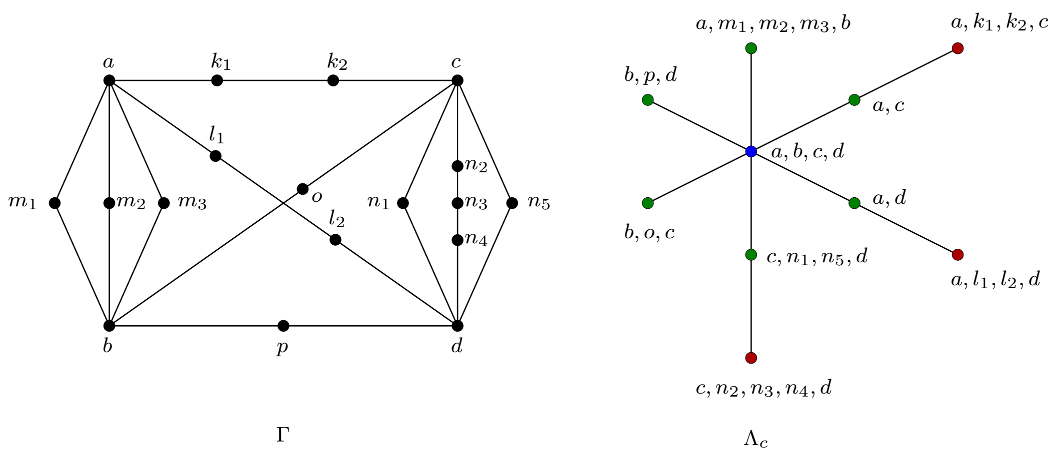

Figure 3.1.5

3.2 QI-invariance

As discussed in Section 1.1, the feature of interest of the JSJ graph of cylinders is that it can give insight on whether two groups can be QI or not. In the case of certain hyperbolic RACGs, we know by Theorem 3.1 that all the two-ended cylinder vertices have finite valence in the JSJ tree of cylinders. Thus, if two groups exhibit different valencies at their cylinder vertices, the JSJ trees of cylinders are not isomorphic, and thus, by Theorem 2.15 the groups are not QI.

However, this argument is not applicable in general, since cylinder vertices with one-ended vertex groups do not have finite valence. Nonetheless, we still might be able to distinguish trees of cylinders with infinite valence cylinder vertices, and thus produce an obstruction for a QI, by taking the additional structure coming from the vertex groups and their interplay via edge groups into account. This idea was formalized by Cashen and Martin in [CM17a] by the introduction of the *structure invariant*. Recall its definition from Section 2.1.1. We first illustrate when it can distinguish two RACGs up to QI and when it cannot. In a second step, we aim to produce a QI between certain groups from identical structure invariants, making the structure invariant a *complete* QI-invariant.

3.2.1 The structure invariant for RACGs

Since the ornaments (recall Definition 2.17) on a JSJ tree of cylinders consisting of vertex type and relative QI-type determine the structure of the group, we can refine our search to decoration-preserving tree isomorphisms. Hence, by Proposition 2.24, the structure invariant is indeed a QI-invariant for RACGs with defining graph satisfying Standing Assumption 1.3 [cf. CM17a, Theorem 3.8].

Example 3.35. The two groups with defining graphs illustrated in Figure 3.2.6 serve as an introductory example as they are easily distinguished as non-QI by use of the structure invariant. While the commensurator of the cut pairs $\{a, b\}$ and $\{c, d\}$ in Γ_1 both give a VFD vertex group, in Γ_2 the commensurator of $\{c, d\}$ corresponds to a VA vertex group. Since both graphs only have those two uncrossed cut collections, the initial decoration consisting of vertex and relative QI-type already shows that the groups cannot be QI.

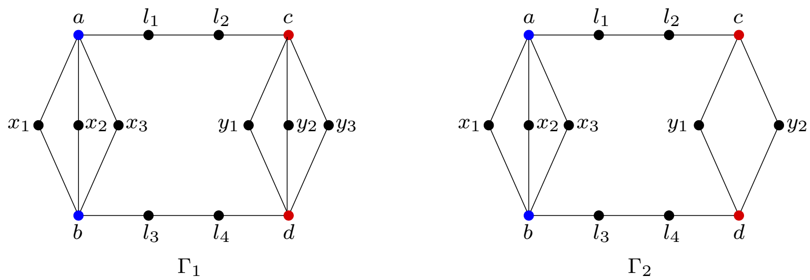

Figure 3.2.6

Example 3.36. To obtain the structure invariants of the JSJ graphs of cylinders Σ_1 and Σ_2 for the two RACGs W_1 and W_2 on the defining graphs $\Gamma_{c,1}$ and $\Gamma_{c,2}$, respectively, illustrated in Figure 3.2.7, we start with the following initial decorations:

$$\begin{array}{rcl} \delta\colon V(\Sigma_1) & \to & \mathcal{O} \\ c & \mapsto & (\text{‘cyl’}, [\![(\text{‘VA’}, \mathcal{P}_c)]\!]) \\ h & \mapsto & (\text{‘hang’}, [\![(\text{‘VF’}, \mathcal{P}_h)]\!]) \\ r & \mapsto & (\text{‘rig’}, [\![(W_{\{a,b,c,d,e,f,g,h\}}, \mathcal{P}_r)]\!]) \end{array} \qquad \begin{array}{rcl} \delta'\colon V(\Sigma_2) & \to & \mathcal{O} \\ c' & \mapsto & (\text{‘cyl’}, [\![(\text{‘VA’}, \mathcal{P}_c)]\!]) \\ h_1,\, h_2 & \mapsto & (\text{‘hang’}, [\![(\text{‘VF’}, \mathcal{P}_h)]\!]) \\ r' & \mapsto & (\text{‘rig’}, [\![(W_{\{a,b,c,d,e,f,g,h\}}, \mathcal{P}_r)]\!]) \end{array}$$

We immediately see that no refinement is possible, the vertices h_1 and h_2 are indistinguishable, and thus, the following structure invariant is the same for both JSJ trees of cylinders:

	(‘cyl’, $[\![(\text{‘VA’}, \mathcal{P}_c)]\!]$)	(‘hang’, $[\![(\text{‘VF’}, \mathcal{P}_h)]\!]$)	(‘rig’, $[\![(W_{\{a,b,c,d,e,f,g\}}, \mathcal{P}_r)]\!]$)
(‘cyl’, $[\![(\text{‘VA’}, \mathcal{P}_c)]\!]$)	0	∞	∞
(‘hang’, $[\![(\text{‘VF’}, \mathcal{P}_h)]\!]$)	∞	0	0
(‘rig’, $[\![(W_{\{a,b,c,d,e,f,g\}}, \mathcal{P}_r)]\!]$)	∞	0	0

Thus, by Proposition 2.24, there is a decoration preserving tree isomorphism between the respective trees of cylinders T_1 and T_2. This leads to the question whether W_1 and W_2 are QI, which we will answer in the negative in Example 3.42.

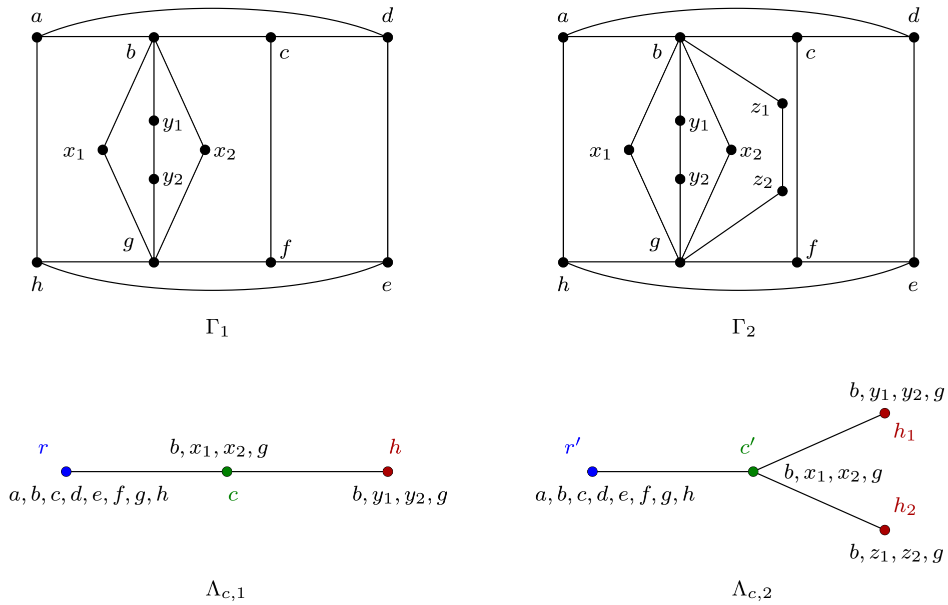

Figure 3.2.7

3.2.2 Promoting to a QI

Given two groups G and G' with identical structure invariants and thus with a decoration-preserving isomorphism between their respective JSJ trees of cylinders, we want to determine when we can promote this isomorphism to a QI of the groups. Since any QI between G and G' needs to restrict to a QI locally at each vertex group by Theorem 2.15, the general idea is the following: Start with any local QI between two cylinder vertex groups with the same entry in the structure invariant, which is bijective on the peripheral structures coming from the incident edge groups (see Example 2.18), and extend it piece by piece from there. By Lemma 3.15, we know that in our setting we can encounter either two-ended, VA or VFD cylinder vertex groups. Hence, we first determine the possible local QIs for these different cases separately and combine the respective results to find a global QI in a next step.

Remark 3.37. All arguments work along the lines of the ones used in [CM17a], where the case of two-ended cylinder vertices is dealt with. However, at this point we need to clarify three technicalities:

- *Rigid vertices need to be handled with special care:*
  - While the relative QI-type of rigid vertices might be hard to determine, it can be the crucial ingredient to distinguish groups. In the case of RACGs for instance, this is illustrated by Cashen, Dani and Thomas in [DT17]. Their Theorem B.1 states that, while all RACGs on 3-convex subdivided complete graphs with at least 4 essential vertices have isomorphic JSJ trees of cylinders, they are pairwise non-QI. The reason for that lies in the relative QI-type of the rigid vertices. To have more control over rigid vertices, Cashen and Martin restrict to those that have the property of being *quasi-isometrically rigid relative to the peripheral structure* [CM17a, Definition 4.1]. For example, free rigid vertex groups have this property by [CM11]. Under this additional assumption, another ornament, the *relative stretch factor*, can be introduced to decorate edges and help distinguish rigid vertices [CM17a, Section 4]. However, whether rigid vertices in JSJ trees of cylinders of

RACGs have this or a similar sufficient property (for example the related *right-angled Artin groups* splitting over cyclic groups do [cf. Mar20, Section 6]) is yet to be determined.

  – Another issue caused by rigid vertices is that they might have adjacent edges whose edge groups are not two-ended as shown in Theorem 3.28. As explained in Remark 1.38, Nguyen and Tran give in [NT19] a complete QI-classification of a class of RACGs with such edge groups, see Theorem 1.37: The defining graphs are connected, triangle-free and planar, have more than 4 vertices, no separating edge or vertex and they are $\mathcal{CFS}$ (see Definition 1.34 and [BFRHS18, Definition 1.3]). It is highlighted in Remark 1.38 that in the proof of Theorem 1.77 they use the *maximal suspension graph*, which is in correspondence with the JSJ graph of cylinders and the decoration consisting of the relative QI-type.
  – Moreover, in [BX20], Bounds and Xie show that RACGs, whose defining graphs are *generalized thick m-gons*, exhibit a strong form of QI-rigidity: They are QI if and only if their defining graphs are isomorphic.

  For simplicity, we focus on groups without any rigid vertices or on pairs of groups which have isomorphic rigid vertex groups as in Examples 3.36 and 3.50.

- *Work on the geometric tree of spaces:* To make technical details more economic, instead of working on graphs of groups, Cashen and Martin state their results for a slightly modified space, the *geometric tree of spaces* X of G over T_c. The construction of X is standard and useful as X is QI to G. Essentially, X is produced from the JSJ graph of cylinders Σ_c by substituting all groups of the same relative QI-type by a uniform *model space* representing the equivalence class. Thus, instead of a subgroup G_t we have a subspace X_t for every $t \in T_c$. Most importantly, if two groups G and G' exhibit subgroups G_t and $G'_{t'}$ with equivalent relative QI-types in their JSJ graphs of cylinders, we choose the same model space X_t for both G_t and $G'_{t'}$. If convenient, we will state results in terms of the geometric tree of spaces X, but spare the bookkeeping, which is done thoroughly in Sections 7.2 and 2.5 of [CM17a].
- *Partial orientations can be omitted:* For the sake of completeness, it should be mentioned that, apart from the neighbor refinement, Cashen and Martin introduce the *cylinder* and the *vertex refinement*, depending on a *partial orientation* chosen essentially on all two-ended spaces. However, since all infinite RACGs, and thus, all edge groups in Σ_c contain an infinite dihedral group D_∞, the orientation can always be reversed. Thus, the refinement processes become trivial and shall therefore be left out of our considerations.

3.2.2.1 Two-ended cylinder vertices In case all cylinder vertex groups are two-ended, like for instance for hyperbolic groups, Cashen and Martin give a structure invariant, which is a complete QI-invariant. Their result, stated for RACGs splitting over two-ended subgroups and thus refining Theorem 2.15, says:

Proposition 3.38. [CM17a, Theorem 7.5] *Let W and W' be two finitely presented, one-ended RACGs with non-trivial JSJ decomposition over two-ended subgroups such that cylinder stabilizers are two-ended and all non-cylinder vertex groups are either hanging or quasi-isometrically rigid relative to the peripheral structure. Define T to be the JSJ tree of cylinders of W and X to be the geometric tree of spaces of W over T. The initial decoration δ_0 on T takes vertex type, relative QI-type and the relative stretch factor into account. Let δ be the neighbor refinement of δ_0. Analogously, we define T', X', δ'_0 and δ' for W'. Then W and W' are QI if and only if there is a bijection $\beta\colon \delta(T) \to \delta'(T')$ such that*

1. $\delta_0 \circ \delta^{-1} = \delta'_0 \circ (\delta')^{-1} \circ \beta$.
2. $S(T, \delta, \mathcal{O}) = S(T', \delta', \mathcal{O}')$ *in the* β*-induced ordering.*
3. *For every ornament* $o \in \mathcal{O}$ *with* $\delta^{-1}(o)$ *containing non-cylinder vertices, there is a vertex* $v \in \delta^{-1}(o)$ *and a vertex* $v' \in (\delta')^{-1}(\beta(o))$ *such that there is a QI between the vertex spaces* X_v *and* $X'_{v'}$*, which is bijective on the peripheral structures* $\mathcal{P}_v$ *and* $\mathcal{P}'_{v'}$ *and respecting the decorations* δ *and* δ'*, respectively.*

The inductive construction of the QI in their proof will serve as a blueprint for the proof of the general Theorem 3.45.

3.2.2.2 VFD cylinder vertices It turns out that VFD cylinder vertex groups have enough flexibility to always find a QI between cylinder vertices with this same entry in the structure invariant. We construct this local QI in the following simplest setting:

Proposition 3.39. *Let* W_1 *and* W_2 *be two RACGs on defining graphs satisfying Standing Assumption 1.3 with identical structure invariants and one single cylinder vertex* v_1 *and* v_2 *in the JSJ graph of cylinders* Σ_1 *and* Σ_2*, respectively. Let the vertex group* V_1 *and* V_2 *of* v_1 *and* v_2*, respectively, be VFD. Then there is a QI between* V_1 *and* V_2 *that is bijective on the respective peripheral structures.*

Proof. The set-up is the following: Both JSJ graphs of cylinders Σ_1 and Σ_2 look like stars, with the cylinder vertex in the middle and their neighbors grouped into $j < \infty$ classes of indistinguishable vertices. Suppose at first that $j = 1$.

Thus, for $i \in \{1, 2\}$, each Σ_i consists of one cylinder vertex v_i, which has a vertex group of the form $V_i = W_{\mathcal{C}_i} \times D_\infty$. The copy of D_∞ is generated by non-adjacent vertices of a cut collection and $W_{\mathcal{C}_i}$ is generated by the set $\mathcal{C}_i$ of their common adjacent vertices. By assumption $W_{\mathcal{C}_i}$ is virtually free, thus $|\mathcal{C}_i| > 2$. At the cylinder vertex v_i in the middle, there is a set $\mathcal{N}_i$ of e_i indistinguishable non-cylinder vertex groups of the same relative QI-type attached along a two-ended edge group. These edge groups are either a copy of D_∞ or of $D_\infty \times \mathbb{Z}_2$ with $\mathbb{Z}_2 = W_{\{c\}}$ for some $c \in \mathcal{C}_i$ by Remark 3.32. Thus, in the corresponding JSJ tree of cylinders, the vertex $1 \cdot V_i$ has infinitely many adjacent vertex groups corresponding to cosets of the form gN: The group N is an element of $\mathcal{N}_i$ and $g \in V_i$ is either any word in $W_{\mathcal{C}_i}$ or a word in $W_{\mathcal{C}_i}$ not ending on c, depending on whether the edge group along which N attaches is D_∞ or $D_\infty \times W_{\{c\}}$.

We want to interpret this set-up in terms of Cayley graphs in order to prove Claim 3.39.1. Before, we need to fix some terminology:

As a *graph* Ω *with tangling edges* E we understand some *base* graph Ω, where at each vertex in $V(\Omega)$ we add some additional neighbors, all of valence 1. Each such additional edge is labelled by an element of E and is called a *tangling edge*. In the new graph, we can think of each tangling edge as the pair (v, i), where $i \in \mathbb{N}$ counts the edges attaching at the base vertex $v \in V(\Omega)$ in Ω. We denote the resulting graph as $\Omega \cup E$, where the union happens via the implicit attaching map. Let $k(v)$ be the number of edges *tangling* at $v \in V(\Omega)$. Then we can interpret the set E as $E = \{t_{v,i} \mid v \in V(\Omega), i \in \{0, \ldots, k(v) - 1\}\}$, where $t_{v,i}$ denotes the i-th tangling edge at vertex $v \in V(\Omega)$.

Claim 3.39.1. *The problem of finding a QI between* V_1 *and* V_2 *that is bijective on the respective peripheral structures can be reduced to finding a QI between two identical infinite, regular trees* T *with tangling edge sets* E_1 *and* E_2 *such that the occurring numbers of tangling edges* $\{k_i(v) \mid v \in V(T)\}$ *in* $T \cup E_1$ *and* $T \cup E_2$ *differ. In addition, this QI must be bijective on the tangling edges.*

Proof of Claim 3.39.1. The idea of the reduction is the following: The Cayley graph of $W_{\mathcal{C}_i}$ reduces to the base tree and the tangling edges are in correspondence with the different cosets gN.

We start the reduction process with the object X_i, illustrated in Figure 3.2.8, constructed as follows: Note first that the Cayley graph of the cylinder vertex group $V_i = W_{\mathcal{C}_i} \times D_\infty$ is the direct product of the Cayley graph T_i of $W_{\mathcal{C}_i}$ and the line D that is the Cayley graph of D_∞. This is true, because with the correct choice of generating sets, the Cayley graph of a direct product is the direct product of the Cayley graphs. Since all the vertices in $\mathcal{C}_i$ are pairwise non-connected in the defining graph, the Cayley graph T_i of $W_{\mathcal{C}_i}$ is a $|\mathcal{C}_i|$-regular tree. We can think of each coset gN adjacent to the vertex $1 \cdot V_i$ as attaching in this Cayley graph. If g can be any word in $W_{\mathcal{C}_i}$, the coset attaches at the vertex g in T_i and along the line D. If g is a word in $W_{\mathcal{C}_i}$ not ending on c, the coset attaches along the edge c starting at the vertex g in T_i and along the line D. Either way, we can think of the e_i different cosets gN as e_i possibly thickened half-planes at the vertex g in T_i attached along the line D. Note that at one vertex g it can happen that there attach both *thick* half-planes along an edge and *thin* half-planes at the vertex. We call the constructed object X_i.

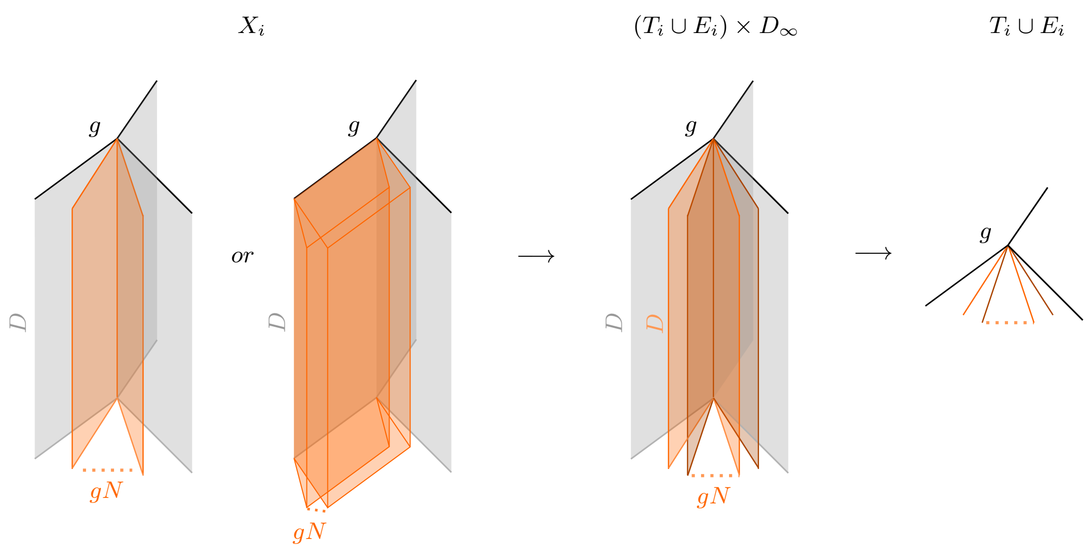

Figure 3.2.8

Since X_i captures the structure of the group, the task of finding a QI between V_1 and V_2 that is bijective on the peripheral structure is done if we can show that there is a QI between X_1 and X_2 that is bijective on the half-planes corresponding to the cosets. For the reduction, we squish for any $N \in \mathcal{N}_i$ attaching along a $D_\infty \times W_{\{c\}}$ the corresponding thick half-plane: Replace each such thick half-plane attaching along an edge by a thin half-plane attached at the terminal vertex of the attaching edge that has less half-planes attached. Then, at each vertex in T_i, there attaches some positive number of thin tangling half-planes corresponding to the cosets gN. We reinterpret this object as $(T_i \cup E_i) \times D$, where E_i is a set of tangling edges. It suffices to find a QI between $(T_1 \cup E_1) \times D$ and $(T_2 \cup E_2) \times D$ that is bijective on the tangling half-planes, because this immediately implies that we can find a QI between the trees with thick tangling half-planes simply by extending the map along the attaching edges via the identity.

However, now it is enough to find a QI between $T_1 \cup E_1$ and $T_2 \cup E_2$ which is bijective on the tangling edges, because again, this immediately implies that we can find a QI between $(T_1 \cup E_1) \times D$ and $(T_2 \cup E_2) \times D$ this time by extending the map to D via the identity.

So, to find this QI, recall the well-known fact that two regular trees are QI to each other by contracting or inserting one edge path of a certain finite length at each vertex of the first tree to turn it into the second. Thus, we start with the tree with smaller regularity, say, without loss of generality, T_1 and perform this operation on the edges to obtain a tree T that is QI to T_1 and isomorphic to T_2. While this contraction and insertion of edges redistributes the tangling edges of T_1, since we started with the tree with smaller regularity, the resulting T also has at least one tangling edge at each vertex. Hence, there is a QI between $T_1 \cup E_1$ and $T_2 \cup E_2$ that is bijective on the tangling edges if we can find a QI between $T \cup E_1$ (with some adjusted tangling edge set E_1) and $T \cup E_2$ that is bijective on the tangling edges.

We could keep track of the exact number $k_i(v)$ of tangling edges at each vertex v in $V(T)$ of $T \cup E_i$. However, since this would require a technical case distinction, we suppress the details. In general, the number $k_i(v)$ of tangling edges at each vertex v varies. However, most importantly, we see from an analysis of the reduction process that all vertices have a bounded number of tangling edges, that is, all vertices have at least $y_i > 0$ and at most $x_i < \infty$ tangling edges, that is, $0 < y_i \leq k_i(v) \leq x_i < \infty$ for all $v \in V(T)$.

With this process, we have reduced the problem of finding a QI between V_1 and V_2 that is bijective on the respective peripheral structures to finding a QI between two copies of an infinite, r-regular tree T with differing occurring numbers $\{k_i(v) \mid v \in V(T)\}$ of tangling edges at its vertices, that is bijective on the tangling edges. □

So, as Claim 3.39.1 suggests, we aim to find a QI q from $T \cup E_1$ to $T \cup E_2$, where the base graph T is an infinite r-regular tree with distinguished base vertex and q is bijective on the tangling edges. Without loss of generality, we set the maximal number x_1 of edges attaching at a base vertex in $T \cup E_1$ to be greater than the maximal number x_2 of edges attaching at a base vertex in $T \cup E_2$.

We define the following notion on $T \cup E_i$: Given an edge $e \in E(T)$ with some tangling edge $t \in E_i$ at the vertex $o(e)$, we call it a *slide* along e if we detach t from $o(e)$ and reattach it at $t(e)$.

Claim 3.39.2. *There is a constant $d \in \mathbb{N}$ such that in $T \cup E_1$ every tangling edge at each vertex of T needs to slide at most along d edges of T away from the distinguished base vertex, such that the resulting graph is isomorphic to $T \cup E_2$.*

Now, we get the desired QI q from $T \cup E_1$ to $T \cup E_2$ which is bijective on the tangling edges: The sliding process in Claim 3.39.2 defines a bijective map $q' : E_1 \to E_2$ mapping each edge in E_1 to the edge in E_2 on whose position it is slid to. We define $q : T \cup E_1 \to T \cup E_2$ to be the map that is the identity on T and q' on the elements of E_1 as “half-open” edges without the endpoint contained in T.

Let $t_{v,j}$ and $t_{v',j'}$ in E_1 be two tangling edges based at v and v' in T, respectively. Since tangling edges are always slid away from the distinguished base vertex, their images can get at most d edges closer to each other than v and v' are. Hence

$$d(t_{v,j}, t_{v',j'}) - d \leq d(q(t_{v,j}), q(t_{v',j'}))\,,$$

which gives the lower QI-bound. For the upper QI-bound, note that, since both tangling edges are slid at most along d edges, their distance can grow at most by $2d$, that is

$$d(q(t_{v,j}), q(t_{v',j'})) \leq d(t_{v,j}, t_{v',j'}) + 2d\,.$$

Since a vertex $w \in V(T)$ is not moved by q, analogous bounds hold for $d(q(t_{v,j}), q(w))$. This implies that q is a quasi-isometric embedding. The bijectivity of q' ensures the quasi-surjectivity of q. Therefore, q is a QI and the only thing left to prove is Claim 3.39.2:

Proof of Claim 3.39.2. For simplicity, we want to define the graphs T_i', which are identical to $T \cup E_i$ with the exception that in T_i', the base vertex of T does not carry any tangling edges. Since the number of tangling edges we remove is bounded by $k < \infty$, it suffices to prove the claim for T_i', because the argument for $T \cup E_i$ works analogously. However, when we work on T_i', we can give d explicitly in terms of the maximal number x_1 of tangling edges at a vertex in $T \cup E_1$, the minimal number y_2 of tangling edges at a vertex in $T \cup E_2$ and the degree r of the regularity of T as

$$d = \left\lceil \frac{\log(\frac{x_1}{y_2})}{\log(r-1)} \right\rceil .$$

If the base vertex carries at most k tangling edges as well, d is bounded by $\left\lceil \frac{\log(\frac{x_1}{y_2})}{\log(r-1)} \right\rceil + k$, a complication we avoid by working with T_i' instead of $T \cup E_i$. Now, the key feature of the proof is the following algorithm:

Algorithm 3.39.3. Since we always need to slide away from the base vertex, we can reduce the problem by dividing the tree T into r subtrees by removing the base vertex, which does not have any tangling edges. Then we have r rooted trees, where the root has $r-1$ outgoing edges. We consider one rooted tree R, which is oriented away from the root $*$. A vertex is at level l of R if it has distance l to the root $*$. Note that every vertex in R, with exception of the root, has one incoming, $(r-1)$ outgoing and at least y_i and at most x_i tangling edges. The root has $r-1$ outgoing and at least y_i and at most x_i tangling edges.

Now, the idea is the following: Every vertex receives some tangling edges via a slide along its incoming edge and superfluous tangling edges leave the vertex via a slide along the outgoing edges. The sliding process follows two rules:

1. The distribution of the superfluous tangling edges along the $r-1$ outgoing edges is uniform.
2. The edges that are kept at each vertex are always the ones that have been slid the furthest.

Without loss of generality, we can assume that all vertices in T_1' have the maximal number of x_1 tangling edges and all vertices in T_2' with the minimal number of y_2 tangling edges. If the bound d works for these special cases, it works for numbers of tangling edges in between.

We show by induction on the level l of R that d satisfying $\left\lceil \frac{x_1}{(r-1)^d} \right\rceil \leq y_2$ works as a uniform bound. Consider the root of R at level 0 as the base case. We need to keep y_2 edges at the root and by rule 2, we keep at most a total of $(r-1) \cdot y_2$ edges coming from the root at level 1. In general we keep at most a total of $(r-1)^i \cdot y_2$ edges coming from the root at level i. But since

$$\sum_{i=0}^{d} (r-1)^i \cdot y_2 \geq (r-1)^d \cdot y_2 \geq x_1 \, ,$$

it is immediate that none of the x_1 edges coming from the root will be slid more than d steps.

For the inductive step, suppose that each edge up to level l will be slid at most along d edges. We consider a vertex v at level $l+1$. If we slide its tangling edges along d edges in R, they are now attached at a vertex at level $l+1+d$. But by hypothesis any edge slid away from a vertex at level l or any level above cannot be attached at level $l+1+d$. Thus, the edges from level $l+1$ are the ones that have been slid the furthest, so by rule 2 they are the ones that need to stay. However, by choice of d, there are at most y_2 edges coming from v per vertex at level $l+1+d$. Therefore, no edges coming from level $k+1$ are slid any further, proving that the chosen d gives a uniform bound. □

The way to interpret Algorithm 3.39.3 is that in the JSJ graphs of cylinders, we can duplicate or collapse the neighboring vertices of v_1 of the same QI-type to match the neighbors of v_2.

In order to produce a QI between V_1 and V_2 when Σ_1 and Σ_2 have $j \geq 2$ classes of indistinguishable vertices attached at the cylinder vertex, we apply the Claims 3.39.1 and 3.39.2 and execute Algorithm 3.39.3 for each class individually. □

3.2.2.3 VA cylinder vertices The flexibility of the VA cylinder vertices lies in between the flexibility of the other two types: In the tree of cylinders, they have infinite valence like the VFD cylinder vertices. However, in order to get a QI from one VA cylinder vertex group to another, the different classes of indistinguishable neighboring vertex groups must occur with matching densities in the respective JSJ graphs of cylinders. This behaviour is similar to the two-ended cylinder vertices. The robustness comes from the fact that the QI cannot be of any type, but it must be bounded distance from scaling by precisely the density. Shepherd and Woodhouse also make use of these densities in [SW22, Section 5.6].

As for the VFD cylinders, we construct the local QI in the following simplest setting:

Proposition 3.40. *Let W_1 and W_2 be two RACGs on defining graphs satisfying Standing Assumption 1.3 with identical structure invariants and one single cylinder vertex v_1 and v_2 in the JSJ graph of cylinders Σ_1 and Σ_2, respectively. Let the cylinder vertex groups $V_1 \cong D_\infty \times D_\infty$ and $V_2 \cong D_\infty \times D_\infty$ at v_1 and v_2, respectively, be VA. Suppose at v_1 attach e_1 and at v_2 attach e_2 neighbors of the same class of indistinguishable vertices. The number e_1 decomposes as the sum of m_1 vertices attaching along a D_∞-edge and n_1 vertices attaching along a $D_\infty \times \mathbb{Z}_2$-edge. Analogously, $e_2 = m_2 + n_2$.*

There is a QI from V_1 to V_2 that is bijective on the respective peripheral structures if and only if there is a QI that is the identity map on the first D_∞-copy of V_1 and V_2 and that scales under the natural identification with $\mathbb{Z}$ the second D_∞-copy of V_1 to the second D_∞-copy of V_2 by

$$\frac{2\,m_1 + n_1}{2\,m_2 + n_2}\,.$$

Furthermore, every QI between V_1 and V_2 that is bijective on the respective peripheral structures is bounded distance from one of the form

$$\begin{array}{rccc} \psi\colon & D \times L & \to & D \times L \\ & (x,y) & \mapsto & (\psi'(x,y), \psi''(x,y))\,, \end{array}$$

where D and L are Cayley graphs of D_∞ and ψ'' is scaling by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$.

Proof. The proof resembles the proof of Proposition 3.39, we have a similar set-up: For $i \in \{1,2\}$, the JSJ graph of cylinders Σ_i looks like a star with one VA cylinder vertex v_i in the middle. The VA vertex group V_i at v_i corresponds to the uncrossed cut collection $\{a_i - b_i\}$ with common adjacent vertex set $\{s_i, t_i\}$, thus $V_i = W_{\{a_i,b_i\}} \times W_{\{s_i,t_i\}} = D_\infty \times D_\infty$.

As before, in Σ_i, at the cylinder vertex v_i, there is a set $\mathcal{N}_i$ of e_i indistinguishable non-cylinder vertex groups of the same relative QI-type. Out of these, m_i are attached along a $W_{\{a_i,b_i\}} = D_\infty$-edge group and n_i are attached along a $W_{\{a_i,b_i,s_i\}} = D_\infty \times \mathbb{Z}_2$-edge group or a $W_{\{a_i,b_i,t_i\}} = D_\infty \times \mathbb{Z}_2$-edge group (cf. Remark 3.32). In the corresponding JSJ tree of cylinders, the vertex $1 \cdot V_i$ has infinitely many adjacent vertex groups corresponding to cosets of the form gN. The group N is an element of $\mathcal{N}_i$ and $g \in V_i$ is either any word in $W_{\{s_i,t_i\}}$ or any word in $W_{\{s_i,t_i\}}$ not ending on s_i or on t_i, depending on whether N attaches along the edge group $W_{\{a_i,b_i\}}$ or $W_{\{a_i,b_i,t_i\}}$ or $W_{\{a_i,b_i,s_i\}}$, respectively.

Again, we want to interpret this set-up in terms of Cayley graphs in order to prove:

Claim 3.40.1. *The problem of finding a QI between V_1 and V_2 that scales $W_{\{s_1,t_1\}}$ to $W_{\{s_2,t_2\}}$ by $\frac{2m_1+n_1}{2m_2+n_2}$ and that is bijective on the respective peripheral structures can be reduced to finding a QI between two copies of the number line with different occurring numbers of tangling edges that scales the number line by $\frac{2m_1+n_1}{2m_2+n_2}$.*

Proof of Claim 3.40.1. Analogous to the procedure for a VFD cylinder vertex in the proof of Proposition 3.39, we use the Cayley graph of V_i. It is given by the $a_ib_i \times s_it_i$-grid. Note that the bi-labelled s_it_i-line L corresponds to T_i in the proof of Proposition 3.39 and the bi-labelled a_ib_i-line corresponds to D. If g can be any word in $W_{\{s_i,t_i\}}$, the coset gN attaches at vertex g in L along D. If g is a word in $W_{\{s_i,t_i\}}$ not ending on s_i, the coset gN attaches along the edge s_i starting at the vertex g in L and along D. If g is a word in $W_{\{s_i,t_i\}}$ not ending on t_i, the coset gN attaches along the edge t_i starting at the vertex g in L and along D. Either way, we think of the e_i different cosets gN as e_i possibly thickened half-planes at the vertex g in L attached along the line D. Using the notation from the proof of Proposition 3.39, we call this object X_i, it is illustrated in Figure 3.2.9.

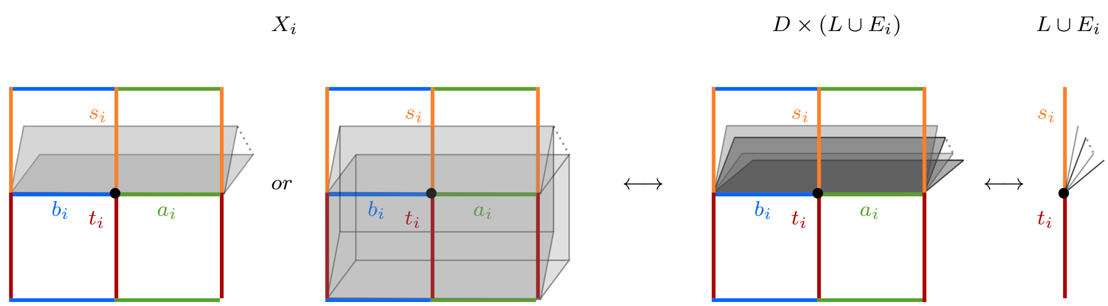

Figure 3.2.9

Again, we obtain a QI between V_1 and V_2 that is bijective on the peripheral structure if we can find a QI between X_1 and X_2 that is bijective on the half-planes corresponding to the cosets. Unlike in the proof of Proposition 3.39, we need to make sure that the following reductions work both ways in order to prove the fact about the scaling, that is, we show that we find the desired QI between the reduced objects if and only if we find one between the original ones.

First, we get rid of the thick half-planes, as in the proof of Claim 3.39.1. For any $N \in \mathcal{N}_i$, we find attaching along a $W_{\{a_i,b_i,s_i\}}$- or $W_{\{a_i,b_i,t_i\}}$-edge, we need to squish the corresponding thick half-planes: Replace half of the thick half-planes attaching along an edge $e \in \{s_i, t_i\}$ by a thin half-plane attached at $o(e)$ and the other half at $t(e)$. Then, at each vertex in L, there attach $\frac{n_i}{2}$ thin tangling half-planes coming from thick ones and m_i originally thin ones. In total, there attach $m_i + \frac{n_i}{2}$ thin half-planes at each vertex. Of course, it can happen that n_i is odd and we have produced "half a half-plane" with this procedure. However, this will not affect the rest of the argument.

We reinterpret this object as $(L \cup E_i) \times D$, where E_i is the set of tangling edges. It suffices to find a QI between $(L \cup E_1) \times D$ and $(L \cup E_2) \times D$ that scales L by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$ that is bijective on the tangling half-planes. The existence of such a QI immediately implies that we can find a QI between the grids with thick tangling half-planes simply by extending the map along the attaching edge via the identity. Conversely, if we find a QI between two grids with thick tangling half-planes which is bijective on all tangling half-planes, this means that the horizontal D-lines are preserved. Thus, we can restrict this QI to obtain the desired QI between the grids with thin tangling edges.

In the second reduction step we check that it is enough to find a QI between $L \cup E_1$ and $L \cup E_2$ that is bijective on the tangling edges and scaling L by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$. Again, given such a QI, we can find the desired QI between $D \times (L \cup E_1)$ and $D \times (L \cup E_2)$ by extending the map via the identity to D.

For the converse, suppose $\psi\colon D \times (L \cup E_1) \to D \times (L \cup E_2)$ is a QI which is bijective on the tangling half-planes and scaling L by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$. Restricted to the grid, ψ is of the form:

$$\begin{array}{cccc} \psi\colon & D \times L & \to & D \times L \\ & (x,y) & \mapsto & (\psi'(x,y), \psi''(x,y))\,. \end{array}$$

However, the bijectivity on the tangling half-planes implies that ψ coarsely preserves the copies of D, that is, the $a_i b_i$-horizontal lines. This means, given the pair (x_0, y_0) in the grid of $D \times (L \cup E_1)$, with image $\psi((x_0, y_0)) = (x_0', y_0')$, any other pair (x, y_0) is mapped to (x', y_0'). This means that $\psi''(x, y)$ is independent of the input of x, hence we can interpret ψ'' as

$$\begin{array}{cccc} \psi''\colon & L & \to & L \\ & y & \mapsto & \psi''(y)\,. \end{array}$$

Via ψ we extend ψ'' again to the tangling edges, that is, we find map $\psi''\colon L \cup E_1 \to L \cup E_2$ that is bijective on the tangling edges. Lastly note, that ψ'' is a QI. Indeed, given two pairs (x_0, y_0),(x_1, y_1) in the grid $D \times L$, we can decompose their distance as: $d_{D\times L}((x_0, y_0), (x_1, y_1)) = d_D(x_0, x_1) + d_L(y_0, y_1)\,.$ This implies that a QI-inequality for ψ also holds for ψ''.

For convenience, we include a third reduction step. As in the proof of Claim 3.39.1 of Proposition 3.39, we contract every other edge of L. This way, we have $2\,m_i + n_i$ edges at each vertex, removing the issue with the "half-edges".

So, we have a QI between V_1 and V_2 which is bijective on the respective peripheral structures and which scales one copy of D_∞ to the other by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$ if and only if we find a QI between two copies of the line L with tangling edge set E_1 and E_2 with $2\,m_1 + n_1$ and $2\,m_2 + n_2$ tangling edges at each vertex, respectively, that scales L by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$ and is bijective on the tangling edges. □

As Claim 3.40.1 suggests, we need to find a QI from $L \cup E_1$ to $L \cup E_2$, where the base graph L is a line whose vertex set we can identify with $\mathbb{Z}$ and the QI is bijective on the tangling edges and scaling by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$. Without loss of generality, we set $2\,m_1 + n_1 > 2\,m_2 + n_2$.

The first step is to define for $i \in \{1, 2\}$ the following map

$$\begin{array}{cccl} \phi_i\colon & L \cup E_i & \to & \mathbb{Z} \\ & l & \mapsto & l \cdot (2m_i + n_i) \\ & t_{z,j} & \mapsto & \begin{cases} z \cdot (2\,m_i + n_i) + j & \text{if } z \geq 0 \\ (z+1) \cdot (2\,m_i + n_i) - j - 1 & \text{if } z < 0\,, \end{cases} \end{array}$$

where $t_{z,j}$ is one of the $2\,m_i + n_i$ tangling edges at $z \in \mathbb{Z}$, that is, $j \in \{0, \ldots, 2\,m_i + n_i - 1\}$. It is easily checked that ϕ_i is bijective on E_i for both $i \in \{1, 2\}$ and by definition ϕ_i scales L by $2\,m_i + n_i$.

Now, it suffices to show the following:

Claim 3.40.2. *Any bijective QI $f\colon \mathbb{Z} \to \mathbb{Z}$ which fixes 0, ∞ and $-\infty$ is bounded distance from the identity map.*

By using Claim 3.40.2, for any isometry i composed with f, we obtain the commuting diagram of the form:

$$\begin{array}{ccccc} & L & \xrightarrow{\varphi} & & L \\ \phi_1|_L & \Big\downarrow & \cdot(2\,m_1+n_1) \qquad & \phi_2|_L & \Big\uparrow \quad : (2\,m_2+n_2) \\ & \mathbb{Z} & \xrightarrow{i\circ f} & & \mathbb{Z} \end{array}$$

This implies that φ is a QI scaling by $\frac{2\,m_1+n_1}{2\,m_2+n_2}$.

Thus, we are left to prove Claim 3.40.2:

Proof of Claim 3.40.2. Let f be a (C, D)-QI satisfying the assumptions and suppose it is not bounded distance from the identity. Then, for any $n \in \mathbb{N}$, we can find a $z_n \in \mathbb{N}$ such that $d(z_n, f(z_n)) > n$.

First we claim that there is a maximal $k \in \mathbb{N}$ such that $f(-k) \geq 0$, implying by surjectivity of f that $[0, \infty) \subseteq f([-k, \infty))$. Suppose this is not true. Then for every $k \in \mathbb{N}$ with $f(-k) \geq 0$, there is a $k' \in \mathbb{N}$ such that $k < k'$ and $f(-k') \geq 0$. However, by the QI-property and the fact that f fixes 0 we have

$$\frac{1}{C} \cdot d(-k, 0) - D \leq d(f(-k), 0) = f(-k)$$

for every $k \in \mathbb{N}$. Thus, with $k \in \mathbb{N}$ tending to ∞, so does $f(-k)$, contradicting the assumption that f fixes $-\infty$.

Now, let $B_R(z_n)$ be a ball of radius R centered at z_n. We want to show that

$$[0, f(z_n) + \frac{R}{C} - D] \subseteq f([-k, z_n + R])$$

for any $R \in \mathbb{N}$ large enough.

Since f is bijective by assumption, some elements must map onto the interval $[0, f(z_n) + \frac{R}{C} - D]$. It is indeed $[-k, z_n + R]$ by the following observations illustrated in Figure 3.2.10 below:

1. By choice of k, there is no element $k' < -k$ such that $f(k') \geq 0$.
2. Since f is a QI, $B_{\frac{R}{C}-D}(f(z_n)) \subseteq f(B_R(z_n))$.
3. Any element $z > z_n + R$ maps to an element $f(z) > f(z_n) + \frac{R}{C} - D$: Pick some $a > C \cdot f(z_n) + R$, for which
$$f(a) = d(f(a), 0) \geq \frac{a}{C} - D > f(z_n) + \frac{R}{C} - D\,.$$
Such an a must exist, since f fixes 0, ∞ and $-\infty$. Thus, a is mapped to the right side of $Z \setminus B_{\frac{R}{C}-D}(f(z_n))$. Now, choose a' such that $d(a, a') = 1$. This implies
$$d(f(a), f(a')) \leq C + D\,.$$
If we choose $R \in \mathbb{N}$ such that $2(\frac{R}{C} - D) > C + D$, then $f(a)$ and $f(a')$ cannot be mapped to different sides of $Z \setminus B_{\frac{R}{C}-D}(f(z_n))$ and not in the ball. Thus, they are both mapped to the right side. Now, for any arbitrary $z > z_n + R$, we pick a sequence $(a_i)_{i=0}^{k}$, where $a_0 = a$, $d(a_i, a_{i+1}) = 1$ for every $i \in \{0, \dots, k-1\}$ and $a_k = z$. Then all $f(a_i)$ with $i \in \{0, \dots, k\}$, in particular $f(z)$, must be on the right side of $Z \setminus B_{\frac{R}{C}-D}(f(z_n))$, that is, $f(z) > f(z_n) + \frac{R}{C} - D$.

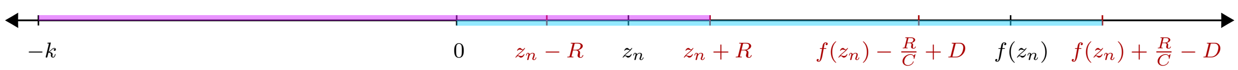

Figure 3.2.10

Hence, we ruled out all the elements outside of $[-k, z_n + R]$ to be mapped to $[0, f(z_n) + \frac{R}{C} - D]$, implying that $[0, f(z_n) + \frac{R}{C} - D] \subseteq f([-k, z_n + R])$. Thus, since f is bijective, we obtain

$$|[0, f(z_n) + \frac{R}{C} - D]| \leq |f([-k, z_n + R])|$$

and thus

$$f(z_n) + \frac{R}{C} - D + 1 \leq z_n + R + k + 1$$
$$f(z_n) - z_n \leq (1 - \frac{1}{C})R + k + D\,.$$

But now choose $n > (1 - \frac{1}{C})R + k + D$, then

$$(1 - \frac{1}{C})R + k + D < n \leq d(f(z_n), z_n) = f(z_n) - z_n \leq (1 - \frac{1}{C})R + k + D\,,$$

which is a contradiction. □

Thus, to conclude, given two RACGs W_1 and W_2 with one cylinder vertex with VA vertex group, and one class of e_1 and e_2 indistinguishable non-cylinder vertices, respectively, there is a QI between W_1 and W_2 and any such QI is bounded distance from scaling by $\frac{2m_1+n_1}{2m_2+n_2}$. □

If at the VA cylinder vertex, there attach $j \geq 2$ classes of indistinguishable neighbors we can apply Proposition 3.40 to each class individually to obtain the following generalization:

Corollary 3.41. *Let W_1 and W_2 be two RACGs on defining graphs satisfying the Standing Assumption 1.3 with the same structure invariant and one single cylinder vertex v_1 and v_2 in the JSJ graph of cylinders Σ_1 and Σ_2, respectively. Let the cylinder vertex groups V_1 and V_2 at v_1 and v_2, respectively, be VA. Suppose at both v_1 and v_2, there attach $j \geq 1$ classes of indistinguishable vertices, respectively, and let $e_{i,k} = m_{i,k} + n_{i,k}$ for $i \in \{1,2\}$ and $k \in \{1,\ldots,j\}$ denote the number of neighbors v_i in class k with $m_{i,k}$ the number of neighbors attaching along a D_∞-edge and $n_{i,k}$ the number of neighbors attaching along a $D_\infty \times \mathbb{Z}_2$-edge. If there is a QI between W_1 and W_2, then the ratio $\frac{2m_{1,k}+n_{1,k}}{2m_{2,k}+n_{2,k}}$ is the same for all $k \in \{1,\ldots,j\}$.*

Example 3.42. In Example 3.36 illustrated by Figure 3.2.7, we have $j = 2$ classes of indistinguishable tangling edges. Since all occurring edge groups are D_∞, we have $e_{i,1} = m_{i,1}$ counting the hanging vertices and $e_{i,2} = m_{i,2}$ counting the rigid. Then we have

$m_{i,k}$	$k=1$	$k=2$
$i=1$	1	1
$i=2$	2	1
ratio	$\frac{1}{2}$	1

and by Corollary 3.41 the VA cylinder vertices of W_1 and W_2 don't have the same relative QI-type. Thus, W_1 and W_2 are not QI by Proposition 2.24.

3.2.3 Refinement of the structure invariant

In Corollary 3.41, we have seen that the number of neighbors per class of indistinguishable vertices at a VA cylinder vertex in the JSJ graph of cylinders is an essential characteristic to determine whether or not two groups are QI. Thus, we aim to alter the structure invariant in a way such that this information is taken into account. For that purpose, we introduce a process we call *density refinement*.

Construction 3.43. We start with an initial decoration δ_0 with an initial set of ornaments consisting of the vertex and the relative QI-type. We perform the neighbor refinement, giving us a stable decoration δ_i.

Now, we define the map $\nu_i \colon V(T) \to \mathbb{N}^{\delta_i(V(T))}/{\sim} \cup \{\#\}$, where $\sim$ is an equivalence relation defined in Step 2 below, as follows:

- For any vertex $v \in V(T)$, whose vertex group is not VA, ν_i maps v to $\#$.
- A vertex $v \in V(T)$, whose vertex group is VA, is mapped to an equivalence class of tuples with entries in $\mathbb{N}$ indexed by the image of the decoration δ_i. We obtain the image $\nu_i(v)$ in two steps:
    1. We associate to v a tuple α obtained as follows: The entry indexed by $o \in \mathcal{O}_i$ is computed from the JSJ graph of cylinders Σ_c. We look at the neighbors of the vertex in Σ_c corresponding to the orbit of v with ornament o. Let m be the number of such neighbors attached along a D_∞-edge and n be the number of such neighbors attached along a $D_\infty \times \mathbb{Z}_2$-edge. Then the entry is $2\,m + n$.
    2. Define the image of v under ν_i as the *projective class* of α, that is, the equivalence class under the relation: $\alpha \sim \beta$ if and only if there is a $k \in \mathbb{R}^+$ such that $k \cdot \alpha = \beta$, where the multiplication $\cdot$ is defined coordinate-wise.

With the map ν_i, we provide a new decoration: The new set of ornaments is

$$\mathcal{O}'_i := \mathcal{O}_0 \times \mathbb{N}^{\delta_i(V(T))}/{\sim} \cup \{\#\} \times \overline{\mathbb{N}}^{\mathcal{O}_i} ,$$

with $\overline{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$. The decoration is $\delta'_i \colon T \to \mathcal{O}'_i$ with

$$\delta'_i(v) := (\delta_0(v), \nu_i(v), f_{v,i})$$

for any $v \in V(T)$. Possibly, δ'_i is a refinement of δ_i, and thus, we can perform the neighbor refinement on it. Again, we obtain a stable decoration δ_j for which we can define a map ν_j as above. We define a new set of ornaments $\mathcal{O}'_j := \mathcal{O}_0 \times \mathbb{N}^{\delta_j(V(T))}/{\sim} \cup \{\#\} \times \overline{\mathbb{N}}^{\mathcal{O}_j}$ and the decoration $\delta'_j \colon T \to \mathcal{O}'_j$ with $\delta'_j(v) := (\delta_0(v), \nu_j(v), f_{v,j})$ for any $v \in V(T)$. We repeat this alternating refinement process. Since there are only finitely many cylinder vertices in Σ_c, this process will eventually stabilize. The resulting decoration is the *density refinement* of δ_0.

Combining Proposition 2.24 and Corollary 3.41 yields that two RACGs can only be QI if their structure invariants, where δ_s is stable with respect to the density refinement, are identical.

Example 3.44. The original structure invariant for the group, illustrated in Figure 3.2.11, with respect to only the neighbor refinement is shown in the following table:

	vertex type	QI type	c_1 c_4 c_6	c_3	c_2 c_5	h_1 h_3 h_4	h_2	r
c_1, c_4, c_6	'cyl'	2-ended	0	0	0	0	0	1
c_3	'cyl'	2-ended	0	0	0	0	1	1
c_2, c_5	'cyl'	'VA'	0	0	0	∞	0	∞
h_1, h_3, h_4	'hang'	'VF'	0	0	∞	0	0	0
h_2	'hang'	'VF'	0	∞	0	0	0	0
r	'rig'		∞	∞	∞	0	0	0

We see that the vertices c_2 and c_5 are indistinguishable. However, when performing the density refinement according to Construction 3.43, the images of c_2 and c_5 under ν_i differ:

$$\nu_i(c_2) = [(0,0,0,2,0,2)] \quad \text{and} \quad \nu_i(c_5) = [(0,0,0,4,0,2)]\,.$$

This makes it possible to further distinguish h_1 from h_3 and h_4. We obtain the following refined structure invariant:

	vertex type	QI type	ν_{stable}	c_1 c_4 c_6	c_3	c_2	c_5	h_1	h_3 h_4	h_2	r
c_1, c_4, c_6	'cyl'	2-ended	#	0	0	0	0	0	0	0	1
c_3	'cyl'	2-ended	#	0	0	0	0	0	0	1	1
c_2	'cyl'	'VA'	$[(0,0,0,0,1,0,0,1)]$	0	0	0	0	∞	0	0	∞
c_5	'cyl'	'VA'	$[(0,0,0,0,0,2,0,1)]$	0	0	0	0	0	∞	0	∞
h_1	'hang'	'VF'	#	0	0	∞	0	0	0	0	0
h_3, h_4	'hang'	'VF'	#	0	0	0	∞	0	0	0	0
h_2	'hang'	'VF'	#	0	∞	0	0	0	0	0	0
r	'rig'		#	∞	∞	∞	∞	0	0	0	0

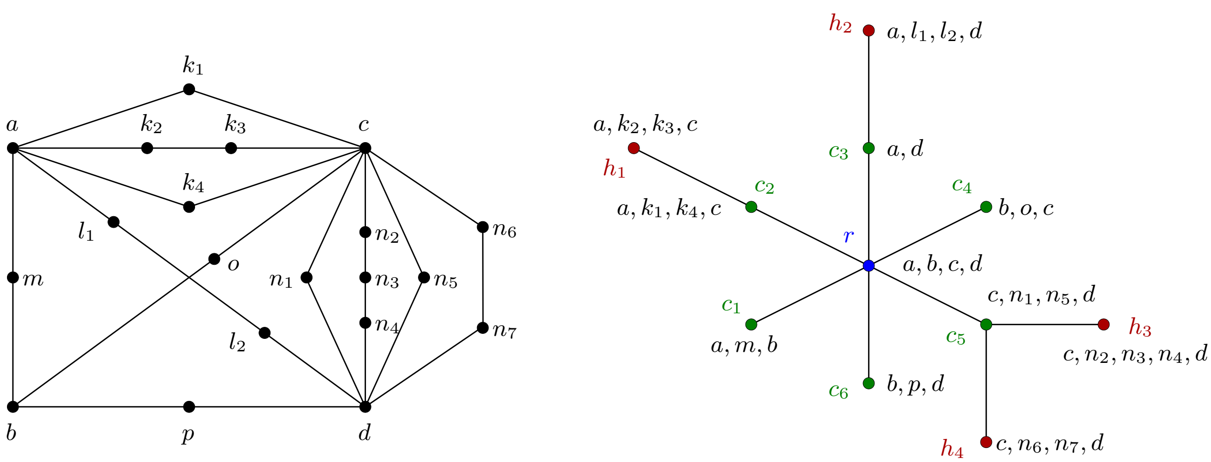

Γ $\qquad$ Λ_c

Figure 3.2.11

3.2.4 Complete QI-Invariant

Now, we aim to put the local QIs between cylinder vertex groups together to obtain a global QI between the groups, and thus, have a structure invariant which is a complete QI-invariant for certain groups. As mentioned in Remark 3.37, we exclude rigid vertices so the only missing piece are the local QIs between the hanging vertices. We see that we can choose them with a lot of flexibility:

Theorem 3.45. *Let W and W' be two finitely presented, one-ended RACGs with non-trivial JSJ decompositions over two-ended subgroups, which both have no rigid vertices. Define T to be the JSJ tree of cylinders of W and X to be the geometric tree of spaces of W over T. The initial decoration δ_0 on T takes vertex type and relative QI-type into account. Let δ be the density refinement of δ_0. Analogously, we define T', X', δ_0' and δ' for W'. Then W and W' are QI if and only if there is a bijection $\beta\colon \delta(T) \to \delta'(T')$ such that*

1. *$\delta_0 \circ \delta^{-1} = \delta_0' \circ (\delta')^{-1} \circ \beta$.*
2. *$S(T, \delta, \mathcal{O}) = S(T', \delta', \mathcal{O}')$ in the β-induced ordering.*
3. *For every ornament $o \in \mathcal{O}$, there is a vertex $v \in \delta^{-1}(o)$ and a vertex $v' \in (\delta')^{-1}(\beta(o))$ such that there is a QI between the vertex spaces X_v and $X'_{v'}$ respecting the decorations δ and δ' and which is bijective on the peripheral structures $\mathcal{P}_v$ and $\mathcal{P}'_{v'}$, respectively.*

Sketch of the Proof. This is an analogue of the proof of [CM17a, Theorem 7.5], with some generalizations and some specializations. The statement is more specialized in the two aspects laid out in Remark 3.37: We assume that the considered groups do not have any rigid vertices, thus, the relative stretch factors do not apply. Moreover, since we restrict to RACGs, partial orientations can be omitted. However, we do not assume the cylinder vertex groups to be two-ended, which makes the statement more general.

The idea is to inductively build a tree isometry $\chi\colon T \to T'$, which respects the decorations by using the local vertex QIs $\phi_v\colon X_v \to X'_{\chi(v)}$ bijective on the respective peripheral structures inducing χ on the link of the vertex $v \in V(T)$. Then χ induces a global QI.

For the base case, we pick some cylinder vertex $c \in V(T)$ and some $c' \in (\delta')^{-1}(\beta(\delta(c)))$ and define $\chi(c) := c'$. Because the initial decoration depends on the relative QI-type, there is a QI between X_c and $X'_{c'}$. Depending on whether c has a two-ended, a VFD or a VA vertex group, we pick such a QI $\phi_c : X_c \to X'_{c'}$ according to Propositions 3.38, 3.39 and 3.40, respectively. By construction, ϕ_c will be bijective on the respective peripheral structures and thus defines how to pick the bijection between the edge spaces incident to c. Thus, we can extend χ to the link of c according to this bijection.

Since the considered trees are bipartite, the inductive step consists of two parts: First we extend χ to a hanging vertex and from there we extend χ to a cylinder vertex.

Suppose there is an edge $e_1 \in E(T)$ such that $o(e_1)$ is a cylinder vertex, $\tau(e_1) =: h$ is a hanging vertex and $\chi(o(e_1))$ is already defined. Then there is a QI $\phi_{o(e_1)}\colon X_{o(e_1)} \to X'_{\chi(o(e_1))}$ respecting the decorations and bijective on the respective peripheral structures. Thus, $\phi_{o(e_1)}|_{X_{e_1}} : X_{e_1} \to X'_{\chi(e_1)}$ defines the QI on X_{e_1}. The QI on X_h can now be produced as suggested in [CM17a, Proposition 7.1], which is guided by [BN08, Theorem 1.2]. The key feature is the following: Pick for any other edge e adjacent to h some real constant σ_e. The only condition is that for all edges in the same orbit the constant needs to be identical. Then we can choose a QI $\phi_h : X_h \to X'_{\chi(h)}$ such that when restricted to X_{e_1} it matches $\phi_{o(e_1)}|_{X_{e_1}}$ and when restricted to X_e for any other e adjacent to h, this $\phi_h|_{X_e}$ is a QI with multiplicative constant σ_e.

Of course, we do not pick the σ_e arbitrarily, but we choose them among the set Σ of multiplicative constants occurring in the QIs produced by the Propositions 3.38, 3.39 and 3.40. Since there are only finitely many orbits of cylinder vertices, this set Σ is finite, and thus, we only pick a finite configuration of σ_e's from Σ. If we later see that our choice of configuration conflicts with the constants forced by the QIs of the adjacent cylinder vertices, we return to h and pick a different configuration. Since the number of such different configurations is finite, we know that eventually we have found the correct QI and extend χ to the link of h accordingly. Thus, without loss of generality, we can assume that we have picked a suitable QI at h satisfying all requirements.

Suppose now that $e_2 \in E(T)$ is an edge such that $o(e_2)$ is a hanging vertex, $\tau(e_2) = c_2$ is a cylinder vertex and $\chi(o(e_2))$ is already defined in the previous step. We repeat the extension process: We know, there is a QI $\phi_{o(e_2)} \colon X_{o(e_2)} \to X'_{\chi(o(e_2))}$ respecting decorations and bijective on the respective peripheral structures, which restricts to a QI on X_{e_2}. We can now extend χ to the link of c_2 and define $\phi_{c_2} \colon X_{c_2} \to X_{\chi(c_2)}$ according to the Propositions 3.38, 3.39 and 3.40 such that it agrees with $\phi_{o(e_2)}$ on X_{e_2}. □

Example 3.46. We make the introductory example of the two groups with defining graphs illustrated in Figure 3.2.12 explicit. By Proposition 3.39, we see that the VFD cylinder vertices coming from the blue uncrossed cut pairs are QI. In both cases, there is one hanging vertex group generated by the l_i's, thus, by Proposition 3.40, the VA cylinder vertices coming from the red uncrossed cut pairs are QI. Hence, the structure invariants are identical and Theorem 3.45 implies that the groups are QI.

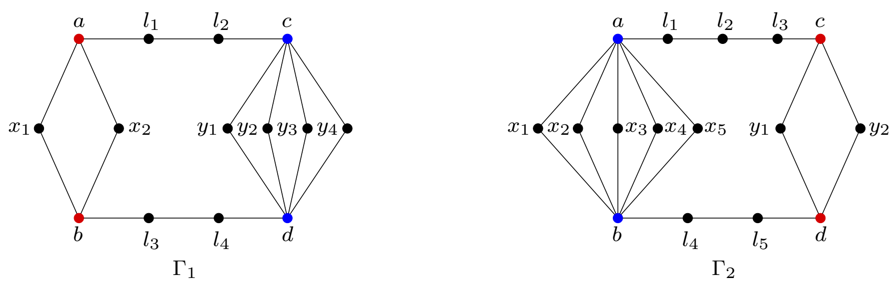

Figure 3.2.12

Remark 3.47. As discussed in Remark 3.37, Theorem 3.45 excludes groups whose JSJ decompositions have rigid vertices. However, in certain cases we can add another induction step to the proof of Theorem 3.45, handling rigid vertices following again the proof of Theorem 7.5 of [CM17a]. For instance, we can consider the subgraphs Σ and Σ' of the graphs of cylinders Σ_c and Σ'_c, respectively, which consist of one rigid vertex and all its adjacent cylinder vertices. If there is a decoration preserving graph isomorphism ϕ between Σ and Σ' and in addition, every vertex and edge group G_t is isomorphic to the image vertex group $G_{\phi(t)}$, then the induction extends also to these rigid vertices. The obvious method to produce such an example is to simply use identical defining graphs for corresponding special subgroups. This is illustrated in Example 3.50.

Alternatively, if the rigid vertices are virtually free, they are quasi-isometrically rigid relative to the peripheral structure by [CM11], and thus, relative stretch factors can be used as introduced in Section 4 of [CM17a].

Outline 3.48. Theorem 3.45 illustrates the flexibility we have to change the defining graph in a way such that the group on the resulting graph is QI to the one on the original graph. The changes happen at the cylinder vertices:

- At a virtually cyclic cylinder vertex coming from the uncrossed cut collection $\{a - b\}$ we can only remove or add a common adjacent vertex such that $|\mathcal{C}| \in \{0, 1\}$ is maintained, since the valencies in the JSJ tree of cylinders need to be preserved.
- At a VFD cylinder vertex coming from the uncrossed cut collection $\{a - b\}$ and its common adjacent vertices $\mathcal{C}$, we can duplicate or remove tangling pieces in the JSJ tree of cylinders that are equivalent up to QI (cf. Algorithm 3.39.3). In the defining graph Γ this corresponds to duplicating or removing any connected component of $\Gamma \setminus \{a - b\}$ disjoint from $\mathcal{C}$, and reattaching the new collection of pieces to a and b. Note that in the reattaching the roles of a and b can be interchanged, thus, this move can be interpreted as a reflection along the subgraph on $\{a, b\} \cup \mathcal{C}$. Additionally the number of vertices in $\mathcal{C}$ can be changed. The only restriction is that $|\mathcal{C}| > 2$.
  Also, within a connected component containing vertices of a set A contributing to a hanging vertex, the number of vertices can be altered while preserving the virtually free QI-type. That means, we can add or remove elements on a branch, as long as the resulting vertex set A still produces a hanging vertex. Thus, by Proposition 3.20, the altered set A should still satisfy conditions (A1), (A2) and (A3) and W_A has to be infinite and not a cylinder vertex group.
- At a VA cylinder vertex coming from the uncrossed cut collection $\{a - b\}$ and its common adjacent vertices $\mathcal{C}$ we can perform changes similar to the ones at VFD cylinder vertices. There are only two differences: We perform the duplication or removal of pieces with a fixed ratio and the number of common adjacent vertices has to stay fixed $|\mathcal{C}| = 2$.

These observations can be used as a method to produce examples of QI RACGs:

Example 3.49. The RACGs on the defining graphs Γ_1 and Γ'_1 with JSJ graphs of cylinders $\Sigma_{c,1}$ and $\Sigma'_{c,1}$ respectively, illustrated in Figure 3.2.13, are QI by Theorem 3.45.

Example 3.50. The RACGs on the defining graphs Γ_2 and Γ'_2 with JSJ graphs of cylinders $\Sigma_{c,2}$ and $\Sigma'_{c,2}$, respectively, illustrated in Figure 3.2.14, are QI by Theorem 3.45 and Remark 3.47.

Remark 3.51. It would be most interesting to produce QIs that do not arise from algebraic considerations. One might guess that a simple graph operation like duplicating the complement of a subgroup corresponding to a cylinder vertex would produce either a group which is a finite index subgroup of the original one or at least produce a group which shares a common finite index subgroup with it, implying that the groups are commensurable and thus QI. However, our construction has much more flexibility than that.

Only partial commensurability results are known, such as the commensurability classification for certain hyperbolic RACGs done by Dani, Stark and Thomas in [DST18], see Section 1.3.1. Their proof is not applicable to our more general setting, as it depends on the fact that the finite valence of cylinder vertices in the JSJ tree of cylinders of hyperbolic RACGs is a QI-invariant. This tool is lost for non-hyperbolic RACGs. In [HST20, Section 4], Hruska, Stark and Tran provide examples of commensurable non-hyperbolic RACGs, whose defining graphs are generalized Θ-graphs. However, a complete classification for some class of non-hyperbolic RACGs is yet to be stated and should be addressed separately. Nonetheless, we can show that the non-hyperbolic Examples 3.49 and 3.50, for which we produced QIs with our methods, are not abstractly commensurable, by application of the following Lemma 3.52, which is guided by Lemma 7.2 of Shepherd and Woodhouse [SW22].

Lemma 3.52. *The two RACGs W_1 and W_1' in Example 3.49 on the defining graphs Γ_1 and Γ_1' in Figure 3.2.13 are not commensurable to each other and the two RACGs W_2 and W_2' in Example 3.50 on the defining graphs Γ_2 and Γ_2' in Figure 3.2.14 are not commensurable to each other.*

Proof. [cf. SW22, Lemma 7.2] Let W and W' be two RACGs whose JSJ graphs of cylinders have cylinder vertices v and v' with vertex groups $W_{\mathcal{C}} \times D_\infty$ and $W_{\mathcal{C}'} \times D_\infty$, respectively, such that $W_{\mathcal{C}}$ and $W_{\mathcal{C}'}$ are both virtually free, that is, $|\mathcal{C}|, |\mathcal{C}'| > 2$. In fact, given $\mathcal{C} = \{c_1, \ldots, c_{i+1}\}$, as per the proof of Theorem B.1 of Cashen, Dani and Thomas in [DT17, Appendix B], $W_{\mathcal{C}}$ has a free subgroup F_i generated by $\langle c_1c_2, \ldots, c_1c_{i+1}\rangle$ of rank i and index 2. Analogously, $W_{\mathcal{C}'}$ has a free subgroup F_j of index 2 and rank $|\mathcal{C}'| - 1 =: j$.

Suppose that W and W' are commensurable, that is, they have isomorphic finite index subgroups. By [GL17, Corollary 7.4], we can assume that the induced JSJ graphs of cylinders of these subgroups are identical. Call this induced JSJ graph of cylinders $\hat{\Gamma}$ with fundamental group $\hat{W}$. The idea is now to compute the degree of a vertex in $\hat{\Gamma}$, using first W and then W', and obtain a contradiction for the groups we are interested in as the computed degrees cannot match.

Suppose there is a $\hat{v} \in V(\hat{\Gamma})$ with vertex group $\hat{G}_{\hat{v}}$ covering v and v'. Then we can embed $\hat{G}_{\hat{v}}$ into both $W_{\mathcal{C}} \times D_\infty$ and $W_{\mathcal{C}'} \times D_\infty$ as a finite index subgroup. Note that such a vertex $\hat{v}$ exists in both the examples we consider here: $\Sigma_{c,1}$ has only one VFD cylinder vertex, the vertex c_2. Thus, any vertex $\hat{v}$ covering some VFD cylinder vertex in $\Sigma'_{c,1}$ has to cover c_2 as well. This argument works also for $\Sigma_{c,2}$ with its only VFD vertex c_3 and $\Sigma'_{c,2}$.

Moreover, in the considered examples all edge groups are the same D_∞ generated by the cut pair. Hence, the number of edges incident to $\hat{v}$ corresponds to the number of double cosets $\hat{G}_{\hat{v}}gD_\infty$ with g an element in the cylinder vertex group, multiplied by the degree of the cylinder vertex. So, we aim to compute $\deg(\hat{v})$ in two ways, first via v, then via v':

$$\begin{aligned}\deg(\hat{v}) &= |\{\hat{G}_{\hat{v}}gD_\infty \mid g \in W_{\mathcal{C}} \times D_\infty\}| \cdot \deg(v)\\ &= |\{\hat{G}_{\hat{v}}gD_\infty \mid g \in W_{\mathcal{C}'} \times D_\infty\}| \cdot \deg(v')\end{aligned}$$

In order to do this, we consider $F_i \times D_\infty \leq W_{\mathcal{C}} \times D_\infty$. This is a subgroup of index 2, thus, the intersection $G_i := F_i \times D_\infty \cap \hat{G}_{\hat{v}} \leq \hat{G}_{\hat{v}}$ is at most of index 2 in $\hat{G}_{\hat{v}}$. For $F_j \times D_\infty \leq W_{\mathcal{C}'} \times D_\infty$, we define analogously $G_j := F_j \times D_\infty \cap \hat{G}_{\hat{v}} \leq \hat{G}_{\hat{v}}$, which is also at most of index 2 in $\hat{G}_{\hat{v}}$. Hence, we have for the intersection $G := G_i \cap G_j$

$$|\hat{G}_{\hat{v}} : G| \leq |\hat{G}_{\hat{v}} : G_i|\,|\hat{G}_{\hat{v}} : G_j| \leq 2 \cdot 2 = 4\,.$$

Now, we decompose the double cosets $\hat{G}_{\hat{v}}gD_\infty$ further into double cosets of G with representatives f_i in $F_i \times D_\infty$. Since D_∞ is central, it suffices to consider $\hat{G}_{\hat{v}}g$: Each such coset $\hat{G}_{\hat{v}}g$ consists of at most 8 cosets of the form Gf_i. Indeed, at most 4 cosets come from the partition of $\hat{G}_{\hat{v}}$ into G-cosets as the index of G in $\hat{G}_{\hat{v}}$ is at most 4 and then we multiply by 2, because $F_i \times D_\infty$ is of index 2 in $W_{\mathcal{C}} \times D_\infty$. This bounds the number of double cosets $\hat{G}_{\hat{v}}gD_\infty$ by:

$$\frac{1}{8}\,|\{Gf_iD_\infty \mid f_i \in F_i \times D_\infty\}| \leq |\{\hat{G}_{\hat{v}}gD_\infty \mid g \in W_{\mathcal{C}} \times D_\infty\}| \leq 2\,|\{Gf_iD_\infty \mid f_i \in F_i \times D_\infty\}|\,.$$

Let $\pi_i : F_i \times D_\infty \to F_i$ be the projection map. Then the image $\pi_i(G)$ is a subgroup of F_i and thus free. This implies that the short exact sequence

$$1 \to G \cap \ker(\pi_i) \to G \to \pi_i(G) \to 1$$

splits, that is, there is a section $\sigma_i : \pi_i(G) \to G$ with image P_i isomorphic to $\pi_i(G)$. But since D_∞ is central in $F_i \times D_\infty$, we know that $G = P_i \times (G \cap \ker(\pi_i))$. Thus, the number of double cosets Gf_iD_∞ is equal to the number of cosets $\pi_i(G)\pi_i(f_i) \times D_\infty$ in $F_i \times D_\infty$. But this number is the index of $\pi_i(G)$ in F_i, which we compute by the Schreier-index formula

$$|F_i : \pi_i(G)| = \frac{\mathrm{rk}(\pi_i(G)) - 1}{i - 1}\,.$$

Analogously, we perform the same argument for $F_j \times D_\infty$ and the projection map $\pi_j : F_j \times D_\infty \to F_j$ to compute the number of double cosets Gf_jD_∞ with $f_j \in F_j \times D_\infty$ via

$$|F_j : \pi_j(G)| = \frac{\mathrm{rk}(\pi_j(G)) - 1}{j - 1}\,.$$

However, we note that

$$\mathrm{rk}(\pi_i(G)) = \mathrm{rk}(G/\ker(\pi_i)) = \mathrm{rk}(G/\mathrm{Z}(G)) = \mathrm{rk}(G/\ker(\pi_j)) = \mathrm{rk}(\pi_j(G))\,,$$

that is, both occurring ranks are identical, call them r. Thus, when computing $\deg(\hat{v})$ via v, we can use the first computation to obtain the bound

$$\frac{1}{8}\,\frac{r-1}{i-1} \cdot \deg(v) \leq \deg(\hat{v}) \leq 2\,\frac{r-1}{i-1} \cdot \deg(v)\,.$$

When computing $\deg(\hat{v})$ via v', we obtain using the second computation

$$\frac{1}{8}\,\frac{r-1}{j-1} \cdot \deg(v') \leq \deg(\hat{v}) \leq 2\,\frac{r-1}{j-1} \cdot \deg(v')\,.$$

This implies that we arrive at a contradiction, whenever

$$2\,\frac{r-1}{j-1} \cdot \deg(v') < \frac{1}{8}\,\frac{r-1}{i-1} \cdot \deg(v)\,,$$

that is, whenever

$$j > 16 \cdot (i-1) \cdot \frac{\deg(v')}{\deg(v)} + 1\,.$$

In case of Example 3.49 this inequality is satisfied for the two VFD cylinder vertices labelled c_2 and c_2': In Γ_1, we have $|\mathcal{C}| = 3$, thus $i = 2$ and $\deg(c_2) = 1$. In Γ_1' we have $|\mathcal{C}'| = 35$, thus $j = 34$ and $\deg(c_2') = 2$. In Example 3.50, the condition is satisfied for the vertices labelled c_3. Hence, W_1 and W_1' in Example 3.49 and W_2 and W_2' in Example 3.50, respectively, are not commensurable. □

Remark 3.53. The proof of Lemma 3.52 works for various other examples. In fact, it can even provide a more sensitive commensurability invariant. Recall that the argument involves computing for an edge e with edge group D_∞ at the vertex v the number of cosets $|\{\hat{G}_{\hat{v}} g D_\infty \mid g \in W_{\mathcal{C}} \times D_\infty\}|$. Then we sum over all such edges e, which is the same as multiplying by the degree $\deg(v)$ of v.

However, instead of summing over all edges e incident to v, we can restrict to a certain subclass of edges. For example, we can restrict to edges whose incident vertices share the same vertex types, because the vertex types of the incident vertices of a covering edge must be the same as the ones of the covered edge. Even finer than just considering the vertex type would be to restrict to edges with incident vertices sharing a particular decoration which has to be preserved by the covering.

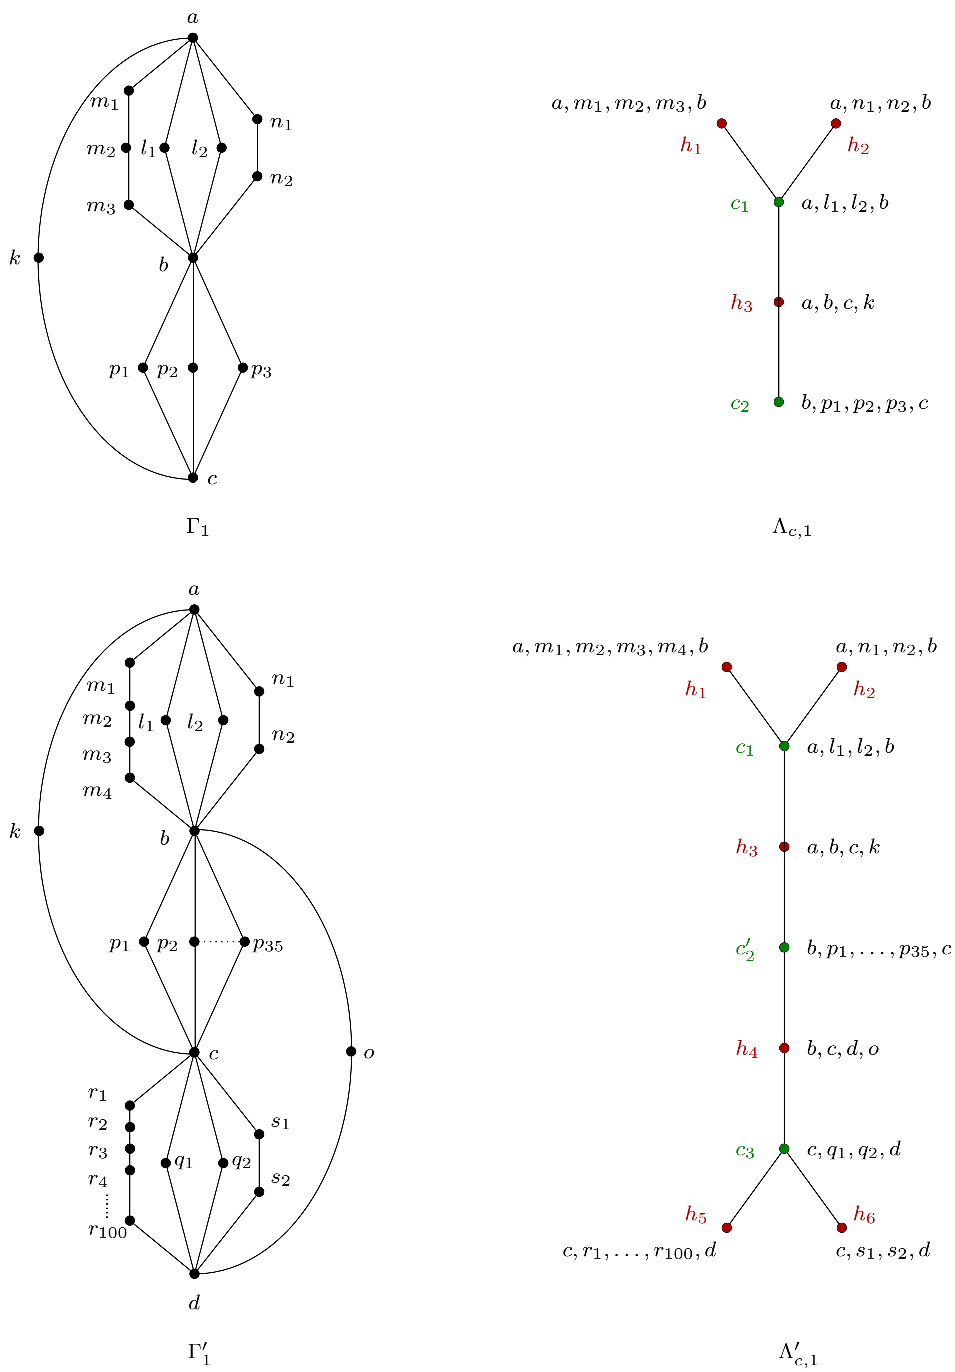

a
m1
m2
m3
l1
l2
n1
n2
k
b
p1
p2
p3
c
Γ1
a, m1, m2, m3, b
a, n1, n2, b
h1
h2
c1
a, l1, l2, b
h3
a, b, c, k
c2
b, p1, p2, p3, c
Λc,1
a
m1
m2
m3
m4
l1
l2
n1
n2
k
b
p1
p2
p35
c
o
r1
r2
r3
r4
r100
q1
q2
s1
s2
d
Γ'1
a, m1, m2, m3, m4, b
a, n1, n2, b
h1
h2
c1
a, l1, l2, b
h3
a, b, c, k
c'2
b, p1, . . . , p35, c
h4
b, c, d, o
c3
c, q1, q2, d
h5
h6
c, r1, . . . , r100, d
c, s1, s2, d
Λ'c,1

Figure 3.2.13

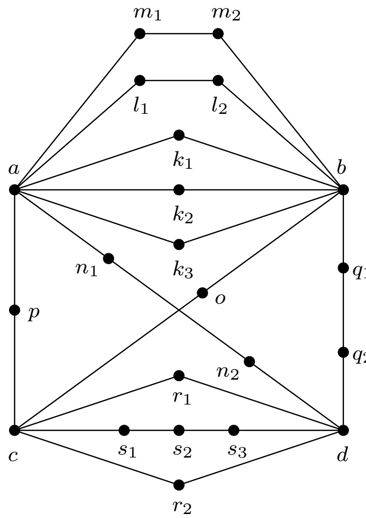

Γ_2

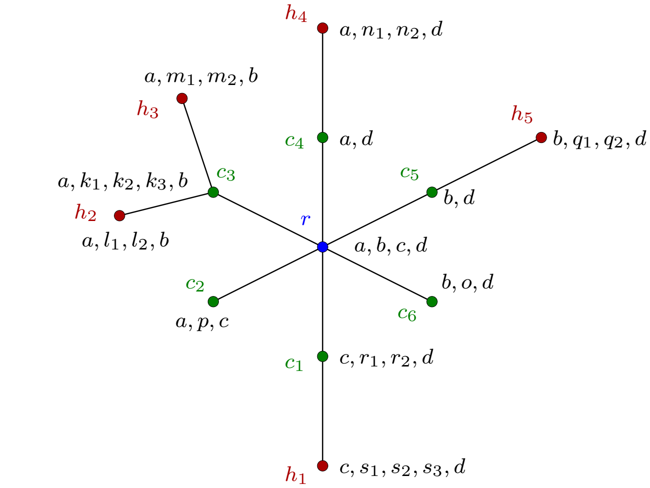

$\Lambda_{c,2}$

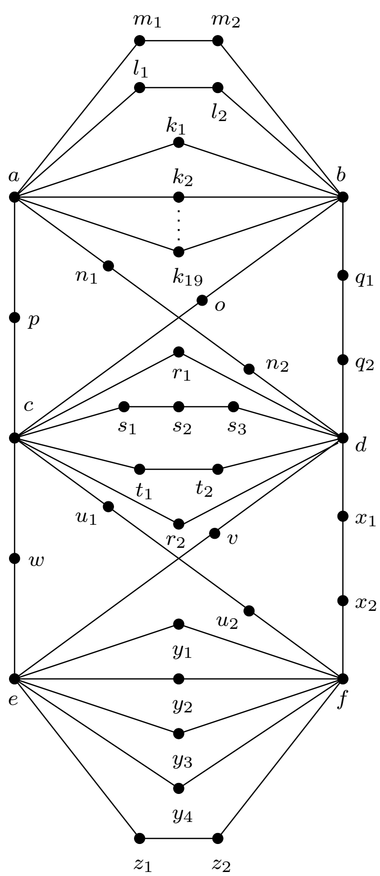

Γ'_2

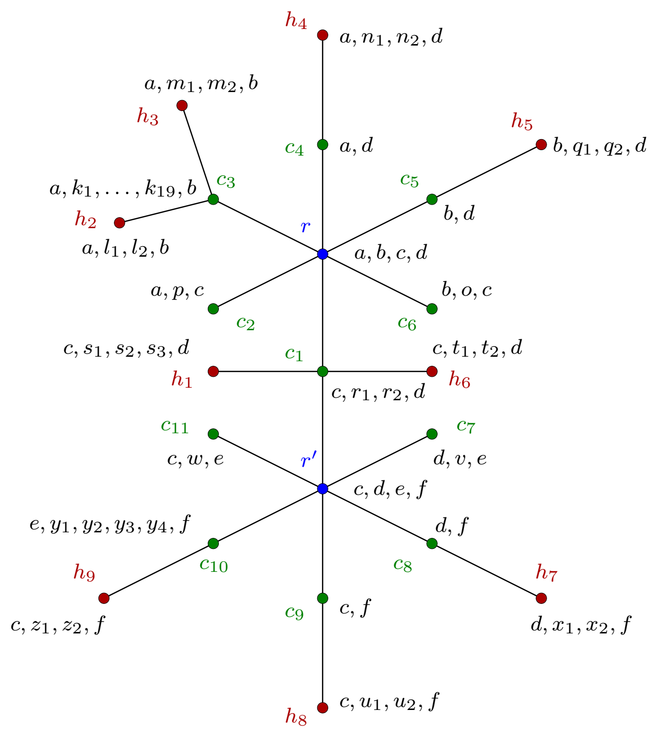

$\Lambda'_{c,2}$

Figure 3.2.14

4 QIs between RACGs and RAAGs

4.1 DL-Algorithm

As outlined in the introductory Section 1.4, one idea to find RAAGs QI to a given RACG is by finding finite index RAAG subgroups. In a RACG W_Γ, every pair of vertices $a, b \in V(\Gamma)$ which is not connected by an edge generates an infinite dihedral subgroup $W_{\{a,b\}} \cong D_\infty$. Thus, the infinite order element $ab \in W_\Gamma$ generates an infinite cyclic subgroup $\langle ab \rangle \leq W_{\{a,b\}} \leq W_\Gamma$ of index 2 in $W_{\{a,b\}}$.

Based on initial results of LaForge in [LaF17], Dani-Levcovitz investigate in [DL20] whether and how we can choose a collection of such infinite order elements as generators of a RAAG and generate a finite index RAAG subgroup by them. In this section, we introduce their main graph theoretical algorithm on Γ, the *DL-Algorithm*, and further develop its range of applications.

Since the infinite order elements we aim to choose as generators for the RAAG correspond to missing edges in the defining graph Γ, we can visualize them by introducing an induced subgraph of the complement graph Γ^c of Γ:

Definition 4.1. Let W_Γ be a RACG and let $\Lambda \leq \Gamma^c$ be an induced subgraph of the complement graph Γ^c of Γ.

- The *commuting graph* Δ *associated to* Λ has a vertex v_{ab} for every edge $(a, b) \in E(\Lambda)$ and two vertices $v_{a_1b_1}$ and $v_{a_2b_2}$ in Δ are connected by an edge if the elements a_1b_1 and a_2b_2 in W_Γ commute.
- The graph $\Theta = \Theta(\Gamma, \Lambda)$ is the graph on the vertex set $V(\Gamma)$ and with edge set $E(\Gamma) \cup E(\Lambda)$. We depict $\Theta(\Gamma, \Lambda)$ with the edges $E(\Gamma)$ in black and $E(\Lambda)$ in color.

Remark 4.2. Observe the following about a commuting graph Δ associated to $\Lambda \leq \Gamma^c$:

(i) Since all graphs, in particular Λ, are undirected, (a, b) and (b, a) refer to the same edge in $E(\Lambda)$. Hence, the vertices v_{ab} and v_{ba} in $V(\Delta)$ are identical. This is consistent with the fact that $ba = (ab)^{-1}$, so ab and ba generate the same subgroup isomorphic to $\mathbb{Z}$. The edge set is still well-defined, because a_1b_1 commutes with a_2b_2 if and only if b_1a_1 commutes with a_2b_2.

(ii) For $a_1, a_2, b_1, b_2 \in V(\Gamma)$, the elements a_1b_1 and a_2b_2 in the RACG W_Γ commute if and only if (a_1, a_2, b_1, b_2) is a square in Γ.

(iii) Define a map $\phi_\Lambda \colon A_\Delta \to W_\Gamma$ by sending a generator $v_{ab} \in V(\Delta)$ to the infinite order element $ab \in W_\Gamma$. Two generators $v_{a_1b_1}, v_{a_2b_2} \in V(\Delta)$ commute if they are adjacent, which by definition of Δ is the case if the elements a_1b_1 and a_2b_2 commute in W_Γ. So, ϕ_Λ is a homomorphism. The difficult part is to determine when ϕ_Λ is injective.

Definition 4.3. Let W_Γ be a RACG with defining graph Γ, induced subgraph $\Lambda \leq \Gamma^c$ of the complement graph Γ^c of Γ and commuting graph Δ associated to Λ. The RAAG A_Δ is called the *visual RAAG subgroup* of W_Γ associated to Λ if the following homomorphism is injective:

$$\begin{array}{rccc} \phi_\Lambda \colon & A_\Delta & \to & W_\Gamma \\ & v_{a,b} & \mapsto & ab\,. \end{array}$$

If Λ has at most two connected components, Dani-Levcovitz give in [DL20] four graph theoretical conditions on $\Theta(\Gamma, \Lambda)$, the *subgroup conditions* to determine that the RAAG A_Δ on the commuting graph Δ associated to Λ is a visual RAAG subgroup. Moreover, they give two conditions, the *index conditions*, to ensure that the visual RAAG subgroup is of finite index in W_Γ. We introduce these conditions and the results using them in this section.

The following two subgroup conditions were first introduced by LaForge in [LaF17]:

Definition 4.4. [DL20, cf. Definition 3.1] The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{R}_1$ if Λ does not contain a cycle.

Definition 4.5. [DL20, cf. Definition 3.3] The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{R}_2$ if each component of Λ is an induced subgraph of Θ.

An example illustrates that conditions $\mathcal{R}_1$ and $\mathcal{R}_2$ are necessary for the injectivity of ϕ_Λ:

Example 4.6. Let W_Γ be the RACG on the defining graph Γ shown in Figure 4.1.1 with induced subgraphs $\Lambda_1, \Lambda_2 \leq \Gamma^c$ of the complement Γ^c of Γ and associated commuting graphs Δ_1 and Δ_2.

- The graph $\Theta(\Gamma, \Lambda_1)$ does not satisfy condition $\mathcal{R}_1$ and indeed the map ϕ_{Λ_1} is not injective: Since A_{Δ_1} is isomorphic to $\mathbb{F}_3$, the element $v_{ab}v_{bc}v_{ca} \in A_{\Delta_1} \setminus \{1_{A_{\Delta_1}}\}$ is non-trivial, but

$$\phi_{\Lambda_1}(v_{ab}v_{bc}v_{ca}) = abbcca = aa = 1_{W_\Gamma} .$$

- The graph $\Theta(\Gamma, \Lambda_2)$ does not satisfy condition $\mathcal{R}_2$ and indeed the map ϕ_{Λ_2} is not injective: Since A_{Δ_2} is isomorphic to $\mathbb{F}_2$, the element $v_{ac}v_{cm_1} \in A_{\Delta_2} \setminus \{1_{A_{\Delta_2}}\}$ is non-trivial, but

$$\phi_{\Lambda_2}(v_{ac}v_{cm_1}v_{ac}v_{cm_1}) = accm_1accm_1 = am_1am_1 = aam_1m_1 = 1_{W_\Gamma} .$$

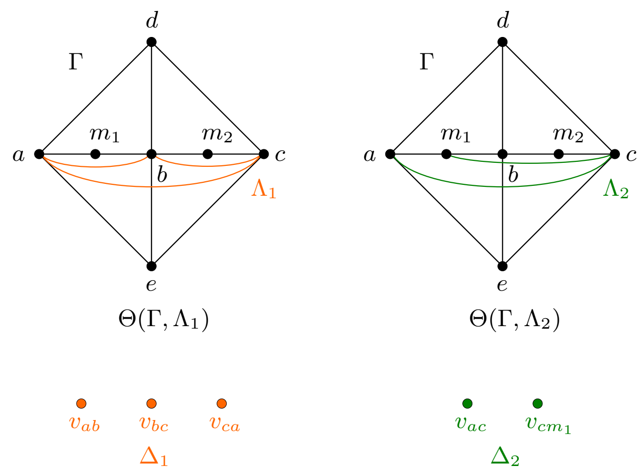

Figure 4.1.1

To define the other two subgroup conditions, we need some graph theoretical terminology:

Definition 4.7. [DL20, cf. Definitions 3.6 and 3.7] Let W_Γ be a RACG with defining graph Γ, induced subgraph $\Lambda \leq \Gamma^c$ of the complement Γ^c of Γ and associated graph $\Theta(\Gamma, \Lambda)$ and let Λ_c and Λ_d be two distinct connected components of Λ.

- A path $\gamma \subseteq \Theta$ is called Γ-*path* or Λ-*path* if all edges in γ are in Γ or Λ, respectively.
- A Γ-path γ is called a *2-component path* if $\gamma = (c_1, d_1, c_2, d_2, \ldots, c_n, d_n)$, where $\{c_1, \ldots, c_n\} \in V(\Lambda_c)$ and $\{d_1, \ldots, d_n\} \in V(\Lambda_d)$ for $n \geq 1$.
- A 2-component path γ is called *2-component cycle* if c_1 and d_n are connected by an edge $(c_1, d_n) \in E(\Theta)$ and a 2-component cycle is called *2-component square* if $n = 2$.
- For vertices $c_1, c_2, \ldots, c_n \in V(\Lambda_c)$, the Λ-*convex hull* $T_\Lambda(c_1, c_2, \ldots, c_n)$ of $c_1, c_2, \ldots, c_n$ is the minimal induced subgraph of Λ_c containing all vertices $c_1, c_2, \ldots, c_n$.

Definition 4.8. [DL20, cf. Definition 3.3] Let Λ_c and Λ_d be two distinct components of Λ. The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{R}_3$ if every 2-component square (c_1, d_1, c_2, d_2) in Θ with $c_1, c_2 \in \Lambda_c$ and $d_1, d_2 \in \Lambda_d$ satisfies the following condition: The graph Γ contains the join of $V(T_\Lambda(c_1, c_2))$ and $V(T_\Lambda(d_1, d_2))$ as a subgraph.

Definition 4.9. [DL20, cf. Definition 3.14] Let Λ_c and Λ_d be two distinct components of Λ. The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{R}_4$ if for every 2-component cycle $\gamma = (c_1, d_1, c_2, d_2, \ldots, c_n, d_n)$ in Θ with $\{c_1, \ldots, c_n\} \in V(\Lambda_c)$, $\{d_1, \ldots, d_n\} \in V(\Lambda_d)$ and $n \geq 2$ satisfies the following condition: Every edge of γ is contained in a 2-component square (t_c, t_d, t'_c, t'_d) of Θ with $t_c, t'_c \in V(T_\Lambda(c_1, \ldots, c_n))$ and $t_d, t'_d \in V(T_\Lambda(d_1, \ldots, d_n))$.

The necessity of these subgroup conditions is less obvious and well illustrated in Examples 3.5 and 3.13 of [DL20], so we refrain from giving an example.

On the contrary, the following two index conditions are rather straightforward:

Definition 4.10. [DL20, cf. Definition 4.2] The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{F}_1$ if given any $s \in V(\Gamma)$ such that s is not a cone vertex, s is contained in some edge $(s, t) \in E(\Lambda)$.

Definition 4.11. [DL20, cf. Definition 4.2] Let Λ_s and Λ_t be two distinct components of Λ. The graph $\Theta(\Gamma, \Lambda)$ satisfies *condition* $\mathcal{F}_2$ if given any $s \in V(\Lambda_s)$ and $t \in \Lambda_t$, then there is a 2-component path $(s, t_1, s_2, t_2, \ldots, s_n, t)$ in Θ with $\{s_2, \ldots, s_n\} \subseteq V(\Lambda_s)$, $\{t_1, \ldots, t_{n-1}\} \subseteq V(\Lambda_t)$ and some $n \in \mathbb{N}$.

It turns out that these subgroup and index conditions suffice to determine RAAG subgroups of W_Γ of index 2 or 4:

Theorem 4.12. [DL20, Theorem C] *Let W_Γ be a RACG satisfying Standing Assumption 1 and let $\Lambda \leq \Gamma^c$ be an induced subgraph of the complement Γ^c of Γ with no isolated vertex, at most two components and associated commuting graph Δ. Then A_Δ is a visual RAAG subgroup of W_Γ if and only if $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$-$\mathcal{R}_4$.*

Theorem 4.13. [DL20, Theorem A] *Let W_Γ be a RACG satisfying Standing Assumption 1 and let $\Lambda \leq \Gamma^c$ be an induced subgraph of the complement Γ^c of Γ with no isolated vertex and associated commuting graph Δ. Then the following are equivalent:*

(1) A_Δ is a finite index visual RAAG subgroup of W_Γ.
(2) A_Δ is a visual RAAG subgroup of W_Γ and has index 2 or 4 in W_Γ.
(3) Λ has at most two components and $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$-$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$.

We refer to a $\Lambda \leq \Gamma^c$ providing a visual RAAG subgroup and to the search for such a Λ, which has 2-components by Theorems 4.12 and 4.13, as follows:

Definition 4.14. For a RACG W_Γ with defining graph Γ, an induced subgraph $\Lambda \leq \Gamma^c$ of the complement graph Γ^c of Γ is called

- *Dani-Levcovitz-*Λ or *DL-*Λ if $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$-$\mathcal{R}_4$.
- *finite index Dani-Levcovitz-*Λ or *FIDL-*Λ if $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$-$\mathcal{R}_4$, $\mathcal{F}_1$ and $\mathcal{F}_2$.

We call the procedure of running through all the induced subgraphs of Γ^c and checking whether one of them is a 2-component FIDL-Λ the *Dani-Levcovitz-Algorithm* or *DL-Algorithm*.

Remark 4.15.

- In general, the choice of a $\Lambda \leq \Gamma^c$ is not unique, usually the DL-Algorithm generates more than one visual RAAG subgroup.
- The input Γ for the DL-Algorithm does not have to be a $\mathcal{CFS}$ graph, but if it satisfies Standing Assumption 1, the existence of a 2-component FIDL-Λ implies by Theorem 4.13 that Γ is $\mathcal{CFS}$.
- The original statement of Theorem 4.13 in [DL20] is slightly more general. Item (2) of Standing Assumption 1 that W_Γ is one-ended is replaced by the statement that Γ does not have any isolated vertex. However, if we assume that W_Γ is one-ended, we can conclude that if two vertices v_1 and v_2 are adjacent in a FIDL-Λ they are the diagonal of a square in Γ. Indeed, suppose they are not, then the vertex $v_{v_1 v_2} \in V(\Delta)$ is isolated, implying that the corresponding RAAG A_Δ is infinitely ended. This cannot happen if the RACG W_Γ it is QI to is one-ended.

Remark 4.16. As mentioned in Remark 2.51, the DL-Algorithm and Theorem 4.13 relate to Theorem 2.50 about finding RAAG subgroups of RAAGs via the extension graph of RAAGs: In the DL-Algorithm, we consider infinite order elements of special D_∞-subgroups of a RACG as generators of a RAAG. These correspond to bi-labelled geodesics in the Davis complex. We could try to consider these as an analogue of standard geodesics in the Salvetti complex of a RAAG used in Definition 2.42 of the extension graph. In an attempt to answer Question 1, we define a potential analogue of the extension graph for a RACG: Draw a vertex for each parallel class of such bi-labelled geodesics and connect two vertices if their corresponding geodesics span a flat. Then we check which of the induced subgraphs of this graph give a (finite index) RAAG subgroup of our original RACG. The ones corresponding to the commuting graph of a FIDL-Λ provide a finite index RAAG subgroup, indeed. Theorem 2.50 suggests that one might find an algorithm on this potential analogue of the extension graph, to find other RAAG subgroups of W_Γ as well.

However, we emphasize that we might not detect all finite index RAAG subgroups of a RACG with this potential analogue: Given a RACG W_Γ with a FIDL-Λ, we can produce a finite index RACG subgroup $W'_\Gamma \leq W_\Gamma$ of W_Γ by doubling over the star of a vertex by Proposition 1.43. If W'_Γ has a FIDL-Λ as well, this produces a RAAG subgroup of finite index in both W'_Γ and W_Γ. It is unclear how this RAAG subgroup is detectable from the potential analogue of the extension graph. Possibly, the definition of a standard geodesic in a Davis complex has to be extended beyond the class of bi-labelled geodesics.

Given the DL-Algorithm, we ask three follow-up questions that guide the rest of this section:

1. When does a RACG W_Γ have a FIDL-Λ?
2. If a RACG W_Γ has a FIDL-Λ, how do we find it efficiently?
3. If a RACG W_Γ does not have a FIDL-Λ, does this mean that W_Γ does not have any finite index RAAG subgroup?

4.1.1 Construction of a FIDL-Λ

Restricting to planar defining graphs, we get a full QI-classification:

Theorem 4.17. [DL20, Theorem 5.5] *Let W_Γ be a RACG on a planar defining graph Γ satisfying Standing Assumption 1. Then W_Γ has a 2-component FIDL-Λ if and only if W_Γ is QI to a RAAG.*

Idea of the Proof. One direction of the theorem is obvious: If W_Γ has a FIDL-Λ, it has a finite index RAAG subgroup by Theorem 4.13 and is thus QI to a RAAG.

The proof of the other direction of the theorem uses the QI-classification of [NT19] between RACGs and RAAGs on planar defining graphs given in Theorem 1.77. It says that there is a QI if and only if every maximal suspension subgraph is spacious (see Definition 1.36). In [DL20, Proof of Theorem 5.5], Dani-Levcovitz use these spacious suspension subgraphs by adding a Λ-edge in the space of the suspension to construct an explicit FIDL-Λ. □

There is a variety of examples that are QI to a RAAG by Theorem 4.17. The following graph was brought to the author's attention by Pallavi Dani:

Example 4.18. The planar defining graph Γ shown on the left of Figure 4.1.2 has many FIDL-Λ. One of them is illustrated in the middle of Figure 4.1.2 with the components Λ_1 and Λ_2 highlighted in blue and red. By Theorem 4.13, the RACG W_Γ has a visible finite index RAAG subgroup A_Δ, for Δ illustrated on the right of Figure 4.1.2, where a vertex of the form v_{xy} is denoted by xy.

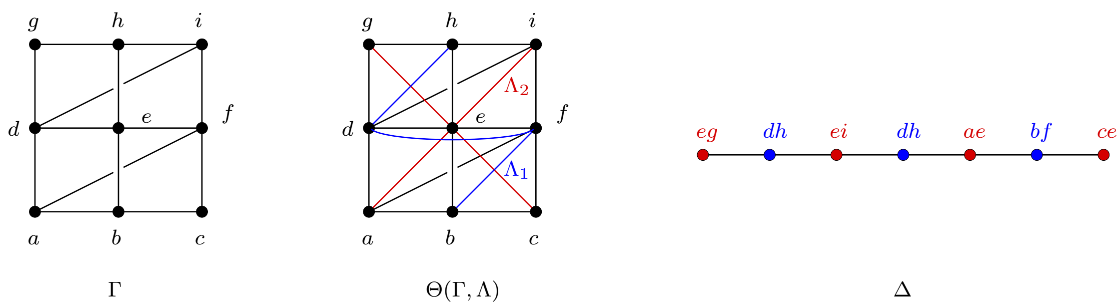

Figure 4.1.2

Example 4.19. The planar defining graph Γ from Figure 4.1.1 in Example 4.6 has a 2-component FIDL-Λ, shown on the left of Figure 4.1.3 in red and blue. By Theorem 4.17, the RACG W_Γ is QI to the RAAG A_Δ, for Δ on the right of Figure 4.1.3, where a vertex of the form v_{xy} is denoted by xy.

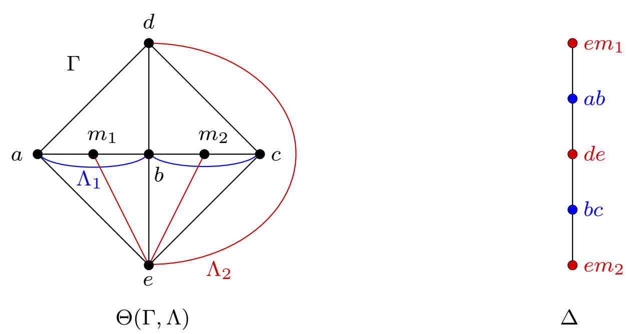

Figure 4.1.3

Naturally, we are not content with excluding non-planar graphs, as they form the much larger and richer class of examples. In [DL20, Section 5.1], Dani-Levcovitz give two explicit families of non-planar graphs which have a FIDL-Λ. The following example is contained in one of them:

Example 4.20. The graph Γ that is the 1-skeleton of a 3-cube with one space diagonal as illustrated on the left of Figure 4.1.4, has a FIDL-Λ whose components Λ_1 and Λ_2 are highlighted in red and blue. By Theorem 4.13, the RACG W_Γ has a visible finite index RAAG subgroup A_Δ, for Δ depicted on the right of Figure 4.1.4, where a vertex of the form v_{xy} is denoted by xy.

In fact, we characterize all graphs with a 2-component FIDL-Λ in terms of a structural property:

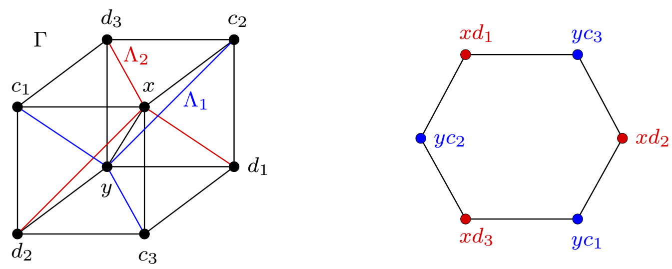

$\Theta(\Gamma, \Lambda)$ Δ

Figure 4.1.4

Definition 4.21. We perform the following inductive procedure to build a graph Γ with an associated graph $\Lambda \leq \Gamma^c$ with two connected components:

Algorithm 4.21.1.

1. The initial graph Γ_0 is a square, the associated graph $\Lambda_0 = \Gamma_0^c$ is the complement graph of Γ_0.
2. Build a sequence of pairs $((\Gamma_0, \Lambda_0), (\Gamma_1, \Lambda_1), \ldots, (\Gamma_n, \Lambda_n))$ by applying the induction step:
   (∗) Given (Γ_i, Λ_i), pick a vertex v_i such that there is a set $N_i \subseteq nbs_\Gamma(v_i)$ that is connected as an induced subgraph of Λ_i. Define $(\Gamma_{i+1}, \Lambda_{i+1})$ by adding a new vertex x_{i+1} and edges as follows:

$$\begin{aligned} V(\Gamma_{i+1}) &= V(\Gamma_i) \cup \{x_{i+1}\} \\ E(\Gamma_{i+1}) &= E(\Gamma_i) \cup \{(n, x_{i+1}) \mid n \in N_i\} \\ V(\Lambda_{i+1}) &= V(\Gamma_{i+1}) \\ E(\Lambda_{i+1}) &= E(\Lambda_i) \cup \{(v_i, x_{i+1})\} \end{aligned}$$

3. Stop the inductive procedure after some $n \in \mathbb{N}$ steps and set $\Gamma = \Gamma_n$ and $\Lambda = \Lambda_n$.

- We refer to Algorithm 4.21.1 as the *Coning Algorithm*.
- For a given RACG W_Γ with an induced subgraph $\Lambda \leq \Gamma^c$ of the complement, we say that the pair (Γ, Λ) is *constructed by the Coning Algorithm* if there is an initial square $\Gamma_0 \leq \Gamma$ and an induced $\Lambda_0 \leq \Lambda$ with $\Lambda_0 = \Gamma_0^c$ such that the pair (Γ, Λ) can be built by the Coning Algorithm 4.21.1 with a finite sequence of pairs $((\Gamma_0, \Lambda_0), (\Gamma_1, \Lambda_1), \ldots, (\Gamma_n, \Lambda_n))$ such that $(\Gamma_n, \Lambda_n)) = (\Gamma, \Lambda)$.
- The *Coning Sequence* of (Γ, Λ) is the sequence $((\Gamma_0, \Lambda_0), (\Gamma_1, \Lambda_1), \ldots, (\Gamma_n, \Lambda_n))$ such that $(\Gamma_n, \Lambda_n) = (\Gamma, \Lambda)$.
- The *Coning Decomposition of* Γ is the sequence $(\Gamma_0, \Gamma_1, \ldots, \Gamma_n)$.
- The *Coning History* of (Γ, Λ) is the sequence

$$((\Gamma_0, \Lambda_0, v_0, N_0), (x_1, \Gamma_1, \Lambda_1, v_1, N_1), \ldots, (x_{n-1}, \Gamma_{n-1}, \Lambda_{n-1}, v_{n-1}, N_{n-1}), (x_n, \Gamma_n, \Lambda_n)).$$

Remark 4.22.

- In the inductive step (∗) of the Coning Algorithm 4.21.1 in Γ_{i+1}, all vertices in the collection $N_i \subseteq nbs_{\Gamma_i}(v_i)$ are being connected to the same new vertex x_{i+1} and $nbs_{\Gamma_{i+1}}(x_{i+1}) = N_i$. Thus, x_{i+1} is the cone on N_i, hence the choice of terminology. We also say that we add the vertex x_{i+1} by *coning off* N_i.
- The Coning Algorithm 4.21.1 is dependent on the initial choice of (Γ_0, Λ_0) as well as the choice of v_i and x_{i+1} for all $i \in \{0, \ldots, n-1\}$. Thus, the Coning Decomposition of Γ is not unique.

Theorem 4.23. *Let W_Γ be RACG satisfying Standing Assumption 1 and let $\Lambda \leq \Gamma^c$ be an induced subgraph of the complement with two connected components. If the pair (Γ, Λ) can be constructed by the Coning Algorithm 4.21.1, then Λ is a FIDL-Λ.*

Proof. Suppose first that the pair (Γ, Λ) is constructed by the Coning Algorithm 4.21.1 with the Coning Sequence $((\Gamma_0, \Lambda_0), (\Gamma_1, \Lambda_1), \ldots, (\Gamma_n, \Lambda_n))$ such that $(\Gamma_n, \Lambda_n) = (\Gamma, \Lambda)$. We show that $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$ - $\mathcal{R}_4$ and conditions $\mathcal{F}_1$ and $\mathcal{F}_2$ by induction on $i \in \mathbb{N}$: It is easy to check that the base case $\Theta(\Gamma_0, \Lambda_0)$ satisfies all the conditions. So, we can suppose by the induction hypothesis that $\Theta(\Gamma_i, \Lambda_i)$ satisfies all the conditions and show that $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$ does as well:

condition $\mathcal{R}_1$: By assumption, $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{R}_1$, so Λ_i does not contain a cycle. By construction, $E(\Lambda_{i+1}) = E(\Lambda_i) \cup \{(x_{i+1}, v_i)\}$ and $x_{i+1} \notin V(\Lambda_i)$, so, Λ_{i+1} does not contain a cycle.

condition $\mathcal{R}_2$: By construction, the elements in N_i are all adjacent in Γ_i to v_i. Since by assumption, $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{R}_2$, this implies that $N_i \subseteq V(\Lambda_{i,1})$ and $v_i \in V(\Lambda_{i,2})$ for $\Lambda_{i,1}$ and $\Lambda_{i,2}$ two distinct connected components of Λ_i. So, the edge $\{(x_{i+1}, v_i)\} \in E(\Lambda_{i+1})$ is added to the connected component $\Lambda_{i,2}$, while in Γ_{i+1}, the vertex x_{i+1} is only connected to vertices in $\Lambda_{i,1}$. Hence, the connected components of Λ_{i+1} are induced in $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$.

condition $\mathcal{R}_3$: Recall Definition 4.8 of condition $\mathcal{R}_3$: For every 2-component square (c_1, d_1, c_2, d_2) in Θ, the graph Γ contains the join of $V(T_\Lambda(c_1, c_2))$ and $V(T_\Lambda(d_1, d_2))$.

Since $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{R}_3$, this is true for every 2-component square in $\Theta(\Gamma_i, \Lambda_i)$ and we only need to show that this holds for every 2-component square in $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$ containing the new vertex x_{i+1}.

The following observations are illustrated in Figure 4.1.5: By construction of Λ_{i+1}, any such 2-component square is of the form (x_{i+1}, n, l, n') with $n, n' \in N_i$ and $l \in \Lambda_{i+1, x_{i+1}}$, where $\Lambda_{i+1, x_{i+1}}$ is the component of Λ_{i+1} containing x_{i+1}. Since N_i is connected as an induced subgraph of Λ_i and by condition $\mathcal{R}_1$, Λ_i does not contain a cycle, there is a unique geodesic $p = (n, n_1, \ldots, n_k, n') \subseteq \Lambda_i$ with $\{n, n_1, \ldots, n_k, n'\} \subseteq N_i$ connecting n and n'. Thus, by construction of Γ_{i+1}, every vertex on the path p is connected with x_{i+1}, and thus, the join of x_{i+1} and $T_{\Lambda_{i+1}}(n, n')$ is in Γ_{i+1}.

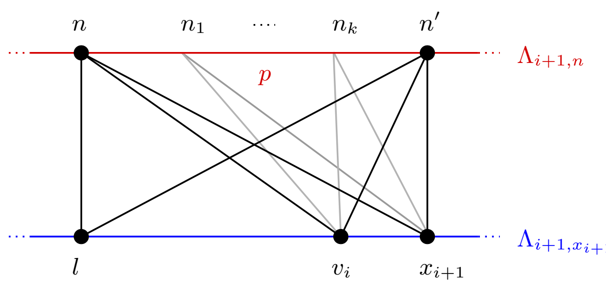

Figure 4.1.5

But the vertex x_{i+1} is only connected to v_i in Λ_{i+1}. Hence, $T_{\Lambda_{i+1}}(x_{i+1}, l) \setminus \{x_{i+1}\} = T_{\Lambda_{i+1}}(v_i, l)$. Consider now the 2-component square $(v_i, n, l, n') \subseteq \Gamma_i$. By hypothesis, the join of $T_{\Lambda_{i+1}}(v_i, l) = T_{\Lambda_{i+1}}(x_{i+1}, l) \setminus \{x_{i+1}\}$ and $T_{\Lambda_{i+1}}(n, n')$ is in Γ_{i+1}. If we include x_{i+1} again and merge the joins, in total we have that the join of $T_{\Lambda_{i+1}}(x_{i+1}, l)$ and $T_{\Lambda_{i+1}}(n, n')$ is in Γ_{i+1}, as required.

condition $\mathcal{R}_4$: Recall Definition 4.9 of condition $\mathcal{R}_4$: For every 2-component cycle $\gamma = (c_1, d_1, c_2, d_2, \ldots, c_n, d_n)$ in Θ, every edge of γ is contained in a 2-component square (t_c, t_d, t'_c, t'_d) of Θ with $t_c, t'_c \in V(T_\Lambda(c_1, \ldots, c_n))$ and $t_d, t'_d \in V(T_\Lambda(d_1, \ldots, d_n))$.

Since $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{R}_4$, we only need to show that this holds for every 2-component cycle in $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$ containing the new vertex x_{i+1}.

Let γ be a 2-component cycle containing x_{i+1}. Since by construction, in Γ_{i+1} the new vertex x_{i+1} is only adjacent to vertices in N_i, γ is of the form $\gamma = (x_{i+1}, n, l_1, \ldots, l_k, n')$ for $n, n' \in N_i$. But $N_i \subseteq nhs_{\Gamma_i}(v_i)$, so both n and n' are also adjacent to v_i. Thus, $\gamma' = (v_i, n, l_1, \ldots, l_k, n')$ is another 2-component cycle in $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$. But since γ' does not contain x_{i+1}, γ' is also a 2-component cycle in $\Theta(\Gamma_i, \Lambda_i)$. Thus, by hypothesis, every edge of γ' is contained in a 2-component square in $\Theta(\Gamma_i, \Lambda_i)$ and hence in $\Theta(\Gamma_{i+1}, \Lambda_{i+1})$. This implies that every edge of γ is contained in a 2-component square, except possibly the two edges (x, n) and (x, n'). But these edges are contained in the 2-component square (x_{i+1}, n, v_i, n'). Hence, the condition is satisfied.

condition $\mathcal{F}_1$: Since by assumption $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{F}_1$, we know that every vertex of Γ_{i+1} (that is not a cone vertex) is contained in a Λ_i-edge and thus in a Λ_{i+1}-edge, except possibly x_{i+1}. But x_{i+1} is in the Λ_{i+1}-edge (x_{i+1}, v_i) by construction. Thus, the condition is satisfied.

condition $\mathcal{F}_2$: Recall Definition 4.11 of condition $\mathcal{F}_2$: Any two vertices in different components of Λ are connected by a 2-component path.

Since $\Theta(\Gamma_i, \Lambda_i)$ satisfies condition $\mathcal{F}_2$, this is true for every pair of vertices in $V(\Gamma_i)$ and we only need to show that there is a path between the new vertex x_{i+1} and some vertex $l \in V(\Gamma_{i+1})$. Since x_{i+1} is only adjacent to vertices in N_i in Γ_{i+1}, any such 2-component path connecting l and x_{i+1} passes through some $n \in N_i$, before reaching x_{i+1}. As $N_i \subseteq nbs_{\Gamma_i}(v_i)$, we know that v_i is adjacent to n. But since $v_i, l \in V(\Gamma_i)$, by hypothesis, there is a 2-component path between v_i and l. This path extends via the edges (v_i, n) and (n, x_{i+1}) to the desired 2-component path between x_{i+1} and l.

Hence, we have shown that given a pair (Γ, Λ) constructed by the Coning Algorithm 4.21.1, $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$ - $\mathcal{R}_4$ and the conditions $\mathcal{F}_1$ and $\mathcal{F}_2$ and thus Λ is a FIDL-Λ. ☐

With Theorem 4.23, we can always consider coning off a collection of vertices in a given graph Γ which has a 2-component FIDL-Λ:

Corollary 4.24. *Given a graph Γ satisfying Standing Assumption 1 with $\Lambda \leq \Gamma^c$ a 2-component FIDL-Λ, define the pair (Γ', Λ') as follows: Choose a vertex $v \in V(\Gamma)$ and a set $N \subseteq nbs_\Gamma(v)$ such that N is connected as an induced subgraph of Λ and define $V(\Gamma') = V(\Gamma) \cup \{x\}$ with $E(\Gamma') = E(\Gamma) \cup \{(x, n) \mid n \in N\}$ and $V(\Lambda') = V(\Gamma')$ with $E(\Lambda') = E(\Lambda) \cup \{(x, v)\}$. Then Λ' is a 2-component FIDL-Λ for Γ'.*

Example 4.25. The graph Γ that is the 1-skeleton of a 3-cube with one space diagonal has a FIDL-Λ, as discussed in Example 4.20 and illustrated in Figure 4.1.4. The pair (Γ, Λ) can be constructed by the Coning Algorithm 4.21.1. The Coning History

$$((\Gamma_0, \Lambda_0, x, \{y, c_2\}), (d_1, \Gamma_1, \Lambda_1, y, \{x, d_3\}), (c_1, \Gamma_2, \Lambda_2, x, \{y, c_1\}), (d_2, \Gamma_3, \Lambda_3, y, \{x, d_1, d_2\}), (c_3, \Gamma_4, \Lambda_4))$$

is illustrated in Figure 4.1.6.

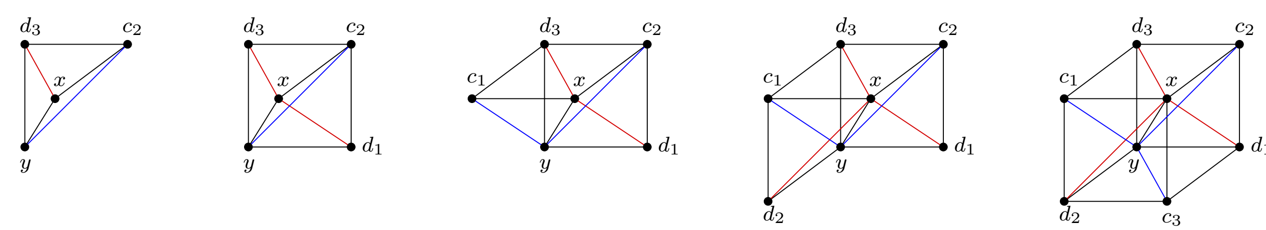

$(\Gamma_0, \Lambda_0, x, \{y, c_2\})$ $(d_1, \Gamma_1, \Lambda_1, y, \{x, d_3\})$ $(c_1, \Gamma_2, \Lambda_2, x, \{y, c_1\})$ $(d_2, \Gamma_3, \Lambda_3, y, \{x, d_1, d_2\})$ $(c_3, \Gamma_4, \Lambda_4)$

Figure 4.1.6

With the Coning Algorithm 4.21.1, we can extend known examples and give new, in particular non-planar, examples of RACGs, not mentioned in [DL20], that have finite index visual RAAGs:

Example 4.26. We start with the planar graph Γ from Example 4.18 shown in Figure 4.1.7 on the left and construct a 2-component FIDL-Λ, as illustrated in the middle. We use Corollary 4.24 to produce a non-planar example Γ' with FIDL-Λ Λ': Observe that the set of vertices adjacent to the vertex f contains $\{a, e, i\} =: N \subseteq nbs_\Gamma(f)$. Since N is connected as an induced subgraph of Λ, we can apply the Coning Algorithm 4.21.1. We obtain a new graph Γ' by adding a new vertex x adjacent to the vertices in N and for a corresponding 2-component FIDL-Λ Λ' we attach x to f.

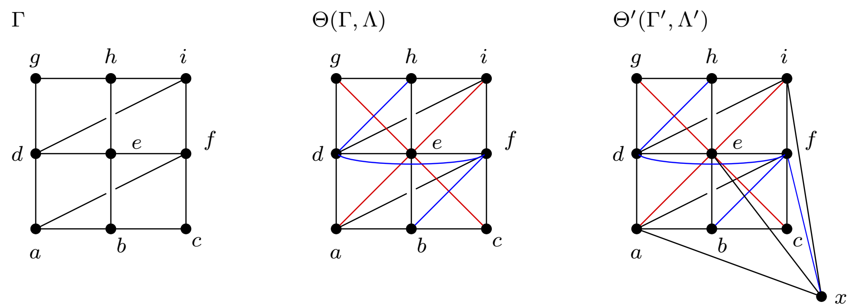

Figure 4.1.7

We can also use the Coning Decomposition in Theorem 4.23 to alter a graph Γ, whose RACG is known not to be QI to a RAAG. This way, we produce a graph Γ' on the same vertex set as Γ and with some additional edges that is QI to a RAAG. Such a pair of graphs might be helpful to find properties to distinguish the RACGs QI to a RAAG from the ones not QI to a RAAG.

Example 4.27. In Example 4.49 we will see that the Diamond graph Γ_D in Figure 4.2.15 on the left is not QI to a RAAG. However, we can try to use the Coning Algorithm 4.21.1 to add some edges and obtain a graph that has a 2-component FIDL-Λ. For instance, if we add the edges (b_5, b_2) and (c_2, b_2) to Γ_D the resulting graph Γ'_D in Figure 4.1.8 has a Coning Decomposition. The Coning

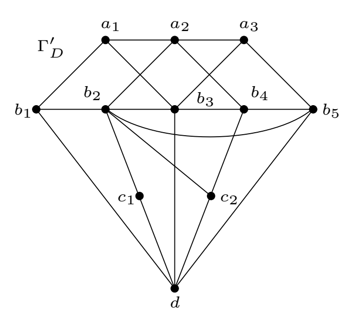

Figure 4.1.8

History of Γ' is illustrated in Figure 4.1.9.

Observe that until the construction of the pair (Γ_4, Λ_4), the Coning Decomposition would also work for the original Diamond graph Γ_D. However, to add the vertex b_5, we want to cone off the vertices $\{a_3, b_4, d\}$, but these vertices are not connected as an induced subgraph of Λ_4. Thus, we need to add the vertex b_2 to this collection to satisfy this condition. The same problem occurs on the pair (Γ_6, Λ_6). We cannot cone off $\{b_4, d\}$, but have to add b_2 to ensure that the collection is connected as an induced subgraph of Λ_6. Now, the resulting RACG $W_{\Gamma'_D}$ is QI to a RAAG.

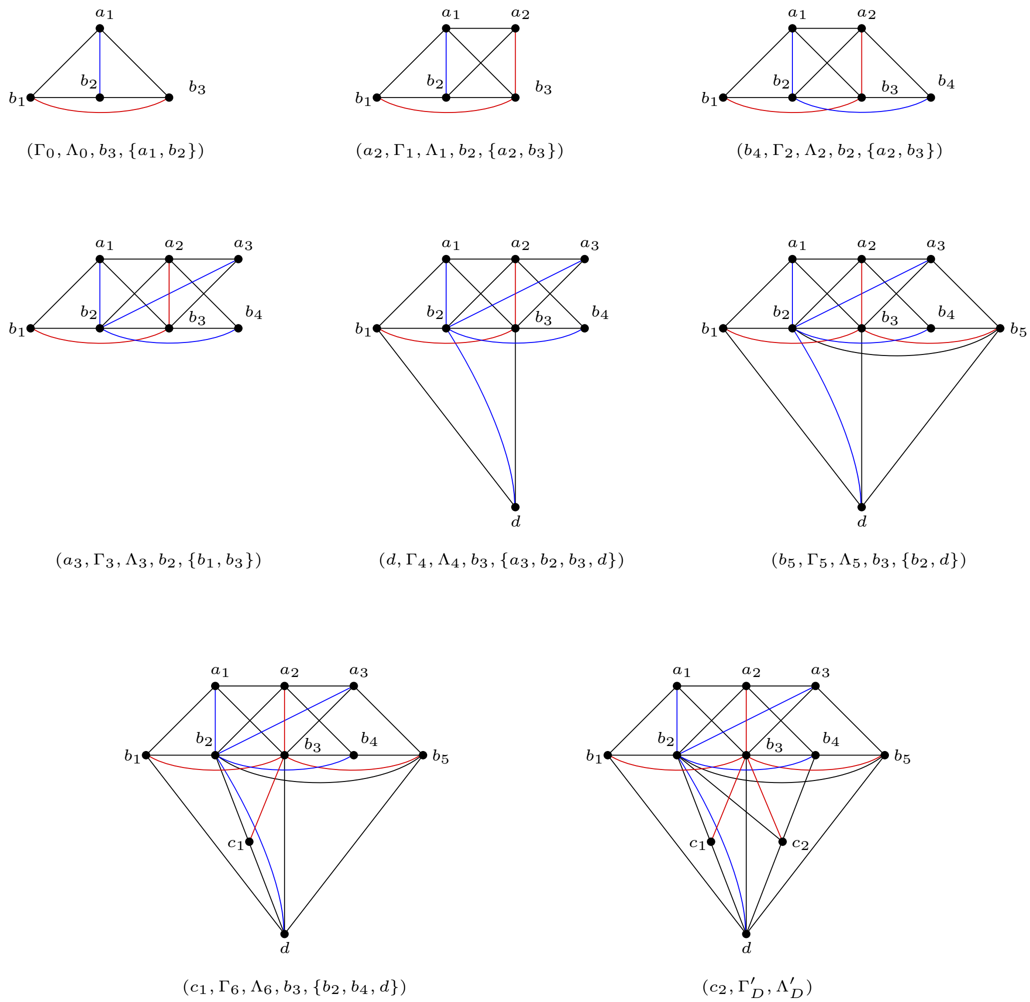

Figure 4.1.9

4.1.2 Splittings and the DL-Algorithm

One of the downsides of the DL-Algorithm is that it is computationally quite expensive to run through all possible induced subgraphs of the complement Γ^c and check whether they satisfy the conditions $\mathcal{R}_1$-$\mathcal{R}_4$ and the conditions $\mathcal{F}_1$ and $\mathcal{F}_2$ to find a FIDL-Λ. The idea investigated in this section is to make the process more efficient by considering a JSJ graph of cylinders decomposition of the RACG W_Γ and applying the DL-Algorithm to each vertex group individually. If we find local 2-component FIDL-Λs, we aim to patch them together to a global one.

First, we show that the occurrence of a cut collection has consequences for a FIDL-Λ:

Lemma 4.28. *Let $\{a - b\}$ be a cut collection of the graph Γ satisfying Standing Assumption 1 with common adjacent vertices $\mathcal{C} = \{c_1, \ldots, c_k\}$. If there is a FIDL-$\Lambda$ with two components Λ_1 and Λ_2, then, up to renumbering:*

1. *$(a, b) \in E(\Lambda_1)$.*
2. *The induced subgraph of Λ_2 on the vertices in $\mathcal{C}$ is connected.*

Proof. Suppose $(a, b) \notin E(\Lambda)$. Since by assumption Γ has a FIDL-Λ, W_Γ is QI to a RAAG by Theorem 4.13, and thus, Γ is $\mathcal{CFS}$ by Theorem 1.59 and Remark 4.15. Hence, since a and b are contained in a cut collection, they have at least two common adjacent vertices c_1 and c_2 in different connected components of $\Gamma \setminus \{a, b\}$, that is, $|\mathcal{C}| \geq 2$. Thus, by conditions $\mathcal{F}_1$ and $\mathcal{R}_2$, a and b lie, without loss of generality, in Λ_1, they are connected via a path $p = (a, x_1, \ldots, x_n, b)$ in $T_\Lambda(a, b)$, and all their common adjacent vertices in $\mathcal{C}$ lie in Λ_2. For the square (a, c_1, b, c_2), condition $\mathcal{R}_3$ then implies that the join of $T_\Lambda(a, b)$ and $T_\Lambda(c_1, c_2)$ is contained in Γ. However, this implies that c_1 and c_2 are both connected to $x_1 \in T_\Lambda(a, b)$. But then c_1 and c_2 are not separated by the cut collection $\{a - b\}$, in contradiction to the assumption. Thus, such a vertex x_1 cannot exist and $(a, b) \in E(\Lambda_1)$.

Suppose now that induced subgraph of Λ_2 on the collection $\mathcal{C}$ of common adjacent vertices of a and b is not connected. Then there are vertices $c_i, c_j \in \mathcal{C}$ such that every path in Λ connecting c_i with c_j passes through some vertex $x \in T_\Lambda(c_i, c_j) \setminus \mathcal{C}$. However, a and b are both connected to c_i and c_j, thus there is a square (a, c_i, b, c_j). Hence, by condition $\mathcal{R}_3$, every vertex in $T_\Lambda(c_i, c_j)$, in particular x, is connected to a and b. Thus, $x \in \mathcal{C}$, in contradiction to the assumption. Thus, the induced subgraph of Λ_2 on all vertices in $\mathcal{C}$ is connected. $\square$

Remark 4.29. Lemma 4.28 implies that Γ does not contain a cycle of cut collections if W_Γ has a FIDL-Λ: Suppose there is a set of vertices $\{a_1, a_2, \ldots, a_n\} \subseteq V(\Gamma)$ such that $\{a_i - a_{i+1}\}$ for every $i \in \{1, \ldots, n-1\}$ and $\{a_n - a_1\}$ are cut collections of Γ. Then, by item 1 of Lemma 4.28, Λ would contain a cycle, contradicting condition $\mathcal{R}_1$.

Under certain conditions, we can patch FIDL-Λs together. We use for the convenience of the reader and the illustration of the argument the most basic setting that can be easily generalised for any JSJ graph of cylinders:

Proposition 4.30. *Let Γ be a graph satisfying Standing Assumption 1 with one single uncrossed cut collection $\{a - b\} \subseteq V(\Gamma)$ and let Σ_c be the corresponding graph of cylinders with one cylinder vertex c with defining graph Γ_c associated to the cut collection $\{a - b\}$ and two rigid vertices r and r' with defining graphs Γ_r and $\Gamma_{r'}$. If for Γ_r, $\Gamma_{r'}$ and Γ_c, there are 2-component FIDL-Λs Λ_r, $\Lambda_{r'}$ and Λ_c, respectively, such that on the vertex set $V(\Gamma_c) \cap V(\Gamma_r)$ the graphs Λ_c and Λ_r are identical and on the vertex set $V(\Gamma_c) \cap V(\Gamma_{r'})$ the graphs Λ_c and $\Lambda_{r'}$ are identical, then Γ has a 2-component FIDL-Λ consisting of $\Lambda_r \cup \Lambda_{r'} \cup \Lambda_c$.*

Proof. By Lemma 4.28, Λ_c has one component $\Lambda_{c,1}$ which consists only of the edge $\{(a, b)\} = E(\Lambda_{c,1})$. We denote the other component on $\mathcal{C}$ as $\Lambda_{c,2}$. Let $\Lambda_{r,1}$ and $\Lambda_{r,2}$, and $\Lambda_{r',1}$ and $\Lambda_{r',2}$ be the two connected components of Λ_r and $\Lambda_{r'}$, respectively, where $E(\Lambda_{c,1}) = \{(a, b)\} \subseteq E(\Lambda_{r,1} \cap \Lambda_{r',1})$. Then the union $\Lambda_{r,1} \cup \Lambda_{r',1}$ is a connected induced subgraph of Γ^c.

Let us now consider $\Lambda_{r,2} \cup \Lambda_{r',2} \cup \Lambda_{c,2}$: There exist $c \in \mathcal{C} \cap V(\Lambda_{r,2})$ and $c' \in \mathcal{C} \cap V(\Lambda_{r',2})$. But by Lemma 4.28, c and c' are connected in $\Lambda_{c,2}$. Thus, $\Lambda_{r,2} \cup \Lambda_{r',2} \cup \Lambda_{c,2}$ is connected.

We show that $\Lambda = \Lambda_r \cup \Lambda_{r'} \cup \Lambda_c$ is a FIDL-Λ of Γ with the two components $\Lambda_1 = \Lambda_{r,1} \cup \Lambda_{r',1}$ and $\Lambda_2 = \Lambda_{r,2} \cup \Lambda_{r',2} \cup \Lambda_{c,2}$, see Figure 4.1.10: By construction, $\Lambda \leq \Gamma^c$ is indeed an induced subgraph of the complement with these two components. So, we need to show that $\Theta(\Gamma, \Lambda)$ satisfies the conditions $\mathcal{R}_1$ - $\mathcal{R}_4$ and conditions $\mathcal{F}_1$ and $\mathcal{F}_2$.

condition $\mathcal{R}_1$: By construction, Λ does not contain a cycle, as Λ_r, $\Lambda_{r'}$ and Λ_c do not.

condition $\mathcal{R}_2$: Since $\Theta(\Gamma_r, \Lambda_r)$, $\Theta(\Gamma_{r'}, \Lambda_{r'})$ and $\Theta(\Gamma_c, \Lambda_c)$ satisfy condition $\mathcal{R}_2$, we suppose, there are vertices $v \in V(\Gamma_r) \setminus V(\Gamma_{r'})$ and $v' \in V(\Gamma_{r'}) \setminus V(\Gamma_r)$ such that $v, v' \in V(\Lambda_1)$. Then every path in Γ connecting v and v' passes through the cut collection $\{a - b\}$, implying $d_\Gamma(v, v') \geq 2$, thus $(v, v') \notin E(\Gamma)$. We argue similarly for Λ_2.

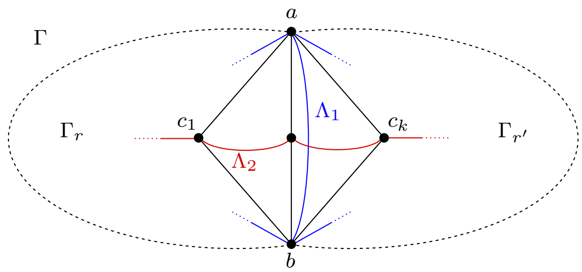

Figure 4.1.10

condition $\mathcal{R}_3$: Since $\Theta(\Gamma_r, \Lambda_c)$, $\Theta(\Gamma_r, \Lambda_r)$ and $\Theta(\Gamma_{r'}, \Lambda_{r'})$ satisfy condition $\mathcal{R}_3$, it suffices to consider a 2-component square γ not contained in Γ_c, Γ_r or $\Gamma_{r'}$. As $\{a - b\}$ is a cut collection, this implies that $\gamma = (a, c, b, c')$ with $c \in V(\Gamma_r)$ and $c' \in V(\Gamma_{r'})$. But then $c, c' \in \mathcal{C}$ and thus $\gamma \subseteq \Gamma_c$.

condition $\mathcal{R}_4$: Since $\Theta(\Gamma_r, \Lambda_r)$, $\Theta(\Gamma_{r'}, \Lambda_{r'})$ and $\Theta(\Gamma_c, \Lambda_c)$ satisfy condition $\mathcal{R}_4$, again it suffices to consider a 2-component cycle γ not contained in Γ_r, $\Gamma_{r'}$ or Γ_c. So, suppose γ is a simple 2-component cycle passing back and forth between Γ_r and $\Gamma_{r'}$, respectively. Then γ passes twice through the cut collection $\{a - b\}$. Cut γ into two 2-component cycles, one component in Γ_r and one contained in $\Gamma_{r'}$ and apply condition $\mathcal{R}_4$ on each cycle. This implies that every edge in γ is contained in a 2-component square, as required.

condition $\mathcal{F}_1$: As $V(\Lambda) = V(\Lambda_r) \cup V(\Lambda_{r'}) \cup V(\Lambda_c) = V(\Gamma)$, every vertex of Γ occurs in a Λ-edge.

condition $\mathcal{F}_2$: It suffices to consider two vertices $v \in V(\Gamma_r)$ and $v' \in V(\Gamma_{r'})$, since $\Theta(\Gamma_r, \Lambda_r)$, $\Theta(\Gamma_{r'}, \Lambda_{r'})$ and $\Theta(\Gamma_c, \Lambda_c)$ satisfy condition $\mathcal{F}_2$ and the case with one vertex in a rigid defining graph and one in the cylinder defining graph is analogous. Let $\gamma = (v, \ldots, a)$ be the 2-component-path in Γ_r connecting v and a and let $\gamma' = (a, \ldots, v')$ be the 2-component path in $\Gamma_{r'}$ between a and v'. Then $\gamma \circ \gamma'$ is a 2-component path connecting v and v' in Γ. □

Proposition 4.30 generalizes to the following:

Corollary 4.31. *The graph Γ satisfying Standing Assumption 1 has a 2-component FIDL-Λ if the following holds: For every defining graph Γ_v of a vertex group in the graph of cylinders Σ_c of the RACG W_Γ, there is a FIDL-Λ Λ_v such that for any pair $c, r \in V(\Sigma_c)$ of adjacent vertices, the FIDL-Λs Λ_c and Λ_r, respectively, agree on the vertex set $V(\Gamma_c) \cap V(\Gamma_r)$.*

We illustrate an application of Corollary 4.31 on an example from [DL20, Corollary 5.2]:

Example 4.32. In the JSJ graph of cylinders Σ_c of W_Γ shown in the middle of Figure 4.1.11 for the graph Γ on the left, the FIDL-Λs for the cylinder vertices c_1 and c_2 already determine the potential FIDL-Λ Λ_r of the rigid vertex r, once we ensure that the condition in Corollary 4.31 is satisfied. So, all we need to do is check that Λ_r is a FIDL-Λ for Γ_r. Since this is true, Corollary 4.31 implies that we can patch the local FIDL-Λs together to a global one, which is illustrated on the right of Figure 4.1.11.

4.1.3 Beyond the DL-Algorithm

Unfortunately, there is nothing we can say just from seeing that the DL-Algorithm fails on some graph Γ. Unless Γ is planar and we can apply Theorem 1.77 of [NT19], it is in general unknown whether the RACG W_Γ has a finite index RAAG subgroup or is QI to a RAAG. Thus, we need to use some other QI-invariant to investigate such examples further. In particular, since the $\mathcal{CFS}$-property is invariant under adding edges, we are curious about the following set-up:

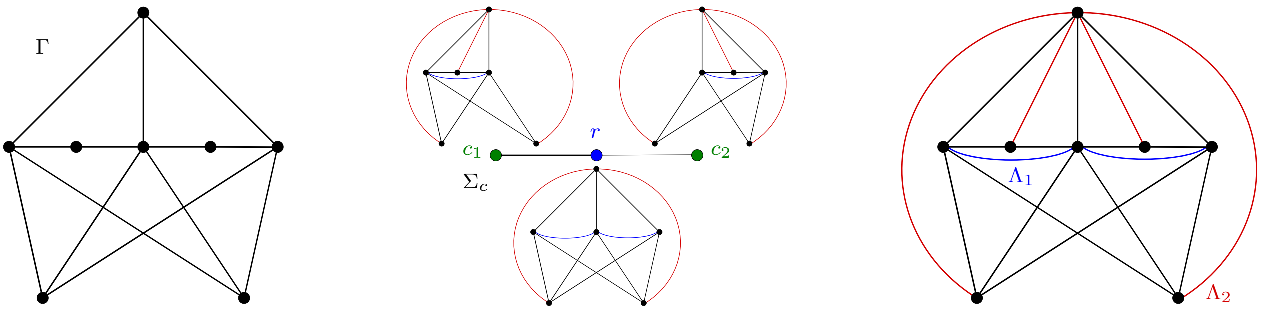

Figure 4.1.11

Question 4. Let Γ be a graph with a 2-component FIDL-Λ. Add one edge to Γ without creating a triangle to obtain a new graph Γ' such that the DL-Algorithm fails on Γ'. What conditions on the new edge determine whether or not Γ' is QI to a RAAG?

To find an answer to Question 4 for a given Γ, we execute the following strategy:

1. Check if the new edge to define Γ' is contained in a square. If yes, check if the resulting graph is still minsquare to potentially apply Corollary 1.70.
2. Try to use the structure invariant (see Section 4.2) to see that $W_{\Gamma'}$ is not QI to a RAAG.
3. Try to use the MPRG Γ'^p, which is the main tool in the Section 4.3.

Moreover, we emphasize that the QIs provided by the Dani-Levcovitz-Algorithm originate from commensurability. In the planar case, being commensurable to a RAAG and being QI to a RAAG are equivalent by Theorem 4.17. This leads to the following question:

Question 5. Is there a RACG that is QI to some RAAG but is not commensurable to any RAAG?

In particular, the JSJ graph of cylinders has potential to be used to advance the distinguishability of QI and commensurability, as demonstrated in the proof of Lemma 3.52 for RACG. It might be used to not only tackle Question 5 but also its relatives:

Question 6. Can the JSJ graph of cylinders be used like in Lemma 3.52 to provide new examples of

- a RACG and a RAAG that are QI but not commensurable?
- two RAAGs that are QI but not commensurable?

4.2 Structure Invariant

We aim to compare the structure invariants determined by the JSJ graphs of cylinders introduced in Section 2.1.1 of RACGs and RAAGs and use them to show that certain RACGs are not QI to any RAAG. The method to describe the JSJ graph of cylinders for RAAGs splitting over two-ended subgroups was introduced by Margolis in [Mar20], using the fact shown by Clay in [Cla14] that any splitting over a two-ended subgroup is in correspondence with a cut vertex of the defining graph. We first recall these relevant results about RAAGs and then highlight the differences between the structure invariants of RAAGs and RACGs and their implications for the QI-classification.

We start with some graph theoretical terminology:

Definition 4.33. An induced subgraph $\Delta' \leq \Delta$ of a graph Δ is *biconnected* if it is connected and does not have any cut vertex.

We describe the JSJ graph of cylinders of a RAAG A_Δ:

Theorem 4.34. [Mar20, Proposition 3.6] *Let A_Δ be a RAAG with Δ satisfying the Standing Assumption 2. Then its JSJ graph of cylinders Σ_c consists of the following vertices:*

- *For every cut vertex $v \in V(\Delta)$, there is a cylinder vertex with vertex group $A_{st_\Delta(v)}$.*
- *For every induced subgraph $\Delta' \leq \Delta$ satisfying the properties (BC1), (BC2) and (BC3), there is a rigid vertex with vertex group $A_{\Delta'}$, where the properties (BC1), (BC2) and (BC3) are the following:*

  (BC1) Δ' is biconnected.
  (BC2) Either Δ' contains two cut vertices of Δ or Δ' is not contained in the star of a cut vertex of Δ.
  (BC3) Δ' is maximal with respect to the properties (BC1) and (BC2).

There is an edge between a cylinder vertex corresponding to the cut vertex v and the rigid vertex corresponding to the induced subgraph Δ' if and only if $v \in V(\Delta')$. The corresponding edge group is $A_{st_{\Delta'}(v)} = A_{\Delta'} \cap A_{st_\Delta(v)} = A_{\{v\}} \times A_{lk_{\Delta'}(v)}$.

Remark 4.35. Observe that the RAAG A_Δ does not have any hanging subgroups: These are essentially the fundamental group of a surface with boundary and thus (virtually) free groups. But a free RAAG has a discrete defining graph, which can never occur as a biconnected component.

For the structure invariant, we want to use the initial decoration consisting of the vertex type (*cylinder* or *rigid*) and the relative QI-type of the vertex. To give the latter, we need:

Definition 4.36. [Mar20, Section 5] Let v be a cut vertex of Δ corresponding to a cylinder vertex c in the JSJ graph of cylinders Σ_c of the RAAG A_Δ satisfying Standing Assumption 2. Then the vertex group $A_{st_\Delta(v)}$ decomposes as

$$A_{st_\Delta(v)} = \langle v \rangle \times (A_{\Delta_1} * \cdots * A_{\Delta_n} * A_{\Delta'_1} * \cdots * A_{\Delta'_m}),$$

where $\Delta_1, \ldots, \Delta_n \leq lk_\Delta(v)$ are induced, connected subgraphs contained in the defining graph of a rigid vertex and $\Delta'_1, \ldots, \Delta'_m \leq lk_\Delta(v)$ are induced, connected subgraphs not contained in any defining graph of a rigid vertex. The groups $A_{\Delta_1}, \ldots, A_{\Delta_n}$ are the *peripheral factors of* c and the groups $A_{\Delta'_1}, \ldots, A_{\Delta'_m}$ are the *non-peripheral factors of* c.

These factors determine the relative QI-types of the cylinder vertices:

Theorem 4.37. [Mar20, Proposition 5.2] *Let A_{Δ_1} and A_{Δ_2} be two RAAGs with cut vertices $v_1 \in V(\Delta_1)$ and $v_2 \in V(\Delta_2)$ in their defining graphs Δ_1 and Δ_2, respectively. Then the corresponding cylinder vertex groups of v_1 and v_2, $A_{st_{\Delta_1}(v_1)}$ and $A_{st_{\Delta_2}(v_2)}$, respectively, are relatively QI if and only if they have the same QI-types of peripheral factors and of one-ended non-peripheral factors.*

Since by Standing Assumption 2 A_Δ is two-dimensional and Δ is triangle-free, we get:

Corollary 4.38. *Every cylinder vertex group in the JSJ graph of cylinders of a RAAG satisfying Standing Assumption 2 is the direct product of the infinite cyclic group and a non-abelian free group. All cylinder vertices have the same relative QI-type.*

Proof. Recall that by Theorem 4.34, a cylinder vertex group is generated by the star $st_\Delta(v)$ of a cut vertex v in the defining graph Δ. Since v is a cut vertex, its link $lk_\Delta(v)$ contains at least two vertices and as Δ is triangle-free, no two vertices in $lk_\Delta(v)$ are adjacent. Hence, all peripheral and non-peripheral factors are isomorphic to $\mathbb{Z}$ and the vertex group is $A_{st_\Delta(v)} \cong \mathbb{Z} \times (\mathbb{Z} * \cdots * \mathbb{Z})$. Thus, by Theorem 4.37, all cylinder vertex groups have the same relative QI-type. □

Example 4.39. The graph Δ on the left of Figure 4.2.12 has two cut vertices c_1 and c_2. There is one biconnected induced subgraph Δ' with vertex set $\{c_1, b_1, c_2, b_2\}$ satisfying the conditions (BC1), (BC2) and (BC3). Thus, we obtain the JSJ graph of cylinders $\Sigma_{c,\Delta}$ on the right of Figure 4.2.12 consisting of two cylinder vertices c and c' adjacent to one rigid vertex r. The vertex group of c is $A_{st_\Delta(c_1)} = \langle c_1 \rangle \times \langle b_1, b_2, a_1 \rangle \cong \mathbb{Z} \times F_3$, the vertex group of c' is $A_{st_\Delta(c_2)} = \langle c_2 \rangle \times \langle b_1, b_2, a'_1, a'_2 \rangle \cong \mathbb{Z} \times F_4$ and the vertex group of r is $A_{\Delta'} = \langle c_1, b_1, c_2, b_2 \rangle = \langle c_1, c_2 \rangle \times \langle b_1, b_2 \rangle \cong F_2 \times F_2$.

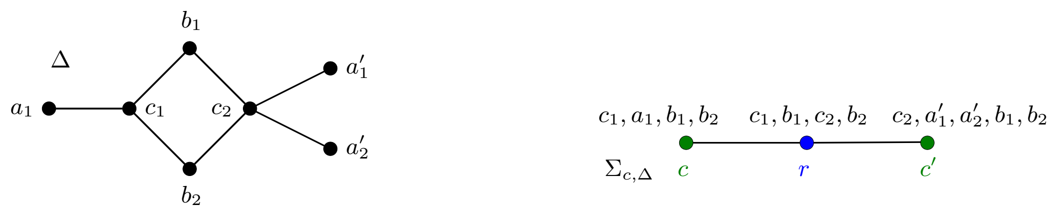

Figure 4.2.12

4.2.1 Comparison

We compare the JSJ decompositions of a RACG W_Γ satisfying Standing Assumption 1 and a RAAG A_Δ satisfying Standing Assumption 2 by use of the structure invariant and Proposition 2.24:

Proposition 4.40. *If W_Γ is a RACG that is QI to a RAAG A_Δ, then every cylinder vertex in the JSJ graph of cylinders $\Sigma_{c,\Gamma}$ of W_Γ is VFD, that is, the direct product of a virtually non-abelian free group and an infinite dihedral group.*

Proof. If W_Γ is QI to a RAAG A_Δ, then every cylinder vertex group of the JSJ graph of cylinders $\Sigma_{c,\Gamma}$ of W_Γ is QI to a cylinder vertex group of the JSJ graph of cylinders $\Sigma_{c,\Delta}$ of A_Δ by Proposition 2.24. But by Corollary 4.38, the cylinder vertex groups of $\Sigma_{c,\Delta}$ are all the direct product $\mathbb{Z} \times F$ of an infinite cyclic group $\mathbb{Z}$ and a non-abelian free group F. By Lemma 3.15, cylinder vertices in $\Sigma_{c,\Gamma}$ are either two-ended, virtually $\mathbb{Z}^2$ or VFD. However, out of these three types, only a VFD cylinder vertex group is QI to a product $\mathbb{Z} \times F$. □

Proposition 4.41. *If W_Γ is a RACG that is QI to a RAAG A_Δ, then every rigid vertex in its JSJ graph of cylinders $\Sigma_{c,\Gamma}$ has a $\mathcal{CFS}$ defining graph and a one-ended vertex group.*

Proof. If W_Γ is QI to a RAAG A_Δ, then a rigid vertex group W_{Γ_r} of the JSJ graph of cylinders $\Sigma_{c,\Gamma}$ of W_Γ is QI to a rigid vertex group A_{Δ_r} of the JSJ graph of cylinders $\Sigma_{c,\Delta}$ of A_Δ by Proposition 2.24. By Theorem 4.34, the defining graph Δ_r of A_{Δ_r} is a maximal biconnected induced subgraph of the connected graph Δ, thus A_{Δ_r} is one-ended. So, by Theorem 1.35 and Corollary 1.59, the defining graph Γ_r is $\mathcal{CFS}$ and W_{Γ_r} is one-ended as well. □

Proposition 4.42. *Let W_Γ be a RACG that is QI to a RAAG A_Δ with rigid vertex r in the JSJ graph of cylinders $\Sigma_{c,\Gamma}$ of W_Γ. The relative QI-type $[\![(G_r, \mathcal{P}_r)]\!]$ of r has one of the following forms:*

1. *G_r and every edge group in $\mathcal{P}_r$ are virtually $\mathbb{Z}^2$.*
2. *G_r is not virtually $\mathbb{Z}^2$ and no edge group in $\mathcal{P}_r$ is virtually $\mathbb{Z}^2$.*

Proof. If W_Γ is QI to A_Δ, then every rigid vertex of $\Sigma_{c,\Gamma}$ has the same relative QI type as some rigid vertex in $\Sigma_{c,\Delta}$ by Proposition 2.24.

Suppose first that G_r is virtually $\mathbb{Z}^2$. This implies that G_r is QI to a rigid vertex group $A_{\{v,w\}}$ of A_Γ that has as defining graph an edge $(v,w) \in E(\Delta)$. But by Theorem 4.34, every adjacent cylinder group contains $A_{\{v,w\}}$. Thus, all adjacent edge groups are of the form $A_{\{v,w\}} \cong \mathbb{Z}^2$.

Assume now that G_r is not virtually $\mathbb{Z}^2$ and thus QI to a rigid vertex r' with vertex group $A_{\Delta'}$ that is not virtually $\mathbb{Z}^2$. Thus, the defining graph Δ' is not a single edge and by Theorem 4.34 biconnected. Let c' be a cylinder vertex adjacent to r'. Then c' corresponds to a cut vertex v of Δ that is contained in Δ'. We determine the edge group: As Δ' contains more than one edge, $st_\Delta(v) \cap \Delta'$ contains v and at least two additional vertices w_1 and w_2 in Δ' adjacent to v. Thus, $A_{\{v,w_1,w_2\}} = A_{\{v\}} \times A_{\{w_1,w_2\}} \cong \mathbb{Z} \times F_2$ is a subgroup of the edge group of the edge connecting r' and c'. Thus, no edge group in the peripheral structure $\mathcal{P}_{r'}$ of $G_{r'}$ is virtually $\mathbb{Z}^2$, and thus, no edge group in $\mathcal{P}_r$ can be virtually $\mathbb{Z}^2$. □

Proposition 4.43. *If W_Γ is a RACG that is QI to a RAAG A_Δ, has a rigid vertex group W_{Γ_r} in its JSJ graph of cylinders $\Sigma_{c,\Gamma}$ that is not virtually $\mathbb{Z}^2$, then the defining graph Γ_r has no cut collection.*

Proof. Suppose the JSJ graph of cylinders $\Sigma_{c,\Gamma}$ of W_Γ has a rigid vertex group W_{Γ_r} whose defining graph Γ_r has a cut collection. Then W_{Γ_r} splits over a two-ended subgroup. But since W_Γ is QI to A_Δ, W_{Γ_r} is QI to a rigid vertex group A_{Δ_r} of A_Δ by Proposition 2.24. Thus, A_{Δ_r} splits over a two-ended subgroup as well. But by [Cla14, Theorem A], this implies that the defining graph Δ_r of A_{Δ_r} has a cut vertex. However, since A_{Δ_r} is rigid, by Theorem 4.34, Δ_r is biconnected. Thus, Δ_r cannot have a cut vertex, which is a contradiction. □

Since in a triangle-free suspension the suspension vertices form a cut pair, we deduce:

Corollary 4.44. *If W_Γ is a RACG that is QI to a RAAG A_Δ, then no defining graph of a rigid vertex group in its JSJ graph of cylinders $\Sigma_{c,\Gamma}$ is a suspension, unless it is a square.*

Remark 4.45. Margolis uses Theorem 4.34 in [Mar20] to give a QI-classification for a class of RAAGs that include RAAGs on trees, which were first classified up to QI in [BN08, Theorem 5.3]. Moreover, the conclusions of Theorem 4.34 in Propositions 4.40 - 4.43 are consistent with the QI-classification between RACGs on planar defining graphs and RAAGs by [NT19] (see Theorem 1.77).

Outline 4.46. To summarize, for a given RACG W_Γ we can check the following properties on its defining graph Γ and its JSJ graph of cylinders $\Sigma_{c,\Gamma}$ to determine whether W_Γ can be QI to a RAAG:

- By Proposition 4.40: Is every cylinder vertex group VFD?
- By Proposition 4.41: Is the defining graph of every rigid vertex $\mathcal{CFS}$?
- By Proposition 4.42: Does every rigid vertex have relative QI type $[\![(G_r, \mathcal{P}_r)]\!]$ with either G_r and all edge groups in $\mathcal{P}_r$ virtually $\mathbb{Z}^2$ or neither G_r nor any edge group in $\mathcal{P}_r$ virtually $\mathbb{Z}^2$?
- By Proposition 4.43: Does no defining graph of a rigid vertex group have a cut collection?

If at least one of the answers to these questions is *NO*, then W_Γ is not QI to a RAAG.

In addition, if we have a given RACG W_{Γ_0} and an explicit given RAAG A_{Δ_0} and we want to check whether they are QI, we can compare their structure invariants in more detail. Suppose all cylinder vertices of W_{Γ_0} are VFD. Since then by Proposition 4.40 and Corollary 4.38 all cylinder vertices of the RACG and the RAAG are QI to each other, the rigid vertices make the difference. Thus, if for instance the number of QI-types of rigid vertices of W_{Γ_0} and A_{Δ_0} differ, we can conclude that W_{Γ_0} and A_{Δ_0} are not QI to each other.

However, we emphasize that even if W_{Γ_0} and A_{Δ_0} have (up to reordering) identical structure invariants for some chosen initial decoration, this does not imply that W_Γ is QI to A_Δ!

Nonetheless, the structure invariant might be refineable like in Section 3.2 for the QI-classification of certain RACGs to answer:

Question 7. Can we choose an initial decoration for the structure invariant to make it a complete QI-invariant between, at least a certain class of, RACGs and RAAGs?

4.2.2 New Examples of RACGs not QI to a RAAG

With the structure invariant, one can produce a variety of new examples of RACGs on non-planar defining graphs that are not QI to a RAAG. We present a selection of $\mathcal{CFS}$ and minsquare examples:

Example 4.47. Consider the RACG on the defining graph Γ_A on the left of Figure 4.2.13. By Theorem 1.77 of [NT19], the RACG $W_{\Gamma_A \setminus \{m_4\}}$ on the planar defining graph $\Gamma_A \setminus \{m_4\}$ is not QI to a RAAG, while the RACG $W_{\Gamma_A \setminus \{m_3\}}$ on the planar defining graph $\Gamma_A \setminus \{m_3\}$ is QI to a RAAG. However, W_{Γ_A} is not QI to a RAAG by Proposition 4.43 (or Corollary 4.44): The vertex group of the rigid vertex r in the JSJ graph of cylinders Σ_{c,Γ_A} of W_{Γ_A}, shown in Figure 4.2.13 on the right, is the special subgroup generated by $\{a, b, c, x, y\}$, whose induced defining subgraph is a suspension and has the cut pair $\{x, y\}$. Thus, the rigid vertex group splits over a two-ended subgroup.

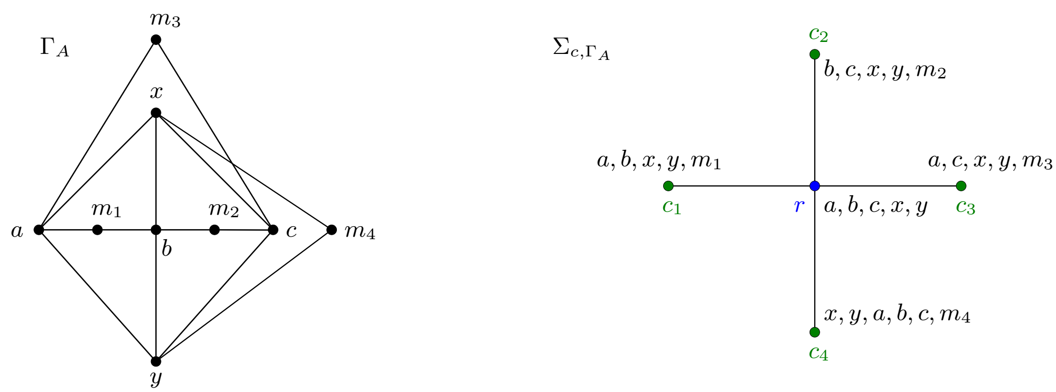

Figure 4.2.13

However, in general a rigid vertex can have a defining graph that is not a suspension and split over a two-ended subgroup:

Example 4.48. For the RACG W_{Γ_B} on the graph Γ_B, illustrated in Figure 4.2.14 on the left, the JSJ graph of cylinders Σ_{c,Γ_B} on the right has a rigid vertex group that is the special subgroup of W_{Γ_B} generated by $\{a, b, c, d, x, y, z\}$. The induced subgraph on these vertices has the two cut pairs $\{x, z\}$ and $\{y, z\}$ and the cut collection $\{a - c\}$. Thus, the rigid vertex group splits over a two-ended subgroup and by Proposition 4.43 W_{Γ_B} is not QI to a RAAG.

The graphs in the following examples were pointed out to the author by Christopher Cashen:

Example 4.49. The *Diamond graph* Γ_D shown on the left of Figure 4.2.15 has only the two cut pairs $\{b_2, d\}$ and $\{b_4, d\}$ with corresponding cylinder vertex groups $\langle b_2, d, b_1, b_3, c_1 \rangle$ and $\langle b_4, d, b_3, b_5, c_2 \rangle$, respectively in the JSJ graph of cylinders Σ_{c,Γ_D} (illustrated in Figure 4.2.15 on the right). Thus, the rigid vertex, whose defining graph contains the nine vertices $\{a_1, a_2, a_3, b_1, b_2, b_3, b_4, b_5, d\}$, is not QI to $\mathbb{Z}^2$ and intersects both cylinder vertex groups only in a group generated by a square. Thus, the edge groups, highlighted in Figure 4.2.15 in pink and orange, are virtually $\mathbb{Z}^2$ and thus, by Proposition 4.42, W_{Γ_D} is not QI to a RAAG.

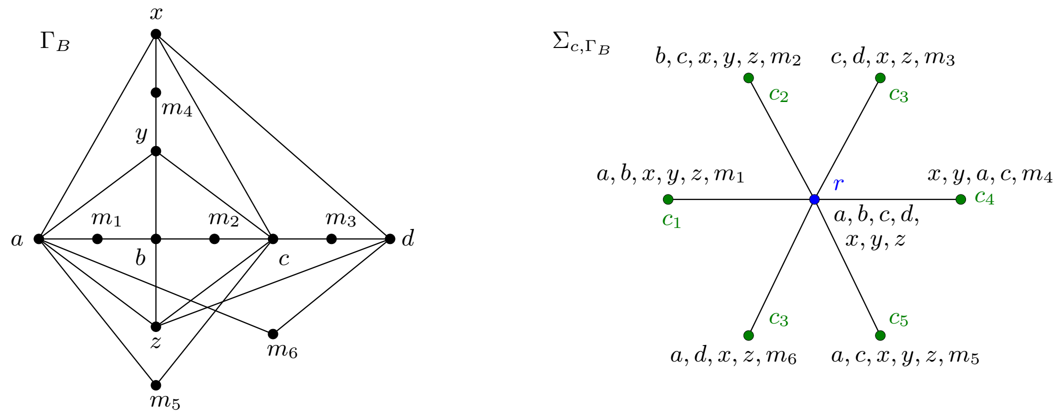

Figure 4.2.14

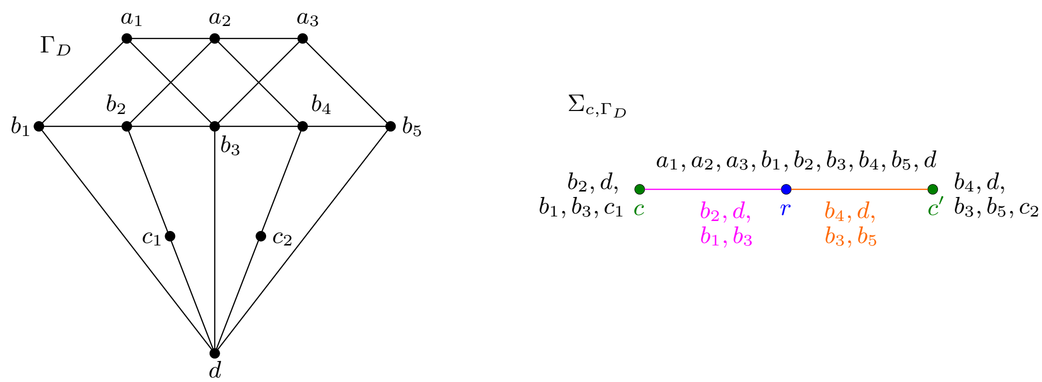

Figure 4.2.15

A similar situation happens for the following *Fox graph*:

Example 4.50. The RACG W_{Γ_F} on the defining graph Γ_F shown on the left of Figure 4.2.16 has a rigid vertex group in its JSJ graph of cylinders Σ_{c,Γ_F} shown on the right of Figure 4.2.16 generated by the vertices $\{a, b, c, d, f, h, i, k\}$ and thus is not QI to $\mathbb{Z}^2$. However, its adjacent edge groups highlighted in pink and orange are virtually $\mathbb{Z}^2$. So, by Proposition 4.42, W_{Γ_F} is not QI to a RAAG.

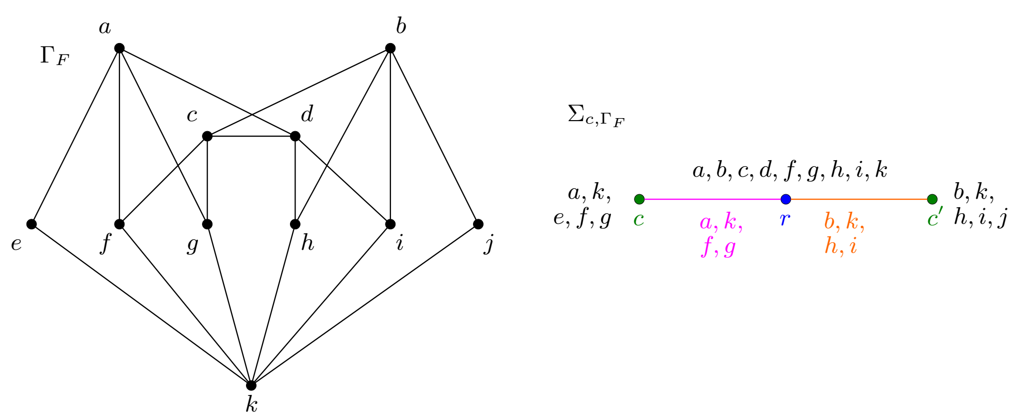

Figure 4.2.16

Also, we can find a min$\mathcal{CFS}$ graph that is not QI to a RAAG:

Example 4.51. The graph Γ_P in Figure 4.2.17 on the left is min$\mathcal{CFS}$, but W_{Γ_P} is not QI to a RAAG by Proposition 4.42: While the rigid vertex group in the JSJ graph of cylinders Σ_{c,Γ_P} depicted on the right of Figure 4.2.17, generated by the five vertices $\{a, b, c, d, e\}$, is not QI to $\mathbb{Z}^2$, all edge groups highlighted in orange are. Therefore, W_{Γ_P} is not QI to a RAAG by Proposition 4.42.

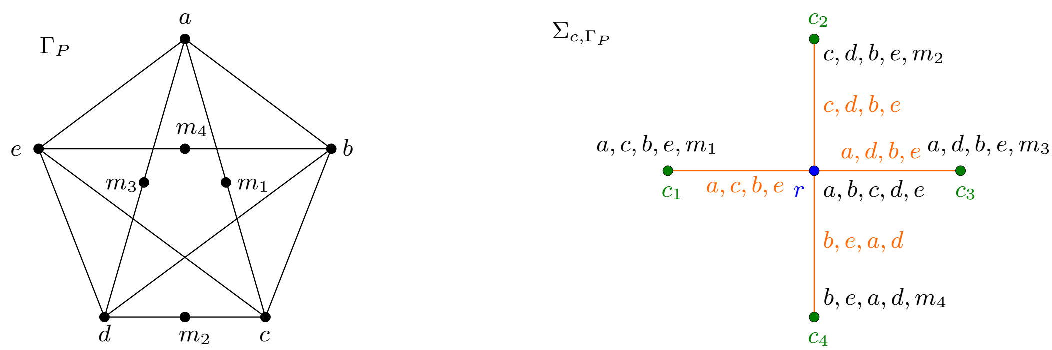

Figure 4.2.17

4.3 MPRG

As described in the introductory Sections 1.1.2 and 2.2 and summarized in Outline 2.40, the aim of this section is to use the properties of the MPRG of RAAGs to distinguish it from the MPRG of certain RACGs and conclude that the RACGs are not QI to any RAAG. The ideas in this Section were developed in collaboration with Christopher Cashen, Jingyin Huang and Annette Karrer. The idea to consider the MPRG as a QI-invariant was introduced to us by Jingyin Huang.

By Proposition 2.33, the MPRG of a RAAG is connected, thus so is the MPRG of a RACG QI to a RAAG by Theorem 2.31. This relates to the property of the defining graph being strongly $\mathcal{CFS}$ (see Definition 1.34): In [RST23], Russell, Spriano and Tran introduce the concept of strongly $\mathcal{CFS}$ graphs and show that RACGs on strongly $\mathcal{CFS}$ defining graphs enjoy some properties similar to those of RAAGs. In their examples, when it is known that a RACG is QI to a RAAG, the defining graph of the RACG is strongly $\mathcal{CFS}$. One might ask if this condition is necessary, we show that it is:

Proposition 4.52. *If there is a QI between a RACG W_Γ satisfying Standing Assumption 1 and a RAAG A_Δ satisfying Standing Assumption 2, then Γ is strongly $\mathcal{CFS}$.*

Proof. Since by assumption W_Γ is QI to a RAAG, by Corollary 1.59 we can assume that Γ is $\mathcal{CFS}$. So, suppose Γ is not strongly $\mathcal{CFS}$, that is, the 4-Cycle-Graph $\square(\Gamma)$ of Γ has more than one connected component. Let $S_1, S_2 \leq \Gamma$ be two induced squares in Γ corresponding to the vertices $s_1, s_2 \in V(\square(\Gamma))$, respectively, that lie in different connected components of $\square(\Gamma)$. Thus, S_1 and S_2 are contained in two distinct essential maximal product subgraphs $M_1, M_2 \leq \Gamma$ with corresponding vertices $m_1, m_2 \in V(\Gamma^p)$ in the MPRG, respectively.

Given two essential maximal product subgraphs of Γ, they are adjacent in the MPRG Γ^p if they intersect in a square. In this case, they correspond to two induced connected subgraphs of $\square(\Gamma)$ intersecting in a vertex. So, the fact that s_1 and s_2 are not connected in $\square(\Gamma)$ implies that m_1 and m_2 are not connected in Γ^p. Thus, the MPRG Γ^p is not connected. However, the MPRG Δ^p of the RAAG A_Δ is connected by Proposition 2.33. So, if W_Γ is QI to A_Δ, by Theorem 2.31, Γ^p is connected as well. This is a contradiction, so Γ is strongly $\mathcal{CFS}$. $\square$

In light of Question 4, we are particularly interested in graphs which have an edge not contained in a square. The following lemma establishes the effects such an edge has on the structure of the MPRG and is a result of a collaboration with Christopher Cashen and Jingyin Huang:

Lemma 4.53. *Let Γ be a triangle-free $\mathcal{CFS}$ graph containing an edge $(s,t) \in E(\Gamma)$ that is not contained in any square. Let $\Gamma' = \Gamma \setminus \{(s,t)\}$ be the graph obtained from Γ by removing the edge (s,t). Then the MPRGs Γ^p and Γ'^p have identical fundamental domains R and R', respectively, but Γ^p contains a cycle that does not occur in Γ'^p.*

Proof. Since the edge (s,t) is not contained in a square, it does not contribute to any essential maximal product subgraph of Γ. Thus, the essential maximal product subgraphs of Γ and Γ' are identical, implying that the fundamental domains R and R' of the MPRGs Γ^p and Γ'^p, respectively, are the same. However, Γ is $\mathcal{CFS}$, thus so is Γ' and s and t are both contained in some square. Therefore, s and t are contained in different essential maximal product subgraphs, and thus, s and t leave at least one vertex v and w, respectively, of R and R' invariant under conjugation. Hence, R overlaps with both ${}^sR$ and ${}^tR$ and analogously R' overlaps with both ${}^sR'$ and ${}^tR'$. Now, consider conjugation with st and ts. Since in Γ the vertices s and t are adjacent, the conjugates ${}^{st}R$ and ${}^{ts}R$ are the same. Thus, the four conjugates R, ${}^sR$, ${}^{ts}R = {}^{st}R$ and ${}^tR$ form a cycle in Γ^p. However, in Γ', s and t do not commute, thus, ${}^{st}R'$ is not equal to ${}^{ts}R'$ and so no cycle is created. See also Figure 4.3.18 for a schematic picture. □

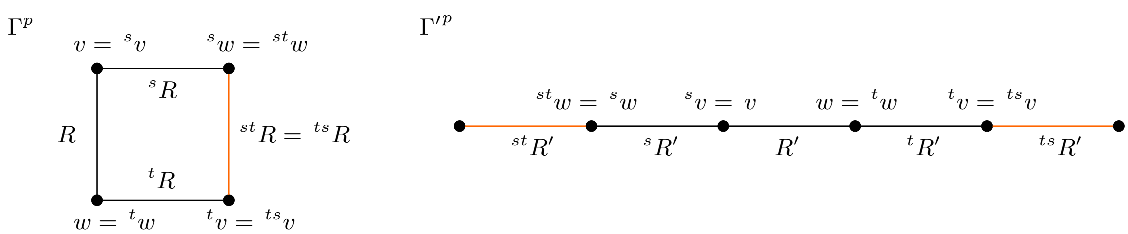

Figure 4.3.18: In Γ^p the conjugates ${}^{st}R$ and ${}^{ts}R$ are identical, in Γ'^p they are different.

Remark 4.54. Figure 4.3.18 is indeed only a schematic picture as the fundamental domain is not always a line or tree and the vertices v and w left invariant under conjugation by s and t, respectively are not necessarily leaves of the fundamental domain. However, since the vertices v and w are contained in two conjugates of the fundamental domain and thus mimic *corners*, we refer to the cycle in the MPRG Γ^p caused by the edge $(s,t) \in E(\Gamma)$ by Lemma 4.53 that is not contained in a square of Γ as the (s,t)*-square* $S_{(s,t)} \subseteq \Gamma^p$.

The following theorem was proved in joint work with Christopher Cashen and Annette Karrer.

Theorem 4.55. *Let Γ be a triangle-free $\mathcal{CFS}$ graph with an edge $(s,t) \in \Gamma$ not contained in any square. Assume the following:*

1. *The (s,t)-square $S_{(s,t)}$ has side length at least 3, that is, the distance between its corner vertices v and w in the fundamental domain R left invariant under conjugation with s and t, respectively is at least 3.*
2. *The (s,t)-square $S_{(s,t)}$ is induced in the MPRG Γ^p.*
3. *There is no vertex in the MPRG Γ^p whose star separates the (s,t)-square $S_{(s,t)}$.*
4. *There exist a vertex $x \in V(\Gamma) \setminus \{st(s) \cup st(t)\}$ and a segment $L \subseteq R \cap S_{(s,t)}$ contained in the side of the (s,t)-square $S_{(s,t)}$ that is also contained in the fundamental domain R of Γ^p of at least length 3, such that ${}^xL = L$.*

Then the RACG W_Γ is not QI to any RAAG.

Proof. The following situation is illustrated in Figure 4.3.19: Consider the (s,t)-square $S_{(s,t)}$ which exists by Lemma 4.53. By assumption *4.*, its side contained in the fundamental domain R of Γ^p contains a segment L that is at least of length 3 and is left invariant under conjugation by some element x that does not commute with either s or t. Thus, the side of $S_{(s,t)}$ contained in ${}^{st}R$ contains the segment ${}^{st}L$, which is equivalent to ${}^{stx}L$. We build a new square with ${}^{stx}L$ as the base. Its other sides are contained in ${}^{stxt}R$, ${}^{stxs}R$ and ${}^{stxst}R$. Now, we can use ${}^{stxstx}L$ as the new base and continue this process of building squares.

On the other hand, we can also build a square with base ${}^{x}L$ (which is equivalent to L). Its other sides are contained in ${}^{xs}R$, ${}^{xst}R$ and ${}^{xt}R$. ${}^{xst}R$ contains ${}^{xst}L$ which is equivalent to ${}^{xstx}L$. Again, we can use this as the base for a new square and continue the process.

In the end, we obtain a chain of squares with sides contained in conjugates of the form ${}^{p_i}R$, ${}^{p_is}R$, ${}^{p_{i\pm1}}R$ and ${}^{p_it}R$ which are attached along the segments ${}^{p_i}L$ with

$$p_i = \begin{cases} (xst)^{-i}x & i \in \mathbb{Z}_- \\ x & i = 0 \\ (stx)^i & i \in \mathbb{Z}_+ \,. \end{cases}$$

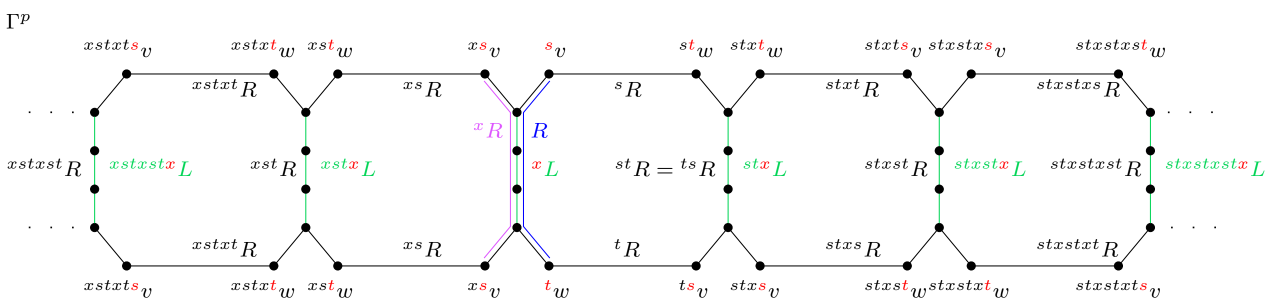

Figure 4.3.19: The letters in red can be added or left out, because they leave the element invariant.

By assumption *2.*, the (s,t)-square $S_{(s,t)}$ is an induced subgraph of Γ^p and by assumption *3.*, $S_{(s,t)}$ is not separated by the star of a vertex. By assumption *4.*, L is a segment of at least length 3. So, no star of any vertex separates this chain of squares. Thus, we can build a sequence $(v_i)_{i=0}^{\infty} \subseteq V(\Gamma^p)$, as suggested in item 1 of Outline 2.40: Choose $v_i \in {}^{p_i}L$ for $i \in \mathbb{Z}_+^0$. Then v_0 and v_i are connected in $\Gamma^p \setminus \{st(v)\}$ for every $v \in \Gamma^p \setminus \{st(v_i) \cup st(v_0)\}$ and for every $i \in \mathbb{Z}_+$. If W_Γ is QI to a RAAG A_Δ, then Γ^p is isomorphic to Δ^p by Theorem 2.31. Thus, there is a fundamental domain R_Δ of Δ^p. However, since R_Δ is finite, there is a $j \in \mathbb{Z}_+$ such that v_0 and v_j are in different conjugates of R_Δ. But then v_0 and v_j must be disconnected by the star of some vertex by Theorem 2.36. By the construction of the sequence in the chain of squares, this is not the case. Thus, the RACG W_Γ is not QI to any RAAG. □

Remark 4.56. Motivated by the shape of the picture in Figure 4.3.19, we refer to the chain of squares constructed in the proof of Theorem 4.55 as *ladder*. Accordingly, the conjugates of the segment L along which the different squares attach are the *rungs*, the trivial conjugate L is the *base rung* and the two connected, infinite, induced subgraphs not intersecting the interior of any rung are the *rails*.

Remark 4.57. Even if the MPRG Γ^p of a RACG W_Γ contains a ladder, it can still be a quasi-tree. A ladder only gives a lower bound on bottleneck constants. Thus, to use the existence of a ladder in an MPRG to argue about the existence of QIs, it is critical that by Theorem 2.31, the MPRG is preserved up to isomorphism by QIs of the group and not just preserved up to QI-type.

4.3.1 New Examples of RACGs not QI to a RAAG

Theorem 4.55 gives a new tool to investigate examples in the class of RACGs discussed in Section 4.1.3, in particular in Question 4: Start with a Γ' that has a DL-Λ and add an edge that is not contained in any square to create a graph Γ. We aim to show by application of Theorem 4.55 that W_Γ is not QI to a RAAG.

For ease of notation in the figures in this section, we fix:

Convention. Whenever an essential maximal product subgraph of a defining graph Γ is a suspension $S = \{a, b\} \circ S'$ with suspension points a and b, we denote the corresponding vertex in Γ^p by the identifying element ab. Note that while a and b do not commute, the element ba corresponds to the same suspension as the element ab. They both generate the same $\mathbb{Z}$ subgroup of $W_{\{a,b\}}$ that gives the *central direction* in the maximal product subcomplex of W_S.

The following example was introduced to the author by Pallavi Dani. It is $\mathcal{CFS}$, does not contain any stable or eccentric subgroup and does not split over a two-ended subgroup, so the techniques introduced in the previous sections are not applicable. The example was first understood not to be QI to a RAAG in collaboration with Christopher Cashen and Jingyin Huang and the proof was simplified by use of Theorem 4.55 in collaboration with Christopher Cashen and Annette Karrer:

Example 4.58. The graph Γ illustrated in Figure 4.3.20 is the graph of Figure 4.1.2 with one additional edge $(c, g) \in E(\Gamma)$. By Example 4.18, the RACG on the defining graph without this edge is QI to a RAAG by the DL-Algorithm. However, with this additional edge, the DL-Algorithm fails.

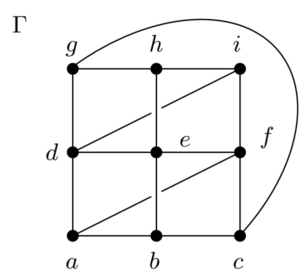

Figure 4.3.20

By application of Theorem 4.55, we show that in fact W_Γ is not QI to any RAAG: Γ has five essential maximal product subgraphs, which are all suspensions with the following suspension identifiers: dh, ei, df, ae and bf. Per the Convention, we use these elements also to denote the vertices in the MPRG Γ^p and their conjugates, see Figure 4.3.21. Since the edge (c, g) is not contained in any square, we obtain by Lemma 4.53 a (c, g)-square $S_{(c,g)}$.

With Remark 2.29 and the fact that all essential maximal product subgraphs are suspensions, we conclude that two vertices are connected by an edge in Γ^p if and only if the corresponding conjugated identifying elements commute. Thus, we immediately see that $S_{(c,g)}$ is an induced subgraph of Γ^p. In addition, we can explicitly check that there are no shortcuts of length two between any two vertices in $S_{(c,g)}$. Moreover, the vertex $e \in V(\Gamma)$ is not contained in $st(c) \cup st(g)$ and leaves all elements in the fundamental domain R invariant under conjugation. Thus, the fundamental domain R is the segment $L = {}^eL$ of length $5 \geq 3$ producing the rungs of the ladder illustrated in Figure 4.3.21. Thus, W_Γ is not QI to a RAAG.

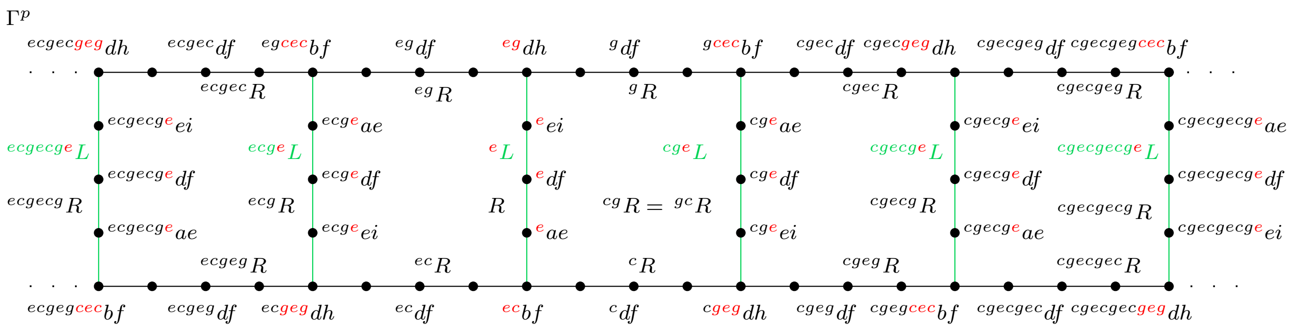

Figure 4.3.21: The letters in red leave the element invariant.

The following example, to which Theorem 4.55 was applied with Christopher Cashen and Annette Karrer, illustrates that even if the fundamental domain R is not a tree, a segment L of R can suffice to create an induced square in the MPRG.

Example 4.59. Consider the graph Γ illustrated on the left of Figure 4.3.22.

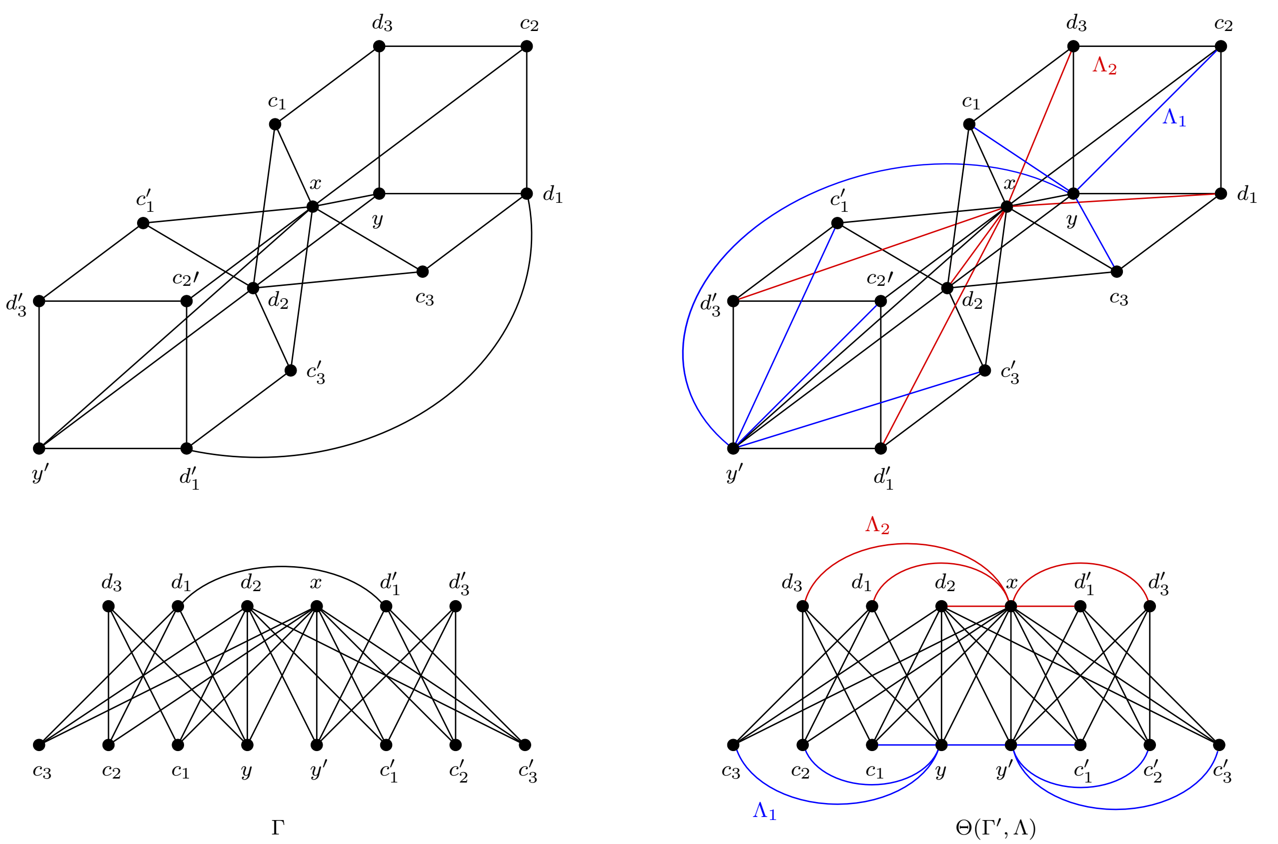

Figure 4.3.22

Note that Γ consists of two copies of the 1-skeleton of the 3-cube from Example 4.20, glued together along the diagonal x and d_2 plus the additional edge between the vertices d_1 and d_1', which is not contained in any square. The graph $\Gamma' = \Gamma \setminus \{(d_1, d_1')\}$ without this edge has a FIDL-Λ illustrated on the right of Figure 4.3.22, while the DL-Algorithm fails on Γ.

By Theorem 4.55, we can show that the RACG W_Γ is not QI to a RAAG: The edge (d_1, d_1') creates a (d_1, d_1')-square in the MPRG Γ^p by Lemma 4.53. On the left of Figure 4.3.23, we see how the four translates of the fundamental domain R of Γ^p, R, ${}^{d_1}R$, ${}^{d_1 d_1'}R$, ${}^{d_1'}R$ attach in Γ^p.

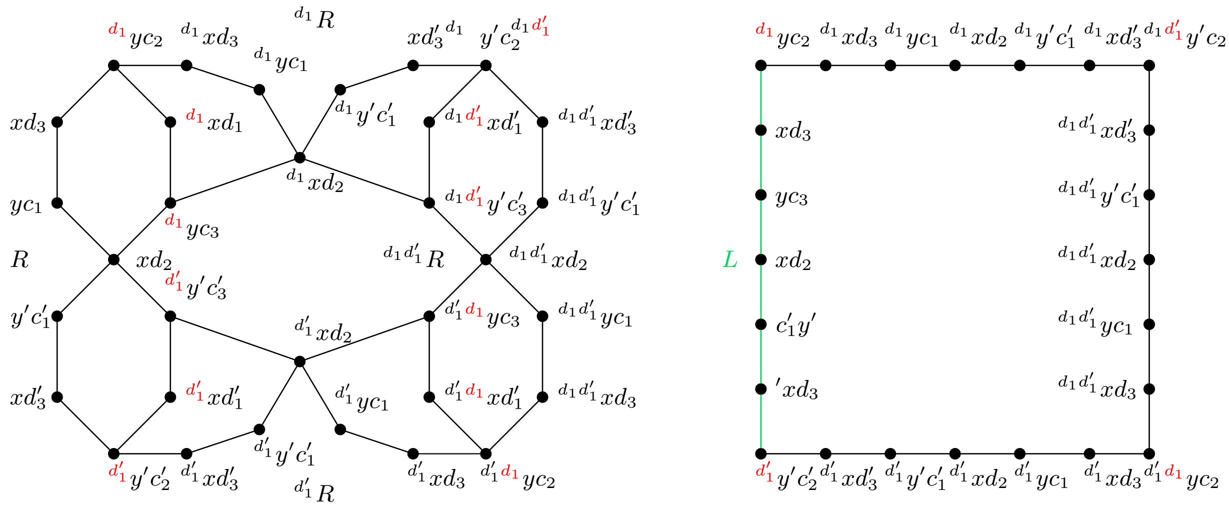

Figure 4.3.23: The letters in red leave the element invariant.

On the right of Figure 4.3.23 we see an induced (d_1, d_1')-square $S_{(d_1, d_1')}$ which exists in Γ^p but not in Γ'^p. Its side of length $6 \geq 3$ that is contained in R is what we define to be the base rung L. It is invariant under conjugation with the element $x \in V(\Gamma) \setminus \{st(d_1) \cup st(d_1')\}$ and thus produces the ladder in Figure 4.3.24. Thus, W_Γ is not QI to a RAAG.

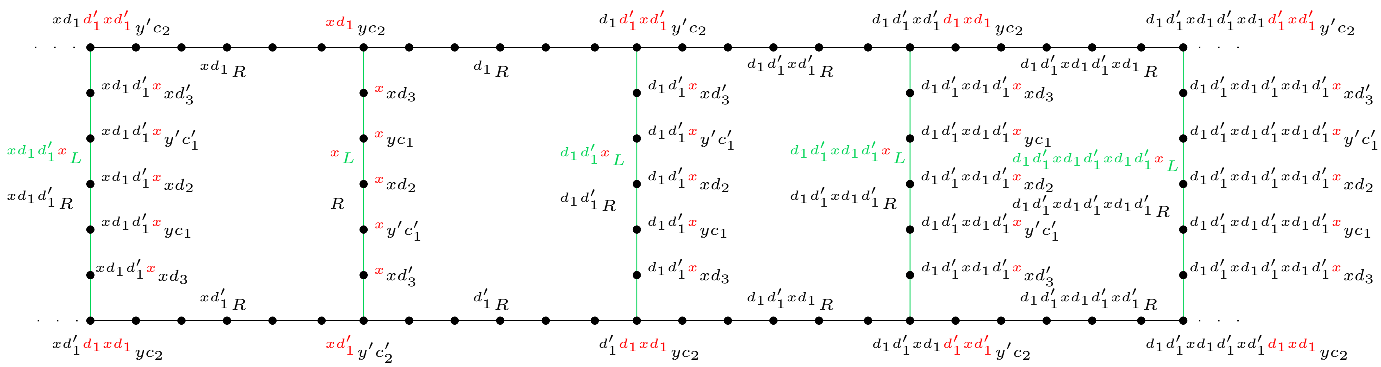

Figure 4.3.24: The letters in red leave the element invariant.

The following class of examples was introduced to the author by Pallavi Dani:

Example 4.60. The graph in Figure 4.3.20 can be generalized: We use the graph B on the left of Figure 4.3.25 as the building block that we stack n times. We add one edge not contained in a square connecting the two vertices of degree 2 to produce the graph Γ_n. Analogously to Example 4.58, an application of Theorem 4.55 shows that the RACG W_{Γ_n} is not QI to a RAAG for any $n \in \{2, 3, \dots\}$.

The following example investigated in collaboration with Christopher Cashen and Annette Karrer illustrates that the assumption in Theorem 4.55 that the base rung is at least of length 3 is necessary:

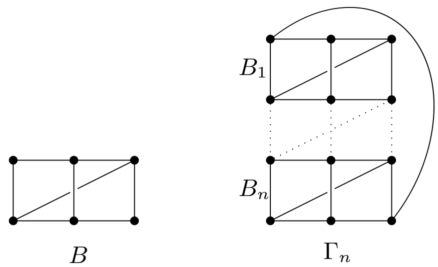

Figure 4.3.25

Example 4.61. The graph Γ, illustrated in Figure 4.3.26, is a generalized Θ-graph with one additional edge $(s,t) \in E(\Gamma)$ that is not contained in any square and the DL-Algorithm fails on Γ.

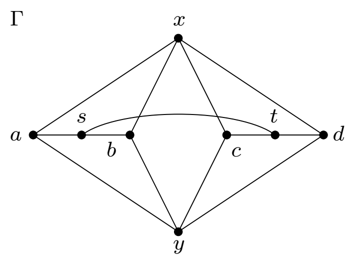

Figure 4.3.26

By Lemma 4.53 the edge produces an induced (s,t)-square $S_{(s,t)}$ in the MPRG Γ^p. However, its side is only of length 2. So, while it is invariant under conjugation with the element $x \in V(\Gamma) \setminus \{st(s) \cup st(t)\}$ and thus creates a ladder, see Figure 4.3.27, the rungs are too short. Removing the star of the vertex xy in Γ^p disconnects it.

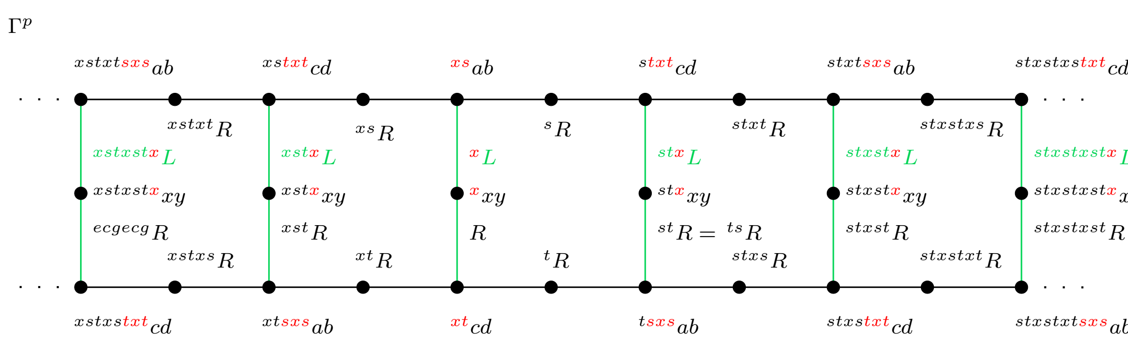

Figure 4.3.27: The letters in red leave the element invariant.

In fact, the RACG W_Γ is QI to a RAAG: We can double Γ, as illustrated in Figure 4.3.28, first along s to obtain the graph Γ_s and then Γ_s along t to obtain the graph Γ_{st}. So, by Lemma 1.43, $W_{\Gamma_{st}}$ provides an index 4 subgroup of W_Γ. However, the graph $W_{\Gamma_{st}}$ is also what we obtain by constructing a RACG commensurable to the RAAG A_{C_8}, where C_8 is the cycle graph on 8 vertices, following the Davis-Januszkiewicz construction in [DJ00].

A natural next step is the following:

Question 8. Determine a class of graphs $\mathcal{L}$ such that every $\Gamma \in \mathcal{L}$ satisfies the assumptions of Theorem 4.55, and thus, Γ^p contains a ladder with long enough rungs.

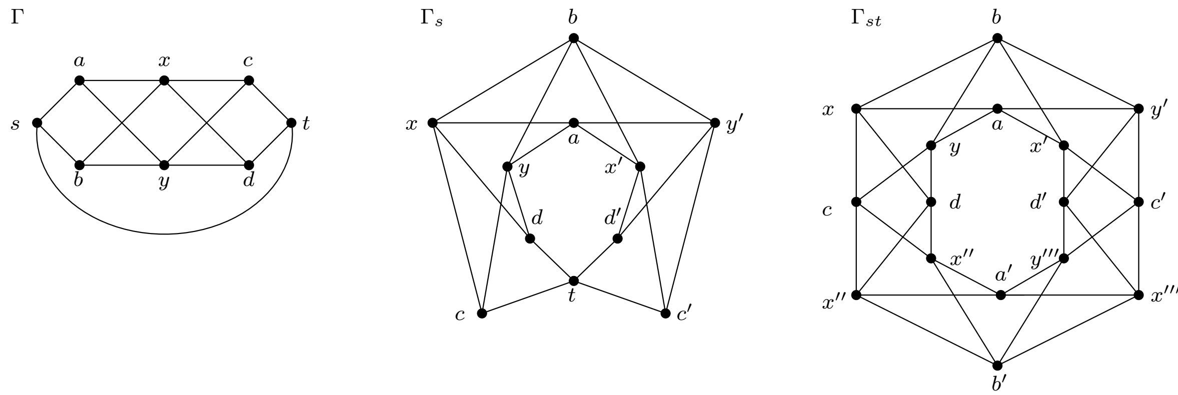

Figure 4.3.28: We double Γ first along s and then along t.

Outline 4.62*.* An edge in the defining graph Γ can cause an obstruction: By Lemma 4.53, it creates a cycle in the MPRG Γ^p that in some cases can be used to show that the Γ^p cannot be isomorphic to the MPRG of a RAAG. However, recall that by Example 4.61 such an edge in Γ does not necessarily prevent the existence of a QI between W_Γ and a RAAG.

So, this Outline 4.62 leads to the following:

Question 9. If Γ is triangle-free and has an edge not contained in a square, when does this imply that the RACG W_Γ is not QI to a RAAG?

4.3.2 Is the MPRG a quasi-tree?

Recall that by Theorem 2.39, the MPRG of a RAAG is always a quasi-tree. Thus, as highlighted in item 3 of Outline 2.40, if we can show for a given RACG W_Γ that its MPRG is not a quasi-tree, we can conclude by Theorem 2.31 that W_Γ is not QI to any RAAG. However, to the author's understanding, no results in this direction are known.

Question 10. When is the MPRG of a RACG W_Γ a quasi-tree?

If the MPRG of a RACG is not a quasi-tree, we would like to know:

Question 11. Is the MPRG of a RACG W_Γ always hyperbolic?

Examples investigated with Christopher Cashen and Pallavi Dani suggest that in certain cases, there is a connection between the connectedness of the Morse boundary of W_Γ and its MPRG:

Question 12. Does the Morse boundary of W_Γ determine whether its MPRG is a quasi-tree? In particular:

1. If Γ has an n-cycle creating a circle in the Morse boundary of W_Γ, can the MPRG Γ^p be a quasi-tree?
2. Suppose W_Γ has a circle in its Morse boundary coming from the 3-paths condition on Γ developed in [GKLS21] (see Section 1.4.2.2). When is the MPRG Γ^p a quasi-tree?
3. If W_Γ has totally disconnected Morse boundary, is the MPRG Γ^p always a quasi-tree?

In addition, an n-cycle creating a circle in the Morse boundary of the RACG W_Γ provides a stable special subgroup by Corollary 1.69. Thus, Question 12 has a follow-up:

Question 13. When is the orbit map $W_{\Gamma'} \to \Gamma^p$ of a stable special subgroup $W_{\Gamma'} \leq W_\Gamma$ of the RACG W_Γ acting on the MPRG Γ^p a QI embedding?

References

[Bar18] Benjamin Barrett, *Computing JSJ decompositions of hyperbolic groups*, J. Topol. **11** (2018), no. 2, 527–558.

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

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

[BFRHS18] Jason Behrstock, Victor Falgas-Ravry, Mark F. Hagen, and Tim Susse, *Global structural properties of random graphs*, Int. Math. Res. Not. IMRN (2018), no. 5, 1411–1441.

[BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups*, Geom. Topol. **21** (2017), no. 3, 1731–1804.

[BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula*, Pacific J. Math. **299** (2019), no. 2, 257–338.

[BM01] Noel Brady and John Meier, *Connectivity at infinity for right angled Artin groups*, Trans. Amer. Math. Soc. **353** (2001), no. 1, 117–132.

[BN08] Jason A. Behrstock and Walter D. Neumann, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141** (2008), no. 2, 217–240.

[Bog08] Oleg Bogopolski, *Introduction to group theory*, EMS Textbooks in Mathematics, European Mathematical Society (EMS), Zürich, 2008, Translated, revised and expanded from the 2002 Russian original.

[Bow98] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.

[BX20] Jordan Bounds and Xiangdong Xie, *Quasi-isometric rigidity of a class of right-angled Coxeter groups*, Proc. Amer. Math. Soc. **148** (2020), no. 2, 553–568.

[CCS23] Ruth Charney, Matthew Cordes, and Alessandro Sisto, *Complete topological descriptions of certain Morse boundaries*, Groups Geom. Dyn. **17** (2023), no. 1, 157–184.

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

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

[Cla14] Matt Clay, *When does a right-angled Artin group split over $\mathbb{Z}$?*, Internat. J. Algebra Comput. **24** (2014), no. 6, 815–825.

[CM11] Christopher H. Cashen and Nataša Macura, *Line patterns in free groups*, Geom. Topol. **15** (2011), no. 3, 1419–1475.

[CM17a] Christopher H. Cashen and Alexandre Martin, *Quasi-isometries between groups with two-ended splittings*, Math. Proc. Cambridge Philos. Soc. **162** (2017), no. 2, 249–291.

[CM17b] Matt Clay and Dan Margalit (eds.), *Office hours with a geometric group theorist*, Princeton University Press, Princeton, NJ, 2017.

[Cor17] Matthew Cordes, *Morse boundaries of proper geodesic metric spaces*, Groups Geom. Dyn. **11** (2017), no. 4, 1281–1306.

[Cox34] H. S. M. Coxeter, *Discrete groups generated by reflections*, Ann. of Math. (2) **35** (1934), no. 3, 588–621.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev, *Rank rigidity for CAT(0) cube complexes*, Geom. Funct. Anal. **21** (2011), no. 4, 851–891.

[Dan20] Pallavi 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] ©2020, pp. 107–141.

[Dav08] Michael W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008.

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

[DL20] Pallavi Dani and Ivan Levcovitz, *Right-angled Artin subgroups of right-angled Coxeter and Artin groups*, arXiv:2003.05531 (2020), 1–40.

[DST18] Pallavi Dani, Emily Stark, and Anne Thomas, *Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups*, Groups Geom. Dyn. **12** (2018), no. 4, 1273–1341.

[DT14] Pallavi Dani and Anne Thomas, *Quasi-isometry classification of certain right-angled Coxeter groups*, arXiv:1402.6224v1 (2014), 1–28.

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

[DT17] ———, *Bowditch's JSJ tree and the quasi-isometry classification of certain Coxeter groups*, J. Topol. **10** (2017), no. 4, 1066–1106.

[Dun85] M. J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), no. 3, 449–457.

[Edl21] Alexandra Edletzberger, *Quasi-isometries for certain right-angled Coxeter groups*, arXiv:2112.10463 (2021), 1–49.

[Efr53] Vadim A. Efremovic, *The proximity geometry of Riemannian manifolds*, Uspekhi Mat. Nauk **8** (1953), no. 5, 1–57.

[Fio22] Elia Fioravanti, *Automorphisms of contact graphs of* CAT(0) *cube complexes*, Int. Math. Res. Not. IMRN (2022), no. 5, 3278–3296.

[FP06] K. Fujiwara and P. Papasoglu, *JSJ-decompositions of finitely presented groups and complexes of groups*, Geom. Funct. Anal. **16** (2006), no. 1, 70–125.

[Gen22] Anthony Genevois, *Quasi-isometrically rigid subgroups in right-angled Coxeter groups*, Algebr. Geom. Topol. **22** (2022), no. 2, 657–708.

[Ger94] S. M. Gersten, *Quadratic divergence of geodesics in* CAT(0) *spaces*, Geom. Funct. Anal. **4** (1994), no. 1, 37–51.

[GKLS21] Marius Graeber, Annette Karrer, Nir Lazarovich, and Emily Stark, *Surprising circles in Morse boundaries of right-angled Coxeter groups*, Topology Appl. **294** (2021), Paper No. 107645, 3.

[GL11] Vincent Guirardel and Gilbert Levitt, *Trees of cylinders and canonical splittings*, Geom. Topol. **15** (2011), no. 2, 977–1012.

[GL17] ———, *JSJ decompositions of groups*, Astérisque (2017), no. 395, vii+165.

[Gor04] C. McA. Gordon, *Artin groups, 3-manifolds and coherence*, Bol. Soc. Mat. Mexicana (3) **10** (2004), 193–198.

[Gro84] Mikhael Gromov, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Warsaw, 1983), PWN, Warsaw, 1984, pp. 385–392.

[Hag14] Mark F. Hagen, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topol. **7** (2014), no. 2, 385–418.

[Har81] W. J. Harvey, *Boundary structure of the modular group*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), Ann. of Math. Stud., vol. No. 97, Princeton Univ. Press, Princeton, NJ, 1981, pp. 245–251.

[HST20] G. Christopher Hruska, Emily Stark, and Hung Cong Tran, *Surface group amalgams that (don't) act on 3-manifolds*, Amer. J. Math. **142** (2020), no. 3, 885–921.

[Hua16] Jingyin Huang, *Quasi-isometry classification of right-angled Artin groups II: several infinite out cases*, arXiv:1603.02372 (2016), 1–53.

[Hua17a] ———, *Quasi-isometric classification of right-angled Artin groups I: the finite out case*, Geom. Topol. **21** (2017), no. 6, 3467–3537.

[Hua17b] ———, *Top-dimensional quasiflats in* CAT(0) *cube complexes*, Geom. Topol. **21** (2017), no. 4, 2281–2352.

[Kar23] Annette Karrer, *Right-angled Coxeter groups with totally disconnected Morse boundaries*, Geom. Dedicata **217** (2023), no. 4, Paper No. 71, 40.

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

[KK14] Sang-Hyun Kim and Thomas Koberda, *The geometry of the curve graph of a right-angled Artin group*, Internat. J. Algebra Comput. **24** (2014), no. 2, 121–169.

[Kra09] Daan Krammer, *The conjugacy problem for Coxeter groups*, Groups Geom. Dyn. **3** (2009), no. 1, 71–171.

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

[Löh17] Clara Löh, *Geometric group theory*, Universitext, Springer, Cham, 2017, An introduction.

[Mar20] Alexander Margolis, *Quasi-isometry classification of right-angled Artin groups that split over cyclic subgroups*, Groups Geom. Dyn. **14** (2020), no. 4, 1351–1417.

[Mil68] J. Milnor, *A note on curvature and fundamental group*, J. Differential Geometry **2** (1968), 1–7.

[Mos73] G. D. Mostow, *Strong rigidity of locally symmetric spaces*, Annals of Mathematics Studies, vol. No. 78, Princeton University Press, Princeton, NJ; University of Tokyo Press, Tokyo, 1973.

[Mou88] Gabor Moussong, *Hyperbolic Coxeter groups*, ProQuest LLC, Ann Arbor, MI, 1988, Thesis (Ph.D.)–The Ohio State University.

[MT09] Michael Mihalik and Steven Tschantz, *Visual decompositions of Coxeter groups*, Groups Geom. Dyn. **3** (2009), no. 1, 173–198.

[NT19] Hoang Thanh Nguyen and Hung Cong Tran, *On the coarse geometry of certain right-angled Coxeter groups*, Algebr. Geom. Topol. **19** (2019), no. 6, 3075–3118.

[Oh22] Sangrok Oh, *Quasi-isometry invariants of weakly special square complexes*, Topology Appl. **307** (2022), Paper No. 107945, 46.

[Pap05] Panos Papasoglu, *Quasi-isometry invariance of group splittings*, Ann. of Math. (2) **161** (2005), no. 2, 759–830.

[Par97] Luis Paris, *Commensurators of parabolic subgroups of Coxeter groups*, Proc. Amer. Math. Soc. **125** (1997), no. 3, 731–738.

[Pra73] Gopal Prasad, *Strong rigidity of* $\mathbf{Q}$*-rank* 1 *lattices*, Invent. Math. **21** (1973), 255–286.

[PW02] Panos Papasoglu and Kevin Whyte, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv. **77** (2002), no. 1, 133–144.

[Rad01] David Glenn Radcliffe, *Unique presentation of Coxeter groups and related groups*, ProQuest LLC, Ann Arbor, MI, 2001, Thesis (Ph.D.)–The University of Wisconsin - Milwaukee.

[RS97] E. Rips and Z. Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. (2) **146** (1997), no. 1, 53–109.

[RST23] Jacob Russell, Davide Spriano, and Hung Cong Tran, *Convexity in hierarchically hyperbolic spaces*, Algebr. Geom. Topol. **23** (2023), no. 3, 1167–1248.

[Ser80] Jean-Pierre Serre, *Trees and amalgams*, Springer, 1980.

[Sta71] John Stallings, *Group theory and three-dimensional manifolds*, Yale Mathematical Monographs, vol. 4, Yale University Press, New Haven, Conn.-London, 1971, A James K. Whittemore Lecture in Mathematics given at Yale University, 1969.

[Šva55] A. S. Švarc, *A volume invariant of coverings*, Dokl. Akad. Nauk SSSR (N.S.) **105** (1955), 32–34.

[SW22] Sam Shepherd and Daniel J. Woodhouse, *Quasi-isometric rigidity for graphs of virtually free groups with two-ended edge groups*, J. Reine Angew. Math. **782** (2022), 121–173.

[Tra19] Hung Cong Tran, *On strongly quasiconvex subgroups*, Geom. Topol. **23** (2019), no. 3, 1173–1235.