NON-QUASICONVEX SUBGROUPS OF HYPERBOLIC GROUPS VIA STALLINGS-LIKE TECHNIQUES

PALLAVI DANI AND IVAN LEVCOVITZ

ABSTRACT. We provide a new method of constructing non-quasiconvex subgroups of hyperbolic groups by utilizing techniques inspired by Stallings' foldings. The hyperbolic groups constructed are in the natural class of right-angled Coxeter groups (RACGs for short) and can be chosen to be 2-dimensional. More specifically, given a non-quasiconvex subgroup of a (possibly non-hyperbolic) RACG, our construction gives a corresponding non-quasiconvex subgroup of a hyperbolic RACG. We use this to construct explicit examples of non-quasiconvex subgroups of hyperbolic RACGs including subgroups whose generators are as short as possible (length two words), finitely generated free subgroups, non-finitely presentable subgroups, and subgroups of fundamental groups of square complexes of nonpositive sectional curvature.

Mathematics Subject Classification: 20F65, 57M07.

1. INTRODUCTION

Given a group G with generating set T, a subgroup H of G is *quasiconvex* (with respect to T) if some finite neighborhood of H in the Cayley graph C of (G, T) contains every geodesic that connects a pair of vertices in $H \subset C$. If G is hyperbolic, then whether or not H is quasiconvex does not depend on the finite generating set chosen. In this article, we provide a simple construction of finitely generated non-quasiconvex subgroups of hyperbolic groups using techniques inspired by Stallings' foldings. The hyperbolic groups we construct are additionally in the natural class of right-angled Coxeter groups and can be chosen to admit a geometric action on a CAT(0) square complex. This adds to the limited list of known techniques for constructing non-quasiconvex subgroups of hyperbolic groups.

Given a simplicial graph Γ with vertex set V and edge set E, the corresponding *right-angled Coxeter group* (RACG for short) W_Γ is the group with presentation:

$$\langle s \in V \mid s^2 = 1 \text{ for all } s \in V, st = ts \text{ for all } (s, t) \in E \rangle$$

We refer to the generators in this presentation as the *standard generating set* for W_Γ, and quasiconvexity of subgroups of W_Γ will always be assumed to be with respect to this generating set. We say that W_Γ is *2-dimensional* if Γ does not contain a 3-cycle. Such RACGs act geometrically on a CAT(0) cube complex that is at most 2-dimensional (their Davis complex).

Let W_Γ be a RACG and $H < W_\Gamma$ a subgroup generated by a finite set of words T in W_Γ. Let R be the "rose graph" associated to T, i.e., a bouquet of $|T|$ circles in which each circle is subdivided and labeled by an element of T. In [DL21], we define a completion Ω of H as the direct limit of a sequence of fold, cube attachment and cube identification operations performed on R, and we show that many properties of $H < W_\Gamma$ are reflected by those of Ω. Notably, H is quasiconvex in W_Γ if and only if Ω is finite.

Given any simplicial graph Γ, we define in this article the notion of Γ-partite graphs (see Section 3). These graphs have a large-scale structure which reflects the graphical structure

of Γ. We show, via an explicit construction, that given any such Γ, there are infinitely many Γ-partite graphs Δ, such that W_Δ is hyperbolic and 2-dimensional (see Theorem 3.1).

The importance of Γ-partite graphs in our setting is explained by the following construction. Given a Γ-partite graph Δ, and a subgroup $H < W_\Gamma$ generated by a finite set of words T in W_Γ, we define a corresponding finite set of words $\overline{T}$ in W_Δ which generates a subgroup $\overline{H} < W_\Delta$. Moreover, fold and cube attachment operations that are performed on the rose graph associated to T have direct analogues to operations performed on the rose graph associated to $\overline{T}$. In fact, we show there is a finite-to-one map from the 1-skeleton of a completion of $\overline{H}$ to the 1-skeleton of a completion of H that is defined by simply collapsing certain bigons to edges. In particular, we deduce that H is quasiconvex in W_Γ if and only if $\overline{H}$ is quasiconvex in W_Δ. Using this construction, we show:

Main Theorem (Theorem 4.5 and Corollary 4.6)**.** *Given any finitely generated, non-quasiconvex subgroup H of a (possibly non-hyperbolic) RACG W_Γ, there exist infinitely many Γ-partite graphs Δ, such that W_Δ is hyperbolic, 2-dimensional and contains the finitely generated, non-quasiconvex subgroup $\overline{H}$. Moreover, if H is torsion-free (resp. not normal in W_Γ) then $\overline{H}$ is torsion-free (resp. not normal in W_Δ).*

Non-quasiconvex subgroups of *non-hyperbolic* RACGs are plentiful. They can, for instance, be found in RACGs that split as a product of infinite groups. Consequently, our construction allows for many easy, explicit examples of non-quasiconvex subgroups of hyperbolic RACGs.

There are other, but not many, known methods of constructing non-quasiconvex subgroups of hyperbolic groups, and we give a non-extensive summary of such constructions. Many of the known examples come from kernels of homomorphisms: infinite-index, infinite, normal subgroups of hyperbolic groups are known to always be non-quasiconvex [ABC+91]. A classical such example is that of a hyperbolic 3-manifold, fibering over the circle, which contains a normal, non-quasiconvex surface subgroup (see [Thu97]). A celebrated early construction of Rips [Rip82], gives infinitely presented, normal subgroups using small cancellation theory. Another source of such kernels is Bestvina-Brady Morse theory developed in [BB97]. This is utilized, for instance, in [Bra99] to give examples of non-hyperbolic (therefore non-quasiconvex) finitely presented subgroups of hyperbolic groups. Bestvina-Brady Morse theory is also applied to RACGs in [JNW19] giving another source of examples in these groups.

An approach to producing non-quasiconvex examples which are not necessarily normal has been to use HNN extensions. For instance, in [Kap99] this is used to show that any non-elementary, torsion-free, hyperbolic group appears as a non-quasiconvex subgroup of a different hyperbolic group. In a similar vein [Mit98, BBD07, BDR13] use iterated HNN extensions to construct free non-quasiconvex subgroups of hyperbolic groups with varying distortion functions.

1.1. **Applications.** We use our main theorem to construct several explicit examples of non-quasiconvex subgroups in hyperbolic groups with additional desirable properties. In a first such application, we construct hyperbolic, 2-dimensional RACGs, each of which contains a finitely generated, non-quasiconvex subgroup whose generators are all length 2 words in the standard generating set of the RACG (see Example 4.8). As any subgroup of a RACG that is generated by length *one* words is convex (and therefore quasiconvex), this construction gives minimal non-quasiconvex subgroups, in the sense that the generators are as short as possible.

In another application (see Theorem 5.1), we construct hyperbolic, 2-dimensional, one-ended RACGs, each of which contains a finitely generated, non-quasiconvex subgroup F that is free. Furthermore, $F \cap K$ is not normal in K for any finite-index subgroup K of the RACG.

By the latter property, these subgroups cannot be obtained, even virtually, as kernels of any homomorphism (which, as discussed earlier, is a known strategy of producing non-quasiconvex examples).

The completions of our subgroups have an explicit description, being analogues of well-understood, simple completions of subgroups of non-hyperbolic groups. This, together with the fact that the fundamental group of a completion of a torsion-free subgroup is the subgroup itself, provides an elementary proof that the examples from the previous paragraph are free.

In Theorems 6.3 and 6.2 respectively, we demonstrate how our method can be used to explicitly construct hyperbolic groups with subgroups that are not finitely presentable and with non-quasiconvex subgroups containing closed surface subgroups.

Wise defines a notion of sectional curvature for 2-complexes [Wis04] and shows that the fundamental group of a compact Euclidean 2-complex with negative sectional curvature is always locally quasiconvex (i.e., all finitely generated subgroups are quasiconvex) and is consequently coherent (i.e., all finitely generated subgroups are finitely presented). On the other hand, fundamental groups of 2-complexes with nonpositive sectional curvature provide a fringe case where coherence is conjectured [Wis19, Conjecture 12.11] and yet local quasiconvexity can fail. This is witnessed by hyperbolic free-by-cyclic groups, as these groups are non-locally quasiconvex (they contain a normal, infinite, infinite-index subgroup) and can be realized as fundamental groups of 2-complexes of nonpositive sectional curvature by [Wis04, Theorem 11.3]. We give explicit examples (see Theorem 5.3) of this phenomenon by constructing non-locally quasiconvex hyperbolic RACGs that are virtually fundamental groups of right-angled square complexes of nonpositive sectional curvature. These examples are likely to be virtually free-by-cyclic.

2. BACKGROUND

2.1. **Completions.** A *cube complex* is a cell complex whose cells are Euclidean unit cubes (of any dimension). Given a graph Γ, a cube complex is Γ-*labeled* if its edges are labeled by vertices of Γ. All Γ-labeled cube complexes considered in this article have the property that edges on opposite sides of squares have the same label. In particular, given the labels of d edges which are all incident to a common vertex and contained in a common d-cube, the labels of the remaining edges of this d-cube are completely determined. This will often be used without mention.

By a *path* in a cube complex Σ, we mean a simplicial path in the 1-skeleton of Σ. The label of a path is a word $s_1 \cdots s_m$ in $V(\Gamma)$, such that s_i is the label of the ith edge traversed by Σ.

Given a Γ-labeled cube complex Σ, we define three operations, each of which produces a new Γ-labeled cube complex (see [DL21] for additional details):

Fold operation: Let e and f be distinct edges of Σ which are incident to a common vertex and have the same label. A *fold operation* produces a new complex in which e and f are identified to a single edge with the same label as e and f.

Cube identification operation: Given two or more distinct d-cubes in Σ with common boundary, a cube identification operation identifies these cubes.

Cube attachment operation: Let $e_1, \ldots, e_d$, with $d \geq 2$, be edges of Σ all incident to a common vertex, with distinct labels corresponding to a d-clique of Γ. A *cube attachment operation* attaches a labeled d-cube c to Σ by identifying a set of d edges of c, all incident to

a common vertex, to the edges $e_1, \ldots, e_d$. Note that the labels of edges of c are completely determined by the labels of $e_1, \ldots, e_d$.

A Γ-labeled cube complex is *folded* if no fold or cube identification operations can be performed on it. It is *cube-full* if, given any set of $d \geq 2$ edges which are incident to a common vertex and whose labels form d-clique of Γ, these edges are contained in a common d-cube.

Let Σ be a Γ-labeled cube complex. Let $\Omega_0, \Omega_1, \ldots$ be a sequence of cube complexes such that $\Omega_0 = \Sigma$ and, for $i > 0$, the complex Ω_i is obtained from Ω_{i-1} by performing one of the three above operations. There is a natural map $\Omega_i \to \Omega_{i+1}$. Let Ω be the direct limit of this sequence of complexes. We say that Ω is a *completion* of Σ if Ω is folded and cube-full. If Σ is finite, one can always construct a completion of Σ [DL21, Proposition 3.3].

Let H be a finitely generated subgroup of a RACG W_Γ, generated by a set of words $T = \{w_1, \ldots, w_m\}$ in W_Γ. The *rose graph* associated to T is the Γ-labeled graph with a base vertex B and a $|w_i|$-cycle attached to B with label w_i for each $1 \leq i \leq m$. For this article, it suffices to define a *completion of H with respect to T* as a completion Ω of the rose graph associated to T. The image of B in Ω is defined to be the base vertex of Ω.

We say that a word w in $V(\Gamma)$ is *reduced* if it has minimal length out of all possible words in $V(\Gamma)$ representing the same element of W_Γ as w. We summarize the key properties of completions that we will need:

Theorem 2.1. [DL21] *Let H be a subgroup of a RACG W_Γ generated by a finite set T of words in W_Γ. Let Ω be a completion of H with respect to T with basepoint B. Then:*

(1) *The completion Ω is finite if and only if H is quasiconvex in W_Γ with respect to the standard generating set of W_Γ.*

(2) *The group H has torsion if and only if there is some loop in the 1-skeleton of Ω with label a reduced word whose letters are contained in a clique of Γ.*

(3) *If H is torsion-free, then Ω is non-positively curved and has fundamental group isomorphic to H.*

(4) *Any loop in Ω based at B has label an element of H, and any reduced word which is an element of H appears as the label of some loop based at B.*

(5) *Suppose additionally that W_Γ does not split as a product with a finite factor (equivalently, there does not exist a vertex of $V(\Gamma)$ which is adjacent to every other vertex of $V(\Gamma)$). Then H is of finite index in G if and only if a) Ω is finite, and b) for every vertex v of Ω and for every $s \in V(\Gamma)$, there is an edge incident to v in Ω which is labeled by s.*

2.2. **Nonpositive sectional curvature.** We review the notion of nonpositive sectional curvature simplified to our setting. See [Wis04] and [Wis19] for additional background. This material will only be needed in Section 5.

Given a graph Θ with vertex set V and edge set E, we define $\kappa(\Theta) := 2 - |V| + |E|/2$. A *spur* of Θ is an edge containing a vertex of valence one, and Θ is *spurless* if it does not contain any spurs.

Let Σ be a right-angled square complex. Following [Wis04], we say that Σ has nonpositive sectional curvature if given any vertex v of Σ and any subgraph Θ of the link of v that is connected, spurless, and has at least one edge, it follows that $\kappa(\Theta) \leq 0$. Similarly, we say that a simplicial graph Γ has *nonpositive sectional curvature* if all connected, spurless subgraphs Θ of Γ having at least one edge satisfy $\kappa(\Theta) \leq 0$.

If a simplicial graph Γ has nonpositive sectional curvature, then the (finite-index) commutator subgroup of the RACG W_Γ is the fundamental group of a right-angled square complex of nonpositive sectional curvature.
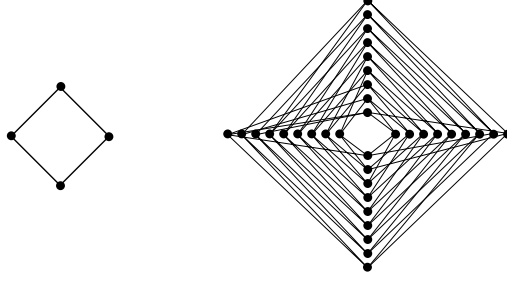
FIGURE 3.1. A graph Γ on the left and, on the right, a Γ-partite graph with cycle connectors, no 3-cycles and no 4-cycles.

3. Γ-PARTITE GRAPHS

Let Γ be a simplicial graph with vertex set $\{s_1, \ldots, s_n\}$. We say that a simplicial graph Δ is Γ-*partite* if its vertex set can be partitioned into n non-empty, disjoint sets $A_1, \ldots, A_n$ such that the following two conditions hold:

(1) No two vertices of A_i are adjacent.
(2) For all $1 \leq i < j \leq n$, the subgraph Δ_{ij} of Δ induced by $A_i \cup A_j$ is connected if s_i is adjacent to s_j in Γ, and is otherwise an edgeless graph, i.e., one with vertices but no edges.

We call $A_1, \ldots, A_n$ the *decomposition* of Δ. We say that Δ has *cycle connectors* if Δ_{ij} is either edgeless or a cycle for every $1 \leq i < j \leq n$. Similarly, we say that Δ has *path connectors* if Δ_{ij} is either edgeless or a simple path for every $1 \leq i < j \leq n$.

The main result of this section is that given any simplicial graph Γ, there are Γ-partite graphs Δ that can be chosen so that W_Δ is hyperbolic and 2-dimensional:

Theorem 3.1. *Given a finite simplicial graph Γ, there exist infinitely many Γ-partite graphs which do not contain any simple 4-cycles or 3-cycles. Additionally, these Γ-partite graphs can be constructed to have either cycle connectors or path connectors.*

Proof. Let Γ have E edges and let $V(\Gamma) = \{s_1, \ldots, s_n\}$. Choose a number $k > 8(3^E)$ which is not divisible by 3. We first construct a Γ-partite graph Δ with decomposition $A_1, \ldots, A_n$, such that $|A_i| = k$ for each $1 \leq i \leq n$, and such that Δ has cycle connectors.

The vertex set of Δ is $A_1 \sqcup \cdots \sqcup A_n$. For each $1 \leq i \leq n$, we fix a labeling $a_0^i, \ldots, a_{k-1}^i$ for the vertices of A_i. To define the edge set of Δ, we arbitrarily label the edges of Γ as $e_1, \ldots, e_E$, and arbitrarily orient each such edge. Fix such an edge e_p, and suppose e_p is incident to the vertices s_x and s_y in Γ, and is oriented from s_x to s_y. Next, we define edges in Δ corresponding to e_p as follows. Each vertex $a_l^x \in A_x$ is adjacent to exactly two vertices of A_y, namely $a_{l+3^p}^y$ and $a_{l+2(3^p)}^y$ (where arithmetic is modulo k). This choice of edges for each edge of Γ defines Δ. The following properties now follow:

(1) A vertex $a_{l'}^y \in A_y$ is adjacent to two vertices of A_x: $a_{l'-3^p}^x$ and $a_{l'-2(3^p)}^x$.
(2) If a_i^x and a_j^x are distinct vertices of A_x which share a neighbor in A_y, then $|i - j| \equiv 3^p$ (mod k). Similarly, if $a_{i'}^y$ and $a_{j'}^y$ are distinct vertices of A_y which share a neighbor in A_x, then $|i' - j'| \equiv 3^p$ (mod k).
(3) As k is not divisible by 3, it follows that 3^p and k are relatively prime. Thus by (2), Δ_{xy} is a simple $2k$-cycle.

It follows that Δ is Γ-partite (by construction) and that Δ has cycle connectors (by (3)).

We now check that there are no simple 4-cycles in Δ. First note that, for each $1 \leq i \leq n$, no simple 4-cycle has all its vertices in A_i, as there are no edges between vertices in A_i. Moreover, no 4-cycle has all its vertices in $A_i \cup A_j$ for some $1 \leq i < j \leq n$, as Δ_{ij} is either a simple $2k$-cycle with $k > 4$ or an edgeless graph.

Suppose there is a simple 4-cycle c with vertices in $A_{i_0} \cup A_{i_1} \cup A_{i_2}$ for some $\{i_0, i_1, i_2\} \subset \{1, \ldots, n\}$. Without loss of generality, we may assume that there are edges e_p, between s_{i_0} and s_{i_1}, and e_q, between s_{i_1} and s_{i_2}, in Γ, and that c contains two vertices $a_u^{i_1}$ and $a_v^{i_1}$ in A_{i_1}. As c is a 4-cycle, $a_u^{i_1}$ and $a_v^{i_1}$ have a common neighbor in A_{i_0} and a common neighbor in A_{i_2}. However, by (2) above, we have that $|u - v| \equiv 3^p \pmod{k}$ and $|u - v| \equiv 3^q \pmod{k}$. This is a contradiction as $p \neq q$ and $k > 8(3^E) > |3^p - 3^q|$. Thus, such a cycle c cannot exist.

Finally, suppose there is a 4-cycle c with vertices in $A_{i_0} \cup A_{i_1} \cup A_{i_2} \cup A_{i_3}$ for some $\{i_0, i_1, i_2, i_3\} \subset \{1, \ldots, n\}$, such that for each $0 \leq j \leq 3$, the cycle c contains exactly one vertex in A_{i_j} and there is an edge e_{p_j} in Γ between s_{i_j} and $s_{i_{j+1}}$ (subscripts taken modulo 4).

To see that such a cycle c cannot exist, we consider an arbitrary length 4 path γ in Δ with startpoint $a_l^{i_0} \in A_{i_0}$, endpoint $a_{l'}^{i_0} \in A_{i_0}$, and exactly one vertex in each of A_{i_1}, A_{i_2} and A_{i_3}. It follows from the definition of edges of Δ and from (1) above that:

$$l' \equiv l + \sum_{m=0}^{3} \epsilon_m 3^{p_m} \pmod{k}$$

where $\epsilon_m \in \{-1, -2, 1, 2\}$ for each $0 \leq m \leq 3$.

We claim that $\sum_{m=0}^{3} \epsilon_m 3^{p_m} \neq 0$. Suppose otherwise, for a contradiction. Let C (resp. C') be the sum of the positive (resp. absolute values of the negative) terms of $\sum_{m=0}^{3} \epsilon_m 3^{p_m}$. It follows that $C = C'$. However, as the numbers p_0, p_1, p_2 and p_3 are all distinct and positive, this implies that we have two ways to represent the number C in base 3. This contradicts the basis representation theorem. We have thus shown that $\sum_{m=0}^{3} \epsilon_m 3^{p_m} \neq 0$.

As $|\sum_{m=0}^{3} \epsilon_m 3^{p_m}| < k$ by our choice of k, and as $\sum_{m=0}^{3} \epsilon_m 3^{p_m} \neq 0$, it follows that $l' \not\equiv l \pmod{k}$. Thus, $a_l^{i_0}$ and $a_{l'}^{i_0}$ are distinct vertices. We have then shown that there cannot be any 4-cycle c as described above.

A similar argument shows that Δ does not contain 3-cycles. This completes the construction of the claimed Γ-partite graph Δ with cycle connectors and no simple 3-cycles or 4-cycles. To produce a Γ-partite graph Δ' with path connectors and no simple 3-cycles or 4-cycles, we can, for each edge (s_i, s_j) of Γ, remove an arbitrary edge from Δ_{ij}. $\qquad\square$

Remark 3.2. The proof of the above theorem gives an explicit, easily implementable construction for Δ, not just an existential one.

Recall that a RACG is not one-ended if and only if its defining graph is either a clique, is disconnected, or is separated by a clique (see [Dav15]). From this and the structure of Γ-partite graphs, we immediately conclude:

Lemma 3.3. *Let Γ be a finite simplicial graph, and let Δ be a Γ-partite graph. If W_Γ is one-ended, then so is W_Δ.*

4. Main Construction

We give our main construction and result (Theorem 4.5 and Corollary 4.6) in this section. Throughout this section, we fix a simplicial graph Γ with vertex set $\{s_1, \ldots, s_n\}$ and a Γ-partite graph Δ with decomposition $A_1, \ldots, A_n$.

Let Π be a Δ-labeled cube complex. Given $1 \leq i \leq n$ and vertices u and v of Π, a *generalized edge between u and v with label A_i* is a collection of $|A_i|$ edges between u and v whose labels are in bijection with the vertices of A_i. Note that u and v could be equal.

Given a Γ-labeled cube complex Σ, we construct an associated Δ-labeled cube complex $\overline{\Sigma}$. We call $\overline{\Sigma}$ the *generalization of Σ with respect to Δ*, and we will always use the bar notation to denote the generalization of a cube complex. When there is no confusion regarding the Γ-partite graph Δ, we simply say that $\overline{\Sigma}$ is the generalization of Σ.

The vertex set of $\overline{\Sigma}$ is in bijective correspondence with that of Σ. The edges of $\overline{\Sigma}$ are defined as follows: for each edge e of Σ there is a corresponding generalized edge with label A_i between the corresponding vertices of $\overline{\Sigma}$, where s_i is the label of e. All edges of $\overline{\Sigma}$ are defined this way.

Note that the graph obtained from the 1-skeleton of $\overline{\Sigma}$ by collapsing generalized edges with label A_i to single edges labeled by s_i, for each $1 \leq i \leq n$, is isomorphic to the 1-skeleton of Σ. We call the resulting map $\overline{\Sigma}^1 \to \Sigma^1$ the *collapsing* map. A *generalized d-cube* in $\overline{\Sigma}$ is a set of 2^d generalized edges whose image under the collapsing map is the 1-skeleton of a d-cube of Σ.

Let c be a d-cube in Σ, with $d \geq 2$, whose edges are labeled by distinct vertices $s_{i_1}, \ldots, s_{i_d}$ of Γ. By construction, there is a corresponding generalized d-cube $\overline{c}$ in $\overline{\Sigma}^1$ with generalized edges labeled by $A_{i_1}, \ldots, A_{i_d}$, such that the image of $\overline{c}$ under the collapsing map is c. For each $(a_1, \ldots, a_d) \in A_{i_1} \times \cdots \times A_{i_d}$ such that $a_1, \ldots, a_d$ span a d-clique in Δ, the complex $\overline{\Sigma}$ contains a single d-cube whose 1-skeleton is the corresponding subset of 2^d edges of $\overline{c}$. Note that even if some other d-cube in Σ has the same boundary as c, there is only one d-cube in $\overline{\Sigma}$ for each such tuple. All d-cubes of $\overline{\Sigma}$, with $d \geq 2$, are defined this way. This completes the construction of $\overline{\Sigma}$. By construction, no cube identification operations can be performed on $\overline{\Sigma}$.

The next lemma is the key observation that allows us to prove our main theorem.

Lemma 4.1. *Let Σ be a finite Γ-labeled cube complex, and let $\overline{\Sigma}$ be the Δ-labeled cube complex which is its generalization. Let Ω be a completion of Σ. Then there is a completion of $\overline{\Sigma}$ that is isomorphic to the generalization $\overline{\Omega}$ of Ω.*

Proof. We first establish three claims, the first of which is obvious by construction.

Claim 4.2. If a cube identification operation can be performed on Σ to obtain a new complex Θ, then $\overline{\Sigma}$ is isomorphic to $\overline{\Theta}$.

Claim 4.3. If a fold operation can be performed on Σ to obtain a new complex Θ, then a sequence of fold and cube identification operations can be performed on $\overline{\Sigma}$ to obtain $\overline{\Theta}$.

Proof. Suppose that there are edges e and f in Σ, each incident to a common vertex v and each labeled by the same $s_i \in V(\Gamma)$. Then there are corresponding generalized edges $\overline{e}$ and $\overline{f}$ in $\overline{\Sigma}$, each incident to a common vertex and each labeled by A_i. By performing $|A_i|$ folds, we identify $\overline{e}$ and $\overline{f}$ into a single generalized edge labeled by A_i. Next, we perform all possible cube identification operations to the resulting complex. This final resulting complex is then isomorphic to $\overline{\Theta}$. $\square$

Claim 4.4. If a cube attachment operation can be performed on Σ to obtain a new complex Θ, then a sequence of cube attachment, fold and cube identification operations can be performed on $\overline{\Sigma}$ to obtain $\overline{\Theta}$.

Proof. Suppose first that such a cube attachment operation attaches a d-cube to Σ with $d = 2$. It follows that there are edges e and f in Σ, incident to a common vertex and with labels s_i and s_j respectively, such that s_i and s_j are adjacent vertices of Γ. In $\overline{\Sigma}$ we see corresponding generalized edges $\overline{e}$ and $\overline{f}$, incident to a common vertex and labeled by A_i and A_j respectively.

Let x be an edge of Δ incident to some $u \in A_i$ and $v \in A_j$. There is an edge in $\overline{e}$ with label u and an edge in $\overline{f}$ with label v. We perform a square attachment operation to $\overline{\Sigma}$ by attaching a square with boundary label $uvuv$ to $\overline{\Sigma}$ along these two edges. We perform such a square attachment for every edge x between A_i and A_j, and we denote the resulting complex by Π.

Next, we perform all possible fold operations between pairs of new edges that were attached in the process of attaching squares to $\overline{\Sigma}$. Let Π' denote the final resulting complex from this sequence of operations. Note that $\overline{\Sigma}$ is naturally a subcomplex of both Π and Π'.

We claim that Π' is isomorphic to $\overline{\Theta}$. To prove this, it is enough to show that all vertices of $\Pi \setminus \overline{\Sigma}$ are identified to a common vertex of Π'. Thus, suppose that z and z' are vertices of $\Pi \setminus \overline{\Sigma}$ which are each contained in a square that was attached to $\overline{\Sigma}$. The labels of these squares are $abab$ and $a'b'a'b'$ for some $a, a' \in A_i$ and $b, b' \in A_j$. As Δ_{ij} is connected (by the definition of Γ-partite graphs), there is a path $a_1 = a, b_1 = b, a_2, b_2 \ldots, a_{m-1}, b_{m-1}, a_m = a', b_m = b'$ in Δ with $a_l \in A_i$ and $b_l \in B_j$ for all $1 \leq l \leq m$. Correspondingly, there are squares with labels $a_l b_l a_l b_l$ and $b_l a_{l+1} b_l a_{l+1}$ in Π for all $1 \leq l \leq m-1$. After folding, all vertices contained in one of these squares which are not contained in $\overline{\Sigma}$ are identified. Thus, Π' is isomorphic to $\overline{\Theta}$.

Suppose now that a cube attachment operation attaches a d-cube to Σ with $d > 2$. From the previous case, it readily follows that there is a sequence of square attachment and fold operations that can be performed to $\overline{\Sigma}$ to produce a complex whose 2-skeleton is isomorphic to that of $\overline{\Theta}$. It is clear that we can obtain $\overline{\Theta}$ by performing a series of cube attachments, folds, and cube identifications to this complex. $\square$

We are now ready to prove the lemma. Let Ω be a completion of Σ obtained as the direct limit of a sequence $\Omega_0 = \Sigma \to \Omega_1 \to \Omega_2 \to \ldots$ of complexes as in the definition of a completion. By the three claims above, there exists a sequence $\Pi_0 = \overline{\Sigma} \to \Pi_1 \to \Pi_2 \to \ldots$ of complexes such that Π_{i+1} is obtained from Π_i by either a cube attachment, cube identification or fold operation, and such that there exists a subsequence $\Pi_0 = \Pi_{i_0}, \Pi_{i_1}, \Pi_{i_2}, \ldots$ such that $\Pi_{i_j} = \overline{\Omega_j}$. Let Π be the direct limit of $\Pi_0 = \overline{\Sigma} \to \Pi_1 \to \Pi_2 \to \ldots$. It follows from the definitions of direct limit and generalization of a complex that Π is isomorphic to $\overline{\Omega}$.

In order to conclude that $\Pi = \overline{\Omega}$ is indeed a completion of $\overline{\Sigma}$, we must show that it is also folded and cube-full. As Ω is a completion, it is folded and cube-full. Let $\phi : \overline{\Omega}^1 \to \Omega^1$ be the collapsing map. To see that $\overline{\Omega}$ is folded, let e and e' be edges of $\overline{\Omega}$ with the same label and incident to a common vertex. Then, by the definition of a collapsing map, $\phi(e)$ and $\phi(e')$ are incident to a common vertex and have the same label. As Ω is folded, $\phi(e) = \phi(e')$. Thus, e and e' must lie in a common generalized edge. As they have the same label, it follows that $e = e'$. Finally, since no cube identification operations can be performed on $\overline{\Omega}$ by definition, we conclude that $\overline{\Omega}$ is folded.

To see that $\overline{\Omega}$ is cube-full, consider a set $e_1, \ldots, e_d$ of edges in $\overline{\Omega}$ with $d \geq 2$, all incident to a common vertex and with labels $t_1, \ldots, t_d$ corresponding to a d-clique of Δ. For $1 \leq i \leq d$, let A_i be the set in the decomposition of Δ that contains the vertex t_i. As there are no edges between vertices of A_i for any i, it follows that $A_i \neq A_j$ for all $1 \leq i \neq j \leq d$. Thus, by the structure of Γ-partite graphs, the labels of $\phi(e_1), \ldots, \phi(e_d)$ span a d-clique in Γ. As Ω is cube-full, $\phi(e_1), \ldots, \phi(e_d)$ are contained in a common d-cube. Consequently, as $\overline{\Omega}$ is the generalization of Ω, there is a d-cube in $\overline{\Omega}$ containing $e_1, \ldots, e_d$. Thus, $\overline{\Omega}$ is cube-full. $\square$

Let H be a subgroup of W_Γ generated by a finite set T of words in W_Γ, and let R be the rose graph associated to T. Let $\overline{R}$ be the generalization of R with respect to Δ. Let $l_1, \ldots, l_m$ be a set of loops in $\overline{R}$, based at the base vertex, which generate $\pi_1(\overline{R})$, and let $\overline{T} = \{w_1, \ldots, w_m\}$

be the set of labels of these loops. We define the *generalization* $\overline{H}$ *of* H *with respect to* Δ *and* T to be the finitely generated subgroup of W_Δ that is generated by $\overline{T}$. We remark that the isomorphism class of $\overline{H}$ very much depends on T and Δ, but not on $\overline{T}$. When there is no confusion regarding T and Δ, we simply say that $\overline{H}$ is the generalization of H. We will always use the bar notation as such to denote the generalization of a group.

Theorem 4.5. *Let* Γ *be a simplicial graph, and let* Δ *be a* Γ-*partite graph. Let* H *be a subgroup of* W_Γ *generated by a finite set* T *of words in* W_Γ, *and let* $\overline{H} < W_\Delta$ *be its generalization with respect to* Δ *and* T. *Then*

(1) H *is quasiconvex in* W_Γ *if and only if* $\overline{H}$ *is quasiconvex in* W_Δ, *where quasiconvexity is with respect to the standard generating sets of* W_Γ *and* W_Δ *respectively.*
(2) *If* H *is torsion-free, then so is* $\overline{H}$.
(3) *If* H *is not normal in* W_Γ, *then* $\overline{H}$ *is not normal in* W_Δ.
(4) *Suppose additionally that* W_Γ *does not split as a product with a finite factor (equivalently, there does not exist a vertex of* $V(\Gamma)$ *which is adjacent to every other vertex of* $V(\Gamma)$). *Then* H *is a finite index subgroup of* W_Γ *if and only if* $\overline{H}$ *is a finite index subgroup of* W_Δ.

Additionally, if Δ *does not contain induced 4-cycles (resp. 3-cycles), then* W_Δ *is hyperbolic (resp. 2-dimensional).*

Proof. The claims regarding W_Δ being hyperbolic or 2-dimensional follow from a theorem of Moussong [Mou88] and by definition respectively.

Let R be the rose graph associated to T. Let $\overline{R}$ be the generalization of R, and let $\overline{T}$ be the set of labels of a finite set of generators of $\pi_1(\overline{R}, B)$, where B is the base vertex. By definition, the generalization $\overline{H}$ is generated by $\overline{T}$. Let R' be the rose graph associated to $\overline{T}$. It readily follows that $\overline{R}$ can be obtained from R' by a sequence of fold operations. In particular, a completion of $\overline{R}$ is a completion of $\overline{H}$.

Let Ω be a completion of R. By Lemma 4.1, $\overline{\Omega}$ is a completion of $\overline{R}$ and is therefore a completion of $\overline{H}$. As Ω is infinite if and only if $\overline{\Omega}$ is infinite, (1) immediately follows from Theorem 2.1(1).

By Theorem 2.1(2), if $\overline{H}$ has torsion, then there is a loop in $\overline{\Omega}$ whose label is a reduced word w with letters in a clique of Δ. It follows from the definition of a Γ-partite graph that w uses at most one letter from each set in the decomposition of Δ. Applying the collapsing map, one obtains a loop in Ω with label a word w' using letters in a clique of Γ, with each letter being used at most once. It follows that w' is reduced. Again by Theorem 2.1(2), this implies that H has a torsion element. Thus, (2) follows.

We define a map $\phi : V(\Delta) \to V(\Gamma)$ by, for each $1 \le i \le n$ and each $a \in A_i$, setting $\phi(a) = s_i$. Moreover, if two generators $s, t \in V(\Delta)$ of W_Δ commute, then it follows from the definition of a Γ-partite graph that $\phi(s)$ commutes with $\phi(t)$ in W_Γ. Thus, ϕ defines a homomorphism $\phi : W_\Delta \to W_\Gamma$ which is readily seen to be surjective.

To see (3), suppose that H is not normal in W_Γ. Let h be a word representing an element of H and g be a word representing an element of W_Γ so that ghg^{-1} does not represent an element of H. By Theorem 2.1(4) and as $\overline{\Omega}$ is a generalization of Ω, there exists a word $\overline{h}$, representing an element of $\overline{H}$, so that $\phi(\overline{h}) = h$. As ϕ is surjective, there is a word $\overline{g}$ representing an element of W_Δ so that $\phi(\overline{g}) = g$. Let $\overline{w}$ be a reduced word equal to $\overline{g}\overline{h}\overline{g}^{-1}$ in W_Δ. Suppose, for a contradiction, that $\overline{w}$ represents an element of $\overline{H}$. By Theorem 2.1(4), we have that there is a loop in $\overline{\Omega}$ with label $\overline{w}$ which is based at the base vertex. By Lemma 4.1, there is a loop with label $\phi(\overline{w})$ in Ω based at the base vertex. However, again by Theorem 2.1(4), this implies
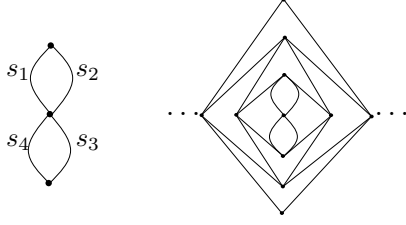
FIGURE 4.1. On the left is the rose graph R of Example 4.8, and on the right is the infinite 1-skeleton of a completion Ω of R. The graph $\overline{R}$ (resp. $\overline{\Omega}^1$) is obtained by replacing edges of R (resp. Ω^1) by generalized edges.

that $\phi(\overline{w}) \in H$. This is a contradiction as $\phi(\overline{w})$ is equal to ghg^{-1} in W_Γ. Thus, $\overline{g}\overline{h}\overline{g}^{-1}$ does not represent an element of $\overline{H}$, and so $\overline{H}$ is not normal in W_Δ.

Claim (4) follows from Theorem 2.1(5). $\qquad \square$

The next corollary follows directly from Theorem 4.5 and Theorem 3.1.

Corollary 4.6. *Let H be a finitely generated non-quasiconvex subgroup of a (possibly non-hyperbolic) RACG W_Γ. Then there exists an infinite set $\{\Delta_i \mid i \in \mathbb{N}\}$ of Γ-partite graphs such that, for each i, W_{Δ_i} is 2-dimensional, hyperbolic and contains the generalization $\overline{H}$ of H (with respect to Δ_i and any finite generating set for H) as a non-quasiconvex subgroup.*

We finish this section by showing how Theorem 4.5 gives easy examples of non-quasiconvex subgroups of hyperbolic groups.

Example 4.7. Let Γ be the graph with vertex set $\{s_1, \ldots, s_{2k}\}$, where $k \geq 2$, and with edges between s_i and s_j for all odd i and even j, i.e., Γ is a complete bipartite graph. Let w be the word $s_1 s_2 \ldots s_{2k}$, and let H be the cyclic subgroup of W_Γ generated by w. The subgroup H can be seen to not be quasiconvex in W_Γ by noting that $W_\Gamma = W_{\{s_1, s_3, \ldots, s_{2k-1}\}} \times W_{\{s_2, s_4, \ldots, s_{2k}\}}$ and that w^n is equal, in W_Γ, to the word $(s_1 s_3 \ldots s_{2k-1})^n (s_2 s_4 \ldots s_{2k})^n$. Alternatively, one can conclude non-quasiconvexity from Theorem 2.1(1) by noting that H has a completion which is an infinite cylinder that is tiled by squares (see for instance Figures 2 and 6 of [DL21]).

By Theorem 3.1, there exists a Γ-partite graph Δ such that W_Δ is 2-dimensional and hyperbolic. Then, by Theorem 4.5, the generalization $\overline{H}$ with respect to Δ and $\{w\}$ is a non-quasiconvex subgroup of W_Δ. Additionally, as H is not normal in W_Γ, Theorem 4.5 implies that $\overline{H}$ is not normal in W_Δ.

The following example gives non-quasiconvex subgroups with generators as short as possible, as discussed in the introduction.

Example 4.8. Let Γ be a 4-cycle with cyclically ordered vertices $\{s_1, \ldots, s_4\}$. Let $w = s_1 s_2$, $w' = s_3 s_4$ and $T = \{w, w'\}$. Note that w and w' are each of order two in W_Γ. The infinite dihedral subgroup $H < W_\Gamma$ generated by T is easily seen to be non-quasiconvex in W_Γ. The rose graph associated to T and the 1-skeleton of a completion for it is shown in Figure 4.1.

By Theorem 3.1, there exists a Γ-partite graph Δ with decomposition $A_1, \ldots, A_4$ such that W_Δ is hyperbolic and 2-dimensional (such as the one in Figure 3.1). Thus, by Theorem 4.5, the generalization $\overline{H}$ of H is a non-quasiconvex subgroup of W_Δ. Moreover, by definition $\overline{H}$ is generated by the labels of all simple loops in $\overline{R}$ based at the base vertex, where R is the rose graph associated to T. Thus, $\overline{H}$ is generated by all words in $V(\Delta)$ of the from aa' with either $\{a, a'\} \subset A_1 \cup A_2$ or $\{a, a'\} \subset A_3 \cup A_4$. In particular, all generators of $\overline{H}$ have length 2.

Examples 4.9 and 4.10 below provide additional examples of non-quasiconvex subgroups of RACGs, which can be used to produce non-quasiconvex subgroups of hyperbolic groups.

Example 4.9. (Products) Given a join graph $\Gamma = \Gamma_1 \star \Gamma_2$, where Γ_1 and Γ_2 are not cliques, the RACG $W_\Gamma = W_{\Gamma_1} \times W_{\Gamma_2}$ contains many non-quasiconvex subgroups. Similarly, let Γ_0 be a 4-cycle, and for integers $i > 0$, let Γ_i be the suspension of Γ_{i-1}. For $i \geq 0$, W_{Γ_i} is virtually isomorphic to $\mathbb{Z}^{i+2}$ and also contains many non-quasiconvex subgroups.

In the following example, W_Γ is not a product. More complicated examples that are not products can be similarly constructed.

Example 4.10. Let Γ be a graph consisting of a 6-cycle with cyclically ordered vertices $s_1, \ldots, s_6$ and an edge between s_1 and s_4. Let $T = \{w = s_1 s_2 s_3 s_4, w' = s_1 s_4 s_5 s_6\}$, and let $H < W_\Gamma$ be generated by T. Let S_1 and S_2 be the 4-cycles of Γ induced by $\{s_1, s_2, s_3, s_4\}$ and $\{s_1, s_4, s_5, s_6\}$ respectively. The group W_Γ is the amalgamation of W_{S_1} and W_{S_2} over a finite group. As $\langle w \rangle$ is not quasiconvex in W_{S_1}, it readily follows that H not quasiconvex in W_Γ. This can also be seen using completions.

5. Non-quasiconvex free subgroups

In this section, we construct non-quasiconvex free subgroups with the additional property that they are not normal in any finite-index subgroup of the parent group (see Theorem 5.1). We also construct non-locally quasiconvex hyperbolic groups with nonpositive sectional curvature (see Theorem 5.3).

The RACG W_Δ in the following theorem can be chosen to be hyperbolic and 2-dimensional by Theorem 3.1.

Theorem 5.1. *As in Example 4.7, let $\Gamma = \{s_1, s_3, \ldots, s_{2k-1}\} \star \{s_2, s_4, \ldots, s_{2k}\}$ be the complete bipartite graph with $k \geq 2$ and let $H < W_\Gamma$ be the cyclic subgroup generated by $s_1 s_2 \ldots s_{2k}$. Let Δ be a Γ-partite graph with path connectors and with decomposition $A_1, \ldots, A_{2k}$ satisfying $|A_i| \geq 2$ for each i. Then the generalization $\overline{H}$ is a finitely generated, non-quasiconvex, non-normal, free subgroup of the one-ended group W_Δ. Moreover, when $k > 2$, then $\overline{H} \cap K$ is not normal in K for any finite index subgroup $K < W_\Delta$.*

Proof. Note that W_Γ is one-ended by Lemma 3.3. Let R denote a $2k$-cycle with label $s_1 s_2 \ldots s_{2k}$. Recall that by definition, $\overline{H}$ is the subgroup of W_Δ generated by the labels of a set of generators of $\pi_1(\overline{R}, B)$, where B is the base vertex. Note also that $\pi_1(\overline{R}, B)$ is free of some rank l.

There exists a completion Ω of R which consists of a tiling of $\mathbb{R} \times S^1$ by squares (see Figures 2 and 6 of [DL21]). By Lemma 4.1, the generalization $\overline{\Omega}$ of Ω is a completion for $\overline{H}$. It readily follows from the description of Ω and Theorems 4.5 and 2.1 that $\overline{H}$ is non-quasiconvex, non-normal and torsion-free. Theorem 2.1(3) then implies that $\overline{H} \cong \pi_1(\overline{\Omega})$. Thus, we show that $\overline{H}$ is free by showing that $\pi_1(\overline{\Omega})$ is a free group of rank l.

We first introduce some new complexes. For $1 \leq i \leq 2k$, let R_i denote the subgraph of $\overline{R}$ consisting of the generalized edges with labels A_i and A_{i+1} (taken modulo $2k$), and let v_i denote the vertex incident to both these generalized edges. Choose one edge in each of these two generalized edges to form a maximal tree, leading to a choice of basis loops for $\pi_1(R_i)$. Let H_i be the subgroup of W_Δ generated by the labels of these loops, which are all infinite order elements of length two. By Theorem 3.17 from [DL20], H_i is a free group with this generating set as a free basis. Let Ω_i denote the generalization of a single 2-cube with boundary label $s_i s_{i+1} s_i s_{i+1}$. It can be seen that Ω_i is a completion for H_i. Moreover, since H_i is free, it is torsion-free, and so $\pi_1(\Omega_i) \cong H_i \cong \pi_1(R_i)$.

It is evident from the description of Ω and the definition of a generalization of a complex, that $\overline{\Omega}$ is a direct limit of a sequence of complexes Σ_i, where $\Sigma_0 = \overline{R}$ and Σ_{j+1} is obtained from Σ_j by gluing a copy of Ω_i for some i to Σ_j along a copy of R_i in each of these.

Now for $j \geq 0$, van Kampen's theorem says that $\pi_1(\Sigma_{j+1})$ is the pushout of the diagram $\pi_1(\Sigma_j) \leftarrow \pi_1(R_i) \rightarrow \pi_1(\Omega_i)$. From what we proved above, the map $\pi_1(R_i) \rightarrow \pi_1(\Omega_i)$ is an isomorphism, and so its inverse provides a natural map $\pi_1(\Omega_i) \rightarrow \pi_1(\Sigma_j)$. By the universal property of pushouts, this map, together with the identity map of $\pi_1(\Sigma_j)$ induces a unique map $\phi : \pi_1(\Sigma_{j+1}) \rightarrow \pi_1(\Sigma_j)$. Moreover, the natural map $\pi_1(\Sigma_j) \rightarrow \pi_1(\Sigma_{j+1})$ is an inverse for ϕ. Thus, $\pi_1(\Sigma_j) \cong \pi_1(\Sigma_{j+1})$. This, together with the fact that $\pi_1(\Sigma_0)$ is a free group of rank l completes the proof that $\overline{H}$ is free of rank l.

To prove the claim regarding normality, choose distinct $a, a' \in A_1$ and $b, b' \in A_3$. Recall that there are no edges between two vertices of A_i, and as s_1 is not adjacent to s_3 in Γ, there are no edges between $\{a, a'\}$ and $\{b, b'\}$. It follows by Tits' solution to the word problem (see [Dav15]) that for every $p, q \in \mathbb{Z} \setminus \{0\}$, the words $(aa')^p$ and $(bb')^q(aa')^p(b'b)^q$ are reduced. In Ω, the base vertex is contained in four distinct edges, labeled by s_1, s_2, s_{2k} and s_{2k-1}. Thus, by Theorem 2.1(4) and the structure of $\overline{\Omega}$, we have that, for every $p, q \in \mathbb{Z} \setminus \{0\}$, the word $(aa')^p$ represents an element of $\overline{H}$ and $(bb')^q(aa')^p(b'b)^q$ does not represent an element of $\overline{H}$. Now, given any finite index subgroup $K < W_\Delta$, there exists some $p, q \in \mathbb{Z} \setminus \{0\}$ such that $(bb')^q \in K$ and $(aa')^p \in K \cap \overline{H}$. The claim follows. $\square$

The next lemma constructs Γ-partite graphs with nonpositive sectional curvature (see Section 2.2 for the definition). These graphs are used in Theorem 5.3 below.

Lemma 5.2. *Let Γ be a 4-cycle, and let Δ be a Γ-partite graph with path connectors. Then Δ has nonpositive sectional curvature.*

Proof. Let $s_1, \ldots, s_4$ be the labels of a cyclic ordering of the vertices of Γ, and let $A_1, \ldots, A_4$ be the corresponding decomposition of Δ. Let Θ be a connected, spurless subgraph of Δ which contains at least one edge. Let V_i be the number of vertices of Θ contained in A_i, and let E_i be the number of edges of Θ connecting a vertex of A_i to a vertex of A_{i+1} (taken modulo 4). We need to check that $\kappa(\Theta) = 2 - \sum_{i=1}^{4} V_i + \frac{1}{2} \sum_{i=1}^{4} E_i$ is at most 0.

As the edges in Δ between A_i and A_{i+1} form a simple path, we have that $E_i \leq 2V_i - 1$ whenever $V_i = V_{i+1} \neq 0$ and $E_i \leq 2\min(V_i, V_{i+1})$ whenever $V_i \neq V_{i+1}$.

Suppose first that $V_i \neq 0$ for all i. We then get:

$$\kappa(\Theta) \leq 2 - \sum_{i=1}^{4} V_i + \sum_{i=1}^{4} \epsilon_i(V_i - \frac{1}{2}) + \sum_{i=1}^{4}(1 - \epsilon_i)\min(V_i, V_{i+1})$$

where $\epsilon_i = 1$ if $V_i = V_{i+1}$ and $\epsilon_i = 0$ otherwise. Regrouping terms and noting that $\epsilon_i(V_i - \min(V_i, V_{i+1})) = 0$ for all i, we get:

$$\kappa(\Theta) \leq 2 - \sum_{i=1}^{4} V_i - \frac{1}{2}\sum_{i=1}^{4} \epsilon_i + \sum_{i=1}^{4}\min(V_i, V_{i+1}) + \sum_{i=1}^{4} \epsilon_i(V_i - \min(V_i, V_{i+1})) = 2 - A - B,$$

where

$$A := \sum_{i=1}^{4} V_i - \sum_{i=1}^{4}\min(V_i, V_{i+1}) \text{ and } B := \sum_{i=1}^{4} \frac{\epsilon_i}{2}.$$

Note that $A \geq 0$, as $V_i - \min(V_i, V_{i+1}) \geq 0$ for any i. We prove below that $A + B \geq 2$. From this, it follows that $\kappa(\Theta) \leq 2 - A - B \leq 0$.

Let M and m respectively be the maximum and minimum of $\{V_i \,|\, 1 \le i \le 4\}$. If $M = m$, then we necessarily have that $\epsilon_i = 1$ for all i. It follows that $B = 2$. Thus, $A + B \ge 2$, since $A \ge 0$. Now suppose $M > m$. Without loss of generality, assume that $V_1 = m$. Then

$$A = m + V_2 + V_3 + V_4 - (m + \min(V_2, V_3) + \min(V_3, V_4) + m).$$

We have $V_i = M$ for some $i \in \{2, 3, 4\}$. In each of these cases, one readily verifies that $A \ge M - m$, (using the observation that $V_i - \min(V_i, V_{i\pm 1}) \ge 0$ for any i). From this, it follows that if $M - m \ge 2$, then $A + B \ge M - m + 0 \ge 2$. Finally, suppose $M - m = 1$, so that $A \ge 1$. Observe that if $\epsilon_i = 1$ for some i, then there exists $j \ne i$ such that $\epsilon_j = 1$. In this case, $B \ge 1$, and once again $A + B \ge 2$. In the last remaining case, namely when $M - m = 1$ and $\epsilon_i = 0$ for all i, we necessarily have $V_2 = V_4 = m + 1$ and $V_3 = m$, and it is easily verified that $A \ge 2$. Therefore, $A + B \ge 2$. We have thus verified that the sectional curvature is at most 0 when $V_i \ne 0$ for all i.

If $V_i = V_j = 0$ for some $i \ne j$, then Θ necessarily contains a spur. This follows as Θ is connected, contains at least one edge and the edges in Δ between A_i and A_{i+1} form a simple path. Thus, it cannot be that $V_i = V_j = 0$ for some $i \ne j$.

The only remaining case to check, up to relabeling, is if $V_4 = 0$ and V_1, V_2 and V_3 are non-zero. Suppose this is the case. Each vertex of A_1 (respectively of A_3), is adjacent to at most two vertices of A_2. As each edge of Θ is incident to a vertex of $A_1 \cup A_3$, we see that the total number of edges of Θ must be at most $2(V_1 + V_3)$. We thus get:

$$\kappa(\Theta) \le 2 - \sum_{i=1}^{3} V_i + (V_1 + V_3)$$

If $V_2 > 1$, this implies that $\kappa(\Theta) \le 0$ and we are done. On the other hand, it cannot be that $V_2 = 1$, for then every vertex of $A_1 \cap \Theta$ must be a spur. This establishes the claim. $\qquad\square$

By applying Theorem 5.1 and Lemma 5.2, we now obtain explicit examples of non-locally quasiconvex, hyperbolic groups with nonpositive sectional curvature.

Theorem 5.3. *Let Γ be a 4-cycle, and let Δ be a Γ-partite graph with path connectors, no 3-cycles and no 4-cycles. Then W_Δ has a finite-index, torsion-free, hyperbolic subgroup K that is the fundamental group of a right-angled square complex with nonpositive sectional curvature. Moreover, K contains a finitely generated, non-quasiconvex, free subgroup. In particular, K is not locally quasiconvex.*

Proof. Let K be the commutator subgroup of W_Δ. Then K is hyperbolic, as it is a finite-index subgroup of the hyperbolic group W_Δ. Morever, K is the fundamental group of a right-angled square complex with nonpositive sectional curvature, as Δ has nonpositive sectional curvature by Lemma 5.2. Let $\overline{H}$ be the finitely generated, non-quasiconvex free subgroup of W_Δ given by Theorem 5.1. Then $\overline{H} \cap K$ is a finitely generated, non-quasiconvex, free subgroup of K. $\quad\square$

6. NON-FINITELY PRESENTABLE, NON-QUASICONVEX SUBGROUPS

In this section we show how Γ-partite graphs with cycle connectors provide examples of finitely generated, non-quasiconvex subgroups of hyperbolic RACGs which contain closed surface subgroups (Theorem 6.2). We also give, using our method, a construction of finitely generated, non-finitely presentable subgroups of hyperbolic RACGs (Theorem 6.3).

Lemma 6.1. *Let C be a cycle with cyclically-ordered vertices $p_1, p_2, \ldots, p_{2k}$, where $k \ge 3$. Let*

$$A = \{p_1 p_i \mid i \text{ is odd}, i > 1\} \text{ and } B = \{p_2 p_i \mid i \text{ is even}, i > 2\}$$

Then the subgroup $K < W_C$ generated by $A \cup B$ is isomorphic to the fundamental group of a closed hyperbolic surface.

Proof. The group W_C is well-known to act geometrically by reflections on the hyperbolic plane (see [Dav15] for instance). Let Γ be a graph with exactly two vertices, t_1 and t_2, and with one edge between them. Note that W_Γ is a finite group. Also, note that C is a Γ-partite graph with decomposition $A_1 = \{s_1, s_3, \ldots, s_{2k-1}\}$ and $A_2 = \{s_2, s_4, \ldots, s_{2k}\}$. Let H be the (trivial) subgroup of W_Γ generated by the set of words $T = \{t_1 t_1, t_2 t_2\}$. It follows that K is the generalization of H with respect to C and T. As H is a torsion-free, finite-index subgroup of W_Γ, it follows by Theorem 4.5 that K is a torsion-free and finite-index subgroup of W_C. Thus, K is isomorphic to the fundamental group of a closed hyperbolic surface. $\square$

Theorem 6.2. *Let Γ be a simplicial graph, Δ be a Γ-partite graph with cycle connectors, and H be a non-quasiconvex subgroup of W_Γ. Then $\overline{H}$ is not quasiconvex in W_Δ and contains a closed surface subgroup.*

Proof. Let Ω be a completion of H. The generalization $\overline{\Omega}$ is a completion of $\overline{H}$ by Lemma 4.1. As H is not quasiconvex in W_Γ, Theorem 2.1 implies that Ω is infinite. In particular, some cube attachment operation must have been performed to construct Ω, and it follows that Ω contains a square c with label $s_i s_j s_i s_j$, where s_i and s_j are adjacent vertices of Γ. We see a corresponding generalized square $\overline{c}$ (whose image under the collapsing map is c) in $\overline{\Omega}$ with label $A_i A_j A_i A_j$.

Let v be a vertex of $\overline{\Omega}$ that is incident to both a generalized edge $\overline{e}$ labeled by A_i and a generalized edge $\overline{f}$ labeled by A_j. Note that Δ_{ij} is a cycle by hypothesis, and the subgroup generated by $A_i \sqcup A_j$ (the set of vertices of Δ_{ij}) is isomorphic to $W_{\Delta_{ij}}$ (see [Dav15]). Let K be the subgroup of $W_{\Delta_{ij}}$ (hence of W_Δ) generated by the labels of all simple loops based at v and contained in $\overline{e} \cup \overline{f}$. Then K is a closed hyperbolic surface subgroup by Lemma 6.1. Let w be the label of a path in $\overline{\Omega}$ from the base vertex to v. By Theorem 2.1(4), $\overline{H}$ contains wKw^{-1}. $\square$

Theorem 6.3. *Let Γ be a 4-cycle, and let Δ be a Γ-partite graph with cycle connectors, no 3-cycles and no 4-cycles. Then W_Δ is 2-dimensional, hyperbolic and contains a finitely generated, non-finitely presentable (and in particular, non-quasiconvex) subgroup.*

Proof. We immediately have that W_Δ is 2-dimensional and hyperbolic as Δ has no 3-cycles and no 4-cycles. Let the vertices of Γ be cyclically labeled $s_1, \ldots, s_4$. Let H be the subgroup of W_Γ generated by the word $s_1 s_2 s_3 s_4$.

Let $\overline{H}$ be the generalization of H. Let K be the commutator subgroup of W_Δ, and let $\overline{H}' := \overline{H} \cap K$. As W_Δ is 2-dimensional, K has cohomological dimension at most 2. We will show that $\overline{H}'$ is a normal, infinite-index, infinite, non-free subgroup of K. It will then follow from a well-known theorem of Bieri [Bie81] (see also [Wis19, Theorem 9.1]) that $\overline{H}'$ is not finitely presentable.

As H is torsion-free and has infinite index in W_Γ, we have by Theorem 4.5 that $\overline{H}$ is torsion-free and has infinite index in W_Δ. In particular, $\overline{H}'$ is infinite and has infinite index in K. By Theorem 6.2, $\overline{H}$ contains a closed surface subgroup S. As K has finite index in W_Δ, it follows that $\overline{H}' = \overline{H} \cap K$ contains a finite-index subgroup of S. As any finite-index subgroup of S is also isomorphic to a closed surface group, it follows that $\overline{H}'$ is not free. It remains to show that $\overline{H}'$ is normal in K.

Let Ω be the completion of H shown in Figure 2 of [DL21]. The generalization $\overline{\Omega}$ of Ω is a completion of $\overline{H}$ by Lemma 4.1. The 1-skeleton of the complex Ω is isomorphic to an

infinite bipartite graph. Moreover, it is readily seen that given two vertices u and v (and corresponding vertices $\overline{u}$ and $\overline{v}$ in $\overline{\Omega}$) in a common part of the bipartition of this graph, there is a label preserving automorphism of Ω (resp. $\overline{\Omega}$) sending u to v (resp. $\overline{u}$ to $\overline{v}$).

To see that $\overline{H}'$ is normal in K, let k be a reduced word representing an element of K and h be a reduced word representing an element of $\overline{H}'$. It is enough to show that $khk^{-1} \in \overline{H}$. Let γ be a path in $\overline{\Omega}$ labeled by khk^{-1} and based at the base vertex b. Such a path exists, since for every $s \in V(\Delta)$ and for every vertex v of $\overline{\Omega}$, there is an edge labeled by s incident to v. Let u be the endpoint of the initial subpath subpath of γ labeled by k. As k has even length, b and u are in the same part of the bipartition of $\overline{\Omega}$ and there is an automorphism of $\overline{\Omega}$ taking b to u. In particular, as there is a loop in $\overline{\Omega}$ based at b with label h (by Theorem 2.1(4)), there is also a loop based at v with label h. From this and as $\overline{\Omega}$ is folded, we conclude that γ is a loop. Thus, again by Theorem 2.1(4), khk^{-1} represents an element of $\overline{H}$. We have thus shown that $\overline{H}'$ is normal in K, concluding our proof. $\square$

REFERENCES

[ABC+91] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, Edited by Short, pp. 3–63.

[BB97] Mladen Bestvina and Noel Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), no. 3, 445–470.

[BBD07] Josh Barnard, Noel Brady, and Pallavi Dani, *Super-exponential distortion of subgroups of CAT(−1) groups*, Algebr. Geom. Topol. **7** (2007), 301–308.

[BDR13] Noel Brady, Will Dison, and Timothy Riley, *Hyperbolic hydra*, Groups Geom. Dyn. **7** (2013), no. 4, 961–976.

[Bie81] Robert Bieri, *Homological dimension of discrete groups*, second ed., Queen Mary College Mathematical Notes, Queen Mary College, Department of Pure Mathematics, London, 1981.

[Bra99] Noel Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, J. London Math. Soc. (2) **60** (1999), no. 2, 461–480.

[Dav15] Michael W. Davis, *The geometry and topology of Coxeter groups*, Introduction to modern mathematics, Adv. Lect. Math. (ALM), vol. 33, Int. Press, Somerville, MA, 2015, pp. 129–142.

[DL20] Pallavi Dani and Ivan Levcovitz, *Right-angled Artin subgroups of right-angled Coxeter and Artin groups*, arXiv:2003.05531, 2020.

[DL21] _____, *Subgroups of right-angled Coxeter groups via Stallings-like techniques*, Journal of Combinatorial Algebra (to appear) (2021).

[JNW19] Kasia Jankiewicz, Sergey Norin, and Daniel Wise, *Virtually fibering right-angled coxeter groups*, J. Inst. Math. Jussieu (2019), Published online.

[Kap99] Ilya Kapovich, *A non-quasiconvexity embedding theorem for hyperbolic groups*, Math. Proc. Cambridge Philos. Soc. **127** (1999), no. 3, 461–486.

[Mit98] M. Mitra, *Cannon–Thurston maps for trees of hyperbolic metric spaces*, J. Differential Geometry **48** (1998), 135–164.

[Mou88] Gabor Moussong, *Hyperbolic Coxeter groups*, ProQuest LLC, Ann Arbor, MI, 1988, Thesis (Ph.D.)– The Ohio State University.

[Rip82] E. Rips, *Subgroups of small cancellation groups*, Bull. London Math. Soc. **14** (1982), no. 1, 45–47.

[Thu97] William P. Thurston, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy.

[Wis04] D. T. Wise, *Sectional curvature, compact cores, and local quasiconvexity*, Geom. Funct. Anal. **14** (2004), no. 2, 433–468.

[Wis19] Daniel Wise, *An invitation to coherent groups*, Princeton University Press, Princeton, NJ, 2019.