A STUDY OF SUBGROUPS OF RIGHT-ANGLED COXETER GROUPS VIA STALLINGS-LIKE TECHNIQUES

PALLAVI DANI AND IVAN LEVCOVITZ

Abstract. We associate a cube complex to any given finitely generated subgroup of a right-angled Coxeter group, called the completion of the subgroup. 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 right-angled Coxeter group. We provide an algorithm that determines whether a given one-ended, 2-dimensional right-angled Coxeter group is isomorphic to some finite-index subgroup of another given right-angled Coxeter group. In addition, we answer several algorithmic questions regarding quasiconvex subgroups. Finally, we give a new proof of Haglund's result that quasiconvex subgroups of right-angled Coxeter groups are separable.

1. Introduction

In the highly influential article [Sta83], 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 [KM02], Kapovich–Miasnikov use Stallings' ideas, re-cast in a more combinatorial form, to systematically study the subgroup structure of free groups, and to answer a number of algorithmic questions about such subgroups. The main goal of this article is to better understand subgroups of right-angled Coxeter groups through generalizations of these authors' techniques.

Given a finite simplicial graph Γ, the associated right-angled Coxeter group 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. Right-angled Coxeter groups form a wide class of groups which have become central objects in geometric group theory. We refer to [Dan18] for a survey of recent work on these groups. One interesting feature of right-angled Coxeter groups is that they have a rich variety of subgroups, which includes all free groups, right-angled Artin groups [DJ00] and surface groups. Incredibly, all hyperbolic 3-manifold groups [Ago13] [Wis11] and Coxeter groups [HW10] are virtually subgroups of right-angled Coxeter groups as well.

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

Given a subgroup G of a right-angled Coxeter group W_Γ, we abstractly define edge-labeled cube complexes associated to G, called completions of G. If G is additionally finitely generated, we explicitly build a *standard* completion for G by the following procedure. We first build a subdivided "rose" graph whose petals are labeled by the words generating G. Next, we perform a sequence of operations of three possible types: fold, cube attachment, and cube identification. A completion can always be obtained as the direct limit of the complexes in this sequence.

As a completion Ω has the structure of a cube complex (and not just a graph), one may use techniques from cubical geometry to study it. Many properties of the subgroup G can be characterized in terms of properties of Ω. Additionally, loops in Ω correspond to elements in G. The following theorem summarizes our results relating properties of G to properties of Ω.

Theorem A. *Let G be a subgroup of the right-angled Coxeter group 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 6.9), G being torsion-free (Proposition 4.6), and G being normal (Theorem 5.5) in terms of properties of a completion.*
(3) *If G is torsion-free, then any completion is non-positively curved (Proposition 7.3) and has fundamental group isomorphic to G (Theorem 4.7).*

We remark that characterizations similar to those above could also be obtained by considering the covering spaces perspective and applying results of Haglund [Hag08]. Theorem A also has some overlap with the work of Kharlampovich–Miasnikov–Weil [KMW17], who construct "Stallings graphs" for automatic groups. We further discuss the relation of our results to these and other works in Section 1.1. Our proof strategies for Theorem A seem to be novel and are independent of the results mentioned above. Additionally, our procedure for constructing Ω is new, and this is what allows us to prove structural results regarding subgroups of right-angled Coxeter groups. Furthermore, the approach using completions is particularly well suited for addressing algorithmic questions.

A given subgroup may have multiple completions. Indeed two completions for a given subgroup need not even be homotopy equivalent (see Example 3.7). Nevertheless, most of our structural results do not depend on the specific completions chosen. 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 5.3).

Theorem A provides a tool to show that subgroups of a right-angled Coxeter group are quasiconvex, by showing that their associated completions must be finite. For instance, we prove that finitely generated reflection subgroups, i.e. subgroups generated by reflections, must be quasiconvex:

Theorem B (Theorem 10.5)**.** *Every finitely generated reflection subgroup of a right-angled Coxeter group is quasiconvex.*

We next turn our attention to Coxeter subgroups, i.e. subgroups that are themselves isomorphic to some abstract finitely generated Coxeter group. A result of [Dye90] and [Deo89] shows that every reflection subgroup of a right-angled Coxeter group is a Coxeter subgroup. The converse to this statement is not true in general (see Remark 11.5.1), but it holds under certain hypotheses:

Theorem C (Theorem 11.4, Corollary 11.5). *Every one-ended Coxeter subgroup of a 2-dimensional right-angled Coxeter group 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 right-angled Coxeter groups. For instance, we consider the problem of determining, given two right-angled Coxeter groups, whether one can be embedded as a finite-index subgroup of the other:

Theorem D (Theorem 12.11). *There is an algorithm which, given a one-ended, 2-dimensional right-angled Coxeter group W_Γ, and any right-angled Coxeter group $W_{\Gamma'}$, determines whether or not $W_{\Gamma'}$ is isomorphic to a finite-index subgroup of W_Γ. Moreover, the time-complexity of the algorithm is bounded by a function of the number of vertices of Γ and Γ'.*

The above theorem gives an algorithm that can often determine when two right-angled Coxeter groups are commensurable; thus, it provides a tool for studying commensurability classification. A few specific families of right-angled Coxeter groups have been classified up to commensurability (see [CP08, DST18, HST17]), but not much is known in general. We note that the precise statements of Theorem C and Theorem D use a significantly weaker hypothesis than one-endedness.

When G is quasiconvex, Theorem A implies that every completion of G is finite. This makes it possible to provide finite-time algorithms to check a number of basic properties concerning G, as described in Theorem E below. We remark that the existence of algorithms for (1) and (4) was already known [Kap96, KMW17] (see the discussion in Section 1.1).

Theorem E (Theorem 13.1). *Let G be a quasiconvex subgroup of a right-angled Coxeter group W_Γ given by a finite generating set of words in W_Γ. Then there exist finite-time algorithms to solve the following problems.*

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

In particular, the above theorem may be applied to reflection subgroups as these are quasiconvex by Theorem B. If the ambient right-angled Coxeter group W_Γ is additionally 2-dimensional, then the time-complexity of these algorithms is bounded only by the length of generators of the reflection subgroup and the number of vertices of Γ (see Theorem 10.7).

As another application, we give new proofs for some known results. We give a proof using completions of the following result of Haglund (see Theorem 9.4):

Theorem F ([Hag08, Theorem A]). *Every quasiconvex subgroup of a right-angled Coxeter group is separable and is a virtual retract.*

Theorem F can be thought of as a generalization of Marshall Hall's Theorem from the free group setting to the setting of right-angled Coxeter groups. Also following Stallings' approach, we give a simple proof, using completions, of the well-known result that right-angled Coxeter groups are residually finite (see Theorem 9.3).

Finally, we note that much of the work presented here can be used to study right-angled Artin groups as well. Given any right-angled Artin group A, Davis–Januszkiewicz construct a right-angled Coxeter group W which contains A as a finite-index subgroup [DJ00]. Thus, given a subgroup G of a right-angled Artin group A, one can first embed A as a finite-index subgroup of W and construct a completion for G considered as a subgroup of W. Now Theorem A and Theorem E can be used to understand properties of G as a subgroup of A.

1.1. **Relation to other works.** The ideas in Stallings' paper have led to many applications which are too numerous to list here. We discuss some results in the literature that are more closely related to ours.

Given a quasiconvex subgroup G of a right-angled Coxeter group, Haglund shows that G acts cocompactly on the combinatorial convex hull $\Sigma(G)$ of G in the Davis complex of W_Γ [Hag08]. It turns out that the completion Ω is very close to being equal to the quotient of $\Sigma(G)/G$ (one has to take care, as this quotient is an orbifold in general and the generating set for G used in constructing Ω may not consist of reduced words). If one further developed this point of view, characterizations similar to those in Theorem A could be obtained from results in Haglund's article.

Beeker–Lazarovich use a version of Stallings folds for cube complexes to give a characterization of quasiconvex subgroups of hyperbolic groups that act properly and cocompactly on CAT(0) cube complexes in terms of their hyperplane stabilizers [BL18]. Similar ideas were also used in Brown's thesis in the setting of hyperbolic VH square complexes [Bro]. Sageev–Wise show that relatively quasiconvex subgroups of relatively hyperbolic groups which act on a finite-dimensional, locally finite CAT(0) cube complex admit a convex core [SW15]. This generalizes ideas of Haglund [Hag08]. We note that the results mentioned in this paragraph use some form of hyperbolicity, wheras, in our case, many right-angled Coxeter groups are not even relatively hyperbolic (see [BHS17]).

Kharlampovich–Miasnikov–Weil [KMW17] show that Stallings graphs (we refer the reader to their paper for a definition) can always be constructed for quasiconvex subgroups of automatic groups. They use this to address several algorithmic problems for quasiconvex subgroups of automatic groups, including the membership problem, deciding whether the subgroup is finite-index or finite, and computing intersections of subgroups. Since right-angled Coxeter groups are automatic [BH93], these results apply, thus providing algorithms for (1) and (4) in our Theorem E. Kapovich had already previously shown that the membership problem is solvable for automatic groups [Kap96]. Schupp uses Stallings graphs to show that certain classes of extra-large type Coxeter groups are locally quasiconvex and to answer algorithmic questions regarding their subgroups [Sch03]. We note that right-angled Coxeter groups are not of extra-large type. We refer the reader to [KMW17] for a detailed summary on the literature regarding constructing Stallings graphs for solving algorithmic problems.

In the spirit of Theorem D, Kim–Koberda consider the question of understanding when one right-angled Artin group can be realized as a (not necessarily finite-index) subgroup of another [KK13]. They relate the existence of such embeddings to properties of certain associated graphs, called extension graphs and clique graphs. Using their work, Casals-Ruiz proves a number of algorithmic results about embeddings between right-angled Artin groups [CR15]. In particular, she shows that there is an algorithm which, given a 2-dimensional right-angled Artin group A_Γ and any right-angled Artin group $A_{\Gamma'}$, determines whether or not $A_{\Gamma'}$ is isomorphic to a subgroup of A_Γ.

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

2. Preliminaries

Given a graph Γ, we will always 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 groups which we will need throughout this article. We refer the reader to [Dan18] for a survey on right-angled Coxeter groups and to both [Dav08] and [BB05] 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 *right-angled Coxeter group* 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 $\mathrm{Support}(w)$, 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 right-angled Coxeter group. 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 [Dav08]). Such a subgroup of W_Γ is called a *special subgroup*.

Given a vertex v of Γ, the *link of* v, denoted by $\mathrm{link}(v)$, is the set of all vertices of Γ which are adjacent to v. The *star of* v, denoted by $\mathrm{star}(v)$, is the set $\mathrm{link}(v) \cup \{v\}$. Many times throughout this article, we will consider the special subgroup of W_Γ generated by the link or star of a vertex.

We now recall some classes of graphs and their corresponding right-angled Coxeter groups. Recall that a graph is a *clique* if any pair of distinct vertices of the graph are adjacent. A right-angled Coxeter group W_Γ is finite if and only if Γ is a clique.

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 right-angled Coxeter group 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 if $v_1 \in V(\Gamma_1)$ and $v_2 \in V(\Gamma_2)$, then v_1 and v_2 are adjacent in Γ. 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 describe a characterization of the number of ends of a right-angled Coxeter group in terms of the defining graph. These results were first proven in [MT09]. The group W_Γ is 2-ended if and only if either Γ consists of two non-adjacent vertices or there exists two non-adjacent vertices $s, t \in V(\Gamma)$ and a clique subgraph $K \subset \Gamma$ such that $\Gamma = K \star \{s, t\}$. Now suppose that W_Γ is not finite or 2-ended (each a case we already described in terms of Γ), then W_Γ has infinitely many ends if and only if Γ has a separating clique, i.e. there is some clique subraph $K \subset \Gamma$ such that $\Gamma \setminus K$ has more than one component.

Finally, we say that a subgroup of a Coxeter group is a *Coxeter subgroup* if it is isomorphic to some finitely generated Coxeter group. We refer the reader to [BB05] for a definition of Coxeter groups. We will actually not need this definition, as we show below (the likely well-known fact) that every Coxeter subgroup of a right-angled Coxeter group is itself a right-angled Coxeter group.

Proposition 2.1. *Let G be a Coxeter subgroup of a right-angled Coxeter group, then G is a right-angled Coxeter group.*

Proof. As any finite subgroup of a Coxeter group is conjugate to a subgroup of a finite special subgroup (see [Dav08, Theorem 12.3.4] for instance), it follows that all finite order elements in a right-angled Coxeter group have order 2.

We note that Coxeter groups which are not right-angled must contain elements of order larger than 2. This is the case since such a group must necessarily contain generators s and t which satisfy the relation $(st)^m$ for some $m > 2$ and do not satisfy the relation $(st)^k$ for all $0 < k < m$ [BB05, Proposition 1.1.1]. Thus, G must indeed be a right-angled Coxeter group. $\square$

2.2. **The word problem in right-angled Coxeter groups.** We discuss Tits' solution to the word problem in right-angled Coxeter groups. We again refer the reader to [Dav08] or [BB05] for proofs of these facts.

Let $w = s_1 \dots s_n$ be a word in the right-angled Coxeter group W_Γ. Suppose that $m(s_i, s_{i+1}) = 2$ for some $1 \le i \le n$. 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, suppose that $s_i = s_{i+1}$ (as vertices of Γ) for some $1 \le i \le n$. We can then 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*.

We now discuss an alternative way to obtain a reduced expression for a word $w = s_1 \dots s_n$ in a right-angled Coxeter group W_Γ. Suppose $s_i = s_{i'} = s$ for some $1 \le i < i' \le n$ and that $m(s, s_j) = 2$ for all $i < j < i'$. It follows that the word

$w' = s_1 \dots s_{i-1}s_{i+1} \dots s_{i'-1}s_{i'+1} \dots s_n$ is an expression for w. We say that w' is obtained from w by a *deletion* (as the two occurrences of s have been deleted).

We remark that our definition of a deletion, defined in the setting of right-angled Coxeter groups, is stated in a slightly stronger form than its classical statement for (not necessarily right-angled) Coxeter groups. It is a basic fact that given any word w in a right-angled Coxeter group, a reduced expression for w can be obtained by performing a sequence of deletions. As we could not find a statement of this in the literature using our exact definition of deletion, we later provide a proof of this fact (see Proposition 2.2).

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 [CS11] and to [Wis12] 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. The cube complexes we consider in this article will all be Γ-labeled. 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.

Let Ω be a cube complex. We say Ω is *non-positively curved* if the (simplicial) link of each vertex in Ω is a flag simplicial complex. If Ω is both non-positively curved and simply connected then we say that 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 $1 \leq i < n$. We remark that for the cube complexes considered in this article, it will be the case that hyperplanes do not self-intersect. Thus, it can be checked that in this setting a hyperplane can also be defined to be a connected collection of midcubes in Ω, such that for any cube c in Ω either $H \cap c$ is a midcube or $H \cap c = \emptyset$.

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 and H be a hyperplane in Ω. We will need the following well-known facts (see for instance [Wis12, Chapters 3.2 and 3.3]):

(1) $\Omega \setminus H$ contains exactly two components
(2) $N(H)$ is convex in the combinatorial path metric
(3) 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 which intersect γ.

2.4. **Disk diagrams in cube complexes.** We now recall some basic facts about disk diagrams, which are a useful tool for studying cube complexes. We refer to [Wis11] and [Wis12] for further details.

A *disk diagram* D is a contractible, finite, 2-dimensional cube complex (i.e., a square complex) that is equipped with a given planar embedding $\Psi : D \to \mathbb{R}^2$. The map Ψ gives a natural cellulation of the 2-sphere $S^2 = \mathbb{R}^2 \cup \infty$. We call the path traced out 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 in* Ω is a disk diagram D which admits a map to $\Phi : D \to \Omega$ 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 cube complex Ω and a closed null-homotopic loop $p : \mathbb{S}^1 \to \Omega$, by a lemma of van Kampen, there always exists a disk diagram D in Ω with combinatorial map $\Phi : D \to \Omega$ 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 [Wis12, Lemma 3.1]).

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

The cube complexes we consider in this article will have the additional property that edges are labeled by elements of a simplicial graph Γ. Furthermore, any square in such a cube complex will have opposite sides labeled by the same vertex of Γ and adjacent sides labeled by distinct adjacent vertices of Γ. Thus, given a disk diagram in such a cube complex we can define the *type* of a dual curve to be the label of an edge (equivalently, all edges) dual to the dual curve. Furthermore, if two dual curves intersect, then their types must be adjacent vertices of Γ.

Let W_Γ be a right-angled Coxeter group. There exists a CAT(0) cube complex Σ_Γ associated to W_Γ, known as the Davis complex, whose 1-skeleton is the Cayley graph of W_Γ and whose edges are naturally labeled by vertices of Γ (see for instance [Dav08]). Thus, given any word w in W_Γ which is an expression for the identity element, it follows that there exists a disk diagram D that has edges labeled by vertices of Γ and whose boundary has label w. Furthermore, D has the properties listed in the previous paragraph. We will often use these facts without mention.

Below we use disk diagrams to give a proof of a version of the so-called deletion property. We remark that a similar proof is used in [Bah05] for proving the standard version of the deletion property in Coxeter groups.

Proposition 2.2. *If $w = s_1 \dots s_n$ is a word in a right-angled Coxeter group W_Γ which is not reduced, then a deletion can be applied to w. In other words, for some $1 \le i < i' \le n$, we have that $s_i = s_{i'} = s$, and $m(s, s_j) = 2$ for all $i < j < i'$. Consequently $w' = s_1 \dots s_{i-1}s_{i+1} \dots s_{i'-1}s_{i'+1} \dots s_n$ is an expression for w.*

Proof. Let r be a reduced expression for w. As rw^{-1} is equal to the identity element in W_Γ, it follows that there is a disk diagram D with boundary rw^{-1}. Let α be the path along the boundary of D with label r, and let β be the path along the boundary of D with label w.

As r is reduced, no dual curve is dual to two edges of α. As w is not reduced, we have that $|w| > |r|$. From these two facts we deduce there must be some dual curve H dual to two distinct edges, say e and e', of β. Furthermore, as e and e' are dual to the same dual curve, these edges are labeled by the same letter $s \in \Gamma$.

Let β' be the subpath of β from e to e'. Without loss of generality, we can assume that no dual curve is dual to two edges of β' (or else we could replace H with such a dual curve). It follows that every dual curve dual to an edge of β' must intersect H. Thus, the label $t_1 \dots t_m$ of β' has the property that $m(s, t_i) = 2$ for all $1 \leq i \leq m$. Hence, the claim follows as w must contain the subword $st_1 \dots t_m s$. Note that if e and e' were actually adjacent edges of D then w has a subword of the form ss and the claim still follows. □

The following lemma, which is required later, easily follows from Proposition 2.2.

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

We will need the following lemma which lets us deduce properties of words from properties of a disk diagram with the same words as labels of its boundary. We remark 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 right-angled Coxeter group W_Γ. Suppose $w = w'w''$ and $z = z'z''$ where w', w'', z', z'' are reduced words in W_Γ. Let D be a disk diagram with boundary label wz^{-1}, and let $\alpha_{w'}$ and $\alpha_{z'}$ be the paths in the boundary of D with labels w' and z' respectively. Suppose further that every dual curve dual to $\alpha_{w'}$ is also dual to $\alpha_{z'}$. Then z' has a reduced expression $z' = w'x$ where x is some word in W_Γ.*

Proof. By [Wis11, Lemma 2.3], there is a minimal area disk diagram D' with the same boundary label and property as in the statement. By the proof of [Wis11, Lemma 2.6], if two dual curves dual to $\alpha_{z'}$ intersect, then there is a smaller area disk diagram D'' with boundary label $w(z'')^{-1}y^{-1}$ where y is a word equal to z' in W_Γ. Additionally, dual curves dual to the path in the boundary of D'' with label w' are also dual to the path in the boundary of D'' labeled by y.

By repeatedly applying this argument, we eventually get a disk diagram $\hat{D}$ with boundary label $wz''^{-1}\hat{x}^{-1}$ where $\hat{x}$ is equal to z' in W_Γ. Furthermore, no pair of curves dual to the path in the boundary of $\hat{D}$ labeled by $\hat{x}$ intersect, and every curve dual to the path in the boundary of $\hat{D}$ labeled by w' is also dual to the path in the boundary of $\hat{D}$ labeled by $\hat{x}$. It follows that w' is a prefix of $\hat{x}$, and the claim follows. □

3. A complex for subgroups of a right-angled Coxeter group

The main goal of this section is to define a completion of a subgroup of a right-angled Coxeter group, 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: 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 end up adding 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 operation or cube identification operation can be performed to the complex. 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 $f_i : \Omega_i \to \Omega_{i+1}$ 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

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

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 $\hat{f} : X \to \Omega_X$ as the direct limit of the maps $\{f_i\}$.

Example 3.1. Let Γ_1 be the graph in Figure 1. The bottom 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.

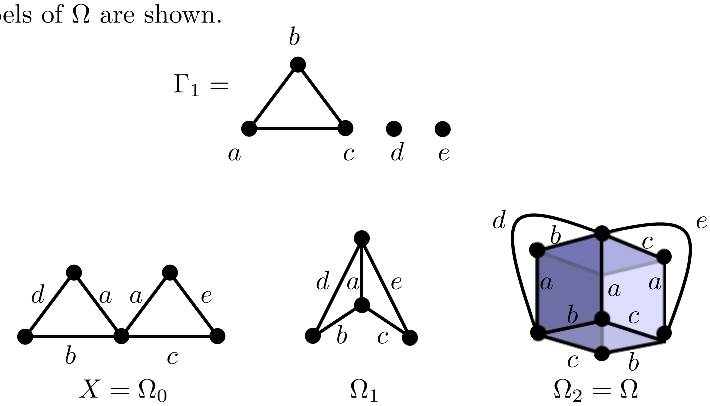

FIGURE 1. A completion Ω for the Γ_1-labeled complex X.

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.

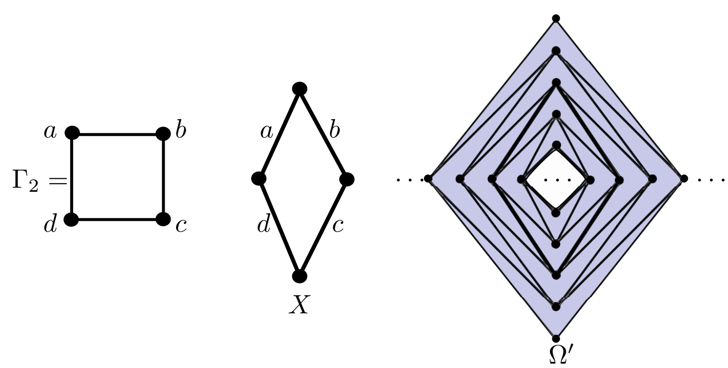

FIGURE 2. A completion Ω' for the Γ_2-labeled complex X.

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 $\Omega_{i+1}, \Omega_{i+2}, ..., \Omega_{i+j}$, 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 $\Omega_{i+j+1}, \Omega_{i+j+2}, ...\Omega_{i+j+k}$, 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 $e_i', \dots, e_n'$ 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. □

Definition 3.4 (Standard Completion)**.** We call the completion algorithm given in the proof of Lemma 3.3 a *standard completion* and the associated sequence:

$$\Omega_0 \to \Omega_1 \to \cdots \to \Omega_X$$

a standard completion sequence. We recall that this is a completion sequence obtained by alternately performing all possible fold and cube identification operations to a complex, then performing all possible cube attachments to the resulting folded complex, and iteratively repeating this procedure whenever possible.

In the next proposition, we show that if a completion is finite, then there is indeed a finite algorithm to obtain the completion.

Proposition 3.5. *Let X be a Γ-labeled complex. Let*

$$X = \Omega_0 \to \Omega_1 \cdots \to \Omega$$

be a standard completion sequence. 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 $f : \Omega_n \to \Omega$ 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 that $\Omega_N = \Omega$. Suppose for a contradiction 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 $f : \Omega_N \to \Omega$ 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 with the same argument, we conclude that $\Omega_N = \Omega$ □

3.2. **Completion of a subgroup.** In this subsection, we define the completion of a subgroup of a right-angled Coxeter group 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 right-angled Coxeter group 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 3.6.1. We sometimes say Ω is a completion of G, dropping the basepoint B from the notation, if the basepoint is not relevant.

Example 3.7. 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 2-cell to X (where we view the 2-cell to be a square). 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 the completion of a certain Γ-labeled graph associated to G, which we now describe. Let G be generated by the finite generating set of words

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

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 resulting based complex by $(X(S_G), B)$ and call it the *S_G-complex.* Note that the label of a circle based at B in $X(S_G)$ always corresponds to a generator in S_G.

Let

$$X(S_G) = \Omega_0 \to \Omega_1 \to \Omega_2 \cdots \to \Omega$$

be a completion of $X(S_G)$. By a slight abuse of notation, we denote by B the vertex $B \in X(S_G)$, its image in Ω_i for any i, and its image in Ω. The next few lemmas show that Ω is a completion of G.

Lemma 3.8. *Let G be a subgroup of a right-angled Coxeter group 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
$$X(S_G) = \Omega_0 \to \Omega_1 \to \Omega_2 \to \cdots \to \Omega$$
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 $w' = h_1...h_k$ 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 $w' = a_1...a_i s t a_{i+1}...a_m$ and $w'' = a_1...a_i t s a_{i+1}...a_m$, 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 equal to w. □

Lemma 3.9. *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 $F : \bar{\Omega} \to \Omega$ 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 right-angled Coxeter group W_Γ.*

Proof. We analyze each type of operation separately:

Cube identification operation: If Ω is obtained from $\bar{\Omega}$ by 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 that Ω is obtained from $\bar{\Omega}$ by a fold operation. 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 every $j < i$.
(4) A path p from $u = u_i$ to $B = u_m$, where u has unique preimage and $u_j = v$ for every $j > i$.
(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 as elements of 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 $\bar{v}_2$.

Let $\bar{z}$ be the path from $\bar{v}_1$ to $\bar{v}_2$ obtained by traversing the edge $\bar{f}_1$ then the edge $\bar{f}_2$. 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.

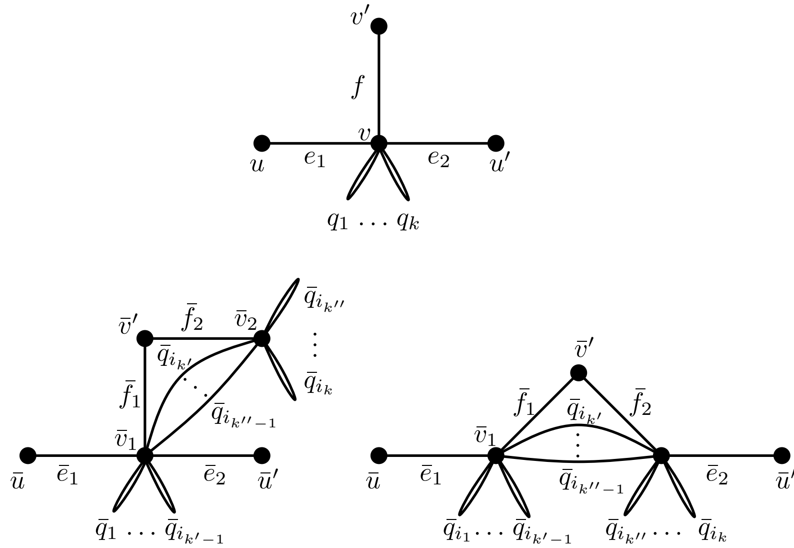

FIGURE 3. The graph on the top demonstrates a path of type 2 in Ω. The two graphs below it show the two possible choices of preimages for p. The graph-loops $\bar{q}_{i_1}, \dots, \bar{q}_{i_{k'}}$ is a subsequence of the graph-loops $\bar{q}_1, \dots, \bar{q}_k$ that consist of graph-loops based at $\bar{v}_1$. Similarly, $\bar{q}_{i_{k'}}, \dots, \bar{q}_{i_{k''-1}}$ are edges between $\bar{v}_1$ and $\bar{v}_2$ and $\bar{q}_{ik''}, \dots, \bar{q}_{i_k}$ 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 $\bar{f}_1 = \bar{e}_1$.

Type 5: In this case $v = B$ necessarily, 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 $q_1, \dots, q_n$. As above, these preimages consist of graph-loops at $\bar{v}_1$, graph-loops at $\bar{v}_2$ 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 $1 \le i < k$, 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 $\bar{B}$.

Type 3 and 4: The analysis of these cases is very similar to the cases done above and are omitted.

Cube attachment operation: Suppose Ω is obtained by attaching a cube c to the union $e_1 \cup \dots \cup e_n \subset \bar{\Omega}$, where for $1 \le i \le n$, 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 the labels $s_1, \dots, s_n$ 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 $s_1, \ldots, s_n$ form a clique in Γ, there is a reduced expression for h given by $h' = s_1^{\epsilon_1} \ldots s_n^{\epsilon_n}$ 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 $q' = e_1^{\epsilon_1} \cup e_2^{\epsilon_2} \cup \cdots \cup e_n^{\epsilon_n}$, 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 (regardless of whether or not 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 $q' = e_1^{\epsilon_1} \cup e_2^{\epsilon_2} \cup \cdots \cup e_n^{\epsilon_n}$ 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 $e_1 e_2^{\epsilon_2} \cup \cdots \cup e_n^{\epsilon_n} e_1$, 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. □

Lemma 3.10. *Let G be a subgroup of a right-angled Coxeter group 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

$$X(S_G) = \Omega_0 \xrightarrow{f_0} \Omega_1 \xrightarrow{f_1} \Omega_2 \xrightarrow{f_2} \cdots \to \Omega$$

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 $\hat{f} : \Omega_n \to \Omega$ is the natural map. By iteratively applying Lemma 3.9 starting with $\hat{l}$, 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. □

The existence of completions for subgroups is now an immediate consequence of Lemma 3.8 and Lemma 3.10:

Theorem 3.11. *Let G be a subgroup of a right-angled Coxeter group 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. Then (Ω, B) is a completion of G.* □

Definition 3.12 (Standard Completion of a Subgroup)**.** Let G be a subgroup the right-angled Coxeter group W_Γ generated by a finite generating set S_G. We call a completion Ω of G obtained by Theorem 3.11 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 3.13. The Γ_1-labeled cube complex Ω in Example 3.1 is a standard completion of the subgroup of W_{Γ_1} generated by the words $w_1 = adb$ and $w_2 = aec$. Similarly, the Γ_2-labeled cube complex Ω' in Example 3.2 is a standard completion of the subgroup of W_{Γ_2} generated by $w = abcd$.

Remark 3.13.1. Although every reduced word in W_Γ representing an element of G appears as a loop in the completion Ω, it is not true that every word in W_Γ representing an element of G appears in Ω as a loop. For instance let s be a vertex in Γ. Let G be a subgroup of W_Γ which is generated by a set of words in W_Γ, 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 right-angled Coxeter group 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 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 $1 \le i < i' \le n$ 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).

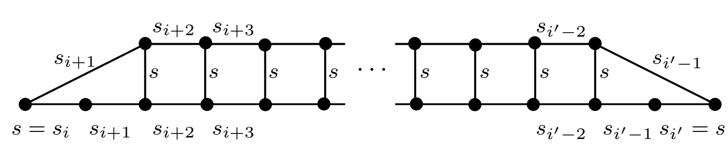

Figure 4.

Thus 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.

On the other hand, if $i' - i = 1$, then as Ω is folded, either α traverses an edge labeled by s twice in opposite directions, or α traverses 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. (However, in graph-loop case the new path is not necessarily 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 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 deletion operations. By Lemma 4.1, there is a path p'' with label w'' and with 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 generators which commute.

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 following definition and proposition allow us to "go backwards" and obtain a subgroup from a Γ-labeled complex.

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 right-angled Coxeter groups using completions:

Proposition 4.6. *Let G be a subgroup of a right-angled Coxeter group 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 [Dav08, 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, Proposition 2.2 implies that k cannot be an expression for the identity element. As k has finite order, it follows that G has torsion. □

For torsion-free subgroups, the following holds:

Theorem 4.7. *Suppose G is a torsion-free subgroup of the right-angled Coxeter group 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 $\phi : \pi_1(\Omega, B) \to G$ 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, we can apply Lemma 4.2(1) to conclude that α is null-homotopic. □

5. Core graphs

In general, a given 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, and we prove that it is unique. This is used in Theorem 5.5 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.

Remark 5.1.1. The core graph of a completion is not necessarily its entire 1-skeleton. For instance, it is straightforward to check that the core graph for the complex Ω of Figure 1 is not the whole 1-skeleton of Ω.

Note that $C(\Omega, B)$ is a connected, Γ-labeled complex, and so it has an associated subgroup as in Definition 4.3. Moreover, the following holds:

Lemma 5.2. *Let Ω be a connected, folded, cube-full, Γ-labeled complex, and let $C = C(\Omega, B)$. Then the subgroup of Γ associated to (C, B) is the same as the subgroup associated to (Ω, B).*

Proof. By Proposition 4.4, every element of the subgroup of W_Γ associated to Ω at B is represented by a loop based at B labeled by a *reduced* word. The lemma follows easily from this observation. □

Core graphs are unique in the following sense:

Proposition 5.3. *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 $f : C_1(\Omega_1, B_1) \to C_2(\Omega_1, B_2)$, such that $f(B_1) = B_2$.*

Proposition 5.3 is a consequence of the following lemma:

Lemma 5.4. *Let $G_1 \leq G_2 \leq W_\Gamma$, and let (Ω_1, B_1) and (Ω_2, B_2) be completions of G_1 and G_2 respectively. Then there is a label preserving combinatorial map $f : C(\Omega_1, B_1) \to C(\Omega_2, B_2)$ such that $f(B_1) = B_2$.*

Proof. Let C_i denote $C(\Omega_i, B_i)$ for $i = 1, 2$. We first define f on the vertices of C_1. Define $f(B_1) = B_2$. Now let $v_1 \neq B_1$ be a vertex in C_1. By the definition of a core graph, there is a loop in Ω_1 based at B_1 which passes through v_1, and whose label w is a reduced word in W_Γ. Let α_1 be the part of this (oriented) loop which goes from B_1 to v_1, and let w_α be the subword of w which labels α_1. Since w labels a loop in Ω_1 based at B_1, by the definition of completion, it represents an element of G_1, which by hypothesis is a subgroup of G_2. Now since w is reduced, again by the definition of a completion, there is a loop in Ω_2 based at B_2 with label w, and by the definition of a core graph, this loop is contained in C_2. Let α_2 be the initial segment of this loop which is labeled by w_α, and let $v_2 \in C_2$ be the endpoint of α_2. Define $f(v_1) = v_2$.

We must show that f is well-defined on vertices, i.e., that it does not depend on the choice of loop containing v_1. Suppose we run the construction in the previous paragraph on a different loop containing v_1, and thus obtain the path β_1 from B_1 to v_1 with label w_β, and corresponding path β_2 in C_2 with label w_β, from B_2 to some vertex v_2'. Then we wish to show $v_2 = v_2'$.

Observe that the concatenation $\alpha_1\beta_1^{-1}$ is a loop based at B_1. Thus its label $w_\alpha w_\beta^{-1}$ represents an element of $G_1 < G_2$. Let u be a reduced word in W_Γ representing $w_\alpha w_\beta^{-1}$. By the definition of completion, there is a loop γ_2 in Ω_2 based at B_2 with label u. Now the concatenation $\gamma_2\beta_2$ is a path in Ω_2 from B_2 to v_2' with label uw_β. Note that w_α is a reduced word which represents uw_β. It follows from Lemma 4.2 that there is a path from B_2 to v_2' with label w_α. On the other hand, α_2 is a path in Ω_2 from B_2 to v_2 with label w_α. However, since Ω_2 is folded, there is at most one path starting at B_2 with a given label. It follows that the two paths are the same, and hence $v_2 = v_2'$ are required.

Let e be an edge in Ω_1 between (possibly non-distinct) vertices u_1 and u_2, and let s be the label of e. In order to extend f to edges, we show there is a unique edge e' in Ω_2 between $f(u_1)$ and $f(u_2)$. Let α be a loop in Ω_1 which contains the edge e and has reduced label h. As before, there must be a loop α' in C_2 based at B_2 with label h. By the construction of the map f on vertices, it readily follows that α' contains the vertices $f(u_1)$, $f(u_2)$ and an edge labeled by s between these vertices. This edge is unique since Ω_2 is folded. □

We now use Lemma 5.4 to establish the uniqueness of core graphs:

Proof of Proposition 5.3. Let C_i denote $C(\Omega_i, B_i)$ for $i = 1, 2$. Lemma 5.4 implies that there are label preserving combinatorial maps $f_1 : C_1 \to C_2$ with $f_1(B_1) = B_2$, and $f_2 : C_2 \to C_1$ with $f_2(B_2) = B_1$. Then $f_2 \circ f_1$ is a label-preserving combinatorial map from C_1 to itself which fixes B_1. Since C_1 is folded and connected, there is exactly one label-preserving combinatorial map which fixes B_1, namely the identity. Thus $f_2 \circ f_1$ is the identity map of C_1, and similarly $f_1 \circ f_2$ is the identity map of C_2. It follows that f_1 is an isomorphism as required. □

Remark 5.4.1. We remark that if C is the core graph of a completion (Ω, B) of some $G < W_\Gamma$, then the complex induced by C in Ω (by including cubes of Ω whose boundaries are contained in C) is a connected, folded, Γ-labeled complex which satisfies properties (2) and (3) of the definition of a completion. However, it may not be cube-full. For instance, this is the case for the completion Ω of Figure 1.

Next we characterize normal subgroups of right-angled Coxeter groups in terms of core graphs. In Section 13, we give a different version of this theorem that is used to give an algorithm that checks whether a quasiconvex subgroup is normal.

Theorem 5.5. *Let G be a subgroup of W_Γ, and let (Ω, B) be a completion of G. Consider the following subset of $V(\Gamma)$:*

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

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 by Lemma 5.2, the group associated with $(C(\Omega, v), v)$ 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. On the other hand, 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. Then by Lemma 4.5, we know that (Ω, v) is a completion for $w^{-1}Gw = G$ (since G is normal). Then by Proposition 5.3, 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 6.9, which states that the index of a subgroup can be computed from a completion. In order to correctly state Theorem 6.9, we define resolved completions and resolved generating sets which address a slight technical issue that arises when Γ contains a vertex that is adjacent to every other vertex.

Definition 6.1 (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 6.2. *Let G be a subgroup of W_Γ, and let Ω be a resolved completion of G. Let $s \in \Gamma$ be such that $V(\Gamma) = star(s)$. 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. ☐

Definition 6.3 (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.

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. For otherwise by Proposition 2.2 there exists a reduced word w', ending with s, which is an expression for w. By Lemma 4.2(1) there is a corresponding 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 6.2, 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 $W_\Gamma = Gg_1 \sqcup Gg_2 ... \sqcup Gg_n$ 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 $M = \max\{|w_1|, ..., |w_n|\}$, 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

$h(h'h'')^{-1} = wkh''^{-1}h'^{-1}$. 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 dual to $p_w p_k$ must necessarily intersect either $p_{h'}$ or $p_{h''}$. Furthermore, a pair of dual curves which are each dual to p_k cannot intersect one another as s and t do not commute. Let C be the dual curve dual to the first edge of p_k. Note that C is of type s.

The dual curve C cannot intersect $p_{h''}$. For if it did, every curve dual to p_k would intersect $p_{h''}$ as well. However, as $|h''| \leq M$ and $|k| = 2M$, this is not possible. Thus, we conclude C intersects $p_{h'}$.

Additionally, no dual curve dual to 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 dual to p_w intersects $p_{h'}$ at an edge occurring before (in the orientation of $p_{h'}$) the edge of $p_{h'}$ dual to C.

By Lemma 2.4, there is a reduced word equal to h' in W_Γ which has ws as a prefix, which contradicts *(iii)* above. We remark that the above argument holds in the case when w is the empty word, i.e. when $B = v$. $\square$

We now define resolved sets. Lemma 6.6 below, whose proof is immediate, states that any generating set can be easily extended to a resolved generating set.

Definition 6.5 (Resolved set)**.** A set $\{w_1, w_2, \dots\}$ of words in W_Γ is *resolved* if given any $s \in V(\Gamma)$ such that $V(\Gamma) = \text{star}(s)$, it follows that $s \in \text{Support}(w_i)$ for some i.

Lemma 6.6. *Every set of words in W_Γ which generate a subgroup G can be extended to a resolved generating set of words for G. More specifically, suppose $T = \{w_1, w_2 \dots\}$ is a generating set of words for a subgroup G of W_Γ. Then $T' = T \cup \{s^2 \mid s \in \Gamma$ and $V(\Gamma) = star(s)\}$ is a resolved generating set for G.*

Lemma 6.7 guarantees that a resolved completion can always be constructed for a finitely generated subgroup.

Lemma 6.7. *Let G be a finitely generated subgroup of W_Γ and let S_G be a finite resolved generating for G (which exists by Lemma 6.6). If Ω is a standard completion for G with respect to S_G, then Ω is resolved.*

Proof. Let $\Omega_0 \to \Omega_1 \to \cdots \to \Omega$ be a standard completion of G with respect to S_G. Let $s \in V(\Gamma)$ be such that $V(\Gamma) = \text{star}(s)$. As S_G is resolved and as Ω_0 is the "rose graph" of words in S_G, it follows that some edge of Ω_0 is labeled by s. Thus, some edge in Ω is labeled by s. $\square$

Lemma 6.8. *Let G be a subgroup of W_Γ, and let (Ω, B) 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_G (which could be infinite).*

Proof. Let $B = v_1, v_2, \dots$ be an enumeration of the vertices of Ω. 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 $w_1, w_2, \ldots$ 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 established 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 $1 \leq i < j \leq n$. It follows that $w_iw_j^{-1}$ 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 and is unique 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. □

Theorem 6.9. *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.*

Proof. First suppose G has finite index in W_Γ. Then by Lemma 6.4, Ω is full valence. Furthermore, Ω is finite by Lemma 6.8. The other direction is an immediate consequence of Lemma 6.8. □

The following corollary for finitely generated subgroups immediately follows from Theorem 6.9 and Lemma 6.7.

Corollary 6.10. *Let G be a finitely generated subgroup of W_Γ and let S_G be a finite resolved generating for G (which exists by Lemma 6.6). Let Ω be a standard completion of G with respect to S_G. Then G has finite index in W_Γ if and only if Ω is finite and full valence.*

7. Nonpositive curvature

This section establishes criteria which guarantee that a completion is non-positively curved or a CAT(0) cube complex.

We begin with the following proposition, which shows that given a completion sequence of a Γ-labeled complex X, any hyperplane in any complex of this sequence is an extension of a hyperplane of X.

Proposition 7.1. *Let X be a Γ-labeled complex. Let*

$$X = \Omega_0 \to \Omega_1 \to \cdots \to \Omega$$

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. The claim that hyperplanes in Ω_n intersect the image of X will be proven 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. Suppose first that Ω_n is obtained by attaching an n-cube c to Ω_{n-1} along edges $e_1, \ldots, e_n$, 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. It is also clear that cube identification and fold operations do not produce new hyperplanes. 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. $\square$

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 7.2. *Let Ω be a folded Γ-labeled cube complex which does not contain a commuting bigon. Suppose further that Ω is either cube-full or that Γ is triangle-free. Then Ω is non-positively curved.*

Proof. Let v be a 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 simplicial. This part of the proof does not require that Ω is cube-full or that Γ is triangle-free.

The complex Δ cannot contain a graph-loop. For if it did, then adjacent sides of some square would be identified in Ω. However, adjacent sides of squares in Ω are always labeled by distinct elements of Γ, so such an identification cannot happen.

We next check that the 1-skeleton of Δ does not contain a bigon. Suppose there is such a bigon. It follows that there is a pair of edges e_1 and e_2 incident to a vertex v and there exist (possibly non-distinct) squares c_1 and c_2 in Ω each containing both e_1 and e_2. If c_1 and c_2 are distinct squares, then by the given identifications, their boundary-labels, read starting from v in the direction of e_1, must be the same. However, this is not possible as Ω is folded.

On the other hand, suppose c_1 and c_2 are the same square, $c = c_1 = c_2$. As c is labeled by $stst$ for some adjacent vertices $s, t \in \Gamma$, the opposite edges of c must be identified in Ω. It is straightforward to check that the only possible such identification producing a bigon in Δ is that of $\mathbb{R}P^2$. In this case, opposite edges of c are identified "with a flip."

It follows that the image of the boundary of c under the attaching map of c is a bigon in Ω whose two edges are labeled s and t. Since s and t are adjacent in Γ, this is a commuting bigon. This contradicts the assumption that Ω does not contain commuting bigons. Hence the 1-skeleton of Δ cannot contain a bigon. We have thus verified that Δ is simplicial.

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 $1 \le i \le n$, let e_i be the edge of Ω incident to v which u_i lies on, and let s_i be the label of e_i. For each $1 \le i < j \le n$, 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 $v \cup \bigcup_{i=1}^n e_i$. 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 7.2 immediately implies:

Proposition 7.3. *Any completion of a torsion-free subgroup of a right-angled Coxeter group is non-positively curved.* $\square$

Our next objective is to show that a completion of a finite Γ-labeled tree is a finite CAT(0) cube complex.

Lemma 7.4. *Let X be a finite Γ-labeled tree, and let*

$$X \to \Omega_1 \to \Omega_2 \to \cdots \to \Omega$$

be a completion sequence for X. Then Ω_i is simply connected for all i, and Ω is a CAT(0) cube complex.

Proof. We first 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 wsw^{-1}. As wsw^{-1} has an odd number of occurrences of the letter s, it represents a nontrivial element of W_Γ. However, by iteratively applying Lemma 3.9 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 $v \cup \bigcup_{i=1}^k e_i$. Hence Ω_{n+1} is simply connected.

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, and hence 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.

Finally, 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.

In order to show Ω is non-positively curved, by Lemma 7.2 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.

Finally, Ω is a CAT(0) cube complex as it is a simply-connected and non-positively curved cube complex. □

Proposition 7.5. *Let X be a Γ-labeled finite tree and let*

$$X = \Omega_0 \to \Omega_1 \to \cdots \to \Omega$$

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. By Lemma 7.4, we know that Ω is CAT(0). 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 Lemma 7.1, 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.

The second part of the claim now follows from Proposition 3.5. □

When Γ is triangle-free, we get the following more precise bound on the length of a standard completion sequence:

Proposition 7.6. *Let X be a Γ-labeled finite tree where Γ is triangle-free. Let*

$$X = \Omega_0 \to \Omega_1 \to \cdots \to \Omega$$

be a standard 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 depending only on the number of edges of X and on $|V(\Gamma)|$.

Proof. By Proposition 7.5, $\Omega = \Omega_N$ 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:

$$\Theta_1 = \Omega_{i_1}, \Theta_2 = \Omega_{i_2}, \ldots, \Theta_n = \Omega_{i_n}$$

By Proposition 7.2, we know that Θ_i is a CAT(0) cube complex. Furthermore, Θ_i has diameter at most E, by the proof of Proposition 7.5.

We claim that the complex Θ_j is not isometric to Θ_k for all $k > i$. 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 and may not be quasiconvex with respect to some other generating set for G. However, if a subgroup is quasiconvex with respect to some generating set, then it is quasi-isometrically embedded with respect to any generating set [BH99, Chapter III.Γ, Lemma 3.5].

When we say a subgroup is quasiconvex in a right-angled Coxeter group, we will always mean with respect to the standard generating set. The main result of this section is that a subgroup of a right-angled Coxeter group is quasiconvex if and only if any 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 of the right-angled Coxeter group. Note that whenever we consider a Cayley graph of a right-angled Coxeter group, it is the Cayley graph associated to the standard Coxeter generating set.

Lemma 8.1. *Let G be a subgroup of the right-angled Coxeter group 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:

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

$\square$

Lemma 8.2. *Let G be a subgroup of the right-angled Coxeter group 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 the endpoints of α are the identity vertex v_{id}, and some vertex labeled by an element g of G. It follows that 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, $d_\Omega(u, B) = d_{\mathcal{C}}(v, G) \leq M$. □

Lemma 8.3. *If G is a quasiconvex subgroup of the right-angled Coxeter group 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 [BH99, Chapter III.Γ Lemma 3.5]. Let (Ω, B) be a standard completion of G. Let Y be the subset of vertices of Ω that are contained in a loop based at B which is labeled by a reduced word, i.e., Y consists of vertices belonging to the core $C(\Omega, B) \subset \Omega$. By Lemma 8.1, $d_\Omega(v, B) \leq M$ for every vertex $v \in Y$.

Suppose Ω is not finite. We will obtain a contradiction by finding a loop based at B in Ω which is labeled by a reduced word, but contains a vertex that has distance greater than M from B.

By Proposition 7.1 and since Ω is a standard completion, there are only finitely many hyperplanes in Ω. Thus, there exists an integer N such that any geodesic in Ω of length N has at least $M+3$ of its edges dual to the same hyperplane. Let v be a vertex in Ω that is at distance N from the basepoint B. Let α be a geodesic from B to v, and let w be the label of α. As α is geodesic, w is a reduced word. Let H be an hyperplane that is dual to at least $M+3$ edges of α, and let $s \in \Gamma$ be the type of H. Let e_1 and e_2 be the last two edges of α (read from B to v) that are dual to H.

Let β_1 be the smallest subpath of α that contains both e_1 and e_2, and let β_2 be the geodesic along the carrier of H from the endpoint of β_1 to the startpoint of β_1. Let b_1 be the label of β_1, and let b_2 be the label of β_2. Note that b_2 is a word in $\text{link}(s)$ and, in particular, has no occurrences of s.

Let β be the subpath of α from B to the startpoint of e_1, and let b be its label. There is some generator $t \in V(\Gamma)$ in the word b which does not commute with s and appears after every occurrence of s in b. This follows since $w = bb_1$ is reduced, b_1 begins with the letter s and b contains the letter s.

Let l be the loop $\beta\beta_1\beta_2\beta^{-1}$ with label $h = bb_1b_2b^{-1}$. Let $\hat{h}$ be a reduced expression for h obtained by a series of deletions. We claim that h and $\hat{h}$ have the same number of occurrences of the generator s. For suppose that a deletion of an s generator occurs at some point. Such a deletion must be between an s generator in bb_1 and one in b^{-1}. This follows since b_2 does not contain any occurrences of the generator s, bb_1 is a reduced word, and b^{-1} is a reduced word. Additionally for this to be possible, the occurrence of t in b^{-1} must first be deleted as well. As t does not commute with s, there is no occurrence of t in b_2. Thus, if this occurence of t were deleted, there must be an occurrence of t also deleted in bb_1. However, this is not possible as bb_1 ends with s and no s occurrence can be deleted before t is deleted.

We will now construct a reduced expression for h. First obtain a reduced expression $\hat{b}\hat{b}_1\hat{b}_2$ for bb_1b_2 by a series of deletion operations (here the hat notation indicates that some generators in the original words have been deleted). By Lemma 4.1, there is a path p with the same endpoints as $\beta\beta_1\beta_2$ and with label $\hat{b}\hat{b}_1\hat{b}_2$. In particular, the endpoint vertex of p is the same as the endpoint vertex of β. Call this vertex v_1. The loop $p\beta^{-1}$ has label $\hat{b}\hat{b}_1\hat{b}_2b^{-1}$ which is an expression for h.

Now let $\hat{h}$ be a reduced expression for h obtained by performing deletion operations to the word $\hat{b}\hat{b}_1\hat{b}_2b^{-1}$. As $\hat{b}\hat{b}_1\hat{b}_2^{-1}$ and b^{-1} are each reduced, each deletion operation must delete a pair of generators: one in b^{-1} and one in $\hat{b}\hat{b}_1\hat{b}_2^{-1}$. As no occurrence of s can be deleted and as b^{-1} has at least $M+1$ occurrences of s, no more than $|b^{-1}| - (M+1) = |\beta| - (M+1)$ deletion operations can occur.

By the definition of a completion of a subgroup, there is a loop $\hat{l}$ in Ω with label $\hat{h}$ based at B. By Lemma 4.1(3), some vertex u of $\hat{l}$ satisfies $d(u, v_1) \leq |\beta| - (M+1)$. As $d(v_1, B) = |\beta|$, it follows that u is distance at least $M+1$ from B. Now $\hat{l}$ is a loop labeled by a reduced word so by definition, $u \in Y$. However this contradicts the fact that all vertices in Y have distance at most M from B. □

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

Theorem 8.4. *Let G be a subgroup of a right-angled Coxeter group. 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.*

9. Residual finiteness and separability

In this section we give a new proof of a result of Haglund, which states that quasiconvex subgroups of right-angled Coxeter groups are separable and are virtual retracts. We additionally give a short proof of the well-known result that right-angled Coxeter groups are residually finite.

We begin with a preliminary result which is also used in later sections. Specifically, in a few arguments throughout the rest of 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 9.1. *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. Further 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 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.

By assumption, Ω' is a folded complex. Thus, 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 $s_1, \dots, s_n$. Further suppose that $v \cup e_1 \cup \dots \cup e_n$ are not all contained in a common n-cube and that n is maximal out of such possible choices. Consider the corresponding cube attachment operation which glues 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 the following holds for some $0 \leq k \leq n$: 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. □

Let G be a finitely generated subgroup of a right-angled Coxeter group W_Γ, and suppose that G has a finite completion, (Ω, B). We define a construction that can be performed to Ω to produce a new complex.

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 9.1. 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 the conclusions of Lemma 9.1, $\mathcal{E}$ is finite, full-valence, and the natural inclusion of Ω into $\mathcal{E}$ is an isometry. The following proposition shows the relation between G and the subgroup associated to $(\mathcal{E}, B)$.

Proposition 9.2. *Let G be a finitely generated subgroup of a right-angled Coxeter group W_Γ, and suppose that G has a finite completion (Ω, B). Let $(\mathcal{E}, B)$ be the full valence extension of (Ω, B). Let H be the subgroup associated to $(\mathcal{E}, B)$ (as in Definition 4.3). Then H is a finite-index subgroup of W_Γ, and there is a retraction from H to G.*

Proof. By Proposition 4.4 and Lemma 6.8, H has finite index in W_Γ. We are left to prove that G is a retract of H.

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. We define a map $\phi : H \to G$ as follows. 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 of Tits moves 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. By the same argument, 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. □

We give a proof using completions that right-angled Coxeter groups are residually finite. This is a well known result. In fact these groups are linear (see for instance [BB05]), and by a theorem of Malcev, every finitely generated linear group is residually finite (see [Weh69, 4.2] for a proof).

Theorem 9.3. *Every right-angled Coxeter group is residually finite.*

Proof. Let W_Γ be a right-angled Coxeter group. Let g be a nontrivial 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 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 7.5, 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 9.1, w is still not the label of a loop in $\mathcal{E}$ based at B. Let H be the subgroup of W_Γ associated to $(\mathcal{E}, B)$. By Proposition 9.2, H 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 $\mathcal{E}$. □

A subgroup G of a group K is a virtual retract if G is a retract of a finite index subgroup of K. The next theorem and corollary show that quasiconvex subgroups of a right-angled Coxeter group are virtual retracts and are separable. These results were first proven by Haglund in [Hag08].

Theorem 9.4. *Let G be a quasiconvex subgroup of a right-angled Coxeter group W_Γ. Then G 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). Let H be the subgroup of W_Γ associated to $(\mathcal{E}, B)$. By Proposition 9.2, H is a finite-index subgroup of W_Γ, and G is a retract of H. □

It is well-known that a virtual retract of a residually finite group is separable. We refer the reader to [Hag08, Proposition 3.8] for a proof. The following corollary thus immediately follows.

Corollary 9.5. *Every quasiconvex subgroup of a right-angled Coxeter group is separable.*

10. Reflection Subgroups

Given a right-angled Coxeter group W_Γ, a *reflection* is an element of W_Γ represented by a word of the form wsw^{-1} where s is a generator in $V(\Gamma)$ and w is a word in W_Γ. In this section, we consider a subgroup G generated by a finite set of reflections. We show there is a constructive algorithm to build a finite completion of G. In particular, such subgroups are always quasiconvex and their index can be computed. We will use these results to study Coxeter subgroups of 2-dimensional right-angled Coxeter groups in the next section.

10.1. **Trimmed sets of reflections.** The following observation guarantees that we can always use a nice generating set of reflections:

Lemma 10.1. *Let W_Γ be a right-angled Coxeter group, and let G be a subgroup generated by a finite set of reflections $\mathcal{R}'$. Then G is generated by a the set of reduced reflections of the form:*

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

where for all $i \neq j$, no reduced expression for w_j begins with $w_i s_i$. 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 $g = wsw^{-1}$ 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 $g' = hgh^{-1} = (w'q)s(w'q)^{-1}$, 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. $\square$

Definition 10.2 (Trimmed reflection set)**.** We say that a set $\mathcal{R}$ of reduced reflections is *trimmed* if it satisfies the conclusion of Lemma 10.1.

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 right-angled Coxeter group 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 10.1 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 3.11, 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 7.5, 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 $\hat{f} : \mathcal{T} \to \mathcal{FT} \to \Omega_{\mathcal{FT}}$ be the natural map. By a slight abuse of notation, we also denote by $\hat{f}$ the natural map $\hat{f} : \mathcal{FT} \to \Omega_{\mathcal{FT}}$. Let $\hat{\mathcal{T}} := \hat{f}(\mathcal{T}) = \hat{f}(\mathcal{FT})$.

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

$$L_{\hat{v}} = \left\{ s \in V(\Gamma) \;\middle|\; \begin{array}{c} \exists\ v \in V(\Omega_0) \text{ incident to a graph-loop labeled } s, \\ \text{such that } \hat{v} = \hat{f}(v) \end{array} \right\}$$

Note that the set of vertices $\{\hat{v} \in \Omega_{\mathcal{FT}} \mid L_{\hat{v}} \neq \emptyset\}$ is exactly the set of vertices of $\Omega_{\mathcal{FT}}$ whose preimage in $\mathcal{T} \subset \Omega_0$ contains a vertex incident to a graph-loop.

We would like to build the completion Ω_G by "adding back" the graph-loops to $\Omega_{\mathcal{FT}}$ and applying Lemma 9.1. 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 9.1 would not be satisfied. The next two lemmas show this situation is not possible. They are the technical results required to prove the main results of this section.

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 a word which is an expression for ks.*

Proof. For a contradiction, suppose there exists such a path α. As $\mathcal{T}$ is a tree, if α were not geodesic then some generator would be consecutively repeated in the label of α. Thus, by possibly passing to a homotopic path, we may assume that α is geodesic in $\mathcal{T}$. 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 in W_Γ, we have $h_2 = ks$ and $r_2 = (ks)(h_2' s' h_2'^{-1}) s k^{-1}$. However, in this case 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 $w_1 s$. The claim follows. $\square$

Lemma 10.4. *For every vertex $\hat{v}$ of $\Omega_{\mathcal{FT}}$ and every $s \in L_{\hat{v}}$, $\hat{v}$ is not incident to an edge labeled by s in $\Omega_{\mathcal{FT}}$.*

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

Let H be the hyperplane (recall that $\Omega_{\mathcal{FT}}$ is a CAT(0) cube complex) dual to d. In particular, H is of type s. By Lemma 7.1, 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 $\Omega_{\mathcal{FT}}$ 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 $\text{link}(s) \subset \Gamma$.

As $\hat{\beta}'$ is geodesic, any hyperplane intersects it at most once. Thus, as hyperplanes separate $\Omega_{\mathcal{FT}}$ 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 of the form 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 right-angled Coxeter group. Then there exists a finite completion of G.*

Proof. As previously discussed, we first obtain the completion $\Omega_{\mathcal{FT}}$ of $\mathcal{T}$ using Proposition 7.5. For each vertex $\hat{v}$ of $\Omega_{\mathcal{FT}}$ and $s \in L_{\hat{v}}$, we attach a graph-loop to $\hat{v} \in \Omega_{\mathcal{FT}}$ labeled by s. Let $\Omega'_{\mathcal{FT}}$ 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 $\Omega'_{\mathcal{FT}}$ can be obtained from Ω_0 by applying the same completion sequence that was applied to $\mathcal{T}$ to obtain $\Omega_{\mathcal{FT}}$, while "ignoring" the graph-loops. We now get a finite completion Ω_G of $\Omega'_{\mathcal{FT}}$ by applying Lemma 9.1. It follows that Ω_G is a finite completion for G. $\square$

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

Corollary 10.6. *Every finitely generated reflection subgroup of a right-angled Coxeter group is quasiconvex.* $\square$

For 2-dimensional right-angled Coxeter groups, 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 result will be important in Section 12.

Theorem 10.7. *Let W_Γ be a 2-dimensional right-angled Coxeter group. 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 10.1, we may assume without loss of generality that $\mathcal{R}$ is trimmed. As before, we first obtain the completion $\Omega_{\mathcal{FT}}$ of $\mathcal{T}$. However, this time we use the more refined Proposition 7.6 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 9.1(3) guarantees the bound. $\square$

Remark 10.7.1. Note that by Lemma 9.1, the completion Ω_G of the reflection group G given by Theorem 10.5 and by Theorem 10.7 contains the complex $\Omega_{\mathcal{FT}}$ as an isometrically embedded subcomplex. Moreover, the inclusion of $\Omega_{\mathcal{FT}}$ into Ω_G satisfies the additional properties given by Lemma 9.1. These facts will be important in Section 12.

11. Coxeter subgroups of 2-dimensional right-angled Coxeter groups

In this section we study Coxeter subgroups of 2-dimensional right-angled Coxeter groups. Recall that such subgroups, by our definition, are always finitely generated.

It is clear that a Coxeter subgroup of a right-angled Coxeter group W_Γ must be generated by a set of involutions (order two elements) in W_Γ. We show in Theorem 11.4 that under mild hypotheses these subgroups are generated by reflections. Consequently, Theorem 10.5 applies to these subgroups.

We first prove three lemmas involving the structure of particular types of words in a right-angled Coxeter group. The first is well-known and addresses involutions:

Lemma 11.1. *Let g be an involution in the right-angled Coxeter group W_Γ. Then there is an expression wkw^{-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 right-angled Coxeter group is contained in a conjugate of a special finite subgroup (see [Dav08, Theorem 12.3.4] for instance). Furthermore, a special subgroup of a right-angled Coxeter group is finite if and only if its defining graph is a clique. The lemma then follows. □

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 right-angled Coxeter group 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 $1 \le i \le m$, either:*

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

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 $wkw^{-1}k^{-1}$. 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 $1 \le i \le m$. 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 $1 \le r \le n$. As dual curves of the same type do not intersect, it follows that for each $1 \le j \le l$, there is a curve in $\mathcal{H}$ dual to both e_j and e'_j. Let H be the dual curve in $\mathcal{H}$ that is bottom-most in R, i.e., that is dual to 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 $1 \le r \le n$. Observe that the label of α has an odd number of occurrences of the letter t. Hence some curve dual to 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 $1 \le r \le n$. Let d be the edge on the bottom of R labeled by k_r, and let d' be the edge on the 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 $\mathcal{H}$ dual to 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 $\mathcal{H}$ dual to 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 jth edge labeled t on the left side is paired with the jth 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, however, 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. □

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 right-angled Coxeter group 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 $bxb^{-1}x^{-1}$. 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 is dual to both e_t and e_b. Let $N(H)$ be the set of cells in R that contain an edge dual to H. It follows that the boundary $\partial N(H)$ of $N(H)$ has label $s_1ys_1y^{-1}$ 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 an 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, $\hat{b}$ and s_1 both commute with $\hat{x}$, 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 $e_1, \ldots, e_l$ be the set of edges labeled by s_1 on the left side of R ordered from bottom to top, and let $e'_1, \ldots, e'_l$ 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 $\mathcal{H}$

are either dual to 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 alternatively are dual to 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 $\mathcal{H}$ has 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 is dual to d. Note that T cannot intersect any dual curve in $\mathcal{H}$ as s_1 and t do not commute. Furthermore, T cannot be dual to 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 is dual to an edge on the right side of R that lies before e'_r. Correspondly, 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. □

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

Theorem 11.4. *Let G be a subgroup of a right-angled Coxeter group 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 right-angled Coxeter group, 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 will inductively construct a sequence of generating sets for G, each of which consists 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 both $k_1 \neq s_1$ 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, $k' = s_1 s_2$. Lemma 11.2 further implies that s_1 and s_2 commute with z. Consequently $x = s_1 s_2 = k'$, 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. $\square$

Now it follows easily from Corollary 10.6 that subgroups of the kind in Theorem 11.4 are quasiconvex:

Corollary 11.5. *Given a 2-dimensional right-angled Coxeter group, every Coxeter subgroup whose defining graph does not have an isolated vertex is quasicovex.* $\square$

Remark 11.5.1. Note that the defining graph of a right-angled Coxeter group 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 right-angled Coxeter group is a finite-index subgroup of a 2-dimensional right-angled Coxeter group

In this section, we provide an affirmative answer to the following question:

Question 12.1. *Is there an algorithm which, given a 2-dimensional, one-ended right-angled Coxeter group W_Γ and any right-angled Coxeter group $W_{\Gamma'}$, determines whether or not W_Γ contains a finite-index subgroup isomorphic to $W_{\Gamma'}$?*

The result we prove, Theorem 12.11, actually assumes a weaker hypothesis than one-endedness of W_Γ. We additionally show that the time-complexity of this algorithm only depends on the number of vertices of Γ and Γ'. Furthermore, when such a subgroup does exist, the output of the algorithm is a set of words in W_Γ which is a standard right-angled Coxeter group generating set for a subgroup isomorphic to $W_{\Gamma'}$.

Given a set of reflections $\mathcal{R}$ which generate a finite-index subgroup of the right-angled Coxeter group W_Γ, under the right hypotheses, Proposition 12.2 below bounds the sizes of elements in $\mathcal{R}$ as a function of $|V(\Gamma)|$ and $|\mathcal{R}|$. This proposition 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.

We define a certain class of graphs in order to state this proposition. Namely, we say a graph Δ is *almost star* if there exist vertices $s, t \in \Delta$ (possibly not distinct) such that $V(\Delta) = \mathrm{star}(s) \cup \{t\}$.

Proposition 12.2. *Let W_Γ be a right-angled Coxeter group such that Γ is triangle-free and is 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 on $|\mathcal{R}|$ 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 in this discussion will be fixed until the proof of Proposition 12.2 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 right-angled Coxeter group. 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 10.7, 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

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

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 12.2. 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 12.3. *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 $v = v'$. $\square$

Lemma 12.4. *Let β be a path in $\Omega_{\mathcal{FT}}$ whose label is 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. $\square$

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 12.5. *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 12.4 $\hat{f}(\eta)$ 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 $\hat{\zeta}$ be a geodesic along the carrier of H_1 from the endpoint to start point of $\hat{\beta}$. Let z be the label of $\hat{\zeta}$. Let D be a disk diagram in $\Omega_{\mathcal{FT}}$ with boundary $\hat{\beta} \cup \hat{\zeta}^{-1}$. The label of the boundary 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 is dual to 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 dual to 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 occurences of $\bar{s}$ appear as “left-most” as possible. Thus, every dual curve dual to 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 a possibly empty word in W_Γ. However, it follows from our choice of expression for w_j (where 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 wk_1 is a prefix of both w_1 and w_j. Hence, it must be the case that $e = e_1$ and so $\hat{e} = \hat{e}_1$. □

Lemma 12.6. *Let Y be the subcomplex of $\Omega_{\mathcal{FT}}$ 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 $\hat{v} \in \Omega_{\mathcal{FT}} \subset \Omega_G$) is in 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 $\mathcal{FT}$ is not contained in γ. By Lemma 12.5 and the fact that every vertex of γ has valence 2 by construction, $\hat{f}(\mathcal{FT}) \cap Y = \gamma$. 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 9.1 and Remark 10.7.1, there is a path η in $\Omega_{\mathcal{FT}}$ 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 12.2. 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 $\mathrm{star}(s) \subsetneq V(\Gamma)$ for any $s \in V(\Gamma)$. In particular, any finite set of words in Γ is resolved. As G is a finite-index subgroup of W_Γ given by a resolved generating set, Ω_G is full valence by Corollary 6.10.

Again, as Γ is not almost star, there exists a vertex $a \in \Gamma$ such that $a \neq t$ and $a \notin \mathrm{star}(\bar{s})$. As Ω_G is full valence, there must exist an edge $\hat{d}$ adjacent to $\hat{v}$ labeled by a. By Lemma 12.6, $\hat{d}$ 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 12.6. 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 $\mathrm{star}(\bar{s})$. Furthermore, $\hat{d}'$ cannot be dual to H_1 either. For otherwise, as $\hat{\gamma}$ is geodesic (by Lemma 12.4), it follows that the hyperplane dual to $\hat{e}$ must intersect H_1, contradicting the fact that t is not in $\mathrm{link}(\bar{s})$. Thus, $H \neq H_1$ and $H \neq H_2$.

By Proposition 7.1, 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 12.5, 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 the main theorem, we need to address a special case and recall some known results. The next lemma describes finite-index subgroups of W_Γ for the case where Γ is a triangle-free join graph.

Lemma 12.7. *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

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

and let $\mathcal{R}_B = \mathcal{R} \setminus \mathcal{R}_A$. Note that if $wsw^{-1} \in \mathcal{R}_A$, then w does not have a letter in $V(B)$. For if it did, wsw^{-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. ☐

Next we recall a theorem proven independently by Deodhar and by Dyer which states that reflection subgroups of Coxeter groups are Coxeter groups themselves.

Theorem 12.8 ([Dye90] [Deo89])**.** *Let W be a Coxeter group and $G < W$ a subgroup generated by reflections. Then G is a Coxeter group.*

The following corollary follows in the setting of right-angled Coxeter groups.

Corollary 12.9. *Let W be a right-angled Coxeter group and $G < W$ a subgroup generated by a set $\mathcal{R}$ of reflections. Then G is a right-angled Coxeter group. Furthermore, if $\mathcal{R}$ is trimmed, then it is a standard Coxeter generating set.*

Proof. By Theorem 12.8 and Proposition 2.1, G is a right-angled Coxeter group.

The second claim follows from the main theorem of [Dye90] and [Dye90, Proposition 3.5]. We also refer the reader to [Dye90, page 69] for an algorithm to determine a standard generating set for a reflection subgroup of a Coxeter group. $\square$

The next proposition describes a trick to produce an index two right-angled Coxeter subgroup of a right-angled Coxeter group. We will use this trick in one of the cases in the proof of Theorem 12.11. This proposition is well known to experts. However, we do not know of a proof in the literature, so we provide one here.

Before stating the proposition, we define a certain graph operation. Namely, given a graph Γ and a vertex $v \in \Gamma$, let Δ be the subgraph of Γ induced by $V(\Gamma) \setminus v$. We define $D(\Gamma, v)$ to be the graph consisting of the union of two copies of Δ which are identified along the subgraph of Δ induced by $\text{link}(v)$.

Proposition 12.10. *Given any graph Γ and vertex $s \in V(\Gamma)$, let $\phi_s : W_\Gamma \to \mathbb{Z}_2 = \{1, a\}$ be the homomorphism defined by $\phi_s(s) = a$ and $\phi_s(t) = 1$ for all $t \in V(\Gamma)$ such that $t \neq s$. Let K be the kernel of ϕ_s. Then K is generated by reflections and is isomorphic to the right-angled Coxeter group $W_{D(\Gamma,s)}$.*

Proof. Fix a vertex $s \in \Gamma$ and set $\phi = \phi_s$. We compute the kernel K of the map ϕ. For any word w in W_Γ, it follows that $\phi(w) = 0$ if and only if w contains an even number of occurrences of the letter s. Let $\mathcal{R}_1 = V(\Gamma) \setminus s$ and $\mathcal{R}_2 = \{sts \mid t \in \mathcal{R}_1\}$. It readily follows that K can be generated by the set of reflections $\mathcal{R} = \{\mathcal{R}_1 \cup \mathcal{R}_2\}$. Thus K is a right-angled Coxeter group by Corollary 12.9.

Let Δ be the graph whose vertices correspond to reflections in $\mathcal{R}$ and whose edges correspond to commuting reflections. By the condition described in [Dye90, Introduction], K is isomorphic to W_Δ. We now leave to the reader to check that Δ is isomorphic, as a graph, to $D(\Gamma, v)$. $\square$

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

Theorem 12.11. *There is an algorithm which, given a 2-dimensional right-angled Coxeter group W_Γ and a right-angled Coxeter group $W_{\Gamma'}$ 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 does exist some subgroup G of W_Γ that is isomorphic to $W_{\Gamma'}$, then by Theorem 11.4 and Lemma 10.1, 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 the graph 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 $wsw^{-1} \in I$. 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 in the first paragraph of the proof. As Γ is not almost star, Proposition 12.2 guarantees that $\mathcal{R} \in \mathcal{I}_M$ where M is as in that proposition 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 10.7 and Corollary 6.10, 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 Corollary 12.9, G is a right-angled Coxeter group and I is a standard Coxeter generating set. As any right-angled Coxeter group is defined by a unique graph [Rad03], it is straightforward to check whether I generates a right-angled Coxeter group isomorphic to $W_{\Gamma'}$. The theorem then follows in this case.

(ii) $|V(\Gamma)| \leq 2$

It easily follows that W_Γ is isomorphic to either $\mathbb{Z}_2$, $\mathbb{Z}_2 \times \mathbb{Z}_2$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. In each case, such a group contains a finite number of finite-index right-angled Coxeter subgroups, up to isomorphism, and one can easily list which these are.

(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 in the first paragraph of this proof.

If G is a finite-index subgroup, then Lemma 12.7 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 correspond to elements of $\mathcal{R}_A$ and vertices of Δ_B correspond to elements of $\mathcal{R}_B$. It readily follows that $W_{\Gamma'}$ is isomorphic to the right-angled Coxeter group W_Δ.

Thus we can assume that Γ' is a join graph, $\Gamma' = A' \star B'$. To prove the claim, again by Lemma 12.7, it is enough to check whether $W_{A'}$ is isomorphic to a finite-index subgroup of W_A and whether $W_{B'}$ 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 $V(\Gamma) \subsetneq \operatorname{star}(v)$ for all $v \in V(\Gamma)$. Let s and t be vertices of Γ such that $V(\Gamma) = \operatorname{star}(s) \cup \{t\}$. Note that s and t must be distinct in this case.

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 $\phi = \phi_u : W_\Gamma \to \mathbb{Z}_2$ be the homomorphism described in Proposition 12.10, let $\Delta = D(\Gamma, u)$ and let $K = \ker(\phi)$. Note that Δ is not almost star and that K is generated by reflections.

Let $i_G : G \to W_\Gamma$ 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.

$$\begin{tikzcd}
K' = \ker(\phi') \arrow[r, "i"'] \arrow[d, "i"] & G \arrow[d, "i_G"] \arrow[dr, "\phi' = \phi \circ i_G"] & \\
K = \ker(\phi) \cong W_\Delta \arrow[r, "i"'] & W_\Gamma \arrow[r, "\phi"'] & \mathbb{Z}_2
\end{tikzcd}$$

Recall that given any group G_1, a subgroup G_2 of G_1 and a subgroup G_3 of G_2 then 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 above 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, suppose that $V(\Gamma) = \mathrm{star}(s)$ for some vertex $s \in V(\Gamma)$. In this case 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. It thus follows by either (i) or (ii) that there is an algorithm to check whether K' is a finite-index subgroup of K, where K' and K are defined as before. $\square$

13. Other algorithmic properties of quasiconvex subgroups

This section is dedicated to the proof of Theorem 13.1 which gives several algorithmic results for quasiconvex subgroups of right-angled Coxeter groups.

We remark that the existence of algorithms for (1) and (4) in Theorem 13.1 below for quasiconvex subgroups of automatic groups (which include right-angled Coxeter groups) are known by work of Kapovich [Kap96, Lemma 2] and Kharlampovich–Miasnikov–Weil [KMW17] respectively. We include proofs here anyway, as it also easily follows from what we have already shown.

Theorem 13.1. *Let G be a quasiconvex subgroup of a right-angled Coxeter group W_Γ given by a finite generating set of words in W_Γ. Then there exist finite-time algorithms to solve the following problems.*

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

Let Ω be a standard completion of G. By Theorem 8.4, Ω is finite and by Proposition 3.5 it can be computed in finite time. We fix this notation throughout the rest of the section.

The proof of (1), (3) and (4) follow from work we have already done:

Proof of (1), (3) and (4): Let w be a reduced word representing some element $g \in W_\Gamma$. It follows from the definition of a completion, that $g \in G$ if and only if there exists a loop in Ω based at B with label w. This shows (1). The claims (3) and (4) follow respectively by Proposition 4.6 and Corollary 6.10. □

Before proving (2), we prove that powers of an element of a right-angled Coxeter group can be represented by words of a special form.

Lemma 13.2. *Let w be a reduced word in the right-angled Coxeter group W_Γ. Then there exist reduced words x, h and k, such that $xhkx^{-1}$ is a reduced expression for w and $xh^n k^{(n \mod 2)} x^{-1}$ 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 $xh^n k^{(n \mod 2)} x^{-1}$ is clearly an expression for the word w^n. Furthermore, the word $xh^n k^{(n \mod 2)} x^{-1}$ must be reduced. For otherwise, it follows by Proposition 2.2 that either h is not reduced or that 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 we know is not possible. □

Proof of (2): 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 $l \le N$. 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 13.2, 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 $1 \le i \le m$, 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 $l \le N$. 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. This proves the claim. □

We would like to use Theorem 5.5 to prove (5). However, the core of a completion may be difficult to algorithmically compute in general. Thus, for finitely generated subgroups, we can give a different characterization of normality which is better suited to the algorithmic approach.

Proposition 13.3. *Let G be a subgroup of W_Γ generated by a finite set of reduced words S_G, and let Ω be a standard completion for G with respect to S_G. Consider the following subset of $V(\Gamma)$:*

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

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 5.5 implies N1 and N2'. On the other hand, suppose N1 and N2' hold. As in the proof of Theorem 5.5, 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 $sGs \subseteq G$. □

Finally, we are ready to show (5):

Proof of (5): Let S_G be a finite set of reduced words in W_Γ which generate G. Since S_G consists of reduced words and generates G, the set Δ from Proposition 13.3 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 now follows that the conditions N1 and N2' can be checked in finite time as well. □

References

[Ago13] Ian Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning. 1

[Bah05] Patrick Bahls, *The isomorphism problem in Coxeter groups*, Imperial College Press, London, 2005. MR 2162146 8

[BB05] Anders Bjorner and Francesco Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, vol. 32, Springer, 2005. 5, 6, 35

[BH93] Brigitte Brink and Robert B. Howlett, *A finiteness property and an automatic structure for Coxeter groups*, Math. Ann. **296** (1993), no. 1, 179–190. MR 1213378 4

[BH99] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999. 30, 31

[BHS17] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto, *Thickness, relative hyperbolicity, and randomness in Coxeter groups*, Algebr. Geom. Topol. **17** (2017), no. 2, 705–740, With an appendix written jointly with Pierre-Emmanuel Caprace. MR 3623669 4

[BL18] Benjamin Beeker and Nir Lazarovich, *Stallings' folds for cube complexes*, Israel J. Math. **227** (2018), no. 1, 331–363. MR 3846326 4

[Bro] Samuel Brown, *Geometric structures on negatively curved groups and their subgroups, PhD thesis*, University College London, 2016. 4

[CP08] John Crisp and Luisa Paoluzzi, *Commensurability classification of a family of right-angled Coxeter groups*, Proc. Amer. Math. Soc. **136** (2008), no. 7, 2343–2349. MR 2390500 (2008k:20090) 3

[CR15] Montserrat Casals-Ruiz, *Embeddability and universal theory of partially commutative groups*, Int. Math. Res. Not. IMRN (2015), no. 24, 13575–13622. MR 3436157 4

[CS11] Pierre-Emmanuel Caprace and Michah Sageev, *Rank rigidity for CAT(0) cube complexes*, Geom. Funct. Anal. **21** (2011), no. 4, 851–891. MR 2827012 7

[Dan18] Pallavi Dani, *The large-scale geometry of right-angled Coxeter groups*, Handbook of Group Actions, vol. V, International Press and Higher Education Press, 2018. 1, 5

[Dav08] Michael W. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008. MR 2360474 5, 6, 8, 20, 39

[Deo89] Vinay V. Deodhar, *A note on subgroups generated by reflections in Coxeter groups*, Arch. Math. (Basel) **53** (1989), no. 6, 543–546. MR 1023969 2, 46

[DJ00] Michael W. Davis and Tadeusz Januszkiewicz, *Right-angled Artin groups are commensurable with right-angled Coxeter groups*, J. Pure Appl. Algebra **153** (2000), no. 3, 229–235. MR 1783167 1, 4

[DST18] Pallavi Dani, Emily Stark, and Anne Thomas, *Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups*, Groups Geom. Dyn. **12** (2018), no. 4, 1273–1341. MR 3874643 3

[Dye90] Matthew Dyer, *Reflection subgroups of Coxeter systems*, J. Algebra **135** (1990), no. 1, 57–73. MR 1076077 2, 46, 47

[Hag08] Frédéric Haglund, *Finite index subgroups of graph products*, Geom. Dedicata **135** (2008), 167–209. MR 2413337 2, 3, 4, 35

[HST17] G. Christopher Hruska, Emily Stark, and Hung Cong Tran, *Surface group amalgams that (don't) act on 3-manifolds*, arXiv:1705.01361, to appear in American Journal of Mathematics, 2017. 3

[HW10] Frédéric Haglund and Daniel T. Wise, *Coxeter groups are virtually special*, Adv. Math. **224** (2010), no. 5, 1890–1903. MR 2646113 1

[Kap96] Ilya Kapovich, *Detecting quasiconvexity: algorithmic aspects*, Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994), DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 25, Amer. Math. Soc., Providence, RI, 1996, pp. 91–99. MR 1364182 3, 4, 49

[KK13] Sang-hyun Kim and Thomas Koberda, *Embedability between right-angled Artin groups*, Geom. Topol. **17** (2013), no. 1, 493–530. MR 3039768 4

[KM02] Ilya Kapovich and Alexei Myasnikov, *Stallings foldings and subgroups of free groups*, J. Algebra **248** (2002), no. 2, 608–668. MR 1882114 1

[KMW17] Olga Kharlampovich, Alexei Miasnikov, and Pascal Weil, *Stallings graphs for quasiconvex subgroups*, J. Algebra **488** (2017), 442–483. MR 3680926 2, 3, 4, 49

[MT09] Michael Mihalik and Steven Tschantz, *Visual decompositions of Coxeter groups*, Groups Geom. Dyn. **3** (2009), no. 1, 173–198. MR 2466022 6

[Rad03] David G. Radcliffe, *Rigidity of graph products of groups*, Algebr. Geom. Topol. **3** (2003), 1079–1088. MR 2012965 48

[Sch03] Paul E. Schupp, *Coxeter groups, 2-completion, perimeter reduction and subgroup separability*, Geom. Dedicata **96** (2003), 179–198. MR 1956839 4

[Sta83] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR 695906 1

[SW15] Michah Sageev and Daniel T. Wise, *Cores for quasiconvex actions*, Proc. Amer. Math. Soc. **143** (2015), no. 7, 2731–2741. MR 3336599 4

[Weh69] B. A. F. Wehrfritz, *Infinite linear groups*, Queen Mary College Mathematical Notes, Queen Mary College, Department of Pure Mathematics, London, 1969. MR 837425 35

[Wis11] Daniel T. Wise, *The structure of groups with a quasiconvex hierarchy*, 2011. 1, 8, 9

[Wis12] ———, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012. MR 2986461 7, 8