

Subgroups of right-angled Coxeter groups via Stallings-like techniques

Pallavi Dani* and Ivan Levcovitz**

Abstract. We associate cube complexes called completions to each subgroup of a right-angled Coxeter group (RACG). A completion characterizes many properties of the subgroup such as whether it is quasiconvex, normal, finite-index or torsion-free. We use completions to show that reflection subgroups are quasiconvex, as are one-ended Coxeter subgroups of a 2-dimensional RACG. We provide an algorithm that determines whether a given one-ended, 2-dimensional RACG is isomorphic to some finite-index subgroup of another given RACG. In addition, we answer several algorithmic questions regarding quasiconvex subgroups. Finally, we give a new proof of Haglund’s result that quasiconvex subgroups of RACGs are separable.

Mathematics Subject Classification (2020). 20F65, 57M07, 20F55.

Keywords. Right-angled Coxeter groups, quasiconvex subgroups, algorithmic problems.

1. Introduction

In the highly influential article [37], Stallings introduced new tools to study subgroups of free groups. A crucial idea in Stallings’ work is that given a finite set of words in a free group, one can associate a labeled graph to this set, and perform a sequence of operations, now known as “Stallings folds,” to this graph. The resulting graph is, in some sense, a canonical object associated to the subgroup generated by the given words. This topological viewpoint provided clean new proofs for many theorems regarding subgroups of free groups. In [29], Kapovich–Miasnikov use Stallings’ ideas, recast in a more combinatorial form, to systematically study the subgroup structure of free groups and to answer a number of algorithmic questions about such free subgroups.

Arzhantseva–Olshanskii were the first to apply Stallings’ techniques to groups that are not free in [7], and Arzhantseva uses these techniques to show quasiconvexity of certain subgroups in [4]. More recently, Kharlampovich–Miasnikov–Weil use similar

*The first author was supported by a grant from the Simons Foundation (#426932, Pallavi Dani) and by NSF Grant No. DMS-1812061.

**The second author was supported by the Israel Science Foundation and in part by a Technion fellowship.

techniques to study several algorithmic questions for automatic groups (which include RACGs) [30]. Beeker–Lazarovich, building on work of Brown [13], define a version of Stallings folds for CAT(0) cube complexes, which they use to give a characterization of finitely presented quasiconvex subgroups of cubulated hyperbolic groups, in terms of hyperplane stabilizers [9]. Other related articles are discussed in Section 1.1.

The main goal of this article is to study subgroups of right-angled Coxeter groups (RACGs) through generalizations of Stallings’ techniques. Restricting to RACGs (rather than considering all automatic groups or CAT(0) cube complex groups) allows for some stronger results. We do not make any hyperbolicity assumption, which is often crucial in more general frameworks. Our approach is inspired by the combinatorial methods used in [29], and is quite different from that of [9]. This is further discussed in Section 1.1.

Given a finite simplicial graph Γ , the associated RACG W_Γ is generated by order two elements corresponding to vertices of Γ , with the additional relations that two such generators commute if there is an edge in Γ between the corresponding vertices. RACGs form a wide class of groups which have become central objects in geometric group theory. We refer to [17] for a survey of recent work on these groups. One interesting feature of RACGs is that they have a rich variety of subgroups, which includes all free groups, RAAGs (right-angled Artin groups) [21] and surface groups. Incredibly, all hyperbolic 3-manifold groups [1, 40] and Coxeter groups [27] are virtually subgroups of RACGs as well.

Given a subgroup G of a RACG W_Γ , we abstractly define a completion of G as an edge-labeled cube complex satisfying certain properties. If G is additionally finitely generated, we explicitly build a *standard* completion for G by the following procedure. Starting with a subdivided “rose” graph whose petals are labeled by the words generating G , we perform a sequence of operations of three possible types: fold, cube attachment, and cube identification. A completion Ω can always be obtained as the direct limit of the complexes in this sequence. Many properties of the subgroup G can be characterized in terms of properties of Ω , as summarized in the following theorem:

Theorem A. *Let G be a subgroup of the RACG W_Γ . Then*

- (1) *G is quasiconvex in W_Γ if and only if G is finitely generated and every (equivalently, some) standard completion for G is finite (Theorem 8.4);*
- (2) *there exist characterizations of G having finite index (Theorem 6.6), G being torsion-free (Proposition 4.6), and G being normal (Theorem 5.3) in terms of properties of a completion of G ;*
- (3) *if G is torsion-free, then any completion is non-positively curved (Proposition 7.2) and has fundamental group isomorphic to G (Theorem 4.7).*

We note that quasiconvex subgroups are always quasi-isometrically embedded and, in RACGs, are separable [25]. We discuss the relation of our results to [9, 30] and other works in Section 1.1.

Our approach to producing completions is constructive. In particular, given a quasiconvex subgroup, a finite completion can be explicitly constructed in finite time (by Theorem A (1) and Proposition 3.5). Furthermore, as our completions are cube complexes (and not just graphs as is the case in many generalizations of Stallings' methods to non-free settings), powerful tools from cubical geometry can be applied to study them. This enables us to use the characterizations in Theorem A to prove several new results on subgroups of RACGs including showing that particular classes of subgroups of RACGs are always quasiconvex and solving various algorithmic problems (see Theorem B–Theorem E below).

We expect that Theorem A will have other applications. For instance, the characterization of finite-index subgroups in Theorem A is an essential tool in our proof of [18, Theorem B], which states that any 2-dimensional, one-ended RACG with planar defining graph is quasi-isometric to a RAAG if and only if it is commensurable to a RAAG.

A given subgroup may have multiple completions. Indeed two completions for a given subgroup need not even be homotopy equivalent (see Example 3.8). Nevertheless, our results do not depend on the specific completions chosen. In fact, the flexibility of choosing different generating sets and order of operations is useful in our arguments. Despite the non-uniqueness of completions, every completion for G has a 1-dimensional subcomplex called its core graph, and any two core graphs for G are isomorphic (see Proposition 5.2).

Theorem A provides a tool to show that subgroups of a RACG are quasiconvex, by showing that their associated completions must be finite. We apply this approach to subgroups of W_Γ generated by *reflections*, i.e., conjugates of the generators corresponding to vertices of Γ . We call such subgroups *reflection subgroups*. More precisely, we prove:

Theorem B (Theorem 10.5). *Every finitely generated reflection subgroup of a RACG is quasiconvex.*

We next turn our attention to Coxeter subgroups, i.e. subgroups that are isomorphic to some abstract finitely generated Coxeter group. A result proved independently by Dyer [23] and Deodhar [22] shows that every reflection subgroup of a RACG is a Coxeter subgroup. The converse to this statement is not true in general (see Remark 11.6), but we show that it holds under certain hypotheses:

Theorem C (Theorem 11.4, Corollary 11.5). *Every one-ended Coxeter subgroup of a 2-dimensional RACG is a reflection subgroup. Consequently, every such Coxeter subgroup is quasiconvex by Theorem B.*

Completions can be used to answer several algorithmic questions about subgroups of RACGs. For instance, we consider the problem of finite-index embeddability between RACGs:

Theorem D (Theorem 12.8). *There is an algorithm (explicitly constructed in Section 12) which, given a one-ended, 2-dimensional RACG W_Γ , and any RACG $W_{\Gamma'}$, determines whether or not $W_{\Gamma'}$ is isomorphic to a finite-index subgroup of W_Γ .*

The above theorem gives an algorithm that can often determine when two RACGs are commensurable; thus, it provides a tool for studying commensurability classification. A few specific families of RACGs have been classified up to commensurability (see [16, 19, 28]), but not much is known in general. We remark that the precise statements of Theorem C and Theorem D use a significantly weaker hypothesis than one-endedness.

As noted above, when G is quasiconvex, a finite completion for G can be computed in finite time. Consequently, we can provide algorithms to check several properties of G :

Theorem E. *Let G be a quasiconvex subgroup of a RACG W_Γ , given by a finite generating set of words in W_Γ . Then there exist algorithms (explicitly constructed in Section 13) to solve the following problems.*

- (1) *Determine whether or not G is torsion-free.*
- (2) *Determine the index of G in W_Γ (even if infinite).*
- (3) *Given a word w representing an element $g \in W_\Gamma$, determine whether or not a positive power of g is in G .*
- (4) *Determine whether or not G is normal in W_Γ .*

Finally, we use completions to recover some known results. Specifically, we give new proofs of the well known result that RACGs are residually finite (see Theorem 9.3), and of a result of Haglund that every quasiconvex subgroup of a RACG is separable (see Theorem 9.5).

We note that much of the work presented here can also be used to study RAAGs. Any RAAG A embeds as a finite-index subgroup of a RACG W by a construction of Davis–Januszkiewicz [21]. Thus, one can construct a completion for G considered as a subgroup of W , and then use Theorem A and Theorem E to study properties of G as a subgroup of A .

1.1. Relation to other works. Every RACG W_Γ admits an automatic structure whose associated language consists of reduced, i.e. geodesic, words of W_Γ [12]. Thus, the algorithms provided by Kharlampovich–Miasnikov–Weil in [30] for quasiconvex subgroups of automatic groups apply to RACGs. These include computing their intersections and (in some cases including hyperbolic groups) the conjugacy and almost malnormality problems. The algorithmic applications in [30] are distinct from those in our Theorems D and E. We note that, for automatic groups satisfying an additional assumption which does not hold for all RACGs, [30] gives an algorithm for (2) of Theorem E, i.e. determining the index of a quasiconvex subgroup.

To provide the algorithms above, Kharlampovich–Miasnikov–Weil show that finite *Stallings-like graphs* can be constructed for subgroups that are quasiconvex with respect to an automatic structure. We refer to [30] for a precise definition, but note that the 1-skeleton of every completion, in our sense, is a Stallings-like graph.

However, completions have strictly stronger properties than Stallings-like graphs. For instance, completions are *cube-full*, which means they do not have “missing relations” (see Section 3). A priori, there is no guarantee that a Stallings-like graph can be completed to a finite cube-full complex, and this is what we prove can be done for quasiconvex subgroups of RACGs (see Theorem A). The cube-full property is essential to our arguments. It allows us to prove, in many cases, that completions are non-positively curved, which is a crucial step in our proofs of Theorem B and Theorem D.

Another important property of the completion of a subgroup $G < W_\Gamma$ is that, given *any* geodesic word w of W_Γ representing an element of G , there exists a loop with label w in the completion. Stallings-like graphs do not necessarily have this property, which can be traced back to the fact that, in general, the language associated to a geodesically automatic structure is strictly smaller than the set of geodesic words. Additionally, loops in a completion whose labels are reduced expressions for the same group element have Hausdorff distance proportional to the number of “moves” required to get from one expression to the other (see Section 4). These properties are used throughout our paper.

To produce Stallings-like graphs, Kharlampovich–Miasnikov–Weil begin with a construction similar to ours and many others: they start with a rose graph, and alternately fold and “add relations.” For each folded graph Λ in the resultant sequence, they use a separate algorithm, which involves constructing a new automaton associated to Λ , to check if the graph obtained is Stallings-like (and stopping if it is). In contrast, our algorithm only involves checking, at each stage, whether the complex is folded and cube-full (and stopping if it is). Consequently, it is a more practical algorithm for doing explicit computations in RACGs.

The name “completion” is inspired by the terminology in [30], which talks of a “completion process” for constructing Stallings-like graphs. We caution the reader that our completions are not the same as the canonical completions considered in [26].

We mention some other works which use generalizations of Stallings’ methods to study algorithmic and structural properties of subgroups, referring to [30] for a more detailed summary.

Arzhantseva–Olshanskii use Stallings-like graphs to study groups that are generic in a certain statistical sense [7]. They define a series of operations, now known as AO-moves, that, when applied to a Stallings-like graph, preserve the subgroup associated to it. These techniques have been generalized to the free product setting in [38] and [24], and the latter uses a notion conceptually similar to our completions. In a series of articles, Arzhantseva and Arzhantseva–Cherix utilize Stallings-like graphs to prove that several results hold generically in finitely presented groups [2–4, 6]. Notably, Stallings-like graphs are first used to show quasiconvexity of subgroups in [4].

McCammond–Wise use a generalization of Stallings’ construction to study local quasiconvexity and coherence in groups satisfying certain small cancellation

conditions [32]. Schupp uses Stallings graphs to show that certain extra-large type Coxeter groups (which don't include RACGs) are locally quasiconvex and to answer algorithmic questions [36].

Beeker–Lazarovich define a version of Stallings folds for CAT(0) cube complexes and use it to give a characterization of finitely presented quasiconvex subgroups of hyperbolic groups that act geometrically on CAT(0) cube complexes [9]. This generalizes work of Brown on hyperbolic VH square complexes [13]. The Beeker–Lazarovich characterization is different in nature to ours; it is in terms of finiteness and quasiconvexity properties of stabilizers of hyperplanes, conditions which are not necessarily easier to check algorithmically. Despite also using folds, their procedure follows a very different outline. Given a hyperbolic group H acting geometrically on a CAT(0) cube complex X , and a finitely presented subgroup G , they start with an equivariant map $K \rightarrow X$, where K is the universal cover of the presentation complex of G . They then apply a “resolution process” (generalizing work of Dunwoody) which uses the pullback in K of the wall structure of X to produce an unbounded CAT(0) cube complex Y . Next, folds are G -equivariantly applied to Y to obtain the cube complex associated to H . In contrast, our completion process is algorithmic in nature (our starting point is an explicit generating set), and the results we obtain do not require any hyperbolicity and finite-presentation hypotheses. Indeed, many RACGs are not even relatively hyperbolic. On the other hand, we strongly rely on combinatorial properties of RACGs which do not hold in the generality of cubulated groups.

Other authors have studied “cores” associated to quasiconvex subgroups. For instance, Sageev–Wise show that relatively quasiconvex subgroups of relatively hyperbolic groups acting on a finite-dimensional, locally finite CAT(0) cube complex admit a convex core [35]. This generalizes ideas of Haglund [25], who shows that a quasiconvex subgroup G of a RACG acts cocompactly on the combinatorial convex hull $\Sigma(G)$ of G in the Davis complex. Indeed the quotient $\Sigma(G)/G$ is very close to being a completion (one has to take care, as cubes may be folded along midcubes under this quotient). We note that characterizations of quasiconvexity using this approach cannot immediately be used to answer algorithmic questions.

In the spirit of Theorem D, Kim–Koberda find conditions for a RAAG to be realized as a (not necessarily finite-index) subgroup of another in terms of properties of the associated extension and clique graphs [31]. Using their work, Casals-Ruiz gives an algorithm which determines if a 2-dimensional RAAG is isomorphic to a subgroup of another RAAG [15].

Acknowledgements. The authors would like to thank Jason Behrstock, Anthony Genevois, Ilya Kapovich, Sang-hyun Kim, Thomas Koberda, Ignat Soroko, and the referee for comments. The second author would like to thank Nir Lazarovich and Michah Sageev for helpful conversations.

2. Preliminaries

Given a graph Γ , we denote the vertex and edge sets of Γ by $V(\Gamma)$ and $E(\Gamma)$, respectively.

2.1. Right-angled Coxeter groups. We summarize some well known facts regarding right-angled Coxeter group (RACGs) which are needed throughout this article. We refer the reader to [17] for a survey on RACGs and to [20] and [10] as references on Coxeter groups.

Let Γ be a simplicial graph with finite vertex set $S = V(\Gamma)$ and edge set $E = E(\Gamma)$. The RACG W_Γ associated to Γ is the group given by the presentation:

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

We say that S is a *standard Coxeter generating set* for W_Γ . Given $s, t \in V(\Gamma)$, we write $m(s, t) = 1$ if $s = t$, $m(s, t) = 2$ if s is adjacent to t and $m(s, t) = \infty$ otherwise.

We refer to the elements of S as *letters*. A word w in W_Γ is a (possibly empty) sequence of letters in S . Let $w = s_1 \dots s_n$ be a word in W_Γ , where $s_i \in S$ for $1 \leq i \leq n$. We let $|w| = n$ denote the *length* of w . If w' is another word in W_Γ such that w and w' are equal as elements of W_Γ , then we say that w' is an *expression* for w . We say that w is *reduced* if $|w| \leq |w'|$ for any expression w' for w . Finally, we define the support of w , denoted by $\text{Support}(w)$, to be the set of vertices of Γ which appear as a letter in w .

Given a graph Γ and a subset V' of $V(\Gamma)$, the graph Δ *induced* by V' is the graph which has vertex set V' and an edge between two vertices of V' if and only if there is an edge between them in Γ . We also say that Δ is an induced subgraph of Γ .

Throughout this article, given any simplicial graph Γ , we will always denote by W_Γ the corresponding RACG. If Δ is an induced subgraph of a graph Γ , then W_Δ is naturally isomorphic to the subgroup of W_Γ generated by the generators corresponding to vertices of Δ (see for instance [20]). Such a subgroup of W_Γ is called a *special subgroup*.

Given a vertex v of Γ , the *link of v* , denoted by $\text{link}(v)$, is the set of all vertices of Γ which are adjacent to v . The *star of v* , denoted by $\text{star}(v)$, is the set $\text{link}(v) \cup \{v\}$. We will often consider the special subgroup of W_Γ generated by the link or star of a vertex.

Recall that a graph is a *clique* if any pair of distinct vertices of the graph are adjacent. A RACG W_Γ is finite if and only if Γ is a clique. Furthermore, W_Γ is one-ended if and only if Γ is connected and does not contain a clique which separates Γ [33]. We say a graph is *triangle-free* if it does not contain a subgraph that is a clique with three vertices, i.e. a triangle. If Γ is triangle-free, we say that the RACG W_Γ is *2-dimensional*.

A graph Γ decomposes as a *join graph* $\Gamma = \Gamma_1 \star \Gamma_2$ if there are induced subgraphs Γ_1 and Γ_2 of Γ such that $V(\Gamma) = V(\Gamma_1) \cup V(\Gamma_2)$ and v_1 and v_2 are adjacent in Γ for all $v_1 \in V(\Gamma_1)$ and $v_2 \in V(\Gamma_2)$. The graph Γ decomposes as a join $\Gamma = \Gamma_1 \star \Gamma_2$ if and only if $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$.

We say that a subgroup of a Coxeter group is a *Coxeter subgroup* if it is isomorphic to some finitely generated Coxeter group. In our setting, it follows from [20, Theorem 12.3.4] and [10, Proposition 1.1.1] that Coxeter subgroups are themselves RACGs:

Proposition 2.1. *Let G be a Coxeter subgroup of a RACG, then G is a RACG.*

We will need the following definition in Section 12.

Definition 2.2. Given a graph Γ and $s \in V(\Gamma)$, let Λ be the subgraph of Γ induced by $V(\Gamma) \setminus \{s\}$. We define $D(\Gamma, s)$ to be the graph consisting of the union of two copies of Λ which are identified along the subgraph of Λ induced by $\text{link}(s)$.

It is well known that $W_{D(\Gamma,s)}$ is isomorphic to the kernel K of the homomorphism $\phi_s: W_\Gamma \rightarrow \mathbb{Z}_2$ defined by $\phi_s(s) = 1$ and $\phi_s(t) = 0$ for all $t \in V(\Gamma) \setminus \{s\}$. To see this, first observe that K is generated by

$$T = \{t, sts \mid t \in V(\Gamma) \setminus \{s\}\}.$$

Since T consists of reflections, a result proved independently by Dyer [23] and Deodhar [22] implies that K is a Coxeter group, and K is therefore a RACG by Proposition 2.1. By the criterion in the first paragraph of [23], the set T' is a Coxeter generating set, where T' consists of all reduced reflections r in K with the property that for every reflection $r' \neq r$ in K , the word $r'r$ cannot be represented by a word of length less than $|r|$. It is easily seen from Tits' solution to the word problem that $T \subset T'$. Since both these sets generate K and as no subset of T' can generate K , it follows that $T = T'$, and consequently, that K is isomorphic to $W_{D(\Gamma,s)}$.

2.2. The word problem in RACGs. We discuss Tits' solution to the word problem and the deletion property in RACGs. We refer the reader to [20] or [10] for more details.

Let $w = s_1 \dots s_n$ be a word in the RACG W_Γ . Suppose that $m(s_i, s_{i+1}) = 2$ for some $1 \leq i \leq n$. Then we may "swap" the letters s_i and s_{i+1} to obtain another expression $w' = s_1 \dots s_{i-1} s_{i+1} s_i s_{i+2} \dots s_n$ for w . We say that w' is obtained from w by a *swap move* or by *swapping* s_i and s_{i+1} . On the other hand, if $s_i = s_{i+1}$ (as vertices of Γ) for some $1 \leq i \leq n$, then we can obtain an expression $w' = s_1 \dots s_{i-1} s_{i+2} \dots s_n$ for w by *cancelling* s_i and s_{i+1} .

Let w be a word in W_Γ , and let w' be a reduced expression for w . There exists a sequence of words $w = w_1, \dots, w_m = w'$ such that w_{i+1} is obtained from w_i by either a swap move or a cancellation. This is known as Tits' solution to the word problem. We call such a sequence of expressions for w a *sequence of Tits moves*.

Given a word $w = s_1 \dots s_n$ in W_Γ , suppose $s_i = s_{i'} = s$ for some $1 \leq i < i' \leq n$ and that $m(s, s_j) = 2$ for all $i < j < i'$. It follows that

$$w' = s_1 \dots s_{i-1} s_{i+1} \dots s_{i'-1} s_{i'+1} \dots s_n$$

is an expression for w . We say that w' is obtained from w by a *deletion* (as two occurrences of s have been deleted). We remark that this definition is slightly stronger than the classical definition of a deletion in general Coxeter groups. The *deletion property* states that if w is a word in W_Γ such that w is not reduced, a deletion can be applied to w . The proof of this fact is similar to that of the corresponding statement for the standard definition of a deletion [8]. The deletion property guarantees that a reduced expression for w can be obtained by performing a sequence of deletions.

The following lemma, which is required later, easily follows from the deletion property.

Lemma 2.3. *Let h and k be reduced words in a RACG W_Γ . Then there is a reduced expression $\widehat{h}\widehat{k}$ for the word hk such that $\widehat{h}s_1 \dots s_m$ is a reduced expression for h and $s_m \dots s_1\widehat{k}$ is a reduced expression for k where $s_i \in V(\Gamma)$ for $1 \leq i \leq m$.*

2.3. Cube complexes. A *cube complex* is a cell complex whose cells are Euclidean unit cubes, $[-\frac{1}{2}, \frac{1}{2}]^n$, of varying dimension. We refer the reader to [14] and to [41] for a background on cube complexes.

Let Γ be a simplicial graph. A cube complex is Γ -*labeled* if every edge in its 1-skeleton is labeled by a vertex of Γ . In this article, the labels of a Γ -labeled cube complex will be the generators of the RACG W_Γ . Given a simplicial path α in the 1-skeleton of a Γ -labeled complex, the *label of α* is the word formed by the sequence of labels of consecutive edges in α .

We say that a cube complex Ω is *non-positively curved* if the (simplicial) link of each vertex in Ω is a flag complex. (Recall that a simplicial complex is called a *flag complex* if any finite clique in its 1-skeleton spans a simplex.) If Ω is both non-positively curved and simply connected then Ω is a CAT(0) cube complex.

A *path* in the cube complex Ω is a simplicial path in its 1-skeleton. Given a path p , we denote by $|p|$ the number of edges in p . Given two paths, p and p' , such that the endpoint of p is equal to the startpoint of p' , we let pp' denote their concatenation. A *loop* in a complex is defined to be a closed path (possibly with backtracking). We define a *graph-loop* to be an edge in a complex that connects a vertex to itself.

We will work with the combinatorial path metric on cube complexes. Namely, given two vertices of Ω , we define their distance to be the length of a shortest path in Ω between them.

A *midcube* of a cube $c = [-\frac{1}{2}, \frac{1}{2}]^n$ is the restriction of one of the coordinates of c to 0. A *hyperplane* H in Ω is a maximal collection of midcubes in Ω , such that for any two midcubes m and m' in H , it follows there is a sequence of midcubes $m = m_1, \dots, m_n = m'$ in H such that $m_i \cap m_{i+1}$ is a midcube in Ω for all $1 \leq i < n$.

The *carrier* of a hyperplane H , denoted by $N(H)$, is the set of all cubes which have non-empty intersection with H . If H intersects an edge e , then we say that e is dual to H .

Let Ω be a CAT(0) cube complex. A path γ in Ω is geodesic if and only if every hyperplane is dual to at most one edge of γ . Thus, if γ is geodesic then $|\gamma|$ is equal to the number of hyperplanes intersecting γ . Given a hyperplane H in Ω , $\Omega \setminus H$ contains exactly two components, and the carrier $N(H)$ is convex in the combinatorial path metric. We refer the reader to [41, Chapters 3.2 and 3.3] for proofs of these well known facts.

2.4. Disk diagrams in cube complexes. We now recall some basic facts about disk diagrams, and refer to [40] and [41] for further details.

In the setting of cube complexes, a *disk diagram* D is a contractible, finite, 2-dimensional cube complex (i.e., a square complex) equipped with a planar embedding $\Psi: D \rightarrow \mathbb{R}^2$. The map Ψ gives a natural cellulation of the 2-sphere $S^2 = \mathbb{R}^2 \cup \infty$. We call the path traced by an attaching map of the cell containing ∞ in this cellulation the *boundary of D* and denote it by ∂D .

Given a cube complex Ω , a *disk diagram over Ω* is a disk diagram D which admits a map $\Phi: D \rightarrow \Omega$ mapping n -cubes isometrically onto n -cubes (i.e., a combinatorial map). As the edges of the cube complexes we consider in this article will be labeled, we accordingly further require the edges of D to be labeled and the map from D to Ω to respect this labeling.

Given a closed null-homotopic loop $p: \mathbb{S}^1 \rightarrow \Omega$, the van Kampen Lemma says that there exists a disk diagram in Ω given by some $\Phi: D \rightarrow \Omega$, and an identification of ∂D with \mathbb{S}^1 such that Φ restricted to ∂D is equal, as a map, to p (see, for instance, [41, Lemma 3.1]).

Given a disk diagram D in Ω and an edge e of D , the *dual curve intersecting e* is the hyperplane in D dual to e . As D is planar, a dual curve in D can intersect at most two edges along ∂D .

Let W_Γ be a RACG. The Cayley graph of W_Γ (with the usual presentation) is the 1-skeleton of a CAT(0) cube complex (the *Davis complex*) whose 2-cells correspond to the commuting relations between generators (see for instance [20]). In this setting, a Γ -labeled disk diagram over the Davis complex of W_Γ has the property that opposite sides of squares must have the same label. In particular, we can define the *type* of a dual curve to be the label of an edge (equivalently, all edges) intersecting the dual curve. Furthermore, if two dual curves intersect, then their types are adjacent vertices of Γ . These observations are used in the following technical lemma which is useful in later sections. Note that any word considered in the lemma below could be the empty word.

Lemma 2.4. *Let w and z be words that are equal as elements of a RACG W_Γ . Suppose that $w = w'w''$ and $z = z'z''$, where w', w'', z' and z'' are words in W_Γ , and suppose that z' is reduced. Let D be a disk diagram whose boundary label,*

starting from a base vertex v , is wz^{-1} . Let $\alpha_{w'}$ and $\alpha_{z'}$ be the oriented paths starting at v in ∂D whose labels, read in the direction of the orientation, are w' and z' , respectively. Suppose further, that every dual curve intersecting $\alpha_{w'}$ also intersects $\alpha_{z'}$. Then z' has a reduced expression $w'x$, where x is some word in W_Γ .

Proof. Let H_1, \dots, H_k , with $k = |\alpha_{w'}|$, be the dual curves intersecting $\alpha_{w'}$ ordered by the orientation of $\alpha_{w'}$ (i.e. for $i < j$, $H_i \cap \alpha_{w'}$ occurs before $H_j \cap \alpha_{w'}$ along this orientation). This order is well-defined, since each H_i necessarily also intersects $\alpha_{z'}$ by hypothesis, and therefore intersects $\alpha_{w'}$ exactly once.

Let e_1, \dots, e_k be the first k edges of $\alpha_{z'}$ ordered by its orientation. If H_i intersects e_i for each $1 \leq i \leq k$, then we are done. For then w' is precisely the initial subword of z' of length k . Otherwise, let c be the smallest integer such that H_c does not intersect e_c . In particular, the initial subword of w' of length $c - 1$ is equal to the initial subword of z' of length $c - 1$.

We claim that we can construct a new disk diagram D' with boundary label $w'w''(yz'')^{-1}$, such that y is a reduced word equal in W_Γ to z' , every dual curve intersecting the part of $\partial D'$ with label w' also intersects the part of $\partial D'$ labeled y , and the initial subword of y of length c is equal to the initial subword of w' of length c . The lemma then follows by iteratively applying this claim.

Orient the edges of $\alpha_{z'}$ with the induced orientation from $\alpha_{z'}$. Let e be the edge of $\alpha_{z'}$ that intersects H_c . Define β to be the subpath of $\alpha_{z'}$ from v to the initial vertex of e if $c = 1$ and from the terminal vertex of e_{c-1} to the initial vertex of e otherwise. If H is a dual curve intersecting β , then $H \neq H_i$ for $1 \leq i \leq c$ by construction, and H cannot intersect $\alpha_{z'}$ twice, since z' is reduced. Thus, any such H must intersect H_c . It follows that the label s of e commutes with the label b of β . We may now attach a $|b| \times 1$ rectangle with label $bsb^{-1}s$ to the ∂D by identifying the sides labeled b and s of the rectangle with β and e respectively, to get a new diagram D' . Define y to be the word obtained from z' by replacing the subword labeled bs by sb . It is clear by construction that D' has the properties in the claim. \square

3. A complex for subgroups of a RACG

The main goal of this section is to define a completion of a subgroup of a RACG, and to construct completions for finitely generated subgroups. We begin by defining a completion of a Γ -labeled complex as the direct limit of a certain sequence of Γ -labeled complexes. We then show that there is a natural labeled graph associated to any finite generating set, such that a completion of this graph is also a completion of the group generated by the set.

3.1. Completion of a complex. Let Γ denote a simplicial graph. In this paper, we only consider Γ -labeled cube complexes whose labeled cubes have two additional

properties. Firstly, any pair of edges dual to a common mid-cube have the same label. As a result, hyperplanes in the cube complex have a well-defined label. Secondly, given a cube in such a complex and a set of edges of this cube which are all incident to a common vertex, no two edges in this set have the same label, and the full subgraph of Γ induced by the vertices of Γ corresponding to the labels of the edges in this set is a clique. When we mention a Γ -labeled cube complex, it will be implicit that the labeling has these additional properties.

Let C be a Γ -labeled cube complex. We describe three operations that can be applied to C to produce a new Γ -labeled cube complex.

Fold operation. A *fold operation* corresponds to collapsing a pair of adjacent edges with the same label into a single edge. More precisely, for $i = 1, 2$, let e_i be an edge in C with endpoints v and v_i , where $e_1 \neq e_2$, but two or more of the vertices v, v_1, v_2 could be equal. Furthermore, suppose that e_1 and e_2 have the same label. Temporarily orient the edge e_i from v to v_i (choosing the orientation arbitrarily if $v = v_i$, i.e. if e_i is a graph-loop). Then the fold operation consists of forming a quotient of C by identifying e_1 and e_2 so that their orientations agree, and then forgetting the orientation.

We remark that although the fold map corresponding to e_1 and e_2 is not unique when one of these edges is a graph-loop, this does not affect any of our applications.

Cube identification operation. This is a higher dimensional analogue of a fold operation. Consider a collection of two or more distinct i -cubes in C , with $i \geq 2$, whose boundaries are equal. A cube identification operation consists of forming the quotient of C in which all of the i -cubes in the collection have been identified to a single cube. Note that the 1-skeleton does not change in this process.

Cube attachment operation. Consider an i -tuple e_1, \dots, e_i of edges in C , with labels s_1, \dots, s_i , which are all incident to a single vertex v . Suppose furthermore, that the vertices corresponding to s_1, \dots, s_i in Γ form an i -clique. A cube attachment operation consists of adding an i -cube c to C by identifying the edges e_1, \dots, e_i to i edges in c which are all incident to a single vertex of c . In the process, we add some vertices and edges to C . Each new edge added is dual to a mid-cube of c which is also dual to a unique edge in the set $\{e_1, \dots, e_i\}$. This induces a labeling on the newly added edges, making the resultant complex Γ -labeled.

We say a complex is *folded* if no fold or cube identification operation can be performed to it. As fold operations and cube identification operations reduce the number of cells, any finite complex C can be transformed into a folded complex through finitely many such operations.

We say a complex is *cube-full* if for any i -tuple of edges all incident to the same vertex such that the vertices corresponding to their labels form an i -clique in Γ , there exists an i -cube of C whose boundary contains these i edges.

Given a connected finite Γ -labeled complex X , consider a possibly infinite sequence:

$$\Omega_0 = X \xrightarrow{f_0} \Omega_1 \xrightarrow{f_1} \Omega_2 \cdots$$

where for each i , the map $f_i: \Omega_i \rightarrow \Omega_{i+1}$ is either a fold, cube identification or cube attachment operation. Let Ω_X be the direct limit of this sequence. If Ω_X is folded and cube-full, we call Ω_X a *completion* of X . We say

$$\Omega_0 = X \xrightarrow{f_0} \Omega_1 \xrightarrow{f_1} \Omega_2 \cdots \rightarrow \Omega_X$$

is a completion sequence for X . We sometimes leave the maps f_i out of the notation when these maps are not relevant. We also set $\hat{f}: X \rightarrow \Omega_X$ as the direct limit of the maps $\{f_i\}$.

Example 3.1. Let Γ_1 be the graph in Figure 1. The right of the figure shows a completion sequence for the Γ_1 -labeled complex X . The completion Ω is obtained from X by a fold operation followed by a cube attachment operation. Note that not all labels of Ω are shown.




Figure 1. A completion Ω for the Γ_1 -labeled complex X .

Example 3.2. Figure 2 shows a graph Γ_2 and a Γ_2 -labeled complex X . A standard completion Ω' of X (see Definition 3.4) is shown on the right (the labels of Ω' are omitted). The cube complex Ω' is topologically a bi-infinite cylinder.




Figure 2. A completion Ω' for the Γ_2 -labeled complex X .

Additional examples of completions can be found in Examples 6.8, 8.5 and 10.6.

Proposition 3.3 (Existence of a completion). *Given any finite Γ -labeled complex X , there exists a completion Ω_X of X .*

Proof. We set $\Omega_0 = X$ and build Ω_i inductively. Suppose a finite complex Ω_i was obtained from Ω_0 by a sequence of fold, cube identification and cube attachment operations. We iteratively perform fold and cube identification operations to Ω_i to obtain the complexes $\Omega_{i+1}, \Omega_{i+2}, \dots, \Omega_{i+j}$, where Ω_{i+j} is folded. (This includes the case $i = 0$.) As Ω_i is finite, we conclude that j (and hence Ω_{i+j}) is finite as well.

Next we describe a sequence of operations to be performed to the finite folded complex Ω_{i+j} . Choose a vertex v of Ω_{i+j} . Consider a maximal tuple of edges incident to v such that their labels form a clique in Γ . If there is no cube in Ω_{i+j} whose boundary contains the tuple of edges, then attach an appropriately labeled cube of the appropriate dimension along the tuple of edges. Do this for each such maximal tuple at v , and then proceed to do the same for all the vertices of Ω_{i+j} . The result is a sequence of complexes $\Omega_{i+j+1}, \Omega_{i+j+2}, \dots, \Omega_{i+j+k}$, such that each one is obtained from the previous one by attaching an n -cube for some n to an n -tuple of edges of Ω_{i+j} which are all incident to a single vertex of Ω_{i+j} . As Ω_{i+j} is finite we conclude that k (and hence Ω_{i+j+k}) is finite. We then repeat the above procedure starting with the finite complex Ω_{i+j+k} .

Let Ω_X be the direct limit of these complexes. Consider a pair of edges, say e and f , in Ω_X incident to the same vertex v . It follows that some Ω_i contains preimages of e and f which are incident to the same vertex. Consequently, there is some folded $\Omega_{i'}$, with $i' \geq i$, which contains preimages of e and f that are incident to the same vertex. Thus, if e and f have the same label, then their preimages in $\Omega_{i'}$ must be identified. A similar argument shows that two cubes with the same boundary in Ω_X must be identified. It follows that Ω_X is folded.

Let e_1, \dots, e_n be edges all incident to a common vertex $v \in \Omega_X$ whose labels form an n -clique in Γ . There is some folded Ω_i which contains a preimage v' of v and preimages e'_1, \dots, e'_n of e_1, \dots, e_n so that e'_1, \dots, e'_n are all incident to v' . As a result of the procedure described above, some Ω_{i+j} is the complex resulting from a cube attachment operation where an n -cube, say c , is attached to the image of the edges e'_1, \dots, e'_n in Ω_{i+j-1} . Thus, the image of c in Ω_X is an n -cube containing the edges e_1, \dots, e_n . This shows Ω_X is cube-full. \square

Definition 3.4 (Standard Completion). We call the completion algorithm given in the proof of Proposition 3.3 a *standard completion* and the associated sequence $\Omega_0 \rightarrow \Omega_1 \rightarrow \dots \rightarrow \Omega_X$ a *standard completion sequence*.

We show that if a completion is finite, then there is indeed an algorithm to obtain it that terminates.

Proposition 3.5. *Let X be a Γ -labeled complex. Consider a standard completion sequence*

$$X = \Omega_0 \rightarrow \Omega_1 \rightarrow \cdots \rightarrow \Omega.$$

If Ω is finite, then the completion sequence is finite, i.e. $\Omega = \Omega_N$ for some N .

Proof. As Ω is finite, by the definition of a direct limit, for some M and all $n \geq M$, Ω_n contains an isometrically embedded subcomplex Y_n isometric to Ω , and the natural map $f: \Omega_n \rightarrow \Omega$ is a label-preserving isometry when restricted to Y_n . As we have a standard completion, there exists an $N \geq M$ such that Ω_N is folded.

We claim $\Omega_N = \Omega$. Suppose instead that $Y_N \subsetneq \Omega_N$. If $\Omega_N \setminus Y_N$ contains a vertex or an edge, then since Ω_N is connected, it follows that some vertex $v \in Y_N$ is incident to an edge e that is not contained in Y_N . Let s be the label of e . As Ω_N is folded, no edge in Y_N that is incident to v has label s . However, the continuous map $f: \Omega_N \rightarrow \Omega$ sends e to an edge incident to $f(v)$ labeled by s . This is a contradiction. Thus Y_N and Ω_N have the same 1-skeleton.

Suppose there is a 2-cell c in $\Omega_N \setminus Y_N$. Then the boundary ∂c of c is contained in Y_N . Since Ω_N folded, there is no cube in Y_N with boundary ∂c . However, f sends ∂c to the boundary of a cube in Ω , and $f(c)$ is a cube in Ω with boundary $f(\partial c)$, leading to a contradiction. Thus Ω_N and Ω have the same 2-skeleton. Proceeding inductively, we conclude that $\Omega_N = \Omega$ □

3.2. Completion of a subgroup. In this subsection, we define the completion of a subgroup of a RACG and show a completion is guaranteed to exist for any finitely generated subgroup.

Definition 3.6 (Completion of a subgroup). Let G be a subgroup of a RACG W_Γ , and let Ω be a connected Γ -labeled cube complex with basepoint the vertex $B \in \Omega$. We say that (Ω, B) is a *completion of G* if

- (1) Ω is folded and cube-full;
- (2) given any loop in Ω based at B , its label is a word which represents an element of G ;
- (3) for any *reduced* word w in W_Γ which represents an element of G , there is a loop l based at B with label w .

Remark 3.7. When B is not relevant, we may simply say Ω is a completion of G .

Example 3.8. Let G be a finite subgroup of W_Γ generated by adjacent vertices a and b in Γ . Let X be the rose graph consisting of one vertex, one graph-loop labeled by a and one graph-loop labeled by b . Let Ω_1 and Ω_2 respectively be the torus and Klein bottle obtained by attaching a (square) 2-cell to X . Then both Ω_1 and Ω_2 are completions for G .

When the group $G < W_\Gamma$ is finitely generated, we can construct a completion for G as follows. Let G be generated by the finite generating set of words

$$S_G = \{w_i = s_{i_1}s_{i_2} \dots s_{i_{m_i}} \mid 1 \leq i \leq n\},$$

where $s_{i_j} \in V(\Gamma)$ for each i, j . We associate to S_G the following Γ -labeled complex. We begin with a single base vertex B . For each generator w_i , we attach a circle subdivided to have m_i edges, such that edges of this circle are sequentially labeled, beginning at B , by the letters s_{i_j} for $1 \leq j \leq m_i$. We denote this based complex by $(X(S_G), B)$ and call it the S_G -complex.

Let Ω be a completion of $X(S_G)$ with completion sequence

$$X(S_G) = \Omega_0 \rightarrow \Omega_1 \rightarrow \Omega_2 \rightarrow \dots \rightarrow \Omega.$$

By a slight abuse of notation, we use B to denote the image of the base point B in Ω_i for any i , as well as in Ω . The next few lemmas show that Ω is a completion of G .

Lemma 3.9. *Let G be a subgroup of a RACG W_Γ , given by a finite generating set S_G . Let Ω be any completion of $X(S_G)$, where $(X(S_G), B)$ is the S_G -complex. If w is a reduced word in W_Γ which represents an element of G , then w is the label of some loop in Ω based at B .*

Proof. Let $X(S_G) = \Omega_0 \rightarrow \Omega_1 \rightarrow \Omega_2 \rightarrow \dots \rightarrow \Omega$ be a completion sequence for $X(S_G)$, and let w be a reduced word in W_Γ which represents an element of G .

As S_G is a generating set of G , it follows that w is equal in W_Γ to a word $w' = h_1 \dots h_k$, where $h_i \in S_G$ for each i . By construction, for each $1 \leq i \leq k$, there is a loop l_i based at B in Ω_0 , with label h_i . Let l be the loop in Ω_0 formed as a concatenation of loops: $l_1 l_2 \dots l_k$. Let \hat{l} be the image of l in Ω . Then \hat{l} has the same label as l .

As w' and w are equal as elements of W_Γ , the word w can be obtained from w' through a sequence of Tits moves. Suppose the first Tits move in this sequence is a swap performed to w' to obtain a new word w'' .

We claim w'' is the label of a loop in Ω as well. Note that there are adjacent edges e and f of \hat{l} , labeled by s and t where $s, t \in V(\Gamma)$ and $m(s, t) = 2$, such that $w' = a_1 \dots a_i s t a_{i+1} \dots a_m$ and $w'' = a_1 \dots a_i t s a_{i+1} \dots a_m$, with $a_j \in V(\Gamma)$. As $m(s, t) = 2$ and Ω is cube-full, there must be a square Q in Ω whose boundary contains ef . We now obtain the desired loop by replacing ef in \hat{l} with the opposite path in the boundary of Q .

On the other hand, suppose the first Tits move is a cancellation. In other words, $w' = a_1 \dots a_i s s a_{i+1} \dots a_m$ is replaced by $w'' = a_1 \dots a_m$ where $s \in V(\Gamma)$ and $a_i \in V(\Gamma)$ for each i . As Ω is folded, \hat{l} must traverse an edge e , labeled by s , twice consecutively. It follows that either e is a graph-loop or that \hat{l} traverses e in one direction and immediately backtracks in the other direction. In either case, we can simply remove these two occurrences of the edge e from \hat{l} to obtain a new loop based at B with label w'' .

By repeating this procedure for each Tits move, we obtain a loop in Ω with label w . □

Lemma 3.10. *Let Ω be a Γ -labeled complex obtained by applying either a fold, cube identification or cube attachment operation to the Γ -labeled complex $\bar{\Omega}$. Let $F: \bar{\Omega} \rightarrow \Omega$ be the natural map. Let B be a vertex of Ω and let \bar{B} be a vertex that is in the preimage under F of B . Let w be the label of a loop l in Ω based at the vertex B . Then there exists a loop \bar{l} in $\bar{\Omega}$ based at \bar{B} , with label \bar{w} , such that \bar{w} and w represent the same element of the RACG W_Γ .*

Proof. We analyze each type of operation separately:

Cube identification operation. If F is a cube identification operation, then Ω and $\bar{\Omega}$ have the same 1-skeleton. Thus l is the image of a loop \bar{l} in $\bar{\Omega}$ with the same label as l .

Fold operation. Suppose F is a fold, and let $B = u_1, \dots, u_m = B$ be the vertices of l listed sequentially by the orientation of l .

Suppose some vertex, say v , of Ω has preimage $F^{-1}(v) = \{\bar{v}_1, \bar{v}_2\}$. As only a single edge is folded in a fold operation, there is at most one such vertex. Let \bar{f}_1 and \bar{f}_2 be the two edges in $\bar{\Omega}$ which are folded and let f be the edge in Ω which is their image. Let s be the label of f . The endpoints of f must be v and some vertex v' (which is possibly equal to v).

We say a vertex or edge of l has unique preimage if its preimage under F is a single vertex or edge. It is straightforward to check that l can be subdivided into subpaths of the five types described below (though not all the types may be used):

- (1) An edge p from $u = u_i$ to $u' = u_{i+1}$ where u and u' each have unique preimage.
- (2) A path p from $u = u_i$ to $u' = u_{i'}$, where u and u' each have unique preimage, and $u_j = v$ for every $i < j < i'$.
- (3) A path p from $B = u_1$ to $u = u_i$, where u has unique preimage and $u_j = v$ for $j < i$.
- (4) A path p from $u = u_i$ to $B = u_m$, where u has unique preimage and $u_j = v$ for $j > i$.
- (5) $p = l$ and $u_i = v$ for all i .

We claim that for each path p of a type described above, there is a path \bar{p} in $\bar{\Omega}$ such that the label of p and the label of \bar{p} are equal in W_Γ . Additionally, the image under F of the endpoints of \bar{p} are equal to the endpoints of p . Finally, if p is of type 3, then \bar{p} begins at \bar{B} , if p is of type 4 then \bar{p} ends at \bar{B} and if p is of type 5, then \bar{p} begins and ends at \bar{B} . The lemma clearly follows from this claim. We proceed to prove the claim for each type of subpath of l .

Type 1. Let $p = e$ be the edge in l between $u = u_i$ and $u' = u_{i+1}$. Let \bar{u} and \bar{u}' be the unique preimages of u and u' under F . The preimage $F^{-1}(e)$ is either a single edge between \bar{u} and \bar{u}' or is a pair of edges between \bar{u} and \bar{u}' (in this case the fold operation identifies this pair of edges to get e). Let \bar{e} be a choice of edge in $F^{-1}(e)$. We define \bar{p} to the path that traverses \bar{e} . Clearly \bar{p} and p have the same label.

Type 2. In this case p consists of an edge e_1 from u to v , followed by a collection of graph-loops q_1, \dots, q_k based at v , followed by an edge e_2 from v to u' .

Let \bar{u} and \bar{u}' be the unique preimages of u and u' . Let \bar{e}_1 and \bar{e}_2 be edges (not necessarily unique) in the preimage of e_1 and e_2 respectively. Let $\bar{q}_1, \dots, \bar{q}_k$ each be a choice of edge in the preimages of q_1, \dots, q_k . For each $1 \leq i \leq k$, the edge \bar{q}_i is either a graph-loop at \bar{v}_1 , a graph-loop at \bar{v}_2 or an edge between \bar{v}_1 and \bar{v}_2 .




Figure 3. The graph on the top shows a path of type 2 in Ω . The two graphs below it show the two possible choices of preimages for p . The graph-loops $\bar{q}_{i_1}, \dots, \bar{q}_{i_{k'}}$ form a subsequence of the graph-loops $\bar{q}_1, \dots, \bar{q}_k$ consisting of those based at \bar{v}_1 . Similarly, $\bar{q}_{i_{k''}}, \dots, \bar{q}_{i_k}$ are edges between \bar{v}_1 and \bar{v}_2 and $\bar{q}_{i_{k''}}, \dots, \bar{q}_{i_k}$ are graph-loops based at \bar{v}_2 . Note that some of the vertices and edges shown may actually be equal $\bar{\Omega}$. For instance, it could be that $\bar{f}_1 = \bar{e}_1$.

Let \bar{z} be the path from \bar{v}_1 to \bar{v}_2 obtained by traversing \bar{f}_1 and then \bar{f}_2 . Note that the label of \bar{z} is equal to the identity element of W_Γ as \bar{f}_1 and \bar{f}_2 have the same label. Form the path

$$\bar{p} = \bar{e}_1 \bar{p}_0 \bar{q}_1 \bar{p}_1 \bar{q}_2 \bar{p}_2 \dots \bar{q}_k \bar{p}_k \bar{e}_2,$$

where for $1 \leq i < k$, we define \bar{p}_i to either be \bar{z} , \bar{z}^{-1} or the empty word in order for

the endpoint of \bar{q}_i to coincide with the startpoint of \bar{q}_{i+1} . The paths \bar{p}_0 and \bar{p}_k are defined similarly in order for the endpoint of \bar{e}_1 to coincide with the startpoint of \bar{q}_1 , and in order for the endpoint of \bar{q}_k to coincide with the startpoint of \bar{e}_2 . The claim then follows for this case as the label of \bar{p} is equal as an element of W_Γ to the label of p .

Type 5. In this case $v = B$ and l consists of a sequence of graph-loops q_1, \dots, q_k . Let $\bar{q}_1, \dots, \bar{q}_k$ be a choice of edges in the preimages of q_1, \dots, q_k . As above, these preimages consist of graph-loops at \bar{v}_1 or \bar{v}_2 , and edges between \bar{v}_1 and \bar{v}_2 . We may define the path

$$\bar{p} = \bar{p}_0 \bar{q}_1 \bar{p}_1 \dots \bar{q}_k \bar{p}_k,$$

where \bar{p}_i , for $1 \leq i < k$, is defined similarly as in the previous case. Note that \bar{B} is either equal to \bar{v}_1 or \bar{v}_2 . We then define \bar{p}_0 to either be $\bar{z} = \bar{f}_1 \bar{f}_2, \bar{z}^{-1}$ or the empty path in order to guarantee \bar{p} begins at \bar{B} . Similarly define \bar{p}_k to guarantee that \bar{p} ends at \bar{B} .

Types 3 and 4. The proofs in these cases are very similar to those above, and are omitted.

Cube attachment operation. Suppose Ω is obtained by attaching a cube c to $e_1 \cup \dots \cup e_n \subset \bar{\Omega}$, where for $1 \leq i \leq n$, e_i is an edge labeled s_i from a vertex v to a vertex v_i (with $v = v_i$ if e_i is a graph-loop), and the vertices of Γ corresponding to s_1, \dots, s_n form an n -clique.

If l does not intersect $c \setminus \{e_1, \dots, e_n\}$ then l is clearly the image of a loop in $\bar{\Omega}$ with the same label. Otherwise let q be the closure of a maximal connected subpath of l that is contained in $c \setminus \{e_1, \dots, e_n\}$. In particular, q is a path between v_k and $v_{k'}$ for some (not necessarily distinct) k, k' . Let h be the label of q . Since the vertices s_1, \dots, s_n form a clique in Γ , there is a reduced expression for h given by $h' = s_1^{\epsilon_1} \dots s_n^{\epsilon_n}$ where for each $1 \leq i \leq n$, $\epsilon_i = 1$ if there is an odd number of occurrences of the generator s_i in h and $\epsilon_i = 0$ otherwise.

First assume $k \neq k'$. After renaming if necessary, we may assume $k = 1$ and $k' = n$. We claim that if $1 < j < n$ and $\epsilon_j = 1$, then e_j is a graph-loop. To prove this, note that if $1 < j < n$ and e_j is not a graph-loop, then v_1 and v_n are on the same side of the midcube of c dual to e_j . Thus q crosses this midcube an even number of times, and therefore $\epsilon_j = 0$.

Consider the union of edges $q' = e_1^{\epsilon_1} \cup e_2^{\epsilon_2} \cup \dots \cup e_n^{\epsilon_n}$, where $e_i^{\epsilon_i}$ is interpreted as empty if $\epsilon_i = 0$. The claim in the previous paragraph implies that this is in fact a path q' with the same endpoints as q (even if e_1 and e_n are graph-loops). Observe that the label of q' is h' .

By a similar argument, if $k = k' = 1$ and e_1 is a graph-loop, then $q' = e_1^{\epsilon_1} \cup e_2^{\epsilon_2} \cup \dots \cup e_n^{\epsilon_n}$ is a path with label h' and the same endpoints as q . Finally, suppose $k = k' = 1$ and e_1 is not a graph-loop. As before, if $j \neq 1$ and $\epsilon_j = 1$

then e_j is a graph-loop. Moreover, $\epsilon_1 = 0$, because q crosses the mid-cube dual to e_1 an even number of times. Thus in this case we define q' to be the concatenation $e_1 e_2^{\epsilon_2} \cup \dots \cup e_n^{\epsilon_n} e_1$, and note that this is a continuous path with the same endpoints as q , and with label $s_1 h' s_1$, which is equal in W_Γ to h' .

In each case we have produced a path q' in $F(\bar{\Omega})$ with the same endpoints as q and whose label is a word equal in W_Γ to h . We replace q with q' in l . By performing all possible replacements of this sort, we obtain a loop in $F(\bar{\Omega})$ whose label is a word equal to w in W_Γ . Thus the lemma follows for this case. \square

Lemma 3.11. *Let G be a subgroup of a RACG W_Γ given by a finite generating set S_G . Let Ω be any completion of $X(S_G)$ where $(X(S_G), B)$ is the S_G -complex. Given a loop in Ω based at B , its label is a word representing an element of G .*

Proof. Let $X(S_G) = \Omega_0 \xrightarrow{f_0} \Omega_1 \xrightarrow{f_1} \Omega_2 \xrightarrow{f_2} \dots \rightarrow \Omega$ be a completion sequence for $X(S_G)$. Let B denote the basepoint of $X(S_G)$ as well as all its images in this sequence.

Consider a loop l based at B in Ω , with label w . Then there exists n such that $l = \hat{f}(l')$ for some loop l' based at B in Ω_n which also has label w , where $\hat{f}: \Omega_n \rightarrow \Omega$ is the natural map. By iteratively applying Lemma 3.10 starting with \hat{f} , it follows there is a loop in Ω_0 based at B whose label is a word equal to w in W_Γ . As the label of any loop based at B in Ω_0 represents an element of G , the lemma follows. \square

The existence of completions is an immediate consequence of Lemma 3.9 and Lemma 3.11:

Theorem 3.12. *Let G be a subgroup of a RACG W_Γ given by a finite generating set S_G . Then any completion (Ω, B) of the based S_G -complex $(X(S_G), B)$ is a completion of G .*

Definition 3.13 (Standard completion of a subgroup). Let G be a subgroup of the RACG W_Γ generated by a finite generating set S_G . We call a completion Ω of G obtained by Theorem 3.12 a *standard completion of G with respect to S_G* . In cases where it is understood that there is a finite generating set for G , we simply say Ω is a *standard completion of G* .

Example 3.14. The Γ_1 -labeled cube complex Ω in Example 3.1 is a standard completion of the subgroup $\langle adb, aec \rangle$ of W_{Γ_1} . Similarly, the Γ_2 -labeled cube complex Ω' in Example 3.2 is a standard completion of the subgroup $\langle abcd \rangle$ of W_{Γ_2} .

Remark 3.15. Although every reduced word in W_Γ representing an element of G labels a loop in the completion Ω , it is not true that every word in W_Γ representing an element of G labels a loop in Ω . For instance, let $s \in V(\Gamma)$ and suppose $G < W_\Gamma$ is generated by a set of words, none of which contains the letter s . It follows that no edge in Ω is labeled s . Then the word ss , which is equal to the identity element, cannot be the label of any path in Ω .

4. Basic properties of completions

We prove a few facts regarding completions that will be used throughout the rest of the paper. Recall from Section 2.2 that a deletion performed to a word w in a RACG produces an expression for w with a pair of generators of the same type removed.

Lemma 4.1. *Let Ω be a folded, cube-full, Γ -labeled complex. Let p be a path in Ω with label w . Let w' be an expression for w obtained by performing k deletions to w . Then there exists a path p' in Ω such that the following properties hold.*

- (1) *The path p' has label w' .*
- (2) *The paths p and p' have the same endpoints.*
- (3) *The Hausdorff distance between p and p' is at most k .*
- (4) *If p does not traverse any graph-loops, then p and p' are homotopic relative to their endpoints.*

Proof. Let w_1 be the word obtained by performing the first deletion to w . If $w = s_1 \dots s_n$, with $s_i \in V(\Gamma)$, then

$$w_1 = s_1 \dots s_{i-1} s_{i+1} \dots s_{i'-1} s_{i'+1} \dots s_n,$$

where $s_i = s_{i'} = s$ for some $1 \leq i < i' \leq n$ and $m(s, s_j) = 2$ for all $i < j < i'$. Let α be the subpath of p labeled by $s_i s_{i+1} \dots s_{i'}$.

Suppose first that $i' - i > 1$. As Ω is cube-full, there exists a sequence of squares in Ω such that Figure 4 holds (although there may be additional edge or vertex identifications that are not shown). The subpath α of p , which runs along the bottom of Figure 4, can be replaced with the path which runs along the top of Figure 4, to obtain a new path p_1 with label w_1 . Then p_1 is homotopic relative to endpoints to p and is at Hausdorff distance at most 1 from p .




Figure 4.

On the other hand, if $i' - i = 1$, then as Ω is folded, α either traverses an edge labeled by s twice in opposite directions, or a graph-loop labeled by s twice consecutively. In either case we can simply remove both occurrences of α from p to obtain a new path p_1 with the same endpoints as p , such that p_1 is labeled by w_1 and is at Hausdorff distance at most 1 from p . If p does not traverse any graph-loops, then p_1 is homotopic relative to endpoints to p .

Finally, if p does not traverse any graph-loops, then p_1 does not either. This is clear when $i' - i = 1$. Now suppose $i' - i > 1$, and suppose p has no graph-loops,

but p_1 has a graph-loop e . Then e is an edge of one of the squares in Figure 4, and since Ω is folded, the two edges of the square incident to e are identified in Ω . It follows that the edge opposite to e in the square, which is a part of p , is also a graph-loop. This is a contradiction.

By repeating this process of obtaining p_1 from p inductively, we obtain the result. □

Lemma 4.2. *Let Ω be a folded, cube-full, Γ -labeled complex. Let p be a path in Ω with endpoints the vertices u and v , and let w be its label.*

- (1) *Any reduced word w' equal to w in W_Γ is the label of some path p' in Ω from u to v . Furthermore, if p does not traverse any graph-loops, then p and p' are homotopic relative to their endpoints.*
- (2) *If p has minimal length then w is reduced.*
- (3) *Let p and p' be paths in Ω with the same endpoints. If p and p' are homotopic relative to their endpoints, then their labels are equal as elements of W_Γ .*

Proof. Let w'' be a reduced word equal to w in W_Γ which is obtained by a sequence of deletions. By Lemma 4.1, there is a path p'' with label w'' and the same endpoints as p . Furthermore, if p does not traverse any graph-loops, then p'' is homotopic relative endpoints to p

By Tits' solution to the word problem, there is a sequence of words $w'' = w_0, w_1, \dots, w_n = w'$ so that w_{i+1} is obtained from w_i by swapping a pair of consecutive commuting generators.

Suppose that $w_0 = a_1 \dots a_i s t a_{i+1} \dots a_n$ and $w_1 = a_1 \dots a_i t s a_{i+1}$, where $s, t \in V(\Gamma)$, with $m(s, t) = 2$, and $a_j \in V(\Gamma)$ for each j . It follows there are consecutive edges of p'' with labels s and t . As Ω is cube-full, these edges are in a square with label $stst$. By replacing these two edges of p'' with the other edges in this square, we obtain a new path whose label corresponds to swapping s and t . Furthermore, this new path is homotopic, relative to endpoints, to p'' . By iteratively repeating this process, we obtain the desired path p' . This proves (1).

Let p be a path with minimal length and label w . If w is not reduced, let w' be a reduced expression for w . By (1) there is a path p' having the same endpoints as p , with label w' . However, $|p'| < |p|$, a contradiction. This proves (2).

Let p and p' be as in (3). As the concatenation pp'^{-1} is null-homotopic, there exists a disk diagram D with boundary pp'^{-1} . It readily follows that a homotopy of p to p' in D induces a sequence of Tits moves which show that $p = p'$. This proves (3). □

The next definition allows us to “go backwards” from a Γ -labeled complex to a subgroup.

Definition 4.3 (Associated subgroup). Let Ω be a connected, Γ -labeled complex with base vertex B . Consider the set of all $g \in W_\Gamma$ such that there exists a loop in Ω based at B whose label is a word in W_Γ that represents of g . This set is easily seen to be a subgroup of W_Γ , and is called the *subgroup of W_Γ associated to (Ω, B)* .

Proposition 4.4. *Let (Ω, B) be a connected, folded, cube-full, Γ -labeled complex, and let G be the subgroup of W_Γ associated to (Ω, B) . Then (Ω, B) is a completion of G .*

Proof. Properties (1) and (2) in the definition of a completion of a subgroup (Definition 3.6) are immediate. To check property (3), let w be a reduced word representing an element g of G . By the definition of G , there exists a loop in Ω based at B whose label, say w' , is a representative of g . Since w is a reduced representative of w' , Lemma 4.2 (1) implies that there is a loop in Ω based at B with label w . \square

The following lemma describes the effect of changing the basepoint in a Γ -labeled complex on the associated subgroup of W_Γ .

Lemma 4.5. *Let Ω be a connected, folded, cube-full, Γ -labeled complex. Let B_1 and B_2 be vertices of Ω , and for $i = 1, 2$, let G_i be the subgroup of W_Γ associated to (Ω, B_i) . Then $G_2 = w^{-1}G_1w$, where w is the label of some path from B_1 to B_2 .*

Proof. Let α be a path from B_1 to B_2 with label w . If β is a loop in Ω based at B_1 representing an element of G_1 , then the concatenation $\alpha^{-1}\beta\alpha$ represents an element of G_2 . It follows that $w^{-1}G_1w \subseteq G_2$. Similarly $wG_2w^{-1} \subseteq G_1$. \square

It is easy to detect torsion in subgroups of RACGs using completions:

Proposition 4.6. *Let G be a subgroup of a RACG W_Γ and let (Ω, B) be a completion for G . Then G has torsion if and only if there exists a loop in Ω (not necessarily passing through B) whose label is a reduced word representing an element in a finite special subgroup of W_Γ .*

Proof. If $g \in G$ has finite order, then g is conjugate into a finite special subgroup of W_Γ (see [20, Theorem 12.3.4] for instance). Write $g = wuw^{-1}$, where u and w are reduced and w is the shortest word for which such an expression for g exists. We claim wuw^{-1} is reduced. If not, then a deletion is possible. It follows from our choices that some letter, say s , occurring in w or w^{-1} cancels with an occurrence of s in u . Since u belongs to a finite special subgroup, s commutes with u . Thus we can write $g = w_1uw_1^{-1}$, where $w = w_1s$, a contradiction.

Since wuw^{-1} is reduced, there is a loop α in Ω based at B with label wuw^{-1} . Let v be the vertex along α such that the label of α between B and v is w . As Ω is folded, the subpaths of α with labels w and w^{-1} are identified, and there is a loop based at v with label u . This proves one direction of the claim.

For the other direction, suppose that there is a loop based at some vertex x of Ω with label a reduced word r representing an element in a finite special subgroup

of W_Γ . As finite special subgroups of W_Γ correspond to clique subgraphs of Γ , it follows that the support of r is contained in a clique of Γ . Let $s \in V(\Gamma)$ be a letter in r . As r is reduced, there is exactly one occurrence of s in r . Let h be the label of a path from B to x in Ω . It follows that there is a loop in Ω based at B with label $k = hrh^{-1}$. Furthermore, as there are an odd number of occurrences of the letter s in the word k , the deletion property implies that k cannot be an expression for the identity element. As k has finite order, it follows that G has torsion. \square

For torsion-free subgroups, the following holds:

Theorem 4.7. *Suppose G is a torsion-free subgroup of the RACG W_Γ . Then the fundamental group of any completion Ω of G is isomorphic to G .*

Proof. Let B be a vertex of Ω . We define the isomorphism $\phi: \pi_1(\Omega, B) \rightarrow G$ as follows. Let α be a loop in Ω based at B representing an element of $\pi_1(\Omega, B)$. We may assume that α is contained in the 1-skeleton of Ω . By property (2) of the definition of a completion, the label of α represents an element of G , and we define $\phi(\alpha)$ to be this element.

To see that ϕ is well-defined, let α and α' be loops based at $B \in \Omega$ that are homotopic relative basepoint. Then by Lemma 4.2 (3), the labels of α and α' are equal as elements of G .

It is clear that ϕ is a surjective homomorphism. To check that ϕ is injective, suppose that $\phi(\alpha)$ is a word in W_Γ equal in G to the identity element. Note that Ω cannot contain a graph-loop by Proposition 4.6. Thus, by Lemma 4.2 (1) we conclude that α is null-homotopic. \square

The next result, which in particular implies that any standard completion of a finite Γ -labeled complex has finitely many hyperplanes, is used in several proofs in this article.

Proposition 4.8. *Let X be a Γ -labeled complex. Let*

$$X = \Omega_0 \rightarrow \Omega_1 \rightarrow \dots \rightarrow \Omega$$

be a completion sequence for X . Then, for all $i \geq 0$, every hyperplane in Ω_i intersects the image of X in Ω_i . Consequently, every hyperplane in Ω intersects the image of X in Ω .

Proof. We prove the first claim by induction on n . The base case for Ω_0 is trivially true. Suppose every hyperplane in Ω_{n-1} intersects the image of X . It is clear that cube identifications and folds do not produce new hyperplanes. Now suppose that Ω_n is obtained by attaching a k -cube c along edges e_1, \dots, e_k , all incident to a common vertex of Ω_{n-1} . Since each midcube of c extends a hyperplane dual to one of the e_i 's, it follows that no new hyperplanes are created in Ω_n . Hence, the claim also holds for Ω_n . The claim follows for the completion Ω as any hyperplane in Ω contains the image of some hyperplane in Ω_n for some n . \square

In a few arguments throughout this article, we will have a finite, cube-full, folded Γ -labeled complex, and we will want to attach certain additional graph-loops to this complex. We will then need the original complex to be isometrically embedded in the completion of the new complex. The following lemma guarantees this property.

Lemma 4.9. *Let Ω be a finite, cube-full, folded Γ -labeled complex. Let Ω' be a complex obtained by attaching a set \mathcal{L} of labeled graph-loops to vertices of Ω . Suppose that the label of an attached graph-loop is distinct from the labels of every other edge incident to the vertex it is attached to, i.e., Ω' is a folded complex. Then there exists a completion Ω'' of Ω' such that*

- (1) *the natural inclusion $i: \Omega \hookrightarrow \Omega''$ is an isometry;*
- (2) *every edge of Ω'' that is not in $i(\Omega)$ is a graph-loop attached to a vertex $v \in i(\Omega)$. Let l be such a graph-loop and let s be its label. Then there exists a graph-loop in \mathcal{L} with label s attached to a vertex $u \in \Omega$ and a path in $i(\Omega)$ from $i(u)$ to v whose label is a word in $\text{link}(s) \subset V(\Gamma)$;*
- (3) *the number of operations performed to obtain Ω'' from Ω' is finite and only depends on the number of edges of Ω' .*

Proof. We build the completion Ω'' by alternately performing a single cube attachment operation followed by all possible fold and cube identification operations. Since Ω' is folded by assumption, each cube attachment operation is only done to a folded complex. We show that each folded complex in this completion sequence satisfies the conclusion of the lemma.

Let v be a vertex of Ω' incident to edges e_1, \dots, e_n with distinct pairwise commuting labels s_1, \dots, s_n , such that $v \cup e_1 \cup \dots \cup e_n$ are not contained in a common n -cube and that n is maximal out of such possible choices. Consider the operation which attaches a labeled cube to $v \cup e_1 \cup \dots \cup e_n$. Let c denote the image of this cube in the resulting complex.

By possibly relabeling, we may assume that there exists $0 \leq k \leq n$ such that if $i \leq k$, then e_i is not a graph-loop (and is therefore necessarily in Ω), while if $i > k$, then e_i is a graph-loop (and may or may not be in Ω). Note that if $k = 0$, then e_i is a graph-loop for all i .

If $k > 0$, it follows (since Ω is cube-full) that e_1, \dots, e_k are contained in a common k -cube q of Ω . We perform fold and cube identification operations to identify q to a face of c . Otherwise if $k = 0$, define q to be the 0-cube v .

Next, starting at v , we perform all possible fold operations to pairs of edges which are both in c . It readily follows that the 1-skeleton of this resulting complex will consist of q and a graph-loop with label s_i , for each $k + 1 \leq i \leq n$ and each vertex of q . By a slight abuse of notation, we call this resulting complex c as well.

We now check what other fold operations are possible. As Ω' is folded, the only type of possible additional fold operation would consist of an edge e in Ω'

and a graph-loop f in c such that e and f have the same label, s , and share an endpoint $u \in q$.

We claim that e must be a graph-loop. This is clear if e is in $\Omega' \setminus \Omega$. Suppose $e \in \Omega$. There is a path p from u to v in q with label a word consisting only of generators that are distinct from and commute with s . Let e_i be the edge at v with label s . Then e_i must be a graph-loop since otherwise e would have already been folded onto c . As Ω is cube-full and $p \cup e \subset \Omega$, we conclude that $e_i \in \Omega$. Thus since e_i is a graph-loop, e is a graph-loop as well.

Thus, we simply fold e onto f . After performing all such folds to $\Omega' \cup c$ we obtain the complex Ω'_1 . After possibly applying some cube identification operations, it follows that Ω'_1 is folded. Furthermore, the 1-skeleton of Ω'_1 is the same as the 1-skeleton Ω' with the possible addition of some new graph-loops. Thus, Ω is isometrically embedded in this new complex. The second conclusion of the lemma also readily follows from the construction.

We then iteratively repeat such cube attachments followed by such a sequence of fold and cube identification operations. After each iteration we have a folded complex satisfying the first two claims of the lemma. As the number of such operations is bounded by a function of the number of edges of Ω' , the third claim of the lemma follows. □

5. Core graphs

In general, a subgroup G of W_Γ does not have a unique completion. However, we now use completions to define a certain graph associated to G called a core, which is unique. This is used in Theorem 5.3 to obtain a characterization for normality of a subgroup.

Definition 5.1 (Core graph). Given a Γ -labeled complex (Ω, B) , define its *core graph at B* , denoted $C(\Omega, B)$, to be the 1-dimensional subcomplex consisting of the union of all the loops in Ω based at B whose labels are reduced words in W_Γ .

The core graph of a completion is not necessarily its entire 1-skeleton. For instance, in Figure 1, the core graph is the part of $\Omega^{(1)}$ (the 1-skeleton of Ω) that is shown in black. That is, $C(\Omega, B)$ is $\Omega^{(1)}$ minus the open star of the vertex v diametrically opposite to B . The paths based at B with (reduced) labels $adabec, aeacdb$ and $adcbea$ show that $C(\Omega, B)$ contains this subcomplex. On the other hand, let v_a, v_b and v_c denote the vertices adjacent to B in $\Omega^{(1)}$. Any loop at B which passes through v contains a subpath which passes through v , starts and ends in $\{v_a, v_b, v_c\}$, and, apart from its endpoints, remains in the component of $\Omega^{(1)} \setminus \{v_a, v_b, v_c\}$ containing v . We leave it to the reader to check that any such subpath γ must cross some midcube twice. This, together with the fact that the label of γ only uses the letters a, b , and c , which pairwise commute, implies that γ cannot have reduced label.

The following proposition states that core graphs are unique in a certain sense. We omit the proof as it very closely follows the proof of [29, Theorem 5.2], which is the corresponding result in the free group setting.

Proposition 5.2. *Let G be a subgroup of W_Γ , and let (Ω_1, B_1) and (Ω_2, B_2) be completions of G . Then there is an isomorphism $f: C_1(\Omega_1, B_1) \rightarrow C_2(\Omega_2, B_2)$, such that $f(B_1) = B_2$.*

We now characterize normal subgroups of RACGs in terms of core graphs. In Section 13, we give a different characterization of normality that is more useful for algorithmic applications.

Theorem 5.3. *Let G be a subgroup of W_Γ , with completion (Ω, B) . Consider*

$$\Delta = \{s \in V(\Gamma) \mid s \text{ commutes with every element of } G\}.$$

Then G is normal if and only if the following conditions are satisfied.

- (N1) *Given any $s \in V(\Gamma) \setminus \Delta$, there is an edge in Ω incident to B with label s .*
- (N2) *For every vertex v of Ω , there is an isomorphism from $C(\Omega, B)$ to $C(\Omega, v)$ which takes B to v .*

Proof. First suppose (N1) and (N2) are satisfied. To show that G is normal, it is enough to show that $sGs \subset G$ for all $s \in V(\Gamma)$. This is obvious when $s \in \Delta$, so consider $s \in V(\Gamma) \setminus \Delta$. By (N1) there is an edge incident to B with label s . Let v be its other endpoint. Then by Lemma 4.5, (Ω, v) is a completion for sGs . Thus, given any element g in sGs , there is a loop in Ω whose label is a reduced word representing g . It follows that the group associated to $(C(\Omega, v), v)$ (in the sense of Definition 4.3) is sGs . On the other hand, since $C(\Omega, v) \cong C(\Omega, B)$ by (N2), the group associated to it is G . Thus $sGs = G$.

Now suppose G is normal. We first show (N1) is satisfied. Let $s \in V(\Gamma) \setminus \Delta$, and let w be a reduced word representing an element of G which does not commute with s . If w has a reduced expression w' which either begins or ends with s , then since Ω is a completion, there is a loop based at B with label w' . It follows that there is an edge incident to B in Ω labeled s . If no expression for w begins or ends with s , then sws is reduced by Lemma 2.3. Moreover, since G is normal, $sws \in G$, and consequently there is a loop in Ω based at B with label sws . Once again, B is incident to an edge labeled by s . Thus (N1) holds in all cases.

To prove (N2), let $v \neq B$ be a vertex of $C(\Omega, B)$, and let α be a path in $C(\Omega, B)$ from B to v with label w . By Lemma 4.5, we know that (Ω, v) is a completion for $w^{-1}Gw = G$. Then by Proposition 5.2, there is an isomorphism from $C(\Omega, B)$ to $C(\Omega, v)$ which takes B to v . □

6. Index of a subgroup

The main result of this section is Theorem 6.6, which states that completions characterize the index of a subgroup of a RACG. This is analogous to a result for subgroups of free groups, first proved in [5].

Definition 6.1 (Full valence). We say a vertex v of a Γ -labeled complex has *full valence* if for each $s \in V(\Gamma)$ there is an edge with label s incident to v . We say a Γ -labeled complex Ω is *full valence* if every vertex of Ω has full valence.

We would like to be able to state that a subgroup has finite index if and only if all its completions are finite and have full valence, and moreover, that the number of vertices of a completion determines the index of a subgroup. However, due to a slight technical issue, this is not quite true in general. Namely, suppose that Γ contains a vertex s that is adjacent to every other vertex. Then $W_\Gamma = W_\Delta \times \mathbb{Z}_2$, where Δ is the subgraph of Γ induced by $V(\Gamma) \setminus s$. We can then build a completion for $W_\Delta \subset W_\Gamma$ by using the generating set $V(\Delta)$. Such a completion has no edges labeled by s and so does not have full valence. Furthermore, it contains only one vertex, and yet, the index of W_Δ is not one in W_Γ . To remedy this, we define resolved completions below, for which the desired statements do hold (see Theorem 6.6 and Lemma 6.5). Any finitely generated subgroup admits a resolved completion (see Remark 6.7), so resolved completions do indeed remedy this technical issue.

Definition 6.2 (Resolved Completion). Let Ω be a completion of a subgroup G of W_Γ . We say that Ω is a *resolved completion* if given any $s \in \Gamma$ such that $V(\Gamma) = \text{star}(s)$, it follows that some edge of Ω is labeled by s .

Lemma 6.3. *Let G be a subgroup of W_Γ , and let Ω be a resolved completion of G . Let $s \in \Gamma$ be such that $V(\Gamma) = \text{star}(s)$. Then every vertex of Ω is incident to an edge labeled by s .*

Proof. As Ω is resolved, let e be an edge of Ω labeled by s . Let v be a vertex of Ω incident to e . Let u be any vertex adjacent to v , and let t be the label of the edge e' between u and v . Either $t = s$ or t is adjacent to s in Γ . As Ω is cube-full, in the latter case there must be a square with label $stst$ in Ω that contains both the edge e and e' . In either case u is incident to an edge labeled by s as well. Proceeding in this manner, since Ω is connected, we conclude every vertex in Ω is incident to an edge labeled by s . \square

Lemma 6.4. *Let G be a subgroup of W_Γ , and let (Ω, B) be a resolved completion of G . If Ω is not full valence, then G has infinite index in W_Γ .*

Proof. Suppose there exists a vertex v in Ω that is not incident to an edge labeled by s , for some $s \in V(\Gamma)$. Let α be a minimal length path in Ω from the base vertex B to v , and let w be the label of this path. We assume $|w|$ is minimal among the possible choices for w and v . By Lemma 4.2 (2), the word w is reduced.

We begin by establishing a few facts, which will be used later in the proof:

(i) *The word ws is reduced.* If not, by the deletion property there exists a reduced word w' , ending with s , which is an expression for w . By Lemma 4.2 (1) there is a path in Ω from B to v with label w' . However, this is not possible as v is not incident to an edge labeled by s .

(ii) *No reduced expression for w ends in a generator that commutes with s .* For suppose $w' = s_1 \dots s_n$, with $s_i \in \Gamma$, is a reduced word such that s_n commutes with s and w is equal to w' in W_Γ . Let α' be the path from B to v with label w' and let $\hat{\alpha}$ be the subpath of α' with label $s_1 \dots s_{n-1}$. Such paths exist by Lemma 4.2 (1). Let \hat{v} be the endpoint of $\hat{\alpha}$. No edge incident to \hat{v} is labeled by s . For if there were such an edge, the fact that Ω is cube-full would imply that there is a square with label ss_nss_n containing both v and \hat{v} , contradicting the fact that v is not incident to an edge labeled by s . However, it now follows that \hat{v} is a vertex that is not incident to an edge labeled by s and $|\hat{w}| < |w|$, contradicting the minimality of our choice of w . Thus no expression for w can end with a generator that commutes with s .

(iii) *No reduced word representing an element of G begins with the label ws .* To see this, note that every reduced word representing an element of G labels a loop in Ω based at B . As Ω is folded, α is the only path beginning at B with label w . The claim follows, since the endpoint v of α is not incident to an edge labeled by s .

We now proceed with the proof. As Ω is resolved and by Lemma 6.3, there exists a vertex t of Γ that is not adjacent to s in Γ . By Tits' solution to the word problem, it follows that $(st)^n$ is reduced for all integers $n \geq 1$. Similarly, $w(st)^n$ is reduced for all integers $n \geq 1$, since ws is reduced by (ii) above.

Suppose now, for a contradiction, that G is a finite-index subgroup of W_Γ . In particular, as a set we have $W_\Gamma = Gg_1 \sqcup Gg_2 \dots \sqcup Gg_n$ for finitely many elements $g_1, \dots, g_n \in W_\Gamma$. Let w_1, \dots, w_n be reduced words representing g_1, \dots, g_n . Let $M = \max\{|w_1|, \dots, |w_n|\}$, and let $k = (st)^M$. Consider the word $h = wk$ which we know to be reduced. It follows that h is equal to $h'h''$ in W_Γ , where $h'' = w_i$ for some i and h' is a reduced word in G . Form a disk diagram D with boundary label $h(h'h'')^{-1} = wk h''^{-1} h'^{-1}$. Let $p_w, p_k, p_{h'}$ and $p_{h''}$ be the paths along the boundary of D with labels respectively w, k, h' and h'' . Thus, $p_w p_k p_{h''}^{-1} p_{h'}^{-1}$ is a path tracing the boundary of D .

As h is reduced, every dual curve intersecting $p_w p_k$ must necessarily intersect either $p_{h'}$ or $p_{h''}$. Furthermore, a pair of dual curves which each intersect p_k cannot intersect one another as s and t do not commute. Let C be the dual curve intersecting the first edge of p_k . Note that C is of type s . Now, C cannot intersect $p_{h''}$. For if it did, every dual curve intersecting p_k would intersect $p_{h''}$ as well. However, as $|h''| \leq M$ and $|k| = 2M$, this is not possible. Thus, C intersects $p_{h'}$.

Additionally, no dual curve intersecting p_w intersects C . For suppose there is such a dual curve, and suppose that it is the furthest such dual curve along p_w .

It follows that the type of such a dual curve commutes with s and commutes with every label of an edge appearing further along p_w . However, this implies that w has an expression ending with a generator that commutes with s , which contradicts (ii) above. Thus, every dual curve intersecting p_w intersects $p_{h'}$ at an edge occurring before (in the orientation of $p_{h'}$) the edge of $p_{h'}$ intersecting C .

By Lemma 2.4, there is a reduced word equal to h' in W_Γ with prefix ws . This contradicts (iii) above. Note that the argument above also holds when w is empty, i.e. when $B = v$. □

Lemma 6.5. *Let G be a subgroup of W_Γ , and let Ω be a completion of G which is full valence. Then the index of G in W_Γ is equal to the number of vertices in Ω (which could be infinite).*

Proof. Let $B = v_1, v_2, \dots$ be an enumeration of the vertices of Ω (where B denotes the basepoint). For each i , choose a minimal length path α_i from B to v_i and let w_i be its label. We will show that the words w_1, w_2, \dots are expressions for right coset representatives for G .

Let w be a reduced word in W_Γ . As every vertex of Ω has full valence, there is a path α in Ω beginning at the vertex B with label w . Then, for some i , the concatenation $\alpha\alpha_i^{-1}$ is a loop based at B , and its label ww_i^{-1} represents an element of G . Thus, w can be represented by the coset $(ww_i^{-1})w_i$.

Let n be the number of vertices in Ω (where n could be infinite). We have shown that the index of G is at most n . We now show it is exactly n . Suppose, to the contrary, that there exist words h and h' representing elements of G such that hw_i is an expression for $h'w_j$, for some $1 \leq i < j \leq n$. Then $w_iw_j^{-1}$ is an expression for an element of G . Now consider the path β in Ω with initial vertex B and label $w_iw_j^{-1}$. This path exists as Ω is full valence.

We claim β is a loop. For suppose not. Then β ends in some vertex $v_k \neq B$ and it follows that $w_iw_j^{-1}w_k^{-1}$ is a loop. By the definition of a completion, the word $w_iw_j^{-1}w_k^{-1}$ represents an element of G . Consequently, as $w_iw_j^{-1}$ is a word representing an element of G , we conclude that w_k represents an element of G as well. However, this contradicts Ω being a completion, as w_k is a reduced word and is not the label of a loop in Ω based at B . Thus β must be a loop. However, since β is labeled by $w_iw_j^{-1}$, this implies that $v_i = v_j$, a contradiction. □

Combining Lemma 6.4 and Lemma 6.5, we obtain:

Theorem 6.6. *Let G be a subgroup of W_Γ , and let Ω be a resolved completion of G . The subgroup G has finite index in W_Γ if and only if Ω is finite and full valence.*

Remark 6.7. For finitely generated subgroups, a resolved completion can always be constructed. For let $T = \{w_1, w_2 \dots w_n\}$ be a generating set of words for $G < W_\Gamma$. Note that G is still generated by $T' = T \cup \{s^2 \mid s \in \Gamma \text{ and } V(\Gamma) = \text{star}(s)\}$, and furthermore, every $s \in V(\Gamma)$ occurs in some word in T' . Thus, a standard completion for G built using T' is necessarily resolved.

Example 6.8. Let Γ be the graph in Figure 5, and let $G = \langle ca, cb \rangle < W_\Gamma$. The right of this figure shows a standard completion Ω for G . Since Γ does not have a vertex adjacent to every other vertex, Ω is automatically resolved. Note that Ω is finite and full-valence. Thus, G has finite index in W_Γ by Theorem 6.6.




Figure 5. Figure illustrating Example 6.8.

7. Nonpositive curvature

This section establishes criteria which guarantee that a completion is non-positively curved or a CAT(0) cube complex, which we rely on in Sections 9 and 10.

Recall that a graph-loop is an edge that connects a vertex to itself. A *bigon* in a CW complex is a pair of edges e_1 and e_2 , such that the set of vertices that are endpoints of e_1 is the same as the set of vertices that are endpoints of e_2 . Note that a bigon could consist of two graph-loops based at the same vertex. A *commuting bigon* in a Γ -labeled complex is a bigon whose edges are labeled by adjacent vertices of Γ . Next, we show that the presence of a commuting bigon is the only obstruction to a Γ -labeled complex being non-positively curved.

Proposition 7.1. *Let Ω be a folded Γ -labeled cube complex which does not contain a commuting bigon. If either Ω is cube-full or Γ is triangle-free, then Ω is non-positively curved.*

Proof. Let v be any vertex of Ω , and let Δ denote the link of v in Ω . We verify that Ω is non-positively curved by checking that Δ is a flag simplicial complex. We first check that Δ is a simplicial complex. This part of the proof does not require that Ω is cube-full or that Γ is triangle-free.

Since Δ is the link of a vertex in a cube complex, to check that it is a simplicial complex, it is enough to show that its 1-skeleton $\Delta^{(1)}$ does not contain any graph-loops or bigons. Suppose Δ contains a graph-loop, and let c be the square of Ω contributing this graph loop to Δ . It follows that a pair of adjacent sides of c are identified in Ω . However, adjacent sides of squares in Ω are always labeled by distinct elements of Γ , so such an identification cannot occur. Thus, Δ cannot contain a graph-loop.

We next check that $\Delta^{(1)}$ does not contain a bigon. Suppose there is such a bigon, whose edges come from (possibly non-distinct) squares c_1 and c_2 of Ω , and whose vertices come from distinct edges e_1 and e_2 of Ω , each incident to v . Then e_1 and e_2

are adjacent sides of c_i , for $i = 1, 2$. Suppose c_1 and c_2 are distinct. Since opposite sides of squares in Ω have the same label, it follows that c_1 and c_2 have the same boundary label, read starting from v in the direction of e_1 . However, this is not possible as Ω is folded.

On the other hand, suppose $c_1 = c_2 = c$. Then the bigon in Δ comes from two corners of c . Since adjacent sides of a square in Ω are never identified, it follows that these are necessarily opposite corners of c , and furthermore, that opposite sides of c are identified with orientation reversing isometries. Consequently, the attaching map of c maps the boundary of c to a bigon in Ω , consisting of the edges e_1 and e_2 . Since e_1 and e_2 are adjacent sides of a square in Ω , their labels commute. This contradicts the assumption that Ω does not contain commuting bigons. Hence, $\Delta^{(1)}$ cannot contain a bigon. We have thus verified that Δ is a simplicial complex.

We now check that Δ is flag. Let u_1, \dots, u_n be the vertices of a complete graph contained in the 1-skeleton of Δ . For $1 \leq i \leq n$, let e_i be the edge of Ω incident to v upon which u_i lies, and let s_i be the label of e_i . For each $1 \leq i < j \leq n$, we know that s_i and s_j are adjacent in Γ since u_i is adjacent to u_j in Δ . If Ω is cube-full, it follows that there is some cube in Ω containing $v \cup \bigcup_{i=1}^m e_i$. Thus, Δ is flag. On the other hand suppose that Γ is triangle-free. As Ω is folded, the labels s_1, \dots, s_n are distinct. Furthermore, as Γ is triangle-free and s_1, \dots, s_n as vertices of Γ form a complete graph, we have that $n \leq 2$. Thus, the flag condition holds under the triangle-free assumption as well. \square

If G is a torsion-free subgroup of W_Γ , then by Proposition 4.6, a completion of G cannot have commuting bigons. Then Proposition 7.1 immediately implies:

Proposition 7.2. *Any completion of a torsion-free subgroup of a RACG is non-positively curved.*

Next, we show that a completion of a finite Γ -labeled tree is a finite CAT(0) cube complex.

Proposition 7.3. *Let X be a Γ -labeled finite tree and let*

$$X = \Omega_0 \rightarrow \Omega_1 \rightarrow \dots \rightarrow \Omega$$

be a completion sequence for X . Then Ω is a finite CAT(0) cube complex. Furthermore, there is a finite bound on the length of the completion sequence.

Proof. We begin by proving that each complex in the completion sequence is simply connected.

As a first step, we show that Ω_n does not contain any graph-loops for $n \geq 0$. Note that the label of every loop in Ω_0 based at B represents the trivial element in W_Γ . For a contradiction, suppose that for some n , Ω_n contains a graph-loop l with label s . Suppose l is incident to a vertex $v \in \Omega_n$. Let p be a geodesic in Ω_n from B to v , and let w be the label of p . It follows that the loop plp^{-1} in Ω_n has label $ws w^{-1}$.

As $ws w^{-1}$ has an odd number of occurrences of the letter s , it represents a non-trivial element of W_Γ . However, by iteratively applying Lemma 3.10 we conclude that the label of some loop in Ω_0 based at B represents a non-trivial element of W_Γ . This is a contradiction. Thus, Ω_n does not contain a graph-loop for any n .

Next, we show by induction that Ω_n is simply connected for all $n \geq 0$. The base case is true by hypothesis. Now assume that Ω_n is simply connected.

Suppose Ω_{n+1} is obtained from Ω_n by attaching a k -cube c to the edges e_1, \dots, e_k of Ω_n which are all incident to the same vertex v . Then Ω_{n+1} can be homotoped onto Ω_n by homotoping c onto $v \cup \bigcup_{i=1}^k e_i$. If Ω_{n+1} is obtained from Ω_n by identifying a collection $\{c_i\}$ of k -cubes ($k \geq 2$) with identical boundary to a single cube c , then any null homotopy using the c_i 's can be replaced with one that only uses c . In both cases, Ω_{n+1} is simply connected.

Now suppose Ω_{n+1} is obtained from Ω_n by a fold operation. Specifically, suppose that the edges e_i (with endpoints v and v_i , for $i = 1, 2$) in Ω_n are identified to get the edge e in Ω_{n+1} . By the first paragraph these edges are not graph-loops.

If $v_1 = v_2$, then $e_1 \cup e_2$ is a loop, which is null homotopic because Ω_n is simply connected. Since identifying e_1 and e_2 is equivalent to attaching a disk to this loop, it follows that Ω_{n+1} is simply connected. If $v_1 \neq v_2$, then let Ω'_n and Ω'_{n+1} be the complexes obtained by collapsing the contractible subspaces $e_1 \cup e_2$ and e in Ω_n and Ω_{n+1} respectively to points. Then Ω'_n is homotopy equivalent to Ω_n , Ω'_{n+1} is homotopy equivalent to Ω_{n+1} , and Ω'_n is homeomorphic to Ω'_{n+1} . Again, we conclude that Ω_{n+1} is simply connected. We have established that Ω_n is simply connected for all $n \geq 0$, and it readily follows that Ω is simply-connected as well.

Next, we show that Ω is non-positively curved. By Proposition 7.1 it is enough to show that Ω does not contain a commuting bigon. Since Ω is simply-connected, Lemma 4.2 (3) implies that the group associated to (Ω, B) is trivial, and therefore torsion-free. By Proposition 4.6 there are no commuting bigons.

It follows that Ω is CAT(0) as it is simply-connected and non-positively curved.

We claim that the diameter of Ω is at most E , where E is the number of edges of X . For consider a geodesic α in Ω . By Proposition 4.8, every hyperplane that intersects α must also intersect the image of X in Ω . Furthermore, as Ω is CAT(0) (and not just non-positively curved), no hyperplane intersects α twice. Thus, the length of α is at most E , and, as α was an arbitrary geodesic, the diameter of Ω is also at most E . It follows that Ω is finite, as it is locally finite (since it is folded) and has finite diameter. Finally, Proposition 3.5 implies that there is a finite bound on the length of the completion sequence. \square

When Γ is triangle-free, we get the following more precise bound on the length of a standard completion sequence, which is used in Theorem 10.8.

Proposition 7.4. *With the set-up of Proposition 7.3, if Γ is additionally triangle-free, then there is a finite bound on the length of the completion sequence depending only on the number of edges of X and on $|V(\Gamma)|$.*

Proof. By Proposition 7.3, $\Omega = \Omega_N$ for some N and Ω is a finite CAT(0) cube complex. Let E be the number of edges of X . We are left to prove that N only depends on E and on $|V(\Gamma)|$. Consider the subsequence of all *folded* complexes of the given standard completion:

$$\Theta_1 = \Omega_{i_1}, \quad \Theta_2 = \Omega_{i_2}, \dots, \Theta_n = \Omega_{i_n}.$$

By Proposition 7.1, we know that Θ_i is a CAT(0) cube complex. Furthermore, Θ_i has diameter at most E , by the proof of Proposition 7.3.

We claim that the complex Θ_j is not isometric to Θ_k for all $k > j$. Suppose otherwise for a contradiction. Consider the sequence of operations performed to Θ_j in order to obtain Θ_k . We can repeat this same sequence of operations to Θ_k in order to obtain another folded complex isometric to Θ_j . By iteratively repeating this process, we obtain a standard completion sequence which is infinite. Furthermore, the direct limit Ω' of this new standard completion sequence must be a finite complex. To see this, note that given m distinct cells in Ω' , there is some complex isometric to Θ_j in the completion sequence, which contains m distinct preimages of the cells (since there are infinitely many such complexes in the sequence). Thus the size of Ω' is bounded by the size of Θ_j . However, this contradicts Proposition 3.5.

Let F be the number of all possible CAT(0) cube complexes of diameter at most E and with at most $|V(\Gamma)|$ edges incident to each vertex. As Θ_j is not isometric to Θ_k for all $j \neq k$, it follows that $n \leq F$. For each $1 \leq j \leq n$, the number of cube attachments that can be applied to Θ_j , and the number of fold and cube identification operations that can be applied to the resulting complex is bounded by a number K which depends only on E . Thus, $N \leq KF$ where K and F depend only on E and on $|V(\Gamma)|$. \square

8. Quasiconvexity

Let H be a group with fixed generating set. Recall that a subgroup G of H is M -*quasiconvex*, for $M \geq 0$, if any geodesic path in the Cayley graph of H with endpoints in G lies in the M -neighborhood of G . We say G is *quasiconvex* if it is M -quasiconvex for some M . In general, G may be quasiconvex with respect to one generating set for G but not another. However, if G is quasiconvex with respect to some generating set, then it is quasi-isometrically embedded with respect to any generating set [11, Chapter III.Γ, Lemma 3.5].

When we say a subgroup is quasiconvex in a RACG, we will always mean with respect to the standard generating set. The main result of this section is that a subgroup of a RACG is quasiconvex if and only if any standard completion of the subgroup is finite.

We first prove a lemma relating distances in a completion of a subgroup to distances in the Cayley graph (associated to the standard generating set) of the RACG.

Lemma 8.1. *Let G be a subgroup of the RACG W_Γ and let (Ω, B) be a completion of G . Let w be the label of a path in Ω from the basepoint B to some vertex $v \in \Omega$. Let \mathcal{C} be the Cayley graph of W_Γ , and let v_w be the vertex in \mathcal{C} which represents the element of W_Γ corresponding to w . Then $d_\Omega(B, v) = d_{\mathcal{C}}(G, v_w)$. (Here G is naturally identified with the vertices in \mathcal{C} which represent elements of G .)*

Proof. By Lemma 4.2, there is a path in Ω from B to v with label a reduced expression for w . Thus, without loss of generality, we may assume that w is reduced. Let α be a geodesic in \mathcal{C} from v_{id} to v_w with label w , where v_{id} is the vertex in \mathcal{C} labeled by the identity element. Let β be a geodesic in \mathcal{C} from v_w to G which realizes the distance from v_w to G . Let h be the label of β . It follows that h is a reduced word. As $\alpha\beta$ is a path from v_{id} to a vertex of G , the word $k = wh$ represents an element of G . By Lemma 2.3, there is a reduced expression $\hat{k} = \hat{w}\hat{h}$ for k in W_Γ such that $w' = \hat{w}s_1 \dots s_m$ is a reduced expression for w and $h' = s_m \dots s_1\hat{h}$ is a reduced expression for h , where $s_i \in V(\Gamma)$ for $1 \leq i \leq m$.

By Lemma 4.2 (1), there is a path α' with label $w' = \hat{w}s_1 \dots s_m$ from B to v in Ω . Furthermore, by the definition of the completion of a subgroup, there is a loop l with label $\hat{k} = \hat{w}\hat{h}$ in Ω based at B . Since Ω is folded, α' and l overlap on the part labeled \hat{w} . It follows that there is a path from B to v labeled by $h'^{-1} = \hat{h}^{-1}s_1 \dots s_m$.

Let γ be a geodesic in Ω from v to B with label z . Note that z must be a reduced word, and that $|z| \leq |h'| = |h|$. As wz is the label of a loop in Ω , it follows by the definition of a completion that wz represents an element of G . Thus there is a path in \mathcal{C} from v_w to G with label z . By the minimality of β , we have that $|h| \leq |z|$. Hence, $|z| = |h|$. It now follows that

$$d_\Omega(v_w, G) = |\beta| = |h| = |z| = |\gamma| = d_\Omega(B, v_w). \quad \square$$

Lemma 8.2. *Let G be a subgroup of the RACG W_Γ . If some completion (Ω, B) of G is finite, then G is M -quasiconvex in W_Γ , where M is the maximal distance of a vertex in Ω from B .*

Proof. Let α be a geodesic in the Cayley graph of W_Γ between two elements of G . Without loss of generality, we may assume that α goes between the identity vertex v_{id} and some vertex labeled by an element g of G . Then the label w of α is a minimal length word representing g . Let v be any vertex along α . Let w' be the label of the subpath of α from v_{id} to v .

By the definition of a completion, there is a loop l in Ω based at B with label w . Consequently, there is an initial subpath, l' , of l with label w' . Let u be the vertex of Ω that is the endpoint of l' . By Lemma 8.1,

$$d_\Omega(u, B) = d_{\mathcal{C}}(v, G) \leq M. \quad \square$$

Lemma 8.3. *If G is a quasiconvex subgroup of the RACG W_Γ , then G is finitely generated and every standard completion of G is finite.*

Proof. Suppose G is M -quasiconvex in W_Γ . Then G must be finitely generated as it is a quasiconvex subgroup of a finitely generated group [11, Chapter III.Γ, Lemma 3.5]. Let (Ω, B) be a standard completion of G . Lemma 8.1 implies that $d_\Omega(v, B) \leq M$ for every vertex v in the core graph $C(\Omega, B)$ from Definition 5.1.

Suppose Ω is not finite. Since it is folded and locally finite, a standard argument can be used to show that Ω has an infinite geodesic ray ζ based at B . By Proposition 4.8, there are only finitely many hyperplanes in Ω . It follows that we may choose a geodesic subsegment γ of ζ starting at B , a hyperplane H , and an integer $C \geq 0$ such that H is dual to at least $M + C + 2$ edges of γ and such that the distance from B to $H \cap \gamma$ is less than C .

Starting from B , let e_1 and e_2 respectively be the first and last edge along γ dual to H . Let u be the endpoint of e_1 closer to B , and let α_1 be the subpath of γ from B to u . Note that $|\alpha_1| \leq C$. Now at least one endpoint of e_2 is in the same component of $N(H) \setminus H$ as u ; call it v . Let α_2 be the subpath of γ from u to v and let α_3 be a geodesic in $N(H) \setminus H$ from v to u . Let a_i be the label of α_i , for $1 \leq i \leq 3$. Then the concatenation $\alpha = \alpha_1\alpha_2\alpha_3\alpha_1^{-1}$ is a loop in Ω based at B , with label $a = a_1a_2a_3a^{-1}$.

We will complete the proof by using α to produce a loop in $C(\Omega, B)$ which leaves the M -neighborhood of B , a contradiction. We first claim that a reduced expression $\hat{a}_2\hat{a}_3$ for a_2a_3 is obtained by a sequence of at most $|a_2| - M - C - 1$ deletions. As a_2 and a_3 are reduced, every pair of generators deleted consists of a generator in a_2 and one in a_3 . Let s be the type of H . The claim now follows from the fact that a_3 has no occurrences of s (being a word in $\text{link}(s)$) while a_2 has at least $M + C + 1$ occurrences of s .

By Lemma 4.1, there is a closed path $\hat{\alpha}$ in Ω with label $\hat{a}_2\hat{a}_3$ based at u , and of Hausdorff distance at most $|a_2| - M - C - 1$ from $\alpha_2\alpha_3$. Then $\beta = \alpha_1\hat{\alpha}\alpha_1^{-1}$ is a path based at B in Ω with label $a_1\hat{a}_2\hat{a}_3a_1^{-1}$, and β is Hausdorff distance at most $|a_2| - M - C - 1$ from α . As the vertex v of α_1 has distance $|a_1| + |a_2|$ from B , it follows that β contains a vertex v' whose distance from B is at least

$$|a_1| + M + C + 1.$$

Let \hat{a} be a reduced expression for $a_1\hat{a}_2\hat{a}_3a_1^{-1}$ obtained by a sequence of deletion operations. As a_1 and $\hat{a}_2\hat{a}_3$ are each reduced, it follows that each pair of generators deleted in such a sequence contains a generator in a_1 or in a_1^{-1} (possibly in both). Thus, there can be at most $2|a_1|$ such deletions. By Lemma 4.1, there is a path $\hat{\alpha}$ in Ω based at B with label \hat{a} and with Hausdorff distance at most $2|a_1|$ from β . As β contains v' , it follows that $\hat{\alpha}$ contains a vertex of distance at least

$$(|a_1| + M + C + 1) - 2|a_1| \geq M + 1$$

from B . This is a contradiction, as $\hat{\alpha}$ is contained in $C(\Omega, B)$, which is itself contained in the M neighborhood of B . □

We immediately obtain the following theorem from Lemmas 8.2 and 8.3.

Theorem 8.4. *Let G be a subgroup of a RACG. The following are equivalent:*

- (1) G is quasiconvex.
- (2) Some completion for G is finite.
- (3) G is finitely generated and every standard completion for G is finite.

Example 8.5. Let Γ be the graph in Figure 6, and let $G = \langle abcde \rangle < W_\Gamma$. The right of this figure shows a standard completion Ω for G . The complex Ω is a bi-infinite cylinder tiled by squares, as the process of attaching squares never stops. As Ω is infinite, G is not quasiconvex by Theorem 8.4.




Figure 6. Figure illustrating Example 8.5.

9. Residual finiteness and separability

In this section we give a proof using completions of the well known result that RACGs are residually finite. We additionally give a new proof of a result of Haglund which states that quasiconvex subgroups of RACGs are separable. We begin by giving a construction which associates a new complex and associated subgroup to each subgroup of an RACG that has a finite completion. This is similar to the complex used in Stallings’ proof of Marshall Hall’s theorem [37] and to the canonical completions of [26].

Definition 9.1 (Full valence extension and associated subgroup). Let G be a finitely generated subgroup of a RACG W_Γ , and suppose that G has a finite completion (Ω, B) . For each $s \in V(\Gamma)$ and each vertex v of Ω that is not incident to an edge labeled by s , we add a graph-loop labeled by s to v . We call this first resulting complex \mathcal{E}_0 . Let (\mathcal{E}, B) be a completion of \mathcal{E}_0 obtained by applying Lemma 4.9. We call (\mathcal{E}, B) the *full valence extension* of Ω . The complex \mathcal{E} is folded and cube-full as it is a completion. Furthermore, by construction and by Lemma 4.9, \mathcal{E} is finite

and full-valence, and the natural inclusion of Ω into \mathcal{E} is an isometry. Let H be the subgroup associated to (\mathcal{E}, B) (as in Definition 4.3).

The following lemma is an immediate consequence of Proposition 4.4 and Lemma 6.5.

Lemma 9.2. *With the notation of Definition 9.1, the subgroup H has finite index in W_Γ .*

We now use full valence extensions to give a direct proof of the residual finiteness of RACGs. The fact that RACGs are residually finite is well known: RACGs are linear (see for instance [10]), and by a theorem of Malcev, every finitely generated linear group is residually finite (see [39, 4.2] for a proof).

Theorem 9.3. *Every RACG is residually finite.*

Proof. Let W_Γ be a RACG. Let g be a non-trivial element in W_Γ , and let w be a reduced word representing g . Let G be the trivial subgroup of W_Γ given by the generating set $S_G = \{ww^{-1}\}$. The S_G -complex $\Omega_0 = X(S_G)$ consists of a circle, subdivided into labeled edges, whose label, read from a base vertex B , is ww^{-1} . We can iteratively perform fold operations to Ω_0 and obtain a complex Ω_N that is a path labeled by w .

By Proposition 7.3, there is a completion (Ω, B) of Ω_N that is a finite CAT(0) cube complex. The image of Ω_N in Ω is a path, p' , based at B and labeled by w . Furthermore, the path p' is not a loop in Ω . This follows since Ω is a completion of the trivial subgroup, and consequently every loop in Ω based at B must have as label a word that is trivial in W_Γ .

Let (\mathcal{E}, B) be the full valence extension of (Ω, B) . As Ω is isometrically embedded in \mathcal{E} by Lemma 4.9, w is still not the label of a loop in \mathcal{E} based at B . By Lemma 9.2, the subgroup H of W_Γ associated to (\mathcal{E}, B) has finite index in W_Γ . Furthermore, $g \notin H$ as w is a reduced word representing g which is not the label of a loop in \mathcal{E} . □

An additional useful property of full valence extensions is that the subgroup H associated to a full valence extension of a finite completion of $G < W_\Gamma$ has a retraction to G :

Proposition 9.4. *With the notation of Definition 9.1, there is a retraction from H to G .*

Proof. We define a map $\phi: H \rightarrow G$ as follows. Let \mathcal{L} be the set of graph-loops in \mathcal{E} that are not in Ω , i.e. the graph-loops added in the construction of \mathcal{E}_0 or in the completion process. Given an element $h \in H$, since \mathcal{E} is a completion of H by Proposition 4.4, there is a loop l in \mathcal{E} based at B whose label w is a word representing h . We remove from l all graph-loops it traverses which are in \mathcal{L} . Let l' be the resulting loop in Ω based at B , and let w' be its label. It follows that w' represents an element $g \in G$. We set $\phi(h) = g$.

We first check that ϕ is well-defined. Let l_1 and l_2 be loops in \mathcal{E} based at B with labels w_1 and w_2 , such that w_1 and w_2 are distinct words, each representing the same element $h \in H$. Let w be a reduced word representing h in W_Γ . Let l'_1, l'_2 and l' be the loops obtained by removing graph-loops in \mathcal{L} from l_1, l_2 and l respectively. Let w'_1, w'_2 and w' be the labels of l'_1, l'_2 and l' respectively. We must show that w'_1 and w'_2 represent the same element of W_Γ . To do so, we will show that w' represents the same element in W_Γ as both w'_1 and w'_2 .

By Tits' solution to the word problem, there is a sequence of Tits moves that can be performed to w_1 to obtain w . This sequence naturally produces a sequence of corresponding loops $l_1 = q_1, q_2, \dots, q_n = l$ in \mathcal{E} whose labels are the corresponding words obtained by the Tits moves. Furthermore, if a cancellation move is performed to q_i in order to obtain q_{i+1} , then as \mathcal{E} is folded, it readily follows that the edges involved in this cancellation move are either both in \mathcal{L} or both not in \mathcal{L} . Thus, by forgetting the Tits moves performed to generators which are labels of graph-loops in \mathcal{L} , this sequence of Tits moves induces a sequence of Tits moves performed to w'_1 to produce w' . Hence, w'_1 and w' represent the same element of W_Γ . Similarly, w'_2 and w' represent the same element of W_Γ . Consequently, ϕ is well-defined.

It is clear that ϕ is a homomorphism. Furthermore, given an element $g \in G$ and a loop l in \mathcal{E} based at B with label a reduced word representing g , we have that l is contained in the subcomplex $\Omega \subset \mathcal{E}$. Thus, l does not traverse any graph-loops in \mathcal{L} . It follows that ϕ restricted to elements of G is the identity. Hence, ϕ provides the desired retraction. □

The above proposition can now be used to recover a theorem of Haglund. Recall that a subgroup G of a group K is a virtual retract if G is a retract of a finite index subgroup of K .

Theorem 9.5 (Haglund [25]). *Let G be a quasiconvex subgroup of a RACG W_Γ . Then G is separable and is a virtual retract of W_Γ .*

Proof. By Theorem 8.4, there is a finite completion (Ω, B) of G . Let (\mathcal{E}, B) be the full valence extension of (Ω, B) , with associated subgroup H . Then H has finite-index in W_Γ by Lemma 9.2 and G is a retract of H by Proposition 9.4. Thus, G is a virtual retract.

It is well known that a virtual retract of a residually finite group is separable (see [25, Proposition 3.8]). So, G is separable. □

10. Reflection subgroups

A *reflection* in W_Γ is an element represented by a word of the form $ws w^{-1}$ where s is a generator in $V(\Gamma)$ and w is a word in W_Γ . In this section, we give a constructive argument to build a finite completion for any subgroup of W_Γ generated by a finite set

of reflections. In particular, such subgroups are always quasiconvex and their index can be computed.

10.1. Trimmed sets of reflections. In this subsection, we define the notion of a trimmed set of reflections (Definition 10.1). We show in Lemma 10.2 that any subgroup generated by a non-trimmed set of reflections can also be generated by a trimmed set of reflections. Moreover, the lengths of words in the trimmed set sum to a strictly smaller sum than those of the original set. This minimality property will facilitate our later arguments. Another advantage of trimmed reflections is that they form a *standard* generating set for a RACG subgroup by a theorem of Dyer (see Proposition 12.7). On the other hand, a non-trimmed set of reflections usually gives a non-standard generating set for the RACG subgroup.

Definition 10.1 (Trimmed reflection set). Let W_Γ be a RACG. We say that a set

$$\mathcal{R} = \{w_i s_i w_i^{-1} \mid w_i \in W_\Gamma \text{ and } s_i \in V(\Gamma), 1 \leq i \leq m\}$$

of reduced reflections is *trimmed* if for all $i \neq j$, no reduced expression for w_j begins with $w_i s_i$.

Lemma 10.2. *Let W_Γ be a RACG, and let G be a subgroup generated by a finite set of reflections \mathcal{R}' . Then G is generated by a trimmed set of reflections \mathcal{R} . Furthermore, there is a constructive algorithm to obtain \mathcal{R} from \mathcal{R}' , whose time-complexity only depends on the number $\sum_{r \in \mathcal{R}'} |r|$.*

Proof. Without loss of generality, we may assume elements in \mathcal{R}' are reduced. Let $g = w s w^{-1}$ be a reflection in \mathcal{R}' so that w has an expression $w = w' s' q$, where $h = w' s' w'^{-1}$ is another reflection in \mathcal{R}' and q is a word in W_Γ .

In \mathcal{R}' , we replace g with a reduced representative of the shorter length reflection

$$g' = h g h^{-1} = (w' q) s (w' q)^{-1},$$

to obtain a new set \mathcal{R}'' . The set \mathcal{R}'' still generates G , as $g = h^{-1} g' h$. By iteratively performing such replacements, we obtain the desired generating set \mathcal{R} . This process must end since at each step we obtain a set of generators whose lengths sum to a strictly smaller number than those in the previous step. □

10.2. A completion for reflection subgroups. Throughout this subsection, we fix the notation in the discussion below. This notation is also used in Section 12.

Let G be a subgroup of the RACG W_Γ , generated by a finite set of reflections:

$$\mathcal{R} = \{w_i s_i w_i^{-1} \mid w_i \in W_\Gamma \text{ and } s_i \in V(\Gamma), 1 \leq i \leq m\}.$$

By Lemma 10.2 we may assume without loss of generality that \mathcal{R} is trimmed.

Our goal is to give a finite completion (Ω_G, B) of G . We begin by describing the first complex Ω_0 in this completion. For each $1 \leq i \leq m$, we attach a subdivided

circle to the base vertex B , with label $w_i s_i w_i^{-1}$. Next, for each i , we fold the two copies of w_i onto one another, and we call this resulting graph Ω_0 . Thus, the graph Ω_0 has, for each $1 \leq i \leq m$, a path emanating from B and labeled by w_i , with a graph-loop labeled by s_i attached at its endpoint. By Theorem 3.12, any completion of Ω_0 is a standard completion of G .

Let \mathcal{T} denote the tree obtained by removing the graph-loops from Ω_0 . Let \mathcal{FT} be the folded tree obtained by iteratively performing fold operations to \mathcal{T} . Let $\Omega_{\mathcal{FT}}$ be a standard completion of \mathcal{FT} . By Proposition 7.3, we know that $\Omega_{\mathcal{FT}}$ is a finite CAT(0) cube complex. Furthermore, $\Omega_{\mathcal{FT}}$ is also a completion of \mathcal{T} by construction. Let $\hat{f}: \mathcal{T} \rightarrow \mathcal{FT} \rightarrow \Omega_{\mathcal{FT}}$ be the natural map. By a slight abuse of notation, we also denote by \hat{f} the natural map $\hat{f}: \mathcal{FT} \rightarrow \Omega_{\mathcal{FT}}$. Let $\hat{\mathcal{T}} := \hat{f}(\mathcal{T}) = \hat{f}(\mathcal{FT})$.

Given a vertex \hat{v} in $\Omega_{\mathcal{FT}}$, define

$$L_{\hat{v}} = \{s \in V(\Gamma) \mid \exists v \in V(\Omega_0) \text{ incident to a graph-loop labeled } s, \text{ such that } \hat{v} = \hat{f}(v)\}.$$

We would like to build the completion Ω_G by “adding back” the graph-loops to $\Omega_{\mathcal{FT}}$ and applying Lemma 4.9. However, there is a technical issue: when adding back a graph-loop labeled s to a vertex of \hat{v} of $\Omega_{\mathcal{FT}}$, a priori \hat{v} might already be incident to an edge labeled by s . If this were true, then the hypothesis of Lemma 4.9 would not be satisfied. The next two technical lemmas show this situation does not arise.

Lemma 10.3. *Let ks be a reduced word in W_Γ such that $s \in V(\Gamma)$ and k is a (possibly empty) word consisting only of generators that are adjacent to s in Γ . Then given any $v \in \Omega_0$ incident to a graph-loop labeled by s , no path in $\mathcal{T} \subset \Omega_0$ starting at v is labeled by an expression for ks .*

Proof. For a contradiction, suppose such a path α exists. We may assume that α is a geodesic in \mathcal{T} . For if not, then as \mathcal{T} is a tree, some generator would be consecutively repeated in the label of α , and we would be able to pass to a homotopic path. Let u be the endpoint of α .

Let β_1 be the geodesic from the base vertex B to v , with label h_1 . Then there is an element $r_1 = h_1 s h_1^{-1}$ in \mathcal{R} . We first show that u does not lie on β_1 . Suppose it does. Then h_1 has a suffix which is a reduced word equal in W_Γ to sk^{-1} . Since k commutes with s , it follows that the expression $h_1 s h_1^{-1}$ is not reduced, a contradiction.

Now let β_2 be the geodesic from B to u , with label h_2 . Since u does not lie on β_1 , it follows that h_2 is non-empty. Moreover, there is a reflection $r_2 \in \mathcal{R}$, given by $(h_2 h_2') s' (h_2'^{-1} h_2^{-1})$, where h_2' could be empty. Next, we claim that h_1 is non-empty. For if not, then $r_1 = s$, and h_2 is an expression for ks . It follows that r_2 has a reduced expression that begins with s , which is not possible as \mathcal{R} is trimmed.

Thus, h_1 and h_2 are non-empty and ks is a reduced expression for $h_1^{-1} h_2$. By Lemma 2.3 there exist (possibly empty) words x , k' and k'' such that either $k'x$ and $x^{-1}k''s$ are reduced expressions in W_Γ for respectively h_1^{-1} and h_2 ,

or alternatively $sk'x$ and $x^{-1}k''$ are reduced expressions in W_Γ for, respectively, h_1^{-1} and h_2 . Moreover, $k'k''$ is equal to k in W_Γ . The latter case implies that the reflection r_1 has another reduced expression $(sk'x)^{-1}s(sk'x)$. However, this is a contradiction as this word is clearly not reduced. In the former case, we have that r_1 has reduced expression $(k'x)^{-1}s(k'x)$ and r_2 has reduced expression $(x^{-1}k''sh'_2)s'(x^{-1}k''sh'_2)^{-1}$. Set $w_1 = (k'x)^{-1}$ and $w_2 = x^{-1}k''sh'_2$. As the given expression for r_1 is reduced and as s commutes with k' , it must be that k' is the empty word. In particular, $w_1 = x^{-1}$. Furthermore, since s commutes with k'' , w_2 has reduced expression $x^{-1}sk''h'_2$. However, this contradicts our choice of \mathcal{R} since some expression for w_2 begins with w_1s . The claim follows. \square

Lemma 10.4. *For every vertex \hat{v} of $\Omega_{\mathcal{F}\mathcal{T}}$ and every $s \in L_{\hat{v}}$, there is no edge labeled by s incident to \hat{v} in $\Omega_{\mathcal{F}\mathcal{T}}$*

Proof. For a contradiction, suppose that there exists some \hat{v} of $\Omega_{\mathcal{F}\mathcal{T}}$, $s \in L_{\hat{v}}$ and an edge d incident to \hat{v} in $\Omega_{\mathcal{F}\mathcal{T}}$ which is labeled by s . Note that $\hat{v} \in \hat{\mathcal{T}}$.

Let H be the hyperplane (recall that $\Omega_{\mathcal{F}\mathcal{T}}$ is a CAT(0) cube complex) dual to d . In particular, H is of type s . By Proposition 4.8, H intersects $\hat{\mathcal{T}}$ at some edge \hat{e} . Let e be an edge of \mathcal{T} , such that $\hat{f}(e) = \hat{e}$. Note that the label of e must be s . Let $v \in \mathcal{T}$ be such that $\hat{f}(v) = \hat{v}$. Let β be a geodesic in \mathcal{T} from v to e . Let $\hat{\beta} = \hat{f}(\beta)$. It follows that $\hat{\beta}$ is a path in $\Omega_{\mathcal{F}\mathcal{T}}$ from \hat{v} to \hat{e} . By Lemma 4.2, there exists a geodesic $\hat{\beta}'$ with the same endpoints as $\hat{\beta}$ and with label a reduced expression for the label of $\hat{\beta}$. Finally, let γ be a path in the carrier of H from \hat{v} to the endpoint of \hat{e} . Note that the label of γ only consists of vertices in $\text{link}(s)$.

As $\hat{\beta}'$ is geodesic, any hyperplane intersects it at most once. Thus, as hyperplanes separate $\Omega_{\mathcal{F}\mathcal{T}}$ into two components, it follows that any hyperplane that intersects $\hat{\beta}'$ must also intersect γ . It follows that the label of $\hat{\beta}'$ consists only of vertices in $\text{link}(s)$. However, the label of $\hat{\beta}'$ and the label of β are expressions for the same element of W_Γ . This implies that $\beta \cup e$ is a path in \mathcal{T} based at v whose label is an expression for the word ks , where every generator in k is adjacent to s in Γ . This contradicts Lemma 10.3. \square

We are now ready to prove the main results of this section.

Theorem 10.5. *Let G be a finitely generated reflection subgroup of a RACG. Then there exists a finite completion of G .*

Proof. As previously discussed, we first obtain the completion $\Omega_{\mathcal{F}\mathcal{T}}$ of \mathcal{T} using Proposition 7.3. For each vertex \hat{v} of $\Omega_{\mathcal{F}\mathcal{T}}$ and $s \in L_{\hat{v}}$, we attach a graph-loop to $\hat{v} \in \Omega_{\mathcal{F}\mathcal{T}}$ labeled by s . Let $\Omega'_{\mathcal{F}\mathcal{T}}$ be the resulting complex. Note that by Lemma 10.4, such a graph-loop is never attached to a vertex that is incident to an edge with the same label as the graph-loop. Furthermore, note that $\Omega'_{\mathcal{F}\mathcal{T}}$ can be obtained from Ω_0 by applying the same completion sequence that was applied to \mathcal{T} to obtain $\Omega_{\mathcal{F}\mathcal{T}}$,

while “ignoring” the graph-loops. We now get a finite completion Ω_G of $\Omega'_{\mathcal{F}\mathcal{T}}$ by applying Lemma 4.9. It follows that Ω_G is a finite completion for G . \square

Example 10.6. Let Γ be the graph in Figure 7. The right of this figure shows a completion for the reflection subgroup $\langle (dac)b(dac)^{-1}, (dc)d(dc)^{-1} \rangle$.




Figure 7. A completion for a reflection subgroup.

The following corollary immediately follows from Theorem 10.5 and Theorem 8.4.

Corollary 10.7. *Every finitely generated reflection subgroup of a RACG is quasiconvex.*

For 2-dimensional RACGs, we obtain the following stronger result which shows that the time-complexity of the algorithm which builds the completion of a reflection subgroup is bounded by the size of words in the generating set of reflections. This is used in Section 12.

Theorem 10.8. *Let W_Γ be a 2-dimensional RACG. Let G be a subgroup of W_Γ generated by a finite set of reflection words \mathcal{R} . Then there is a finite completion sequence for G whose length only depends on the numbers $\sum_{r \in \mathcal{R}} |r|$ and $|V(\Gamma)|$.*

Proof. By Lemma 10.2, we may assume without loss of generality that \mathcal{R} is trimmed. As before, we first obtain the completion $\Omega_{\mathcal{F}\mathcal{T}}$ of \mathcal{T} . However, this time we use the more refined Proposition 7.4 which guarantees that the number of steps in this completion sequence only depends on $|\Omega_0|$ and $|V(\Gamma)|$. The rest of the proof follows by repeating the proof of Theorem 10.5 and noting that Lemma 4.9 (3) guarantees the bound. \square

Remark 10.9. By Lemma 4.9, the completion Ω_G of the reflection group G given by Theorems 10.5 and 10.8 contains the complex $\Omega_{\mathcal{F}\mathcal{T}}$ as an isometrically embedded subcomplex. Moreover, the inclusion of $\Omega_{\mathcal{F}\mathcal{T}}$ into Ω_G satisfies the additional properties given by Lemma 4.9.

11. Coxeter subgroups of 2-dimensional RACGs

In this section we study Coxeter subgroups of 2-dimensional RACGs. Recall that such subgroups, by our definition, are always finitely generated. It is clear that a Coxeter subgroup of a RACG W_Γ must be generated by involutions (order two elements)

in W_Γ . We show in Theorem 11.4 that under mild hypotheses these subgroups are generated by reflections. Consequently, they are quasiconvex by Theorem 10.5 and Theorem 8.4.

We first prove three lemmas involving the structure of particular types of words in a RACG. The first is well known and addresses involutions:

Lemma 11.1. *Let g be an involution in the RACG W_Γ . Then there is an expression kwk^{-1} for g , such that every generator in the word k is in a common clique of Γ .*

Proof. It is well known that every finite subgroup of a RACG is contained in a conjugate of a special finite subgroup (see [20, Theorem 12.3.4] for instance). The lemma follows, as a special subgroup of a RACG is finite if and only if its defining graph is a clique. \square

The next lemma concerns the structure of particular types of commuting words.

Lemma 11.2. *Let $w = s_1 \dots s_m$ and $k = k_1 \dots k_n$ be reduced commuting words in the RACG W_Γ and suppose that the vertices $k_1, \dots, k_n \in V(\Gamma)$ are all contained in a common clique of Γ . Then for each $1 \leq i \leq m$, either:*

- (1) $s_i = k_r$ for some $1 \leq r \leq n$ and $m(s_i, s_j) = 2$ for all $j \neq i$, or
- (2) $m(s_i, k_j) = 2$ for all $1 \leq j \leq n$.

Proof. Note that as k is reduced and k_1, \dots, k_n pairwise commute, it follows that $k_i \neq k_j$ for all $i \neq j$. As w and k commute, there exists a disk diagram R with boundary label kwk^{-1} . We think of R as a rectangle. The vertical sides are labeled, read from bottom to top, by w , and the horizontal sides of R are labeled, read from left to right, by k . As w and k are reduced, no dual curve intersects the same side of R twice.

Fix $s \in \Gamma$ such that $s = s_i$ for some $1 \leq i \leq m$. Let e_1, \dots, e_l be the set of edges labeled by s on the left side of R , ordered from bottom to top. Let e'_1, \dots, e'_l be the edges labeled by s on the right side of R , ordered from bottom to top. We think of the edge e_i as lying “directly across” from e'_i in R . Consider the set \mathcal{H} of all dual curves in R of type s .

Suppose first that $s \neq k_r$ for all $1 \leq r \leq n$. As dual curves of the same type do not intersect, it follows that for each $1 \leq j \leq l$, there is a curve in \mathcal{H} intersecting both e_j and e'_j . Let H be the dual curve in \mathcal{H} that is bottom-most in R , i.e., that intersects e_1 and e'_1 . Let α be the path along the boundary of R from the bottom of e_1 to the bottom of e'_1 . Let $t = k_r$ for any $1 \leq r \leq n$. Observe that the label of α has an odd number of occurrences of the letter t . Hence some curve intersecting an edge in α labeled by t must intersect H . Thus, $m(s, t) = 2$ and item (2) in the statement of the lemma holds in this case.

On the other hand, suppose that $s = k_r$ for some $1 \leq r \leq n$. Let d and d' respectively be the edges on the bottom and top of R labeled by k_r . The dual curves in \mathcal{H} must take one of two possible forms. The first possibility is that there are

curves in \mathcal{H} intersecting the following pairs of edges: d and e_1, e'_j and e_{j+1} for $1 \leq j \leq l - 1$, and e'_l and d' . Otherwise, there are curves in \mathcal{H} intersecting the pairs d and e'_1, e_j and e'_{j+1} for $1 \leq j \leq l - 1$, and e_l and d' .

Let t be a letter in w such that $t \neq s$. If t also appears in k , then $m(s, t) = 2$. Otherwise, the dual curves labeled by t all go across R , such that for all j , the j th edge labeled t on the left side is paired with the j th edge labeled t on the right side. Now the structure of dual curves in \mathcal{H} (in either case) forces each dual curve labeled t to intersect a curve in \mathcal{H} . It follows that $m(s, t) = 2$. Thus, s commutes with every generator of w that is not equal to s . All that is left to show is that there is only one occurrence of the generator s in w , i.e., that $l = 1$. However, if there were two occurrences of s , these occurrences can be deleted (as s commutes with every generator in w). This is not possible, as w is reduced. Thus, $m(s_i, s_j) = 2$ for all $1 \leq j \leq m$ such that $j \neq i$, and item (1) in the statement of the lemma holds. \square

The following lemma about commuting words will be used in the proof of Theorem 11.4.

Lemma 11.3. *Let b and $x = zs_1s_2z^{-1}$ be reduced commuting words in a RACG W_Γ , where $s_1, s_2 \in V(\Gamma)$ and z is a word in W_Γ . Suppose also that s_1 commutes with both s_2 and z . Then b commutes with $\hat{x} = zs_2z^{-1}$*

Proof. As in the previous lemma, we form a “rectangular” disk diagram R with boundary label $bx b^{-1}x^{-1}$. The vertical sides of R are labeled, read from bottom to top, by b , and the horizontal sides of R are labeled, read from left to right, by x . As b and x are reduced, no dual curve intersects the same side of R twice. Let \mathcal{H} be all dual curves in R of type s_1 .

As x is reduced and s_1 commutes with both s_2 and z , it readily follows that there is only one occurrence of s_1 in x . Let e_t be the unique edge on the top of R labeled by s_1 , and let e_b be the unique edge on the bottom of R labeled by s_1 .

First suppose that b does not contain any occurrences of the generator s_1 . Then \mathcal{H} consists of a single dual curve H which intersects both e_t and e_b . Let $N(H)$ be the set of cells in R that contain an edge intersecting H . It follows that the boundary $\partial N(H)$ of $N(H)$ has label $s_1 y s_1 y^{-1}$ for some word y . We can then excise $(N(H) \setminus \partial N(H)) \cup e_t \cup e_b$ from R and then glue back together the resulting components along their boundary paths labeled by y . What results is a new disk diagram with boundary label $b\hat{x}b^{-1}\hat{x}^{-1}$. Thus, b commutes with \hat{x} .

On the other hand, suppose that b has one or more occurrence of s_1 . We consider two cases. The first case is that s_1 commutes with every generator of b . As b is reduced, it readily follows that there is a unique occurrence of s_1 in b . Furthermore, $s_1\hat{b}$ and $\hat{b}s_1$ are both expressions for b , where \hat{b} is the word obtained from b by removing the generator s_1 . Now we have the following equalities in W_Γ :

$$\hat{b}\hat{x} = \hat{b}s_1s_1\hat{x} = bx = xb = \hat{x}s_1s_1\hat{b} = \hat{x}\hat{b}.$$

Thus, \widehat{b} and s_1 both commute with \widehat{x} , so b does as well, and the lemma follows for this case.

For the second case, suppose there are generators in b that do not commute with s_1 . We will show that this case is actually not possible as we obtain a contradiction.

Let e_1, \dots, e_l be the set of edges labeled by s_1 on the left side of R ordered from bottom to top, and let e'_1, \dots, e'_l be the set of edges labeled by s_1 on the right side of R ordered from bottom to top. As in the previous lemma, the dual curves in \mathcal{H} either intersect the pairs e_b and e_1, e'_j and e_{j+1} for $1 \leq j \leq l - 1$, and e'_l and e_t , or intersect the pairs e_b and e'_1, e_j and e'_{j+1} for $1 \leq j \leq l - 1$, and e_l and e_t . We assume that the dual curves in \mathcal{H} have the first configuration described (the proof in the other possible configuration same).

Consider the first occurrence in b of a generator t that does not commute with s_1 , and let d be the corresponding edge on the left side of R labeled by t . Note that d is the “bottom-most” edge on the left side of R with label t . Let T be the curve in R that intersects d . Note that T cannot intersect any dual curve in \mathcal{H} as s_1 and t do not commute. Furthermore, T cannot intersect an edge on the top or bottom of R as the label of every such edge commutes with s_1 . Thus, it readily follows by the structure of dual curves in \mathcal{H} that d cannot lie before e_1 . Similarly, d cannot occur after e_l , as then the edge on the right side of R labeled by t occurs after e'_l , which is again not possible by the structure of \mathcal{H} .

Thus, d must lie between e_r and e_{r+1} for some $1 \leq r < l$. However, as T cannot intersect a dual curve in \mathcal{H} and cannot intersect the bottom of R , it follows that T intersects an edge on the right side of R that lies before e'_r . Correspondingly, there is an edge on the left side of R lying below e_r with label t . This contradicts the fact that d is the bottom-most such edge on the left side of R . The lemma follows. \square

An *isolated vertex* of a graph is a vertex that is not adjacent to any other vertex. If W_Γ is a 2-dimensional RACG, we show:

Theorem 11.4. *Let G be a subgroup of a RACG W_Γ , where Γ is triangle-free. Suppose that G is isomorphic to the Coxeter group $W_{\Gamma'}$ (which is right-angled by Proposition 2.1), where Γ' does not have an isolated vertex. Then G is generated by a finite set of reflections in W_Γ .*

Proof. As G is a RACG, it is generated by the standard Coxeter generating set corresponding to Γ' . In particular, there exists a finite generating set I_G for G consisting of reduced words representing involutions in W_Γ , such that

- (1) there is a bijective map from $V(\Gamma')$ to the elements of I_G , and
- (2) if there is an edge between two vertices of Γ' then the corresponding elements of I_G commute.

We inductively construct a sequence of generating sets for G , each consisting of $|I_G|$ reduced words representing involutions in W_Γ and satisfies properties (1) and (2).

Furthermore, each generating set in the sequence will contain one more reflection than the previous one.

Let r be the number of reflections in I_G . If $r = |I_G|$, then we are done. Otherwise, let $h \in I_G$ be such that h is not a reflection. By properties (1) and (2) above, and since Γ' does not have isolated vertices, there exists an $h' \in I_G$ which is distinct from and commutes with h . By Lemma 11.1 and since Γ is triangle-free, we conclude that $h = ws_1s_2w^{-1}$ and $h' = w'k'w'^{-1}$, where s_1 and s_2 are adjacent vertices in Γ , k' is a word of length at most two whose generators are in a common clique, and w, w' are words in W_Γ .

Consider the subgroup $H = w'^{-1}Gw'$ of W_Γ . Note that H is generated by $I_H = w'^{-1}I_Gw'$. Let x be a reduced expression for $w'^{-1}hw'$ obtained by applying a sequence of deletions. Then x must be of the form $x = zs_1s_2z^{-1}$ for some word z in W_Γ . Note that $k' = w'^{-1}h'w' \in I_H$ and that k' and x commute.

We claim that either $k' = s_1$ or $k' = s_2$. First suppose for a contradiction that k' is of length two, say $k' = k_1k_2$, for some distinct $k_1, k_2 \in V(\Gamma)$. If $k_1 \neq s_1$ and $k_1 \neq s_2$, then by Lemma 11.2, we have that $m(k_1, s_1) = m(k_1, s_2) = 2$. However, this is not possible as Γ is triangle-free. Thus k_1 is equal to either s_1 or s_2 , and similarly, k_2 is equal to either s_1 or s_2 . It follows that, up to relabeling, $k' = s_1s_2$. Lemma 11.2 further implies that s_1 and s_2 commute with z . Consequently, $x = s_1s_2 = k'$, which is a contradiction since h and h' are distinct.

Suppose now that k' has length one. Again by Lemma 11.2 and the fact that Γ is triangle-free, k' cannot consist of a generator distinct from s_1 and s_2 . Thus $k' = s_1$ or $k' = s_2$. By possibly relabeling, we may assume that $k' = s_1$. Now Lemma 11.2 implies that s_1 commutes with z and s_2 . Let $y = k'x = zs_2z^{-1}$. We replace x with y in I_H to form the new set I'_H . Note that I'_H is still a generating set for H as $x = k'y$.

Let $b \neq s_1$ be any element of I'_H which commutes with x . Then since $x = zs_1s_2z^{-1}$ and s_1 commutes with z and s_2 , Lemma 11.3 implies that b commutes with y . It follows that $w'I'_Hw'^{-1}$ is a generating set for G which satisfies properties (1) and (2) above. Finally, note that the number of reflections in I'_H , and hence in $w'I'_Hw'^{-1}$ is exactly $r + 1$.

By repeating this process enough times, we are guaranteed a finite generating set for G consisting only of reflections. \square

It follows easily from Corollary 10.7 that subgroups as in Theorem 11.4 are quasiconvex:

Corollary 11.5. *Given a 2-dimensional RACG, every Coxeter subgroup whose defining graph does not have an isolated vertex is quasiconvex.*

Remark 11.6. Note that the defining graph of a RACG has an isolated vertex if and only if the group splits as a free product with a \mathbb{Z}_2 factor. Some version of the non-splitting hypothesis is required in Theorem 11.4 and Corollary 11.5. For consider the graph Γ_2 from Figure 2. The subgroup of W_{Γ_2} generated by ab and cd

is isomorphic to the infinite dihedral group. In particular, it is a right-angled Coxeter subgroup of W_{Γ_2} . However, it is straightforward to check that this subgroup cannot be generated by reflections and has an infinite completion (so is not quasiconvex by Theorem 8.4).

12. Deciding when a RACG is a finite-index subgroup of a 2-dimensional RACG

In this section we give an algorithm which, given a 2-dimensional, one-ended RACG W_Γ and any RACG $W_{\Gamma'}$, determines whether or not W_Γ contains a finite-index subgroup isomorphic to $W_{\Gamma'}$ (Theorem 12.8). In fact, the one-endedness hypothesis can be weakened. When such a subgroup does exist, the output of the algorithm is a set of words in W_Γ which is a standard RACG generating set for a subgroup isomorphic to $W_{\Gamma'}$. Furthermore, the time-complexity of this algorithm only depends on the number of vertices of Γ and Γ' .

Given a set of reflections \mathcal{R} which generate a finite-index subgroup of W_Γ , under the right hypotheses, Proposition 12.1 below bounds the sizes of elements in \mathcal{R} as a function of $|V(\Gamma)|$ and $|\mathcal{R}|$. This is a key step in the proof of the main theorem of this section, as it allows us to bound the number of sets of reflections that need to be investigated by our algorithm.

To state the proposition, we require the following definition: we say a graph Δ is *almost star* if there exist vertices $s, t \in \Delta$ (possibly not distinct) such that $V(\Delta) = \text{star}(s) \cup \{t\}$.

Proposition 12.1. *Let W_Γ be a RACG such that Γ is triangle-free and not almost star. Let*

$$\mathcal{R} = \{w_i s_i w_i^{-1} \mid w_i \in W_\Gamma \text{ and } s_i \in V(\Gamma), 1 \leq i \leq N\}$$

be a trimmed set of reflections which generates a finite-index subgroup $G < W_\Gamma$. Then there exists a constant M , depending only on $|V(\Gamma)|$ and $|\mathcal{R}|$ such that $|w_i| \leq M$ for all $w_i s_i w_i^{-1} \in \mathcal{R}$.

In order to prove this proposition, we establish some notation and prove some preliminary lemmas. The notation below will be fixed until the proof of Proposition 12.1 is complete.

Let M be the smallest integer such that if a reduced word w in W_Γ is longer than M , then w contains $2N + 2$ occurrences of some letter of $V(\Gamma)$ (where $N = |\mathcal{R}|$). In particular, M only depends on $|V(\Gamma)|$ and on $|\mathcal{R}|$. This M will be the same as the constant in the proposition.

In order to establish a contradiction, we assume that $|w_1| > M$. By possibly relabeling, we assume that $|w_1| \geq |w_i|$ for all $1 < i \leq N$. By the previous paragraph, we may fix some vertex \bar{s} of Γ which occurs at least $2N + 2$ times in the word w_1 .

We now choose convenient expressions for the elements of \mathcal{R} . Firstly, we assume w_1 is written in an expression where occurrences of \bar{s} appear as far left as possible. More formally, if $w_1 = s_1 \dots s_m$ and $s_i = \bar{s}$, then for all $j < i$ there is no expression for w_1 in W_Γ equal to the word $s_1 \dots s_{j-1} s_{j+1} \dots s_i s_j s_{i+1} \dots s_m$.

Given two words w and w' in W_Γ , let $\phi(w, w')$ denote the length of their largest common prefix. For each $2 \leq i \leq N$, we choose an expression for w_i so that $\phi(w_1, w_i)$ is maximal out of all such possible choices for w_i . Clearly, there is no loss of generality in making these assumptions on w_1, \dots, w_N .

We will now use the notation established in Section 10.2 associated to a reflection subgroup of a RACG. As in that section, we have the based Γ -labeled complex (Ω_0, B) , the labeled tree $\mathcal{T} \subset \Omega_0$ and the associated folded based labeled tree (\mathcal{FT}, B) . Furthermore, $(\Omega_{\mathcal{FT}}, B)$ is a based finite CAT(0) cube complex which is a completion of (\mathcal{FT}, B) . By Theorem 10.8, there is a completion (Ω_G, B) of (Ω_0, B) (which is also a completion for G) whose associated completion sequence has length bounded by a function which depends only on $\sum_{r \in \mathcal{R}} |r|$ and $|V(\Gamma)|$. We again denote by \hat{f} the natural map

$$\hat{f}: \mathcal{FT} \rightarrow \Omega_{\mathcal{FT}} \subset \Omega_G$$

and recall that $\Omega_{\mathcal{FT}}$ is isometrically embedded in Ω_G .

Let \mathcal{V} be the set of vertices in \mathcal{FT} that are the image of a vertex of $\mathcal{T} \subset \Omega_0$ which has a graph-loop attached to it. Observe that $|\mathcal{V}| \leq N$. Also note that at most N vertices in \mathcal{FT} have valence larger than 2.

Let α be the path in \mathcal{FT} based at B with label w_1 . As there are at least $2N + 2$ occurrences of \bar{s} in the word w_1 , there must exist two edges of α , say e_1 and e_2 , each with label \bar{s} such that every vertex between e_1 and e_2 has valence 2 and is not in \mathcal{V} . Let γ be the geodesic in \mathcal{FT} between e_1 and e_2 . By possibly passing to a subpath, we may assume that no edge in γ has label \bar{s} . We also assume that e_1 is closer to B than e_2 .

We now sketch how a contradiction will be established. As G is a finite-index subgroup of W_Γ , it will follow that Ω_G must be full-valence. We then focus on a specific vertex of Ω_G which is contained in $\hat{f}(\gamma)$. We show that the structure of edges and graph-loops incident to this vertex, together with the assumption that Ω_G is full valence, must imply that Γ is an almost star graph, contradicting the hypotheses of Proposition 12.1. In order to carry out this argument, we must first gain a solid understanding of the structure of the subcomplex $\hat{f}(\gamma)$. This is the purpose of the next four lemmas.

Lemma 12.2. *Let v be a vertex of γ , and let v' be any vertex in \mathcal{FT} . If $\hat{f}(v) = \hat{f}(v')$, then $v = v'$.*

Proof. Let β be a path in \mathcal{FT} from the base vertex B to v , and let β' be a path in \mathcal{FT} from B to v' . The label l of β is a prefix of w_1 . Similarly, the label l' of β' is

a prefix of w_i , for some $1 \leq i \leq N$. Let $\hat{\beta} = \hat{f}(\beta)$ and $\hat{\beta}' = \hat{f}(\beta')$ be the images of these geodesics in $\Omega_{\mathcal{FT}}$.

As $\Omega_{\mathcal{FT}}$ is a CAT(0) cube complex, the loop $\hat{\beta}' \cup \hat{\beta}^{-1}$ is homotopic, relative to basepoint, to B . Thus, by Lemma 4.2 (3), the label of $\hat{\beta}' \cup \hat{\beta}^{-1}$ is equal to the identity element in W_Γ . It follows that l' and l represent the same element of W_Γ . However, our choice of w_i guarantees that w_i and w_1 share the largest possible prefix. It follows that l and l' are the same word. As \mathcal{FT} is folded, it follows that $\beta = \beta'$ and that $v = v'$. □

Lemma 12.3. *Let β be a path in $\Omega_{\mathcal{FT}}$ with label a reduced word in W_Γ . Then β is a geodesic.*

Proof. Suppose β is not geodesic. As $\Omega_{\mathcal{FT}}$ is a CAT(0) cube complex, it follows some hyperplane K is dual to two distinct edges k and k' of β . Furthermore, we can choose K, k and k' so that every hyperplane dual to an edge of β between k and k' intersects K . However, from this it readily follows that the label of k (and of k') commutes with the label of any edge of β between k and k' . This implies that the label of β is not reduced, a contradiction. □

For the next two lemmas, let $\hat{e}_1 = \hat{f}(e_1)$ and $\hat{e}_2 = \hat{f}(e_2)$. Furthermore, let H_1 and H_2 be the hyperplanes in $\Omega_{\mathcal{FT}}$ that are dual respectively to \hat{e}_1 and \hat{e}_2 .

Lemma 12.4. *The edge \hat{e}_1 is the only edge of $\hat{f}(\mathcal{FT})$ that is dual to H_1 . Similarly, the edge \hat{e}_2 is the only edge of $\hat{f}(\mathcal{FT})$ that is dual to H_2 .*

Proof. We prove the claim for H_1 . The proof is analogous for H_2 . Let \hat{e} be an edge in $\hat{f}(\mathcal{FT}) \subset \Omega_G$ dual to H_1 . We will show that $\hat{e} = \hat{e}_1$.

As \mathcal{FT} is a tree, if the edge e_1 is removed (but its endpoints are not removed), there are exactly two resulting components. We let C_B denote the component that includes the vertex B , and let \bar{C}_B be the component which does not.

Let e be an edge of \mathcal{FT} such that $\hat{f}(e) = \hat{e}$. We first claim that e cannot be in \bar{C}_B . For suppose otherwise. It follows that e and e_1 are contained in a common path η with reduced label w_j for some $1 \leq j \leq N$. By Lemma 12.3 $\hat{f}(\eta)$ is a geodesic in $\Omega_{\mathcal{FT}}$. However, it now follows that the hyperplane H_1 is dual to two edges of a geodesic, contradicting the fact that a hyperplane in a CAT(0) cube complex is dual to at most one edge of a geodesic.

Thus, we may assume that either $e \in C_B$ or $e = e_1$. Let β be a geodesic in \mathcal{FT} from e_1 to e , which includes these two edges. Due to the tree structure of \mathcal{FT} , it follows that the label of β is $k_1^{-1}k_2$, where wk_1 is a prefix of w_1 and wk_2 is a prefix of w_j for some $1 \leq j \leq N$ and some reduced word w . Note that w, k_1 and k_2 could each be the empty word. Let $\hat{\beta} = \hat{f}(\beta)$ be the corresponding path in $\Omega_{\mathcal{FT}}$.

Let ζ be a geodesic along the carrier of H_1 from the endpoint to start point of $\hat{\beta}$. Let z be the label of ζ . Let D be a disk diagram in $\Omega_{\mathcal{FT}}$ with boundary $\hat{\beta} \cup \hat{\zeta}^{-1}$. The label of ∂D is $k_1^{-1}k_2z$. Let b_1 and b_2 be the paths along ∂D with

labels respectively k_1 and k_2 . Let c the path along ∂D labeled by z . Write the label of b_1 as $k_1 = s_1 s_2 \dots s_m$, where $s_l \in V(\Gamma)$ for $1 \leq l \leq m$. Note that $s_m = \bar{s}$. For $1 \leq l \leq m$, let d_l be the edge of b_1 with label s_l .

We claim that no dual curve intersects both b_1 and c . For suppose there is such a dual curve P . Further suppose that P intersects the edge d_r such that r is maximal out of such possible choices. Note that $r \neq m$ as z only contains letters in $\text{link}(\bar{s})$, being the label of a geodesic in the carrier of a hyperplane of type \bar{s} . It follows that every dual curve intersecting d_l , for $l > r$, intersects P . Let $p \in \Gamma$ be the type of P . As P intersects c , and the label of c is in $\text{link}(\bar{s})$, it follows that $p \in \text{link}(\bar{s}) \subset V(\Gamma)$. Additionally, p commutes with s_l for every $r < l \leq m$. However, this implies that $s_1 \dots s_{r-1} s_{r+1} \dots s_m s_r$ is an expression for k_1 . As k_1 is a subword of w_1 , this contradicts our choice of w_1 having occurrences of \bar{s} appear as “left-most” as possible. Thus, every dual curve intersecting b_1 must intersect b_2 .

By Lemma 2.4, k_2 is an expression in W_Γ for $k_1 k$, where k is possibly empty. However, it follows from our choice of expression for w_j (we chose expressions for words in \mathcal{R} to have maximal common prefix with w_1) that k_2 and $k_1 k$ are actually equal as words. Consequently $w k_1$ is a prefix of both w_1 and w_j . Hence, we conclude that $e = e_1$ and so $\hat{e} = \hat{e}_1$. □

Lemma 12.5. *Let Y be the subcomplex of $\Omega_{\mathcal{F}\mathcal{T}}$ bounded by H_1 and H_2 . Let \hat{v} be a vertex in Y . Then the label of any graph-loop in Ω_G incident to \hat{v} (where we think of $\hat{v} \in \Omega_{\mathcal{F}\mathcal{T}} \subset \Omega_G$) is in $\text{link}(\bar{s}) \subset V(\Gamma)$.*

Proof. Let u be any vertex of Ω_0 that is incident to a graph-loop. By construction, the image of u in $\mathcal{F}\mathcal{T}$ is not contained in γ . By Lemma 12.4 and the fact that every vertex of γ has valence 2 by construction, $\hat{f}(\mathcal{F}\mathcal{T}) \cap Y = \gamma$. Thus either H_1 or H_2 separates \hat{v} from $\hat{f}(u)$.

Suppose the label of the graph-loop attached to \hat{v} is t . By Lemma 4.9 and Remark 10.9, there is a path η in $\Omega_{\mathcal{F}\mathcal{T}}$ from \hat{v} to a vertex \hat{u}' such that $\hat{u}' = \hat{f}(u')$, where $u' \in \Omega_0$ is a vertex that is incident to a graph-loop labeled by t . Furthermore, the label of η is a word in $\text{link}(t)$. It follows that η must intersect either H_1 or H_2 . Hence, the label of η contains the generator \bar{s} . Thus, $\bar{s} \in \text{link}(t)$ and so $t \in \text{link}(\bar{s})$. □

We are now ready to prove the proposition.

Proof of Proposition 12.1. Let e be the edge of γ adjacent to e_2 . Let v be the vertex $e \cap e_2$. Let t be the label of e . By our choice of w_1 (having \bar{s} occurrences appear “left-most”), t and \bar{s} are not adjacent vertices of Γ . Set $\hat{e} = \hat{f}(e)$, $\hat{e}_2 = \hat{f}(e_2)$, $\hat{v} = \hat{f}(v)$ and $\hat{\gamma} = \hat{f}(\gamma)$.

As Γ is not almost star, it readily follows that $\text{star}(s) \subsetneq V(\Gamma)$ for any $s \in V(\Gamma)$. In particular, Ω_G is resolved. As G is a finite-index subgroup, Ω_G is full valence by Theorem 6.6.

Again, as Γ is not almost star, there exists a vertex $a \in \Gamma$ such that $a \neq t$ and $a \notin \text{star}(\bar{s})$. As Ω_G is full valence, there must exist an edge \hat{d} adjacent to \hat{v} labeled by a . By Lemma 12.5, \hat{d} is not a graph-loop. Let \hat{u}' be the vertex of \hat{d} which is not equal to \hat{v} . As Ω_G is full valence, there must exist an edge \hat{d}' adjacent to \hat{u}' with label \bar{s} which is not a graph-loop by Lemma 12.5. Let H be the hyperplane dual to \hat{d}' .

First note that \hat{d}' cannot be dual to H_2 , for otherwise it would follow from the convexity of $N(H_2)$ that $\hat{d} \subset N(H_2)$, contradicting the fact that a is not in $\text{star}(\bar{s})$. Furthermore, \hat{d}' cannot be dual to H_1 either. For otherwise, as $\hat{\gamma}$ is geodesic (by Lemma 12.3), it follows that the hyperplane dual to \hat{e} must intersect H_1 , contradicting the fact that t is not in $\text{link}(\bar{s})$. Thus, $H \neq H_1$ and $H \neq H_2$.

By Proposition 4.8, H must intersect $\hat{f}(\mathcal{FT}) \subset \Omega_G$. As $\hat{\gamma}$ does not have any edges labeled by \bar{s} , it follows that H cannot intersect $\hat{\gamma}$. Thus, by Lemma 12.4, H must intersect either H_1 or H_2 . However, this is a contradiction as H, H_1 and H_2 are all of type \bar{s} . □

Before proving Theorem 12.8, we address a special case. The next lemma describes finite-index subgroups of W_Γ for the case where Γ is a triangle-free join graph.

Lemma 12.6. *Suppose Γ is a triangle-free graph which splits as a join $\Gamma = A \star B$. Let \mathcal{R} be a finite set of reduced reflection words in W_Γ which generates the subgroup $G < W_\Gamma$. Then G is a finite-index subgroup of W_Γ if and only if $\mathcal{R} = \mathcal{R}_A \cup \mathcal{R}_B$ such that*

- (1) \mathcal{R}_A (resp. \mathcal{R}_B) consists only of words in W_A (resp. W_B);
- (2) \mathcal{R}_A (resp. \mathcal{R}_B) generates a finite-index subgroup of W_A (resp. W_B).

Proof. The “if” direction is immediate. For the other direction, suppose that G has finite index in W_Γ . Let

$$\mathcal{R}_A = \{ws w^{-1} \in \mathcal{R} \mid s \in V(A)\}$$

and let $\mathcal{R}_B = \mathcal{R} \setminus \mathcal{R}_A$. Note that if $ws w^{-1} \in \mathcal{R}_A$, then w does not have a letter in $V(B)$. For if it did, $ws w^{-1}$ would not be reduced. Similarly, every reflection in \mathcal{R}_B does not contain a letter of $V(A)$. This shows (1).

Thus, the subgroup G_A generated by \mathcal{R}_A is a subgroup of W_A , and the subgroup G_B generated by \mathcal{R}_B is a subgroup of W_B . As G has finite index in W_Γ , it must be that G_A and G_B are finite-index subgroups respectively of W_A and W_B . □

We also need the proposition below, which follows from known results:

Proposition 12.7. *Let W be a RACG and let $G < W$ be a subgroup generated by a set \mathcal{R} of reflections. Then G is a RACG. Furthermore, if \mathcal{R} is trimmed, then it is a standard Coxeter generating set.*

Proof. Deodhar and Dyer independently proved that reflection subgroups of Coxeter groups are Coxeter groups [22, 23]. This, combined with Proposition 2.1 implies that G is a RACG. The second claim follows from the main theorem of [23] and [23, Proposition 3.5]. We also refer the reader to [23, p. 69] for an algorithm to determine a standard generating set for a reflection subgroup of a Coxeter group. \square

We are now ready to prove the main theorem of the section.

Theorem 12.8. *There is an algorithm which, given a 2-dimensional RACG W_Γ and a RACG $W_{\Gamma'}$ such that Γ' does not have an isolated vertex, determines whether or not $W_{\Gamma'}$ is isomorphic to a finite-index subgroup of W_Γ . The algorithm takes as input the graphs Γ and Γ' , and the time-complexity of this algorithm only depends on the number of vertices of Γ and of Γ' . Furthermore, if $W_{\Gamma'}$ is isomorphic to a finite-index subgroup of W_Γ , then the algorithm outputs an explicit set of words in W_Γ which generate this subgroup.*

Proof. First observe that if there exists some $G < W_\Gamma$ isomorphic to $W_{\Gamma'}$, then by Theorem 11.4 and Lemma 10.2, G is generated by a trimmed set of reflections \mathcal{R} and $|\mathcal{R}| = |V(\Gamma')|$.

We prove the theorem by analyzing a few different cases depending on the structure of Γ .

(i) Γ is not almost star. Let I be a trimmed set of reduced reflections in W_Γ . We say I is M -admissible if $|I| = |V(\Gamma')|$ and $|w| \leq M$ for every reflection $ws w^{-1} \in I$. Let \mathcal{I}_M be the collection of all M -admissible trimmed sets of reflections. Note that there is a bound on $|\mathcal{I}_M|$ depending only on M and $|V(\Gamma')|$.

Suppose G is a finite-index subgroup of W_Γ which is isomorphic to $W_{\Gamma'}$, and let \mathcal{R} be a trimmed generating set for G as described above. As Γ is not almost star, Proposition 12.1 guarantees that $\mathcal{R} \in \mathcal{I}_M$, where M is as in Proposition 12.1, and depends only on $|V(\Gamma)|$ and $|V(\Gamma')| = |\mathcal{R}|$. It follows that there exists a finite-index subgroup of W_Γ isomorphic to $W_{\Gamma'}$ if and only if some $I \in \mathcal{I}_M$ generates a finite-index subgroup that is isomorphic to $W_{\Gamma'}$.

Thus, to prove the theorem we only need to show that there is an algorithm to decide whether a given $I \in \mathcal{I}_M$ generates a finite-index subgroup G isomorphic to $W_{\Gamma'}$. By Theorem 10.8 and Theorem 6.6 there is an algorithm to decide whether or not such a G is a finite-index subgroup, and the time-complexity of this algorithm only depends on $|V(\Gamma')|$ and $|V(\Gamma)|$. By Proposition 12.7, G is a RACG and I is a standard Coxeter generating set. As any RACG is defined by a unique graph [34], it is straightforward to check whether I generates a RACG isomorphic to $W_{\Gamma'}$. The theorem then follows in this case.

(ii) $|V(\Gamma)| \leq 2$. In this case W_Γ is isomorphic to \mathbb{Z}_2 , $\mathbb{Z}_2 \times \mathbb{Z}_2$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. Each of these groups contains finitely many finite-index RACGs, up to isomorphism, and one can easily list them.

(iii) $\Gamma = A \star B$ where $|V(A)|, |V(B)| \geq 2$. Suppose G is a subgroup of W_Γ which is isomorphic to $W_{\Gamma'}$. Let \mathcal{R} be a trimmed generating set of reflections for G as above.

If G is a finite-index subgroup, then Lemma 12.6 tells us that $\mathcal{R} = \mathcal{R}_A \cup \mathcal{R}_B$ where every reflection in \mathcal{R}_A only contains generators in A and every reflection in \mathcal{R}_B only contains generators in B . Furthermore, \mathcal{R}_A generates a finite-index subgroup of W_A and \mathcal{R}_B generates a finite-index subgroup of W_B . As $|V(A)|, |V(B)| \geq 2$, both \mathcal{R}_A and \mathcal{R}_B are non-empty.

Let $\Delta = \Delta_A \star \Delta_B$ be the triangle-free join graph such that vertices of Δ_A (resp. Δ_B) correspond to elements of \mathcal{R}_A (resp. \mathcal{R}_B). It readily follows that $W_{\Gamma'}$ is isomorphic to W_Δ .

Thus, we can assume that $\Gamma' = A' \star B'$, a join graph. To prove the claim, again by Lemma 12.6, it is enough to check whether $W_{A'}$ is isomorphic to a finite-index subgroup of W_A and $W_{B'}$ is isomorphic to a finite-index subgroup of W_B . However, as Γ is triangle-free, A does not contain any edges. It follows that either A is not almost star or A consists of at most two isolated vertices. The same holds for B . Thus, we are done by cases (i) and (ii).

(iv) Γ is not as in (i), (ii) or (iii). As Γ is not as in (i), we may assume that Γ is almost star. Suppose first that $V(\Gamma) \subsetneq \text{star}(v)$ for all $v \in V(\Gamma)$. Let s and t be vertices of Γ such that $V(\Gamma) = \text{star}(s) \cup \{t\}$. Note that s and t must be distinct. As Γ is triangle-free, is not a join as in (iii), is not the star of a vertex, and contains more than two vertices, there must be a vertex u of Γ that is adjacent to s and is not adjacent to any other vertex of Γ .

Let G be a subgroup of W_Γ generated by reflections. Let $\Delta = D(\Gamma, u)$ and $\phi = \phi_u: W_\Gamma \rightarrow \mathbb{Z}_2$ be as in Definition 2.2, and let $K = \ker(\phi)$. Note that Δ is not almost star and that K is generated by reflections.

Let $i_G: G \rightarrow W_\Gamma$ be the inclusion map. Let K' be the kernel of the map $\phi' = \phi \circ i_G$. We get the diagram below where all maps labeled by i are the obvious inclusion homomorphisms.

$$\begin{array}{ccccc}
 K' = \ker(\phi') & \xrightarrow{\quad i \quad} & G & & \\
 \downarrow i & & \downarrow i_G & \searrow \phi' = \phi \circ i_G & \\
 K = \ker(\phi) \cong W_\Delta & \xrightarrow{\quad i \quad} & W_\Gamma & \xrightarrow{\quad \phi \quad} & \mathbb{Z}_2
 \end{array}$$

Recall that given a triple of groups $G_3 < G_2 < G_1$, their indices satisfy the formula

$$[G_1 : G_3] = [G_1 : G_2][G_2 : G_3],$$

where infinite values are interpreted appropriately. If we apply this formula to the

groups in the diagram and note that

$$[G : K'] = [W_\Gamma : K] = 2,$$

we get

$$2[W_\Gamma : G] = [W_\Gamma : G][G : K'] = [W_\Gamma : K'] = [W_\Gamma : K][K : K'] = 2[K : K'].$$

Thus, G is a finite-index subgroup of W_Γ if and only if K' is a finite-index subgroup of K . As Δ is not almost star, it follows by (i) that there is an algorithm to check whether K' is a finite-index subgroup of K . The theorem now follows.

On the other hand, if $V(\Gamma) = \text{star}(s)$ for some $s \in V(\Gamma)$, then we apply the same argument as before but instead take the homomorphism $\phi = \phi_s$. In this case, $\Delta = D(\Gamma, s)$ must be a non-empty graph with no edges. Thus, by either (i) or (ii), there is an algorithm to check whether K' is a finite-index subgroup of K , where K' and K are as before. □

13. Other algorithmic properties of quasiconvex subgroups

This section is dedicated to the proof of Theorem E of the introduction. Throughout this section, we let G be a quasiconvex subgroup of W_Γ given by a finite generating set, and we let Ω be the corresponding completion. By Theorem 8.4, Ω is finite and by Proposition 3.5 it can be computed in finite time.

Note that (1) and (2) of Theorem E immediately follow by Proposition 4.6, and Theorem 6.6 respectively. Before proving (3), we show that powers of an element of a RACG can be represented by words of a special form.

Lemma 13.1. *Let w be a reduced word in the RACG W_Γ . Then there exist reduced words x, h and k , such that $xhkx^{-1}$ is a reduced expression for w and $xh^nk^{(n \bmod 2)}x^{-1}$ is a reduced expression for w^n for all integers $n > 0$.*

Proof. Write $w = xyx^{-1}$ where x and y are reduced words and $|x|$ is maximal out of all such possible expressions. Let $K = \{k_1, \dots, k_n\}$ be the set of vertices in Γ that appear as letters in the word y and which commute with every other letter of y . As w is reduced, each element of K appears as a letter of w exactly once. Define the word $k = k_1 \dots k_n$. By our choice of k , it follows that there exists a reduced word h , such that hk is a reduced expression for y . Note that h has the property that any generator which appears as the last letter of some reduced expression for h , cannot also appear as the first letter in some reduced expression for h . This follows since otherwise either x is not maximal or such a generator should have been in K . The word $xhkx^{-1}$ will be the desired expression for w .

The word $xh^nk^{(n \bmod 2)}x^{-1}$ is clearly an expression for the word w^n . Furthermore, the word $xh^nk^{(n \bmod 2)}x^{-1}$ must be reduced. For otherwise, it follows by the deletion

property that either h is not reduced or some generator appears as the first letter in some reduced expression for h and as the last letter in some reduced expression for h , which is not possible. \square

Proof of Theorem E (3). Without loss of generality, we may assume that w is reduced. Let N be the number of vertices of Ω . We claim that $g^m \in G$ for some positive integer m if and only if w^{2l} represents an element of G for some $l \leq N$. The theorem clearly follows from this claim and the fact that the membership problem is solvable for G .

One direction of the claim is clear. On the other hand, suppose $g^m \in G$ for some positive integer m . By possibly taking a power, we may assume that m is even. By Lemma 13.1, there is a reduced expression for w^m of the form $z = xh^m x^{-1}$. Let β be a loop in Ω based at B with label z . For $1 \leq i \leq m$, let α_i be the first subpath of β with label h^i , and let v_i be the endpoint of α_i . As Ω is folded, the two subpaths of β labeled x are identified. It follows that α_m is a loop based at v_m . As there are at most N vertices of Ω , it follows that the set $\{v_1, \dots, v_m\}$ contains at most N distinct vertices. There must then exist some loop α based at v_m with label h^l for some $l \leq N$. Thus, if we replace α_m with $\alpha\alpha$ in β , we conclude that there is a loop in Ω based at B with label $xh^{2l}x^{-1}$. By the definition of a completion, the word $xh^{2l}x^{-1}$, which is an expression for w^{2l} , represents an element of G . \square

Although Theorem 5.3 characterizes normality, the core of a completion may be difficult to algorithmically compute in general. Thus, we now give a different characterization of normality, for finitely generated subgroups, which is better suited to the algorithmic approach.

Proposition 13.2. *Let $G < W_\Gamma$ be generated by a finite set of reduced words S_G , and let (Ω, B) be a standard completion for G with respect to S_G . Consider $\Delta \subset V(\Gamma)$ defined by:*

$$\Delta = \{s \in V(\Gamma) \mid s \text{ commutes with every element of } G\}.$$

Then G is normal if and only if the following hold:

- (N1) *Given any $s \in V(\Gamma) \setminus \Delta$, there is an edge in Ω incident to B with label s .*
- (N2') *For every generator $w \in S_G$ of G , and for every vertex v of Ω , there exists a loop based at v with label w .*

Proof. If G is normal, then Theorem 5.3 implies (N1) and (N2'). On the other hand, suppose (N1) and (N2') hold. As in the proof of Theorem 5.3, to show that G is normal it is enough to show $sGs \subseteq G$ for $s \in V(\Gamma) \setminus \Delta$. Let v be the vertex which is adjacent to B via an edge labeled s , which exists due to (N1). Then by Lemma 4.5, the subgroup of W_Γ associated to (Ω, v) is sGs , and by (N2'), G is contained in this subgroup. Conjugating by s , it follows that $sGs \subseteq G$. \square

Proof of Theorem E(4). Let S_G be a finite set of reduced words in W_Γ which generate G . In this case, the set Δ from Proposition 13.2 is equal to the following set:

$$\{s \in V(\Gamma) \mid \forall w \in S_G, s \text{ commutes with every letter in the support of } w\}.$$

Thus, Δ can be computed in finite time. It follows that conditions (N1) and (N2') from Proposition 13.2 can be checked in finite time as well. \square

References

- [1] I. Agol, The virtual Haken conjecture. *Doc. Math.* **18** (2013), 1045–1087 [Zbl 1286.57019](#) [MR 3104553](#)
- [2] G. N. Arzhantseva, On groups in which subgroups with a fixed number of generators are free. *Fundam. Prikl. Mat.* **3** (1997), no. 3, 675–683 [Zbl 0929.20025](#) [MR 1794135](#)
- [3] G. N. Arzhantseva, Generic properties of finitely presented groups. PhD thesis, Moscow Lomonosov State University, 1998.
- [4] G. N. Arzhantseva, Generic properties of finitely presented groups and Howson's theorem. *Comm. Algebra* **26** (1998), no. 11, 3783–3792 [Zbl 0911.20027](#) [MR 1647075](#)
- [5] G. N. Arzhantseva, A property of subgroups of infinite index in a free group. *Proc. Amer. Math. Soc.* **128** (2000), no. 11, 3205–3210 [Zbl 0976.20014](#) [MR 1694447](#)
- [6] G. N. Arzhantseva and P.-A. Cherix, On the Cayley graph of a generic finitely presented group. *Bull. Belg. Math. Soc. Simon Stevin* **11** (2004), no. 4, 589–601 [Zbl 1069.05038](#) [MR 2115727](#)
- [7] G. N. Arzhantseva and A. Y. Ol' shanskiĭ, Generality of the class of groups in which subgroups with a lesser number of generators are free. *Mat. Zametki* **59** (1996), no. 4, 489–496, 638 [Zbl 0877.20021](#) [MR 1445193](#)
- [8] P. Bahls, *The isomorphism problem in Coxeter groups*. Imperial College Press, London, 2005 [Zbl 1113.20034](#) [MR 2162146](#)
- [9] B. Beeker and N. Lazarovich, Stallings' folds for cube complexes. *Israel J. Math.* **227** (2018), no. 1, 331–363 [Zbl 06976580](#) [MR 3846326](#)
- [10] A. Björner and F. Brenti, *Combinatorics of Coxeter groups*. Grad. Texts in Math. 231, Springer, New York, 2005 [Zbl 1110.05001](#) [MR 2133266](#)
- [11] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999 [Zbl 0988.53001](#) [MR 1744486](#)
- [12] B. Brink and R. B. Howlett, A finiteness property and an automatic structure for Coxeter groups. *Math. Ann.* **296** (1993), no. 1, 179–190 [Zbl 0793.20036](#) [MR 1213378](#)
- [13] S. Brown, Geometric structures on negatively curved groups and their subgroups. PhD thesis, University College London, 2016.
- [14] P.-E. Caprace and M. Sageev, Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21** (2011), no. 4, 851–891 [Zbl 1266.20054](#) [MR 2827012](#)

- [15] M. Casals-Ruiz, Embeddability and universal theory of partially commutative groups. *Int. Math. Res. Not. IMRN* (2015), no. 24, 13575–13622 [Zbl 1367.20032](#) [MR 3436157](#)
- [16] J. Crisp and L. Paoluzzi, Commensurability classification of a family of right-angled Coxeter groups. *Proc. Amer. Math. Soc.* **136** (2008), no. 7, 2343–2349 [Zbl 1194.20042](#) [MR 2390500](#)
- [17] P. Dani, The large-scale geometry of right-angled Coxeter groups. In *Handbook of group actions. V*, pp. 107–141, Adv. Lect. Math. (ALM) 48, Int. Press, Somerville, MA, 2020 [Zbl 7274139](#) [MR 4237891](#)
- [18] P. Dani and I. Levcovitz, Right-angled Artin subgroups of right-angled Coxeter and Artin groups, 2020. [arXiv:2003.05531](#)
- [19] P. Dani, E. Stark, and A. Thomas, Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups. *Groups Geom. Dyn.* **12** (2018), no. 4, 1273–1341 [Zbl 1456.20039](#) [MR 3874643](#)
- [20] M. W. Davis, *The geometry and topology of Coxeter groups*. London Math. Soc. Monogr. Ser. 32, Princeton Univ. Press, Princeton, NJ, 2008 [Zbl 1142.20020](#) [MR 2360474](#)
- [21] M. W. Davis and T. Januszkiewicz, Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra* **153** (2000), no. 3, 229–235 [Zbl 0982.20022](#) [MR 1783167](#)
- [22] V. V. Deodhar, A note on subgroups generated by reflections in Coxeter groups. *Arch. Math. (Basel)* **53** (1989), no. 6, 543–546 [Zbl 0688.20028](#) [MR 1023969](#)
- [23] M. Dyer, Reflection subgroups of Coxeter systems. *J. Algebra* **135** (1990), no. 1, 57–73 [Zbl 0712.20026](#) [MR 1076077](#)
- [24] D. Gruber, Infinitely presented graphical small cancellation groups. PhD thesis, University of Vienna, 2015. Available at: <http://othes.univie.ac.at/38520>
- [25] F. Haglund, Finite index subgroups of graph products. *Geom. Dedicata* **135** (2008), 167–209 [Zbl 1195.20047](#) [MR 2413337](#)
- [26] F. Haglund and D. T. Wise, Special cube complexes. *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620 [Zbl 1155.53025](#) [MR 2377497](#)
- [27] F. Haglund and D. T. Wise, Coxeter groups are virtually special. *Adv. Math.* **224** (2010), no. 5, 1890–1903 [Zbl 1195.53055](#) [MR 2646113](#)
- [28] G. C. Hruska, E. Stark, and H. C. Tran, Surface group amalgams that (don't) act on 3-manifolds. *Amer. J. Math.* **142** (2020), no. 3, 885–921 [Zbl 07229244](#) [MR 4101334](#)
- [29] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups. *J. Algebra* **248** (2002), no. 2, 608–668 [Zbl 1001.20015](#) [MR 1882114](#)
- [30] O. Kharlampovich, A. Miasnikov, and P. Weil, Stallings graphs for quasi-convex subgroups. *J. Algebra* **488** (2017), 442–483 [Zbl 06768984](#) [MR 3680926](#)
- [31] S.-h. Kim and T. Koberda, Embedability between right-angled Artin groups. *Geom. Topol.* **17** (2013), no. 1, 493–530 [Zbl 1278.20049](#) [MR 3039768](#)
- [32] J. P. McCammond and D. T. Wise, Coherence, local quasiconvexity, and the perimeter of 2-complexes. *Geom. Funct. Anal.* **15** (2005), no. 4, 859–927 [Zbl 1087.20024](#) [MR 2221153](#)
- [33] M. Mihalik and S. Tschantz, Visual decompositions of Coxeter groups. *Groups Geom. Dyn.* **3** (2009), no. 1, 173–198 [Zbl 1196.20046](#) [MR 2466022](#)

- [34] D. G. Radcliffe, Rigidity of graph products of groups. *Algebr. Geom. Topol.* **3** (2003), 1079–1088 [Zbl 1053.20025](#) [MR 2012965](#)
- [35] M. Sageev and D. T. Wise, Cores for quasiconvex actions. *Proc. Amer. Math. Soc.* **143** (2015), no. 7, 2731–2741 [Zbl 1368.20059](#) [MR 3336599](#)
- [36] P. E. Schupp, Coxeter groups, 2-completion, perimeter reduction and subgroup separability. *Geom. Dedicata* **96** (2003), 179–198 [Zbl 1069.20028](#) [MR 1956839](#)
- [37] J. R. Stallings, Topology of finite graphs. *Invent. Math.* **71** (1983), no. 3, 551–565 [Zbl 0521.20013](#) [MR 695906](#)
- [38] M. Steenbock, Rips-Segev torsion-free groups without the unique product property. *J. Algebra* **438** (2015), 337–378 [Zbl 1402.20046](#) [MR 3353035](#)
- [39] B. A. F. Wehrfritz, *Infinite linear groups*. Queen Mary College Mathematical Notes, Queen Mary College, Department of Pure Mathematics, London, 1969 [MR 837425](#)
- [40] D. T. Wise, The structure of groups with a quasiconvex hierarchy, 2011. Available at: <https://www.math.mcgill.ca/wise/papers.html>
- [41] D. T. Wise, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*. CBMS Reg. Conf. Ser. Math. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012 [Zbl 1278.20055](#) [MR 2986461](#)

Received 18 June, 2020

P. Dani, Department of Mathematics, Louisiana State University,
Baton Rouge, LA 70803-4918, USA

E-mail: pdani@math.lsu.edu

I. Levcovitz, Department of Mathematics, Technion – Israel Institute of Technology,
Haifa 32000, Israel

E-mail: ivan.levcovitz@tufts.edu