

Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products

Motiejus Valiunas¹

Abstract. In this paper we study group actions on quasi-median graphs, or “CAT(0) prism complexes,” generalising the notion of CAT(0) cube complexes. We consider hyperplanes in a quasi-median graph X and define the *contact graph* $\mathcal{C}X$ for these hyperplanes. We show that $\mathcal{C}X$ is always quasi-isometric to a tree, generalising a result of Hagen [18], and that under certain conditions a group action $G \curvearrowright X$ induces an *acylindrical* action $G \curvearrowright \mathcal{C}X$, giving a quasi-median analogue of a result of Behrstock, Hagen and Sisto [5].

As an application, we exhibit an acylindrical action of a graph product on a quasi-tree, generalising results of Kim and Koberda for right-angled Artin groups [20, 21]. We show that for many graph products G , the action we exhibit is the “largest” acylindrical action of G on a hyperbolic metric space. We use this to show that the graph products of equationally noetherian groups over finite graphs of girth ≥ 6 are equationally noetherian, generalising a result of Sela [27].

Mathematics Subject Classification (2020). Primary: 20F65; Secondary: 20F70, 20E06.

Keywords. Acylindrically hyperbolic groups, equationally noetherian groups, graph products, quasi-median graphs.

Contents

1	Introduction	144
2	Preliminaries	149
3	Geometry of the contact graph	154
4	Acylindricity	159
5	Application to graph products	169
6	Equational noetherianity of graph products	174
	References	194

¹ The author was partially supported by EPSRC Scholarship 1807335.

1. Introduction

Group actions on CAT(0) cube complexes occupy a central role in geometric group theory. Such actions have been used to study many interesting classes of groups, such as right-angled Artin and Coxeter groups, many small cancellation and 3-manifold groups, and even finitely presented infinite simple groups, constructed by Burger and Mozes in [9]. Study of CAT(0) cube complexes is aided by their rich combinatorial structure, introduced by Sageev in [26].

In the present paper we study *quasi-median graphs*, which can be viewed as a generalisation of CAT(0) cube complexes; see Definition 2.1. In particular, one may think of quasi-median graphs as “CAT(0) prism complexes,” consisting of *prisms*—cartesian products of (possibly infinite dimensional) simplices—glued together in a non-positively curved way. In his Ph.D. thesis [11], Genevois introduced cubical-like combinatorial structure and geometry to study a wide class of groups acting on quasi-median graphs, including graph products, certain wreath products, and diagram products.

In particular, given a quasi-median graph X , we study *hyperplanes* in X : that is, the equivalence classes of edges of X , under the equivalence relation generated by letting two edges be equivalent if they induce a square or a triangle. Two hyperplanes are said to *intersect* if two edges defining those hyperplanes are adjacent in a square, and *osculate* if two edges defining those hyperplanes are adjacent but do not belong to a square; see Definition 2.3. This allows us to define two other graphs related to X , which turn out to be useful in the study of groups acting on X .

Definition 1.1. Let X be a quasi-median graph. We define the *contact graph* $\mathcal{C}X$ and the *crossing graph* ΔX as follows. For the vertices, let $V(\mathcal{C}X) = V(\Delta X)$ be the set of hyperplanes of X . Two hyperplanes H, H' are then adjacent in ΔX if and only if H and H' intersect; hyperplanes H, H' are adjacent in $\mathcal{C}X$ if and only if H and H' either intersect or osculate.

For a CAT(0) cube complex X , Hagen has shown that the contact graph $\mathcal{C}X$ is a quasi-tree—that is, it is quasi-isometric to a tree [18, Theorem 4.1]. Here we generalise this result to quasi-median graphs.

Theorem A. *For any quasi-median graph X , the contact graph $\mathcal{C}X$ is a quasi-tree, and there exists a $(91, 182)$ -quasi-isometry from a simplicial tree to $\mathcal{C}X$. In particular, there exists a universal constant $\delta > 0$ such that $\mathcal{C}X$ is δ -hyperbolic for any quasi-median graph X .*

We prove Theorem A in Section 3.2.

In this paper we study acylindrical hyperbolicity of groups acting on quasi-median graphs.

Definition 1.2. Suppose a group G acts on a metric space (X, d) by isometries. Such an action is said to be *acylindrical* if for every $\varepsilon > 0$, there exist constants $D_\varepsilon, N_\varepsilon > 0$ such that for all $x, y \in X$ with $d(x, y) \geq D_\varepsilon$, the number of elements $g \in G$ satisfying

$$d(x, x^g) \leq \varepsilon \quad \text{and} \quad d(y, y^g) \leq \varepsilon$$

is bounded above by N_ε . Moreover, an action $G \curvearrowright X$ by isometries on a hyperbolic metric space X is said to be *non-elementary* if orbits under this action are unbounded and G is not virtually cyclic.

A group G is then said to be *acylindrically hyperbolic* if it possesses a non-elementary acylindrical action on a hyperbolic metric space.

Acylindrically hyperbolic groups form a large family, including non-elementary hyperbolic and relatively hyperbolic groups, mapping class groups of most surfaces, and $\text{Out}(F_n)$ for $n \geq 3$ [25]. This family also includes “most” hierarchically hyperbolic groups [5, Corollary 14.4], and in particular “most” groups G that act properly and cocompactly on a CAT(0) cube complex with a “factor system”: see [5]. The following result shows that, more generally, many groups acting on quasi-median graphs are acylindrically hyperbolic.

In the following theorem, we say a group action $G \curvearrowright X$ is *special* if there are no two hyperplanes H, H' of X such that H and H' intersect but H^g and H' osculate for some $g \in G$, and there is no hyperplane H that intersects or osculates with $H^g \neq H$ for some $g \in G$. We say a collection \mathcal{S} of sets is *uniformly finite* if there exists a constant $D \in \mathbb{N}$ such that each $S \in \mathcal{S}$ has cardinality $\leq D$.

Theorem B. *Let G be a group acting specially on a quasi-median graph X , and suppose vertices in $\Delta X/G$ have uniformly finitely many neighbours.*

- (i) *If ΔX is connected and $\Delta X/G$ has finitely many vertices, then the inclusion $\Delta X \hookrightarrow \mathcal{C}X$ is a quasi-isometry.*
- (ii) *If stabilisers of vertices under $G \curvearrowright X$ are uniformly finite, then the induced action $G \curvearrowright \mathcal{C}X$ is acylindrical. In particular, if the orbits under $G \curvearrowright \mathcal{C}X$ are unbounded, then G is either virtually cyclic or acylindrically hyperbolic.*

We prove part (i) of Theorem B in Section 3.1, and part (ii) in Section 4.

Note that a large class of examples of group actions on CAT(0) cube complexes with a factor system comes from special actions [5, Corollaries 8.8 and 14.5]. Theorem B (ii) generalises this result to quasi-median graphs. We also show that several other hierarchically hyperbolic space-like results on CAT(0) cube complexes generalise to quasi-median graphs: for instance, existence of “hierarchy paths,” see [5, Theorem A (2)] and Proposition 3.1.

The main application of Theorems A and B we give is to study graph products of groups. In particular, let Γ be a simplicial graph and let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$

be a collection of non-trivial groups. The *graph product* $\Gamma\mathcal{G}$ of the groups G_v over Γ is defined as the group

$$\Gamma\mathcal{G} = \left(\ast_{v \in V(\Gamma)} G_v \right) / \langle\langle g_v^{-1} g_w^{-1} g_v g_w \mid g_v \in G_v, g_w \in G_w, (v, w) \in E(\Gamma) \rangle\rangle.$$

For example, for a complete graph Γ we have $\Gamma\mathcal{G} \cong \prod_{v \in V(\Gamma)} G_v$, while for discrete Γ we have $\Gamma\mathcal{G} \cong \ast_{v \in V(\Gamma)} G_v$. The applicability of the results above to graph products follows from the following result of Genevois.

Theorem 1.3 (Genevois [11, Propositions 8.2 and 8.11]). *Let Γ be a simplicial graph, let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of non-trivial groups, and let $S = \bigcup_{v \in V(\Gamma)} G_v \setminus \{1\} \subseteq \Gamma\mathcal{G}$. Then the Cayley graph X of $\Gamma\mathcal{G}$ with respect to S is quasi-median. Moreover, the action of $\Gamma\mathcal{G}$ on X is free on vertices and special.*

Moreover, if $\Gamma\mathcal{G}$ and X are as in Theorem 1.3, then, as noted before [11, Lemma 8.8], the hyperplanes of X are of the form H_v^g for some $g \in \Gamma\mathcal{G}$ and for a unique $v \in V(\Gamma)$, where H_v is the hyperplane dual to the clique spanned by $G_v \leq \Gamma\mathcal{G} = V(X)$ (meaning H_v is dual to any edge in that clique). Thus, vertices of $\Delta X / \Gamma\mathcal{G}$ are of the form $H_v^{\Gamma\mathcal{G}}$ for $v \in V(\Gamma)$. It follows from [11, Lemma 8.12] that $H_u^{\Gamma\mathcal{G}}$ and $H_v^{\Gamma\mathcal{G}}$ are adjacent in $\Delta X / \Gamma\mathcal{G}$ if and only if u and v are adjacent in Γ , and so we may deduce the following result.

Lemma 1.4 (Genevois [11]). *The graph $\Delta X / \Gamma\mathcal{G}$ is isomorphic to Γ .*

An important subclass of graph products are right-angled Artin groups (RAAGs): indeed, if $G_v \cong \mathbb{Z}$ then $\Gamma\mathcal{G}$ is the RAAG associated to Γ . In this case, a vertex $v \in V(\Gamma)$ is usually identified with a generator of G_v . In [20] Kim and Koberda constructed the *extension graph* Γ^e of a RAAG $G = \Gamma\mathcal{G}$ as a graph with vertex set $V(\Gamma^e) = \{v^g \in G \mid g \in G, v \in V(\Gamma)\}$, where g^v and h^w are adjacent in Γ^e if and only if they commute as elements of G . This graph turns out to be the same as the crossing graph ΔX of the Cayley graph X defined in Theorem 1.3.

In fact, Kim and Koberda showed that, given that $|V(\Gamma)| \geq 2$ and both Γ and its complement Γ^C are connected, Γ^e is quasi-isometric to a tree [20] and the action of G on Γ^e by conjugation is non-elementary acylindrical [21]. In this paper we generalise these results to arbitrary graph products; this follows as a special case of Theorems A and B. As a special case, we recover hyperbolicity of the extension graph Γ^e and acylindricity of the action $\Gamma\mathcal{G} \curvearrowright \Gamma^e$, providing an alternative (shorter and more geometric) argument to the ones presented in [20, 21]. In the following corollary, a graph Γ is said to have *bounded degree* if there exists a constant $D \in \mathbb{N}$ such that each vertex of Γ has degree $\leq D$.

Corollary C. *Let Γ be a simplicial graph, let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of non-trivial groups, and let X be the quasi-median graph defined in Theorem 1.3. Then $\mathcal{C}X$ is a quasi-tree, and if Γ has bounded degree then the induced action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ is acylindrical. Moreover, if $|V(\Gamma)| \geq 2$ and the complement Γ^C of Γ is connected, then either $\Gamma\mathcal{G} \cong C_2 * C_2$ is the infinite dihedral group, or this action is non-elementary.*

The hyperbolicity of $\mathcal{C}X$ and the acylindricity of the action follow immediately from Theorems A, B, 1.3, and Lemma 1.4, while non-elementarity is shown in Section 5.1.

It is worth noting that Minasyan and Osin have already shown in [24] that if $|V(\Gamma)| \geq 2$ and the complement of Γ is connected, then $\Gamma\mathcal{G}$ is either infinite dihedral or acylindrically hyperbolic. However, their proof is not direct and does not provide an explicit acylindrical action on a hyperbolic space. The aim of Corollary C is to describe such an action.

We also show that in many cases the action of $\Gamma\mathcal{G}$ on $\mathcal{C}X$ is, in the sense of Abbott, Balasubramanya and Osin [1], the “largest” acylindrical action of $\Gamma\mathcal{G}$ on a hyperbolic metric space: see Section 5.2. In particular, we show that many graph products are strongly $\mathcal{A}\mathcal{H}$ -accessible. This generalises the analogous result for right-angled Artin groups [1, Theorem 2.18 (c)].

Corollary D. *Let Γ be a finite simplicial graph and let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of infinite groups. Suppose that for each isolated vertex $v \in V(\Gamma)$, the group G_v is strongly $\mathcal{A}\mathcal{H}$ -accessible. Then $\Gamma\mathcal{G}$ is strongly $\mathcal{A}\mathcal{H}$ -accessible. Furthermore, if Γ has no isolated vertices, then the action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$, where X is as in Theorem 1.3, is the largest acylindrical action of $\Gamma\mathcal{G}$ on a hyperbolic metric space.*

We prove Corollary D in Section 5.2.

Remark 1.5. After the first version of this preprint was made available, it has been brought to the author’s attention that most of the results stated in Corollary C follow from the results in [13, 12, 15]. Moreover, a special case of Corollary D (when the vertex groups G_v are hierarchically hyperbolic) follows from the results in [2, 6]. See Remarks 5.4 and 5.5 for details.

As an application, we use Corollary C to study the class of equationally noetherian groups, defined as follows.

Definition 1.6. Given $n \in \mathbb{N}$, let F_n denote the free group of rank n with a free basis X_1, \dots, X_n . Given a group G , an element $s \in F_n$ and a tuple $(g_1, \dots, g_n) \in G^n$, we write $s(g_1, \dots, g_n) \in G$ for the element obtained by replacing every occurrence

of X_i in s with g_i , and evaluating the resulting word in G . Given a subset $S \subseteq F_n$, the *solution set* of S in G is

$$V_G(S) = \{(g_1, \dots, g_n) \in G^n \mid s(g_1, \dots, g_n) = 1 \text{ for all } s \in S\}.$$

A group G is said to be *equationally noetherian* if for any $n \in \mathbb{N}$ and any subset $S \subseteq F_n$, there exists a *finite* subset $S_0 \subseteq S$ such that $V_G(S_0) = V_G(S)$.

Many classes of groups are known to be equationally noetherian. For example, groups that are linear over a field—in particular, right-angled Artin groups—are equationally noetherian [4, Theorem B1]. It is easy to see that the class of equationally noetherian groups is preserved under taking subgroups and direct products; a deep and non-trivial argument shows that the same is true for free products:

Theorem 1.7 (Sela [27, Theorem 9.1]). *Let G and H be equationally noetherian groups. Then $G * H$ is equationally noetherian.*

Using methods of Groves and Hull developed for acylindrically hyperbolic groups [17], we generalise Theorem 1.7 to a wider class of graph products.

Theorem E. *Let Γ be a finite simplicial graph of girth ≥ 6 , and let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of equationally noetherian groups. Then the graph product $\Gamma\mathcal{G}$ is equationally noetherian.*

We prove Theorem E in Section 6.

The paper is structured as follows. In Section 2, we define quasi-median graphs and give several results that are used in later sections. In Section 3, we analyse the geometry of the contact graph and its relation to crossing graph, and prove Theorem A and Theorem B (i). In Section 4, we consider the action of a group G on a quasi-median graph X , and prove Theorem B (ii). In Section 5, we consider the particular case when $G = \Gamma\mathcal{G}$ is a graph product and X is the quasi-median graph associated to it, and deduce Corollaries C and D. In Section 6, we apply these results to prove Theorem E.

Acknowledgments. I am deeply grateful to Anthony Genevois for his PhD thesis filled with many great ideas, for discussions which inspired the current argument and for his comments on this manuscript. I would also like to thank Jason Behrstock, Daniel Groves, Mark Hagen, Michael Hull, Thomas Koberda, Armando Martino and Ashot Minasyan for valuable discussions. A special thanks goes to the anonymous referee for their insightful comments and for pointing out a gap in the proof of the previous version of Theorem E.

2. Preliminaries

Throughout the paper, we use the following conventions and notation. By a graph X , we mean an undirected simple (simplicial) graph, and we write $V(X)$ and $E(X)$ for the vertex and edge sets of X , respectively. Moreover, we write $d_X(-, -)$ for the combinatorial metric on X —thus, we view X as a geodesic metric space. We consider the set \mathbb{N} of natural numbers to include 0.

Given a group G , all actions of G on a set X are considered to be right actions, $\theta: X \times G \rightarrow X$, and are written as $\theta(x, g) = x^g$ or $\theta(x, g) = xg$. Note that this results in perhaps unusual terminology when we consider a Cayley graph $\text{Cay}(G, S)$: in our case it has edges of the form (g, sg) for $g \in G$ and $s \in S$.

2.1. Quasi-median graphs. In this section we introduce quasi-median graphs and basic results that we use throughout the paper. Most of the definitions and results in this section were introduced by Genevois in his thesis [11]. We therefore refer the interested reader to [11] for further discussion and results on applications of quasi-median graphs to geometric group theory.

Definition 2.1. Let X be a graph, let $x_1, x_2, x_3 \in V(X)$ be three vertices, and let $k \in \mathbb{N}$. We say a triple $(y_1, y_2, y_3) \in V(X)^3$ is a k -quasi-median of (x_1, x_2, x_3) if (see Figure 1a):

- (i) y_i and y_j lie on a geodesic between x_i and x_j for any $i \neq j$;
- (ii) $k = d_X(y_1, y_2) = d_X(y_1, y_3) = d_X(y_2, y_3)$; and
- (iii) k is as small as possible subject to (i) and (ii).

We say $(y_1, y_2, y_3) \in V(X)^3$ is a *quasi-median* of $(x_1, x_2, x_3) \in V(X)^3$ if it is a k -quasi-median for some k . A 0-quasi-median is called a *median*.

We say a graph X is a *quasi-median graph* if (see Figure 1b):

- (i) every triple of vertices has a unique quasi-median;
- (ii) $K_{1,1,2}$ is not isomorphic to an induced subgraph of X ; and
- (iii) if $Y \cong C_6$ is a subgraph of X such that the embedding $Y \hookrightarrow X$ is isometric, then the convex hull of Y in X is isomorphic to the 3-cube.

There are many equivalent characterisations of quasi-median graphs: see [3, Theorem 1]. In this paper we think of quasi-median graphs as generalisations of median graphs. Recall that a graph X is called a *median graph* if every triple of vertices of X has a unique median. In particular, every median graph is quasi-median; more precisely, it is known that a graph is median if and only if it is quasi-median and triangle-free: see [11, Corollary 2.92], for instance.

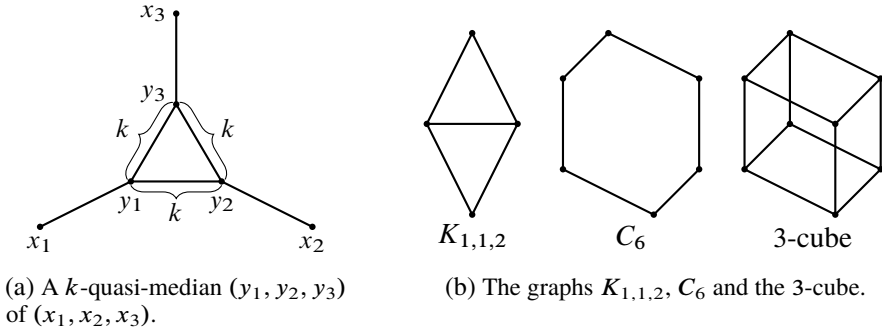


Figure 1. Graphs appearing in Definition 2.1.

Remark 2.2. The class of *quasi-median graphs* is reminiscent of the class of a *bucolic graphs*, introduced in [7], in that both of these generalise median graphs to a class that also includes finite complete graphs. Neither of these two classes includes the other: in essence, bucolic graphs differ from quasi-median ones in that they do not contain infinite complete subgraphs, but may contain an induced subgraph $K_{1,1,2}$ (see Figure 1b); compare [11, Definition 2.1] and [7, Definition 2.11]. We chose to work with quasi-median graphs as they are more suitable for our applications: for instance, the Cayley graph X of a graph product $\Gamma\mathcal{G}$ defined in Theorem 1.3 is quasi-median, but not bucolic unless \mathcal{G} is a collection of *finite* groups.

In what follows, a *clique* is a maximal complete subgraph, a *triangle* is a complete graph on 3 vertices, and a *square* is a complete bipartite graph on two sets of 2 vertices each.

Definition 2.3. Let X be a quasi-median graph. Let \sim be the equivalence relation on $E(X)$ generated by the equivalences $e \sim f$ when e and f either are two sides of a triangle or opposite sides of a square. A *hyperplane* H is an equivalence class $[e]$ for some $e \in E(X)$; in this case, we say H is the hyperplane *dual* to e (or, alternatively, H is the hyperplane dual to any clique containing e). Given a hyperplane H dual to $e \in E(X)$, the *carrier* of H , denoted by $\mathcal{N}(H)$, is the full subgraph of X induced by $[e] \subseteq E(X)$; a *fibre* of H is a connected component of $\mathcal{N}(H) \setminus J$, where J is the union of the interiors of all the edges in $[e]$.

Given two edges $e, e' \in E(X)$ with a common endpoint (p , say) that do not belong to the same clique, let H and H' be the hyperplanes dual to e and e' , respectively. We then say H and H' *intersect* (or *intersect at p*) if e and e' are adjacent edges in a square, and we say H and H' *osculate* (or *osculate at p*) otherwise.

Finally, given two vertices $p, q \in V(X)$ and a hyperplane H , we say H separates p from q if every path between p and q contains an edge dual to H . More generally, we say H separates two subgraphs $P, Q \subseteq X$ if H does not separate any two vertices of P or any two vertices of Q , but it separates a vertex of P from a vertex of Q . Given a path γ in X , we also say H crosses γ if γ contains an edge dual to H .

Another important concept in the study of quasi-median graphs are gated subgraphs. Such subgraphs coincide with *convex subgraphs* for median graphs, but in general form a larger class in quasi-median graphs.

Definition 2.4. Let X be a quasi-median graph, let $Y \subseteq X$ be a full subgraph, and let $v \in V(X)$. We say $p \in V(Y)$ is a *gate* for v in Y if, for any $q \in V(Y)$, there exists a geodesic in X between v and q passing through p . We say a full subgraph $Y \subseteq X$ is a *gated subgraph* if every vertex of X has a gate in Y .

The following result says that the subgraphs of interest to us are gated. Here, by convention, given two graphs Y and Z we denote by $Y \times Z$ the 1-skeleton of the square complex obtained as a cartesian product of Y and Z . The facts that C , $\mathcal{N}(H)$, and F are gated subgraphs follow from Lemmas 2.16, 2.24, and 2.29 (respectively) in [11], while the facts about the isomorphism $\Psi: \mathcal{N}(H) \rightarrow F \times C$ follow from its construction as well as Lemmas 2.28 and 2.29 in [11]. Note that, in our terminology, a “fibre” is a priori different from a “main fibre” defined in [11, Definition 2.27]; however, these two graphs are isomorphic by [11, Lemma 2.29].

Proposition 2.5 (Genevois [11, Section 2.2]). *Let X be a quasi-median graph, H a hyperplane dual to a clique C , and F a fibre of H . Then $\mathcal{N}(H)$, C and F are gated subgraphs of X . Moreover, there exists a graph isomorphism $\Psi: \mathcal{N}(H) \rightarrow F \times C$, and the cliques dual to H (respectively the fibres of H) are precisely the subgraphs $\Psi^{-1}(\{p\} \times C)$ for vertices $p \in V(F)$ (respectively $\Psi^{-1}(F \times \{p\})$ for vertices $p \in V(C)$).*

2.2. Special actions. In this section we describe the hypotheses that we impose on group actions on quasi-median graphs. We first define what it means for an action on a quasi-median graph to be special.

Definition 2.6. Let X be a quasi-median graph, and let G be a group acting on it by graph isomorphisms. We say the action $G \curvearrowright X$ is *special* if

- (i) no two hyperplanes in the same orbit under $G \curvearrowright X$ intersect or osculate; and
- (ii) given two hyperplanes H and H' that intersect, H^g and H' do not osculate for any $g \in G$.

Special actions on CAT(0) cube complexes were introduced by Haglund and Wise in [19]. Notably, there it is shown that, in our terminology, if a group G acts specially, cocompactly and without “orientation-inversions” of hyperplanes on a CAT(0) cube complex X , then the fundamental group of the quotient X/G embeds in a right-angled Artin group.

It is clear from [11, Lemma 2.25] that no hyperplane in a quasi-median graph can self-intersect or self-osculate. The next lemma says that, moreover, the action of the trivial group on a quasi-median graph is special. Recall that two hyperplanes are said to *interosculate* if they both intersect and osculate.

Lemma 2.7. *In a quasi-median graph X , no two hyperplanes can interosculate.*

Proof. Suppose for contradiction that hyperplanes H and H' intersect at p and osculate at q for some $p, q \in V(X)$, and assume without loss of generality that p and q are chosen in such a way that $d_X(p, q)$ is as small as possible. It is clear that $p \neq q$: see, for instance, [11, Lemma 2.13]. On the other hand, since $\mathcal{N}(H)$ and $\mathcal{N}(H')$ are gated (and therefore convex) by Proposition 2.5, and as $p, q \in \mathcal{N}(H) \cap \mathcal{N}(H')$, it follows that a geodesic between p and q lies in $\mathcal{N}(H) \cap \mathcal{N}(H')$. In particular, if r is a vertex on this geodesic, then H and H' either intersect at r or osculate at r ; by minimality of $d_X(p, q)$, it then follows that $d_X(p, q) = 1$.

Let e be the edge joining p and q , and let K be the hyperplane dual to e . It follows from Proposition 2.5 that $K \neq H$ and $K \neq H'$: indeed, if we had $K = H$ (say), then $K = H$ and H' would intersect at q , contradicting the choice of q . Thus K is distinct from H and H' , and so e belongs to a fibre of H and a fibre of H' . It then follows from Proposition 2.5 that K intersects both H and H' at q , and that the graph Y shown in Figure 2 is a subgraph of X .

We now claim that the embedding $Y \hookrightarrow X$ is isometric. Indeed, as H, H' and K are distinct hyperplanes, no two vertices $p', q' \in V(Y)$ with $d_Y(p', q') = 2$ can be joined by an edge in X , as that would create a triangle in X with edges dual to different hyperplanes. It is thus enough to show that if $p', q' \in V(Y)$ and $d_Y(p', q') = 3$, then $d_X(p', q') = 3$. Up to relabelling H, H' and K , we may assume without loss of generality that $p' = s$ and $q' = q$. Now it is clear that $d_X(s, q) \neq 1$: otherwise, p_1s and q_1q are opposite sides in a square in X , contradicting the fact that $H \neq H'$. Thus, suppose for contradiction that $d_X(s, q) = 2$. But then the triple (p_1, s, t) is a quasi-median of (p_1, s, q) for some vertex $t \in V(X)$, and the edges p_1s, p_1t, q_1q are dual to the same hyperplane, again contradicting the fact that $H \neq H'$. Thus the embedding $Y \hookrightarrow X$ is isometric, as claimed.

But now the embedding of the $C_6 \subseteq Y$ formed by vertices s, p_1, q_1, q, q_2 and p_2 into X is also isometric, and so the convex hull of this C_6 in X is a 3-cube. Thus there exists a vertex $u \in V(X)$ joined by edges to s, p_2 and q_2 . This implies that H and H' intersect at q , contradicting the choice of q . Thus H and H' cannot interosculate. \square




Figure 2. Proof of Lemma 2.7: the graph Y (solid edges) and the vertex $t \in V(X)$.

Remark 2.8. We use Lemma 2.7 in the following setting. Let γ be a geodesic in a quasi-median graph X , let e and e' be two consecutive edges of γ , and let H and H' be the hyperplanes dual to e and e' , respectively. Suppose that H and H' intersect. It then follows from Lemma 2.7 that H and H' cannot osculate at the common endpoint p of e and e' , and therefore H and H' must intersect at p . In particular, X contains a square with edges e, e', f and f' , in which f and f' are the edges opposite to e and e' , respectively. We may then obtain another geodesic γ' in X (with the same endpoints as γ) by replacing the subpath ee' of γ with $f'f$. We refer to the operation of replacing γ by γ' as *swapping e and e' on γ* .

2.3. Geodesics in quasi-median graphs. In this section we record two results on geodesics in a quasi-median graph. The first one of these is due to Genevois.

Proposition 2.9 (Genevois [11, Proposition 2.30]). *A path in a quasi-median graph X is a geodesic if and only if it intersects any hyperplane at most once. In particular, the distance between two vertices of X is equal to the number of hyperplanes separating them.* \square

Lemma 2.10. *Let $p, q, r \in V(X)$ be vertices of a quasi-median graph X such that some hyperplane separates q from p and r . Then there exists a hyperplane C separating q from p and r and geodesics γ_p (respectively γ_r) between q and p (respectively q and r) such that q is an endpoint of the edges of γ_p and γ_r dual to C .*

Proof. Let C be a hyperplane separating q from p and r , let γ_p (respectively γ_r) be a geodesic between q and p (respectively q and r), and let c_p and c_r be the edges of γ_p and γ_r (respectively) dual to C . Let q_p and q'_p, q_r and q'_r be the endpoints of c_p, c_r (respectively), labelled so that C does not separate q, q_p and q_r . Suppose, without loss of generality, that γ_p and C are chosen in such a way that $d_X(q, q_p)$ is

as small as possible, and that γ_r is chosen so that $d_X(q, q_r)$ is as small as possible (subject to the choice of γ_p and C). See Figure 3.

We first claim that $q = q_p$. Indeed, suppose not, and let $c'_p \neq c_p$ be the other edge of γ_p with endpoint q_p . Let C'_p be the hyperplane dual to c'_p . Then C'_p does not separate q_p and p (as γ_p is a geodesic), nor q and r (by minimality of $d_X(q, q_p)$), but it separates q_p (and so p) from q (and so r). On the other hand, C separates q_p from p (as γ_p is a geodesic) and q from r (as γ_r is a geodesic). Therefore, C and C'_p must intersect. But then we may swap c_p and c'_p on γ_p (see Remark 2.8), contradicting minimality of $d_X(q, q_p)$. Thus we must have $q = q_p$.

We now claim that $q = q_r$. Indeed, suppose not, and let $c'_r \neq c_r$ be the other edge of γ_r with endpoint q_r . Let C'_r be the hyperplane dual to c'_r . Then C'_r does not separate q and q'_p (as C is the only hyperplane separating $q = q_p$ and q'_p), nor q_r and r (as γ_r is a geodesic), but it separates q (and so q'_p) from q_r (and so r). On the other hand, C separates q_r from r (as γ_r is a geodesic) and q from q'_p . Therefore, C and C'_r must intersect. But then we may swap c_r and c'_r on γ_r , contradicting minimality of $d_X(q, q_r)$. Thus we must have $q = q_r$. \square




Figure 3. Proof of Lemma 2.10.

3. Geometry of the contact graph

Here we analyse the geometry of the contact graph $\mathcal{C}X$ of a quasi-median graph X . In Section 3.1 we show that, under certain conditions, $\mathcal{C}X$ is quasi-isometric to ΔX , and prove Theorem B (i). In Section 3.2 we prove that $\mathcal{C}X$ is a quasi-tree (Theorem A).

3.1. Contact and crossing graphs. The following proposition allows us to lift geodesics in $\mathcal{C}(X)$ back to X . This generalises the existence of “hierarchy paths” in CAT(0) cube complexes [5, Theorem A (2)] to arbitrary quasi-median graphs. Moreover, the same result applies when $\mathcal{C}X$ is replaced by ΔX , as long as ΔX is connected.

Proposition 3.1. *Let $\Gamma = \mathcal{C}X$ or $\Gamma = \Delta X$, and let $A, B \in V(\Gamma)$ be hyperplanes in the same connected component of Γ . Let $p \in V(X)$ (respectively $q \in V(X)$) be a vertex in $\mathcal{N}(A)$ (respectively $\mathcal{N}(B)$). Then there exists a geodesic $A = A_0, \dots, A_m = B$ in Γ and vertices $p = p_0, \dots, p_{m+1} = q \in V(X)$ such that $p_i \in \mathcal{N}(A_{i-1}) \cap \mathcal{N}(A_i)$ for $1 \leq i \leq m$ and $d_X(p, q) = \sum_{i=0}^m d_X(p_i, p_{i+1})$.*

Proof. By assumption, there exists a geodesic $A = A_0, A_1, \dots, A_m = B$ in Γ . For $1 \leq i \leq m$, let $p_i \in V(X)$ be a vertex in the carriers of both A_{i-1} and A_i , and let $p_0 = p, p_{m+1} = q$. Suppose that the A_i and the p_i are chosen in such a way that $D = \sum_{i=0}^m d_X(p_i, p_{i+1})$ is as small as possible. We claim that $D = d_X(p, q)$.

Let γ_i be a geodesic between p_i and p_{i+1} for $0 \leq i \leq m$. Suppose for contradiction that $D > d_X(p, q)$: this means that $\gamma_0 \gamma_1 \cdots \gamma_m$ is not a geodesic. Therefore, there exists a hyperplane C separating p_i and p_{i+1} as well as p_j and p_{j+1} for some $i < j$. Let c_i (respectively c_j) be the edge of γ_i (respectively γ_j) dual to C .

As hyperplane carriers are gated (and therefore convex), any hyperplane separating p_i and p_{i+1} either is or intersects A_i for $0 \leq i \leq m$. Now note that $j - i \leq 2$: indeed, we have $d_\Gamma(A_i, C) \leq 1$ and $d_\Gamma(A_j, C) \leq 1$, so $j - i = d_\Gamma(A_i, A_j) \leq 1 + 1 = 2$. In particular, $j - i \in \{1, 2\}$.

We now claim that $j = i + 1$. Indeed, suppose for contradiction that $j = i + 2$. Let p'_{i+1} (respectively p'_{i+2}) be the endpoint of c_i (respectively c_{i+2}) closer to p_i (respectively p_{i+3}). Then we have

$$\begin{aligned} & d_X(p_i, p_{i+1}) + d_X(p_{i+1}, p_{i+2}) + d_X(p_{i+2}, p_{i+3}) \\ &= d_X(p_i, p'_{i+1}) + d_X(p'_{i+1}, p_{i+1}) + d_X(p_{i+1}, p_{i+2}) \\ & \quad + d_X(p_{i+2}, p'_{i+2}) + d_X(p'_{i+2}, p_{i+3}) \\ & \geq d_X(p_i, p'_{i+1}) + d_X(p'_{i+1}, p'_{i+2}) + d_X(p'_{i+2}, p_{i+3}), \end{aligned} \tag{1}$$

with equality if and only if $\gamma'_i \gamma_{i+1} \gamma'_{i+2}$ is a geodesic, where γ'_i (respectively γ'_{i+2}) is the portion of γ_i (respectively γ_{i+2}) between p'_{i+1} and p_{i+1} (respectively p_{i+2} and p'_{i+2}). But $\gamma'_i \gamma_{i+1} \gamma'_{i+2}$ cannot be a geodesic as it passes through two edges dual to C , and so strict inequality in (1) holds. We may then replace A_{i+1}, p_{i+1} and p_{i+2} with C, p'_{i+1} and p'_{i+2} , respectively, contradicting minimality of D . Thus $j = i + 1$, as claimed.

Therefore, C separates p_{i+1} from p_i and p_{i+2} . By Lemma 2.10, we may assume (after modifying C, γ_i and γ_{i+1} if necessary) that p_{i+1} is an endpoint of both c_i and c_{i+1} . As c_i and c_{i+1} are dual to the same hyperplane, it follows that they belong to the same clique. In particular (as carriers of hyperplanes are gated and so contain their triangles) this whole clique belongs to $\mathcal{N}(A_i) \cap \mathcal{N}(A_{i+1})$. If $r_{i+1} \neq p_{i+1}$ is the other endpoint of c_i , then $d_X(p_i, r_{i+1}) < d_X(p_i, p_{i+1})$ and $d_X(r_{i+1}, p_{i+2}) \leq d_X(p_{i+1}, p_{i+2})$. We may therefore replace p_{i+1} by r_{i+1} , contradicting minimality of D . Thus $D = d_X(p, q)$, as claimed. \square

Taking $\Gamma = \Delta X$ and $p = q$ in Proposition 3.1 immediately gives the following.

Corollary 3.2. *Let $A, B \in V(\Delta X)$ be hyperplanes in the same connected component of ΔX osculating at a point $p \in V(X)$. Then there exists a geodesic $A = A_0, \dots, A_m = B$ in ΔX such that A_{i-1} and A_i intersect at p for $1 \leq i \leq m$. \square*

Lemma 3.3. *Suppose a group G acts on X specially with N orbits of hyperplanes. Let A and B be hyperplanes that osculate and belong to the same connected component of ΔX . Then $d_{\Delta X}(A, B) \leq \max\{2, N - 1\}$.*

Proof. Let $p \in V(X)$ be such that A and B osculate at p . By Corollary 3.2, there exists a geodesic $A = A_0, A_1, \dots, A_m = B$ in ΔX such that A_{i-1} and A_i intersect at p for each i . Let $i_1, \dots, i_k \in \mathbb{N}$, satisfying $0 = i_1 < i_2 < \dots < i_k = m + 1$, be such that $A_{i_j}^G = A_{i_{j+1}-1}^G$ for $1 \leq j \leq k - 1$ (for instance, we may take $i_j = j - 1$). Suppose this is done so that k is as small as possible. Clearly, this implies $A_{i_j}^G \neq A_{i_{j'}}^G$, whenever $1 \leq j < j' \leq k - 1$: otherwise, we may replace i_1, \dots, i_k by $i_1, \dots, i_j, i_{j'+1}, \dots, i_k$, contradicting minimality of k . In particular, $k \leq N + 1$; as $m \geq 2$, note also that $k \geq 2$. We will consider the cases $k = 2$ and $k \geq 3$ separately.

Suppose first that $k \geq 3$. We claim that $i_{j+1} - i_j \leq 1$ whenever $1 \leq j \leq k - 1$. Indeed, note that whenever $1 \leq j \leq k - 2$, $p \in \mathcal{N}(A_{i_j}) \cap \mathcal{N}(A_{i_{j+1}})$, and so A_{i_j} and $A_{i_{j+1}}$ must either intersect or osculate. But $A_{i_{j+1}}$ intersects $A_{i_{j+1}-1}$, and $A_{i_j}^G = A_{i_{j+1}-1}^G$: therefore, as the action $G \curvearrowright X$ is special, it follows that A_{i_j} and $A_{i_{j+1}}$ must intersect. In particular, $i_{j+1} - i_j = d_{\Delta X}(A_{i_j}, A_{i_{j+1}}) \leq 1$ for $1 \leq j \leq k - 2$. For $j = k - 1$, we may similarly note that $\mathcal{N}(A_{i_{j-1}}) \cap \mathcal{N}(A_{i_{j+1}-1}) \neq \emptyset$ and so $A_{i_{j-1}}$ and $A_m = A_{i_{j+1}-1}$ must intersect: thus $i_{j+1} - i_j = d_{\Delta X}(A_{i_{j-1}}, A_{i_{j+1}-1}) \leq 1$ in this case as well. In particular, we get

$$d_{\Delta X}(A, B) = m = i_k - i_1 - 1 = \left(\sum_{j=1}^{k-1} (i_{j+1} - i_j) \right) - 1 \leq k - 2 \leq N - 1,$$

as required.

Suppose now that $k = 2$. Similarly to the case $k \geq 3$, we may note that $p \in \mathcal{N}(A_1) \cap \mathcal{N}(A_m)$, and so, as A_0 and A_1 intersect and as $A_0^G = A_m^G$, it follows that A_1 and A_m intersect. Thus $m - 1 = d_{\Delta X}(A_1, A_m) \leq 1$ and so $d_{\Delta X}(A, B) = m \leq 2$, as required. \square

Proof of Theorem B (i). It is clear that $d_{eX}(A, B) \leq d_{\Delta X}(A, B)$ for any hyperplanes A and B , as ΔX is a subgraph of eX . Conversely, Lemma 3.3 implies that $d_{\Delta X}(A, B) \leq \bar{N}d_{eX}(A, B)$ for any hyperplanes A and B , where $\bar{N} = \max\{2, N - 1\}$. \square

Remark 3.4. We note that all the assumptions for Theorem B (i) are necessary. Indeed, it is clear that ΔX needs to be connected. To show necessity of the other two conditions, consider the following. Let $G_0 = \langle S \mid R \rangle$ be the group with generators $S = \{a_{i,j} \mid (i, j) \in \mathbb{Z}^2\}$ and relators $R = \bigcup_{(i,j) \in \mathbb{Z}^2} \{a_{i,j}^2, [a_{i,j}, a_{i,j+1}], [a_{i,j}, a_{i+1,j}]\}$; this is the (infinitely generated) right-angled Coxeter group associated to a “grid” in \mathbb{R}^2 : a graph Γ with $V(\Gamma) = \mathbb{Z}^2$, where (i, j) and (i', j') are adjacent if and only if $|i - i'| + |j - j'| = 1$. Let X be the Cayley graph of G_0 with respect to S .

Then X is a quasi-median (and, indeed, median) graph by [11, Proposition 8.2]. Furthermore, by the results in [11, Chapter 8], ΔX is connected, and if $H_{i,j}$ is the hyperplane dual to the edge $(1, a_{i,j})$ of X (for $(i, j) \in \mathbb{Z}^2$) then $d_{\mathcal{C}X}(H_{0,0}, H_{i,j}) \leq 1$ but $d_{\Delta X}(H_{0,0}, H_{i,j}) = |i| + |j|$ for all (i, j) . Thus the inclusion $\Delta X \hookrightarrow \mathcal{C}X$ cannot be a quasi-isometry. Moreover, by Theorem A, we know that $\mathcal{C}X$ is a quasi-tree, whereas the inclusion into ΔX of the subgraph spanned by $\{H_{i,j} \mid (i, j) \in \mathbb{Z}^2\}$ (which is isomorphic to the “grid” Γ) is isometric, and so ΔX cannot be hyperbolic—therefore, ΔX and $\mathcal{C}X$ are not quasi-isometric in this case.

It follows from [11, Proposition 8.11] that the usual action of G_0 on X is special—however, there are infinitely many orbits of hyperplanes under this action. On the other hand, let $G = G_0 \rtimes \mathbb{Z}^2$, where the action of $\mathbb{Z}^2 = \langle x, y \mid xy = yx \rangle$ on G_0 is given by $a_{i,j}^{x^n y^m} = a_{i+n, j+m}$; this can be thought of as an example of a *graph-wreath product*, see [22] for details. Then it is easy to see that the action of G on G_0 extends to an action of G on X . This action is transitive on hyperplanes, and therefore not special.

3.2. Hyperbolicity. We show here that $\mathcal{C}X$ is a quasi-tree, proving Theorem A.

Proposition 3.5. *Let $A, B \in V(\mathcal{C}X)$ be two hyperplanes and suppose that $d_{\mathcal{C}X}(A, B) \geq 2$. Then there exists a midpoint M of a geodesic between A and B in $\mathcal{C}X$ and a hyperplane C separating $\mathcal{N}(A)$ and $\mathcal{N}(B)$ such that $d_{\mathcal{C}X}(M, C) \leq 3/2$.*

Proof. By Proposition 2.5, we know that $\mathcal{N}(A)$ and $\mathcal{N}(B)$ are gated. It then follows from [11, Lemma 2.36] that there exist vertices $p \in V(\mathcal{N}(A))$ and $q \in V(\mathcal{N}(B))$ such that any hyperplane separating p from q also separates $\mathcal{N}(A)$ from $\mathcal{N}(B)$. Let $A = A_0, \dots, A_m = B \in V(\mathcal{C}X)$ and $p = p_0, \dots, p_{m+1} = q \in V(X)$ be as given by Proposition 3.1 in the case $\Gamma = \mathcal{C}X$, and let M be the midpoint of the former geodesic. It is clear that $\mathcal{N}(A_i) \cap \mathcal{N}(A_j) = \emptyset$ whenever $|i - j| \geq 2$; in particular, $p_i \neq p_{i+1}$ whenever $1 \leq i \leq m - 1$.

Now let $i = \lfloor m/2 \rfloor \in \{1, \dots, m - 1\}$, and let C be any hyperplane separating p_i and p_{i+1} . By the choice of the p_j , there exists a geodesic between p and q passing through p_i and p_{i+1} : therefore, C separates p and q . Therefore, by the choice

of p and q , C also separates $\mathcal{N}(A)$ from $\mathcal{N}(B)$. Finally, note that as C separates $p_i, p_{i+1} \in \mathcal{N}(A_i)$, we have $d_{\mathcal{C}X}(A_i, C) \leq 1$. Therefore,

$$d_{\mathcal{C}X}(M, C) \leq d_{\mathcal{C}X}(M, A_i) + d_{\mathcal{C}X}(A_i, C) = \left| \frac{m}{2} - i \right| + d_{\mathcal{C}X}(A_i, C) \leq \frac{1}{2} + 1 = \frac{3}{2},$$

as required. \square

Definition 3.6. For a geodesic metric space Y and two points $x, y \in Y$ we say a point $m \in Y$ is a *midpoint* between x and y if $d_Y(m, x) = d_Y(m, y) = \frac{1}{2}d_Y(x, y)$.

Let $D \in \mathbb{N}$. A geodesic metric space Y is said to *satisfy the D -bottleneck criterion* if for any points $x, y \in Y$, there exists a midpoint m between x and y such that any path between x and y passes through a point p such that $d_Y(p, m) \leq D$.

Theorem 3.7 (Manning [23, Theorem 4.6]). *A geodesic metric space Y is a quasi-tree if and only if there exists a constant D such that Y satisfies the D -bottleneck criterion. If this is the case, then there exists a $(26D, 52D)$ -quasi-isometry from a simplicial tree to Y .* \square

Remark 3.8. In [23], the last part of Theorem 3.7 is not explicitly stated, that is, the pair of constants $(26D, 52D)$ is not explicitly computed. However, in [23] Manning constructs a map $\beta: T \rightarrow Y$, where T is a simplicial tree, such that edges in T are mapped to geodesic segments in Y , such that $d_Y(y, \beta(T)) < 20D$ for any $y \in Y$, and such that

$$8Dd_T(v, w) - 16D \leq d_Y(\beta(v), \beta(w)) \leq 26Dd_T(v, w) \quad (2)$$

for any *vertices* $v, w \in V(T)$. It is easy to check, using (2) and the fact that any point in T is within distance $\frac{1}{2}$ from a vertex, that we also have

$$8Dd_T(v, w) - 50D \leq d_Y(\beta(v), \beta(w)) \leq 26Dd_T(v, w) + 52D$$

for any *points* $v, w \in T$, and so β is a $(26D, 52D)$ -quasi-isometry, as required.

Proof of Theorem A. We claim that $\mathcal{C}X$ satisfies the $7/2$ -bottleneck criterion. Thus, let $a, b \in \mathcal{C}X$ be any two points, and let $A, B \in V(\mathcal{C}X)$ be two hyperplanes lying on a geodesic between a and b such that $d_{\mathcal{C}X}(A, a) < 1$ and $d_{\mathcal{C}X}(B, b) < 1$. Note that if M is a midpoint between A and B , then there exists a midpoint m between a and b such that $d_{\mathcal{C}X}(m, M) \leq \frac{1}{2}$.

If $d_{\mathcal{C}X}(a, b) < 7$, then any path between a and b passes through a , and $d_{\mathcal{C}X}(a, m) = d_{\mathcal{C}X}(a, b)/2 < 7/2$ for any midpoint m between a and b , so the $7/2$ -bottleneck criterion is satisfied.

On the other hand, if $d_{\mathcal{C}X}(a, b) \geq 7$, then $d_{\mathcal{C}X}(A, B) > 5$ and so we may define points M and C as given by Proposition 3.5; let also m be a midpoint m between a and b such that $d_{\mathcal{C}X}(m, M) \leq \frac{1}{2}$. Let B_1, \dots, B_n be hyperplanes lying on a path

in $\mathcal{C}X$ between a and b (in order), so that $A = B_r, B_{r+1}, \dots, B_{n+s} = B$ is a path in $\mathcal{C}X$ between A and B , where $r, s \in \{0, 1\}$. Choose vertices $q_1, \dots, q_n \in V(X)$ such that $q_i \in \mathcal{N}(B_{i-1}) \cap \mathcal{N}(B_i)$ for all i . As $q_r \in \mathcal{N}(A)$, $q_{n+s} \in \mathcal{N}(B)$, and as C separates A and B , it follows that C separates q_r and q_{n+s} , and so it separates q_i and q_{i+1} for some i . But as $q_i, q_{i+1} \in \mathcal{N}(B_i)$, it follows that $d_{\mathcal{C}X}(C, B_i) \leq 1$. Then

$$d_{\mathcal{C}X}(M, B_i) \leq d_{\mathcal{C}X}(M, C) + d_{\mathcal{C}X}(C, B_i) \leq \frac{3}{2} + 1 = \frac{5}{2},$$

implying, in particular, that $B_i \notin \{A, B\}$ (as $d_{\mathcal{C}X}(M, A) = d_{\mathcal{C}X}(M, B) = d_{\mathcal{C}X}(A, B)/2 > 5/2$), and so B_i lies on the chosen path between a and b . But then $d_{\mathcal{C}X}(m, B_i) \leq d_{\mathcal{C}X}(m, M) + d_{\mathcal{C}X}(M, B_i) \leq 3 < \frac{7}{2}$, and so $7/2$ -bottleneck criterion is again satisfied, as required.

In particular, Theorem 3.7 implies that $\mathcal{C}X$ is a quasi-tree, and there exists a $(91, 182)$ -quasi-isometry from a simplicial tree T to $\mathcal{C}X$. This implies that there exists a (K, C) -quasi-isometry $\gamma: \mathcal{C}X \rightarrow T$, where $K = 91$ and $C = (3 \cdot 91 + 4) \cdot 182$. The γ -images of edges in a geodesic triangle of $\mathcal{C}X$ are then (K, C) -quasi-geodesics in the tree T , and such quasi-geodesics must be within Hausdorff distance R from geodesics with the same endpoints, where $R = R(K, C)$ is a universal constant: see, for instance, [8, Chapter III.H, Theorem 1.7]. This implies that geodesic triangles in $\mathcal{C}X$ must be δ -slim, where $\delta = K(2R + C)$, proving the last part of the Theorem. \square

4. Acylindricity

In this section we prove Theorem B (ii). To do this, in Section 4.1 we introduce the notion of contact sequences (see Definition 4.2) and show the main technical result we need to prove Theorem B (ii): namely, Proposition 4.3. In Section 4.2 we use this to deduce Theorem B (ii).

Throughout this section, let X be a quasi-median graph.

4.1. Contact sequences. In this subsection, for a gated subgraph $Y \leq X$ and a collection \mathcal{H} of hyperplanes in X , we denote by $Y'_{\mathcal{H}} \subseteq V(X)$ the set of vertices $v \in V(X)$ for which there exists a vertex $p_v \in V(Y)$ such that all hyperplanes separating v from p_v are in \mathcal{H} . Moreover, we write $Y_{\mathcal{H}}$ for the full subgraph of X spanned by $Y'_{\mathcal{H}}$.

Lemma 4.1. *Let $Y \leq X$ be a gated subgraph and let \mathcal{H} be a collection of hyperplanes in X . Then the subgraph $Y_{\mathcal{H}}$ of X is gated.*

Proof. Suppose for contradiction that $Y_{\mathcal{H}}$ is not gated, and let $v \in V(X)$ be a vertex that does not have a gate in $Y_{\mathcal{H}}$. Let p be the gate for v in Y . Let \hat{p} be a vertex of

$Y_{\mathcal{H}}$ on a geodesic between v and p with $d_X(v, \hat{p})$ minimal. By our assumption, \hat{p} is not a gate for v in $Y_{\mathcal{H}}$, and so there exists a vertex $\hat{q} \in V(Y_{\mathcal{H}})$ such that no geodesic between v and \hat{q} passes through \hat{p} . Let q be the gate of \hat{q} in Y . Let $\gamma_p, \gamma_q, \delta, \hat{\delta}, \eta$ be geodesics between \hat{p} and p, \hat{q} and q, p and q, \hat{p} and \hat{q}, v and \hat{p} (respectively), as shown in Figure 4.

Since both η and $\hat{\delta}$ are geodesics, and since $\eta\hat{\delta}$ is not (by the choice of \hat{q}), it follows from Lemma 2.10 that we may assume, without loss of generality, that there exists a hyperplane C and edges c_1, c_2 of $\eta, \hat{\delta}$ (respectively), both of which are dual to C and have \hat{p} as an endpoint. But as p is the gate for v in Y , as $\eta\gamma_p$ is a geodesic by the choice of \hat{p} , and as $q \in Y$, it follows that $\eta\gamma_p\delta$ is a geodesic. Therefore, by Proposition 2.9 H cannot cross $\gamma_p\delta$, and so H does not separate \hat{p} and q . As H separates \hat{p} and \hat{q} , it follows that H separates \hat{q} and q and so crosses γ_q . In particular, since $\hat{q} \in V(Y_{\mathcal{H}})$ and since $q \in V(Y)$ is a gate for \hat{q} in Y , it follows that all hyperplanes crossing γ_q are in \mathcal{H} , and therefore $H \in \mathcal{H}$. But then the endpoint $p' \neq \hat{p}$ of c_1 is separated from $p \in V(Y)$ only by hyperplanes in \mathcal{H} ; this contradicts the choice of \hat{p} . Thus $Y_{\mathcal{H}}$ is gated, as claimed. \square




Figure 4. Proof of Lemma 4.1.

Now let a group G act on a quasi-median graph X . This induces an action of G on the crossing graph ΔX . Let \mathcal{H} be the set of orbits of vertices under $G \curvearrowright \Delta X$ —alternatively, the set of orbits of hyperplanes under $G \curvearrowright X$. We may regard each element of \mathcal{H} as a collection of hyperplanes—thus, for instance, given $\mathcal{H}_0 \subseteq \mathcal{H}$ we may write $\bigcup \mathcal{H}_0$ for the set of all hyperplanes whose orbits are elements of \mathcal{H}_0 .

Let $n \in \mathbb{N}$, and let $\mathcal{H}_1, \dots, \mathcal{H}_n$ be subsets of \mathcal{H} . Pick a vertex (a “basepoint”) $o \in V(X)$, and define the subgraphs $Y_0, \dots, Y_n \subseteq X$ inductively: set $Y_0 = \{o\}$ and $Y_i = (Y_{i-1}) \cup_{\mathcal{H}_i}$ for $1 \leq i \leq n$. By Lemma 4.1, Y_n is a gated subgraph. We denote Y_n as above by $Y(o, \mathcal{H}_1, \dots, \mathcal{H}_n)$, and we denote the gate for $v \in V(X)$ in Y_n by $\mathfrak{g}(v; o, \mathcal{H}_1, \dots, \mathcal{H}_n)$.

Definition 4.2. Let $H, H' \in V(\mathcal{C}X)$, and let $p, p' \in V(X)$ be such that $p \in \mathcal{N}(H)$ and $p' \in \mathcal{N}(H')$. Let $n = d_{\mathcal{C}X}(H, H')$. Given any geodesic $H = H_0, \dots, H_n = H'$ in $\mathcal{C}X$ and vertices $p = p_0, p_1, \dots, p_{n+1} = p' \in V(X)$ such that $p_i, p_{i+1} \in \mathcal{N}(H_i)$ for $0 \leq i \leq n$, we call $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ a *contact sequence* for (H, H', p, p') .

Given a contact sequence $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ for (H, H', p, p') and a vertex $v \in V(X)$, we say $(g_0, \dots, g_n) \in V(X)^{n+1}$ is the v -gate for \mathfrak{S} if g_i is the gate for v in $\mathcal{N}(H_i)$ for $0 \leq i \leq n$. We furthermore denote the tuples $(d_X(p_n, g_n), \dots, d_X(p_0, g_0))$ and $(d_X(p_1, g_0), \dots, d_X(p_{n+1}, g_n))$ by $\mathcal{C}_{/}(\mathfrak{S}, v)$ and $\mathcal{C}_{\backslash}(\mathfrak{S}, v)$, respectively. We say a contact sequence \mathfrak{S} for (H, H', p, p') is v -minimal if for any other contact sequence \mathfrak{S}' for (H, H', p, p') we have either $\mathcal{C}_{/}(\mathfrak{S}, v) \leq \mathcal{C}_{/}(\mathfrak{S}', v)$ or $\mathcal{C}_{\backslash}(\mathfrak{S}, v) \leq \mathcal{C}_{\backslash}(\mathfrak{S}', v)$ in the lexicographical order.

Finally, suppose a group G acts on X . Given a vertex $v \in V(X)$ and a contact sequence $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ for (H, H', p, p') with a v -gate (g_0, \dots, g_n) , we say $(\mathcal{H}_0, \dots, \mathcal{H}_n, \mathcal{H}'_0, \dots, \mathcal{H}'_n)$, where $\mathcal{H}_i, \mathcal{H}'_i \subseteq V(\mathcal{C}X/G)$, is the (v, G) -orbit sequence for \mathfrak{S} if

$$\mathcal{H}_i = \{H^G \mid H \in V(\mathcal{C}X) \text{ separates } p_i \text{ from } g_i\}$$

and

$$\mathcal{H}'_i = \{H^G \mid H \in V(\mathcal{C}X) \text{ separates } p_{i+1} \text{ from } g_i\}$$

for $0 \leq i \leq n$.

It is clear that given any H, H', p and p' as in Definition 4.2, there exists a contact sequence for (H, H', p, p') . As the lexicographical order is a well-ordering of \mathbb{N}^n , it follows that a v -minimal contact sequence exists as well.

Proposition 4.3. *Suppose a group G acts specially on a quasi-median graph X . Let $H, H' \in V(\mathcal{C}X)$, let $p, p' \in V(X)$ be such that $p \in \mathcal{N}(H)$ and $p' \in \mathcal{N}(H')$, and let $v \in V(X)$. Let $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ be a v -minimal contact sequence for (H, H', p, p') with v -gate (g_0, \dots, g_n) and (v, G) -orbit sequence $(\mathcal{H}_0, \dots, \mathcal{H}_n, \mathcal{H}'_0, \dots, \mathcal{H}'_n)$. Write $\mathfrak{g}_i := \mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_i)$ and $\mathfrak{g}'_i := \mathfrak{g}(v; p', \mathcal{H}'_0, \dots, \mathcal{H}'_i)$ for $0 \leq i \leq n$. Then,*

- (i) $\mathfrak{g}_n = \mathfrak{g}'_0$;
- (ii) no hyperplane from $\bigcup \mathcal{H}_i$ osculates with a hyperplane from $\bigcup \mathcal{H}'_j$ whenever $i > j$; and
- (iii) for $1 \leq i \leq n$, the hyperplanes separating \mathfrak{g}_{i-1} from \mathfrak{g}_i (respectively \mathfrak{g}'_i from \mathfrak{g}'_{i-1}) are precisely the hyperplanes separating p_i from g_i (respectively p_i from g_{i-1}).

Proof. Induction on n .

For $n = 0$, we claim that $\mathfrak{g}_0 = g_0$. Indeed, by definition of \mathcal{H}_0 we have $g_0 \in Y(p_0, \mathcal{H}_0)$, and so there exists a geodesic η between p and v passing through g_0 and \mathfrak{g}_0 . Suppose for contradiction $\mathfrak{g}_0 \neq g_0$, let $a \subseteq \eta$ be the edge with endpoint g_0 such that the other endpoint $q_a \neq g_0$ of a satisfies $d_X(v, g_0) > d_X(v, q_a)$, and let A be the hyperplane dual to a ; see Figure 5a. Then $g_0 \in \mathcal{N}(H_0) \cap \mathcal{N}(A)$, and so H_0 and A either coincide, or intersect, or osculate. As A separates p and \mathfrak{g}_0 , we know that A^g separates p and g_0 and so A^g and H_0 either coincide or intersect for some $g \in G$. Thus, as the action $G \curvearrowright X$ is special, it follows that A and H_0 cannot osculate, and therefore they either coincide or intersect. But then we also have $q_a \in \mathcal{N}(H_0)$, contradicting the choice of g_0 . Therefore, $\mathfrak{g}_0 = g_0$, as claimed. A symmetric argument shows that $\mathfrak{g}'_0 = g_0$, and so the conclusion of the proposition is clear.

Suppose now that $n \geq 1$, and let $\hat{\mathfrak{g}}'_i = \mathfrak{g}(v; p_n, \mathcal{H}'_{n-1}, \dots, \mathcal{H}'_i)$ for $0 \leq i \leq n$ (so that $\hat{\mathfrak{g}}'_n = p_n$). Notice that $(H_0, \dots, H_{n-1}, p_0, \dots, p_n)$ is a v -minimal contact sequence for (H, H_{n-1}, p, p_n) . Thus, by the inductive hypothesis we have

$$(i') \quad \mathfrak{g}_{n-1} = \hat{\mathfrak{g}}'_0;$$

(ii') no hyperplane from $\bigcup \mathcal{H}_i$ osculates with a hyperplane from $\bigcup \mathcal{H}'_j$ whenever $n-1 \geq i > j$;

(iii') for $1 \leq i \leq n-1$, the hyperplanes separating \mathfrak{g}_{i-1} from \mathfrak{g}_i (respectively $\hat{\mathfrak{g}}'_i$ from $\hat{\mathfrak{g}}'_{i-1}$) are precisely the hyperplanes separating p_i from g_i (respectively p_i from g_{i-1}).

Moreover, let $\hat{\mathfrak{g}}_i = \mathfrak{g}(v; p_1, \mathcal{H}_1, \dots, \mathcal{H}_i)$ for $0 \leq i \leq n$ (so that $\hat{\mathfrak{g}}_0 = p_1$), and notice that $(H_1, \dots, H_n, p_1, \dots, p_{n+1})$ is a v -minimal contact sequence for (H_1, H', p_1, p') . Thus, by the inductive hypothesis we have

$$(i'') \quad \hat{\mathfrak{g}}_n = \mathfrak{g}'_1;$$

(ii'') no hyperplane from $\bigcup \mathcal{H}_i$ osculates with a hyperplane from $\bigcup \mathcal{H}'_j$ whenever $i > j \geq 1$;

(iii'') for $2 \leq i \leq n$, the hyperplanes separating $\hat{\mathfrak{g}}_{i-1}$ from $\hat{\mathfrak{g}}_i$ (respectively \mathfrak{g}'_i from \mathfrak{g}'_{i-1}) are precisely the hyperplanes separating p_i from g_i (respectively p_i from g_{i-1}).

Finally, the proof of the $n = 0$ case above shows that $\mathfrak{g}(v; p_i, \mathcal{H}_i) = g_i = \mathfrak{g}(v; p_{i+1}, \mathcal{H}'_i)$ for $0 \leq i \leq n$.

Now let $q = \hat{\mathfrak{g}}_{n-1}$ and note that we also have $q = \hat{\mathfrak{g}}'_1$: this is clear if $n = 1$ and follows from the inductive hypothesis if $n \geq 2$. Let $\mathcal{A}, \mathcal{B}, \mathcal{A}', \mathcal{B}' \subseteq V(\mathcal{C}X)$ be the sets of hyperplanes separating q from \mathfrak{g}_{n-1} , q from \mathfrak{g}'_1 , \mathfrak{g}'_1 from \mathfrak{g}_n , \mathfrak{g}_{n-1} from \mathfrak{g}_n , respectively; see Figure 5b. We claim that $\mathcal{A} = \mathcal{A}'$ and $\mathcal{B} = \mathcal{B}'$. We will show this in steps, proving part (ii) of the Proposition along the way.




Figure 5. Proof of Proposition 4.3: general setup.

$\mathcal{A} \cap \mathcal{B} = \emptyset$. Suppose for contradiction that there exists some hyperplane $A \in \mathcal{A} \cap \mathcal{B}$. As $A \in \mathcal{A}$, we know that A separates \hat{g}'_1 from \hat{g}'_0 , and so by (iii') above it also separates p_1 from g_0 : thus $d_{eX}(H_0, A) \leq 1$. Similarly, as $A \in \mathcal{B}$, by (iii'') above we know that A separates p_n from g_n and therefore $d_{eX}(H_n, A) \leq 1$. Hence, $n = d_{eX}(H_0, H_n) \leq 2$, and so either $n = 1$ or $n = 2$.

Let α, β be geodesics between p_1 and g_0 , p_n and g_n , respectively, and let $a \subseteq \alpha$ and $b \subseteq \beta$ be the edges dual to A . As a and b lie on geodesics with endpoint v , we may pick endpoints q_a and q_b of a and b , respectively, such that A does not separate q_a, q_b and v .

Suppose first that $n = 2$: see Figure 6a. Note that in this case H_0, A, H_2 is a geodesic in $\mathcal{C}X$ and that $d_X(p_2, g_2) > d_X(q_b, g_2)$ and $d_X(p_1, g_0) > d_X(q_a, g_0)$. Moreover, since q_a lies on a geodesic between p_1 and g_0 , we have $q_a \in \mathcal{N}(H_0)$; similarly, $q_b \in \mathcal{N}(H_2)$. Furthermore, by the construction we know that $q_a, q_b \in \mathcal{N}(A)$. We may therefore replace p_1, p_2 and H_1 by q_a, q_b and A , respectively, contradicting v -minimality of \mathfrak{S} . Thus $n \neq 2$.

Suppose now that $n = 1$. Then A separates p_1 from both g_0 and g_1 . By Lemma 2.10, we may then without loss of generality assume that p_1 is an endpoint (distinct from q_a and q_b) of both a and b . Now note that both a and b are edges on a geodesic between p_1 and v , so we must have $a = b$, and in particular $q_a = q_b$; see Figure 6b. Since A separates p_1 from both g_0 and g_1 , it intersects or coincides with both H_0 and H_1 , and so $q_a \in \mathcal{N}(H_0) \cap \mathcal{N}(H_1)$. We may therefore replace p_n by q_a ; but we have $d_X(p_1, g_1) > d_X(q_a, g_1)$ and $d_X(p_1, g_0) > d_X(q_a, g_0)$, contradicting v -minimality of \mathfrak{S} . Thus no such hyperplane $A \in \mathcal{A} \cap \mathcal{B}$ can exist and so $\mathcal{A} \cap \mathcal{B} = \emptyset$, as claimed.

$\mathcal{A} \cap \mathcal{B}' = \emptyset$. This is clear, as $g_{n-1} = \hat{g}'_0$ lies on a geodesic between $q = \hat{g}'_1$ and g_n , and so no hyperplane can separate g_{n-1} from both q and g_n .

$\mathcal{A}' \cap \mathcal{B} = \emptyset$. Let \mathcal{C} be the set of hyperplanes separating \mathfrak{g}_n and v . We first claim that $\mathcal{A}' \cap \mathcal{B} = \mathcal{A}' \cap \mathcal{C}$. Indeed, let $B \in \mathcal{A}' \cap \mathcal{B}$. Since $B \in \mathcal{B}$, it separates q and g'_1 ; as $g'_1 = \hat{g}_n$ lies on a geodesic between q and v , B cannot separate g'_1 and v . But as $B \in \mathcal{A}'$, it separates g'_1 and \mathfrak{g}_n , and so B must separate \mathfrak{g}_n and v . Therefore, $B \in \mathcal{A}' \cap \mathcal{C}$. Conversely, let $C \in \mathcal{A}' \cap \mathcal{C}$. Since $C \in \mathcal{C}$, it separates \mathfrak{g}_n and v ; as \mathfrak{g}_n lies on a geodesic between q and v , C cannot separate q and \mathfrak{g}_n . But as $C \in \mathcal{A}'$, it separates g'_1 and \mathfrak{g}_n , and so C must separate q and g'_1 . Therefore, $C \in \mathcal{A}' \cap \mathcal{B}$, and so $\mathcal{A}' \cap \mathcal{B} = \mathcal{A}' \cap \mathcal{C}$, as claimed.

Now suppose for contradiction that there exists a hyperplane $A \in \mathcal{A}' \cap \mathcal{B} = \mathcal{A}' \cap \mathcal{B} \cap \mathcal{C}$. Let γ be a geodesic between \mathfrak{g}_n and v , and let $c \subseteq \gamma$ be the edge dual to A . By Lemma 2.10, we may without loss of generality assume that \mathfrak{g}_n is an endpoint of c : see Figure 6c.

Now let $q_c \neq \mathfrak{g}_n$ be the other endpoint of c . Note that since $A \in \mathcal{B}$, we have $A^G \in \mathcal{H}_n$. Therefore, it follows that q_c is separated from \mathfrak{g}_{n-1} only by hyperplanes in $\bigcup \mathcal{H}_n$; as $d_X(v, \mathfrak{g}_n) > d_X(v, q_c)$, this contradicts the definition of \mathfrak{g}_n . Thus $\mathcal{A} \cap \mathcal{B}' = \emptyset$, as claimed.

$\mathcal{A} \subseteq \mathcal{A}'$ and $\mathcal{B} \subseteq \mathcal{B}'$. As $\mathcal{A} \cap \mathcal{B} = \emptyset = \mathcal{A} \cap \mathcal{B}'$, every hyperplane separating q and \mathfrak{g}_{n-1} does not separate q and g'_1 , nor \mathfrak{g}_{n-1} and \mathfrak{g}_n , thus it separates g'_1 and \mathfrak{g}_n . It follows that $\mathcal{A} \subseteq \mathcal{A}'$. Similarly, as $\mathcal{A} \cap \mathcal{B} = \emptyset = \mathcal{A}' \cap \mathcal{B}$, we get $\mathcal{B} \subseteq \mathcal{B}'$.

Part (ii). By (ii') and (ii'') above, it is enough to show that no hyperplane from $\bigcup \mathcal{H}_n$ osculates with a hyperplane from $\bigcup \mathcal{H}'_0$. Thus, let A (respectively B) be a hyperplane separating p_1 and g_0 (respectively p_n and g_n), so that $A^G \in \mathcal{H}'_0$ and $B^G \in \mathcal{H}_n$. It is now enough to show that A^g and B^h do not osculate for any $g, h \in G$.

But as A separates p_1 from g_0 , we know from (iii') that it also separates $\hat{g}'_1 = q$ from $\hat{g}'_0 = \mathfrak{g}_{n-1}$, that is, $A \in \mathcal{A}$. Similarly, as B separates p_n and g_n , we know from (iii'') that $B \in \mathcal{B}$. But as $\mathcal{A} \cap \mathcal{B} = \emptyset = \mathcal{A} \cap \mathcal{B}'$ and as $\mathcal{B} \subseteq \mathcal{B}'$, it follows that A separates q and g'_1 from \mathfrak{g}_{n-1} and \mathfrak{g}_n , while B separates q from g'_1 and \mathfrak{g}_{n-1} from \mathfrak{g}_n . Therefore, A and B must intersect. But as the action $G \curvearrowright X$ is special, it follows that A^g and B^h do not osculate for any $g, h \in G$. Thus no hyperplane from $\bigcup \mathcal{H}_n$ osculates with a hyperplane from $\bigcup \mathcal{H}'_0$, and so part (ii) holds, as required.

$\mathcal{A}' \cap \mathcal{B}' = \emptyset$. Suppose for contradiction that $A \in \mathcal{A}' \cap \mathcal{B}'$ is a hyperplane. Let α' be a geodesic between g'_1 and \mathfrak{g}_n , let $a \subseteq \alpha'$ be the edge dual to A , and let q_a, q'_a be the endpoints of a so that A does not separate g'_1 and q_a . Suppose, without loss of generality, that α' and A are chosen in such a way that $d_X(g'_1, q_a)$ is as small as possible.

We now claim that $g'_1 = q_a$. Indeed, suppose not, and let $a' \neq a$ be the other edge on α' with endpoint q_a . Let $A' \in \mathcal{A}'$ be the hyperplane dual to a' ; see Figure 6d. Then A' does not separate q and g'_1 (as $\mathcal{A}' \cap \mathcal{B} = \emptyset$), nor \mathfrak{g}_{n-1} and \mathfrak{g}_n

(by minimality of $d_X(\mathfrak{g}'_1, q_a)$), but it separates \mathfrak{g}'_1 (and so q) from \mathfrak{g}_n (and so \mathfrak{g}_{n-1}). In particular, $A' \in \mathcal{A}$, and so $A' \in \bigcup \mathcal{H}'_0$. On the other hand, $A \in \mathcal{B}' \subseteq \bigcup \mathcal{H}_n$, and so A and A' cannot osculate by part (ii). It follows that A and A' must intersect, and therefore we may swap a and a' on α' , contradicting minimality of $d_X(\mathfrak{g}'_1, q_a)$. Thus $\mathfrak{g}'_1 = q_a$, as claimed.

But now q'_a is separated from q just by hyperplanes in $\bigcup \mathcal{H}_n$. Furthermore, A cannot separate \mathfrak{g}_n and v (as \mathfrak{g}_n lies on a geodesic between \mathfrak{g}_{n-1} and v , and as A separates \mathfrak{g}_{n-1} and \mathfrak{g}_n), nor \mathfrak{g}_n and q'_a (as α' is a geodesic), but A separates q'_a (and so \mathfrak{g}_n and v) from \mathfrak{g}'_1 . In particular, $d_X(v, \mathfrak{g}'_1) > d_X(v, q'_a)$, contradicting the fact that $\mathfrak{g}'_1 = \hat{\mathfrak{g}}_n$. Thus $\mathcal{A}' \cap \mathcal{B}' = \emptyset$, as required.

$\mathcal{A} = \mathcal{A}'$ and $\mathcal{B} = \mathcal{B}'$. We have already shown $\mathcal{A} \subseteq \mathcal{A}'$ and $\mathcal{B} \subseteq \mathcal{B}'$. Conversely, as $\mathcal{A}' \cap \mathcal{B} = \emptyset = \mathcal{A}' \cap \mathcal{B}'$, every hyperplane separating \mathfrak{g}'_1 and \mathfrak{g}_n does not separate q and \mathfrak{g}'_1 , nor \mathfrak{g}_{n-1} and \mathfrak{g}_n , thus it separates q and \mathfrak{g}_{n-1} . It follows that $\mathcal{A}' \subseteq \mathcal{A}$ and so $\mathcal{A} = \mathcal{A}'$. Similarly, as $\mathcal{A} \cap \mathcal{B}' = \emptyset = \mathcal{A}' \cap \mathcal{B}'$, we get $\mathcal{B} = \mathcal{B}'$.

Now part (iii) of the Proposition follows immediately. Indeed, given (iii') and (iii''), it is enough to show that the hyperplanes separating \mathfrak{g}_{n-1} from \mathfrak{g}_n (respectively \mathfrak{g}'_1 from \mathfrak{g}'_0) are precisely the hyperplanes separating $\hat{\mathfrak{g}}_{n-1}$ from $\hat{\mathfrak{g}}_n$ (respectively $\hat{\mathfrak{g}}'_1$ from $\hat{\mathfrak{g}}'_0$). But this, and so (iii), follows from the fact that $\mathcal{A} = \mathcal{A}'$ and $\mathcal{B} = \mathcal{B}'$.

Finally, we are left to show part (i). We know that $\mathcal{A}' = \mathcal{A} \subseteq \bigcup \mathcal{H}'_0$, and so $\mathfrak{g}_n \in Y(\mathfrak{g}'_1, \mathcal{H}'_0) \subseteq Y(p', \mathcal{H}'_n, \dots, \mathcal{H}'_0)$. In particular, there exists a geodesic between v and \mathfrak{g}_n passing through $\mathfrak{g}(v; p', \mathcal{H}'_n, \dots, \mathcal{H}'_0) = \mathfrak{g}'_0$. But a symmetric argument can show that there exists a geodesic between v and \mathfrak{g}'_0 passing through \mathfrak{g}_n . Thus $\mathfrak{g}_n = \mathfrak{g}'_0$, proving (i). \square

4.2. Consequences of Proposition 4.3

Corollary 4.4. *Let a group G act specially on a quasi-median graph X . Let $H, H', K, K' \in V(\mathcal{C}X)$, and let $p, p', v, v' \in V(X)$ be such that $p \in \mathcal{N}(H)$, $p' \in \mathcal{N}(H')$, $v \in \mathcal{N}(K)$ and $v' \in \mathcal{N}(K')$. Suppose that $d_{\mathcal{C}X}(H, K) \geq d_{\mathcal{C}X}(H, H') + d_{\mathcal{C}X}(K, K') + 3$.*

If \mathfrak{S} is a v -minimal contact sequence for (H, H', p, p') , then \mathfrak{S} is also v' -minimal. Furthermore, if $(\mathcal{H}_0, \dots, \mathcal{H}_n, \mathcal{H}'_0, \dots, \mathcal{H}'_n)$ is the (v, G) -orbit sequence for \mathfrak{S} , then $\mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_n) = \mathfrak{g}(v'; p', \mathcal{H}_0, \dots, \mathcal{H}_n)$.

Proof. Let $m = d_{\mathcal{C}X}(K, K')$, and let $K = K_0, \dots, K_m = K'$ be a geodesic in $\mathcal{C}X$. For $1 \leq i \leq m$, choose a vertex $v_i \in \mathcal{N}(K_{i-1}) \cap \mathcal{N}(K_i)$; let also $v_0 = v$ and $v_{m+1} = v'$. Let $n = d_{\mathcal{C}X}(H, H')$.

Given a contact sequence $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ for (H, H', p, p') and any $v \in V(X)$, the tuples $\mathcal{C}_{\nearrow}(\mathfrak{S}, v)$ and $\mathcal{C}_{\searrow}(\mathfrak{S}, v)$ only depend on the gates for v in the $\mathcal{N}(H_i)$, $0 \leq i \leq n$. In particular, if for all hyperplanes $A \in V(\mathcal{C}X)$ with




Figure 6. Proof of Proposition 4.3: showing that $\mathcal{A} = \mathcal{A}'$ and $\mathcal{B} = \mathcal{B}'$.

$d_{eX}(H, A) \leq n$ the gates for v and v' in $\mathcal{N}(A)$ coincide, then the set of v -minimal contact sequences for (H, H', p, p') coincides with the set of v' -minimal ones.

Thus, let $A \in \mathcal{V}(eX)$ be a hyperplane with $d_{eX}(H, A) \leq n$, and suppose for contradiction that $g \neq g'$, where g and g' are the gates for v and v' (respectively) in $\mathcal{N}(A)$. Let B be a hyperplane separating g and g' . Since B separates two points in $\mathcal{N}(A)$, we must have $d_{eX}(A, B) \leq 1$, and so $d_{eX}(H, B) \leq n + 1$. On the other hand, as B separates the gates of v and v' in a gated subgraph, B must also separate $v = v_0$ and $v' = v_{m+1}$. Thus B must separate v_i and v_{i+1} for some $i \in \{0, \dots, m\}$. As $v_i, v_{i+1} \in \mathcal{N}(K_i)$, it follows that $d_{eX}(B, K_i) \leq 1$. In particular, $d_{eX}(B, K) \leq d_{eX}(B, K_i) + d_{eX}(K_i, K) \leq i + 1 \leq m + 1$. But then $d_{eX}(H, K) \leq d_{eX}(H, B) + d_{eX}(B, K) \leq n + m + 2$, contradicting our assumption. Thus we must have $g = g'$, and so the set of v -minimal contact sequences for (H, H', p, p') coincides with the set of v' -minimal ones. In particular, \mathfrak{S} is a v' -minimal structural sequence for (H, H', p, p') , and so the conclusion of Proposition 4.3 holds if v is replaced by v' as well.

Now suppose for contradiction that $\mathfrak{g}_n(v) = \mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_n)$ is not equal to $\mathfrak{g}_n(v') = \mathfrak{g}(v'; p, \mathcal{H}_0, \dots, \mathcal{H}_n)$. Let B be a hyperplane separating $\mathfrak{g}_n(v)$ from $\mathfrak{g}_n(v')$. Then B separates gates for v and v' in a gated subgraph, and so as above we get $d_{eX}(B, K) \leq m + 1$. On the other hand, since B separates $\mathfrak{g}_n(v)$ from $\mathfrak{g}_n(v')$, it follows that B separates p from either $\mathfrak{g}_n(v)$ or $\mathfrak{g}_n(v')$: without loss of generality, suppose the former. Then B must separate $\mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_{j-1})$

and $g(v; p, \mathcal{H}_0, \dots, \mathcal{H}_j)$ for some $j \in \{0, \dots, n\}$. By Proposition 4.3 (iii), it then follows that B separates p_j from g_j , and so $d_{\mathcal{C}X}(B, H_j) \leq 1$; in particular, $d_{\mathcal{C}X}(H, B) \leq d_{\mathcal{C}X}(H, H_j) + d_{\mathcal{C}X}(H_j, B) \leq j + 1 \leq n + 1$. Therefore, $d_{\mathcal{C}X}(H, K) \leq d_{\mathcal{C}X}(H, B) + d_{\mathcal{C}X}(B, K) \leq n + m + 2$, again contradicting our assumption. Thus we must have $\mathfrak{g}_n(v) = \mathfrak{g}_n(v')$, as required. \square

Lemma 4.5. *Suppose G acts specially on X . Let $D \in \mathbb{N}$, and suppose every vertex of $\Delta X/G$ has at most D neighbours. If $v, w \in V(X)$, then there exist at most $(D + 1)^2$ hyperplanes $H \in V(\mathcal{C}X)$ such that $w \in \mathcal{N}(H)$ and w is not the gate for v in $\mathcal{N}(H)$.*

Proof. Let $\mathcal{U} \subseteq V(X)$ be the set of vertices $u \in V(X)$ such that $d_X(u, w) = 1$ and $d_X(v, w) = d_X(v, u) + 1$. We claim that $|\mathcal{U}| \leq D + 1$. Indeed, suppose there exist k distinct vertices $u_1, \dots, u_k \in \mathcal{U}$, and let H_i be the hyperplane separating w and u_i for $1 \leq i \leq k$. It is clear that $H_i \neq H_j$ whenever $i \neq j$: indeed, if $H_i = H_j = H$ then by Proposition 2.9 H cannot separate v from either u_i or u_j , and therefore $u_i = u_j$, hence $i = j$. Since $w \in \mathcal{N}(H_i) \cap \mathcal{N}(H_j)$ for every i, j and since the action $G \curvearrowright X$ is special, it also follows that $H_i^G \neq H_j^G$ whenever $i \neq j$.

We now claim that H_i and H_j intersect for every $i \neq j$. Indeed, H_i cannot separate u_i from v (by Proposition 2.9), nor w from u_j (as $H_i \neq H_j$), but it does separate w (and so u_j) from u_i (and so v). On the other hand, a symmetric argument shows that H_j separates w and u_i from u_j and v . Thus H_i and H_j must intersect, as claimed. Therefore, $d_{\Delta X}(H_i, H_j) = 1$ and so, as $H_i^G \neq H_j^G$, we have $d_{\Delta X/G}(H_i^G, H_j^G) = 1$. In particular, $\{H_1^G, \dots, H_k^G\}$ are k vertices of a clique in $\Delta X/G$, and so by our assumption it follows that $k \leq D + 1$. Thus $|\mathcal{U}| \leq D + 1$, as claimed.

Now let $u \in \mathcal{U}$, and let $\mathcal{H} \subseteq V(\mathcal{C}X)$ be the set of hyperplanes $H \in V(\mathcal{C}X)$ such that $u, w \in \mathcal{N}(H)$. It is then enough to show that $|\mathcal{H}| \leq D + 1$. Thus, let $H_1, H_2, \dots, H_k \in \mathcal{H}$ be k distinct hyperplanes, where H_1 is the hyperplane separating u and w . As $w \in \mathcal{N}(H_i) \cap \mathcal{N}(H_j)$ for every i, j and as $G \curvearrowright X$ is special, it is clear that $H_i^G \neq H_j^G$ for any $i \neq j$. Furthermore, it is clear (see, for instance, Proposition 2.5) that H_1 and H_j intersect for every $j \neq 1$. In particular, $d_{\Delta X}(H_1, H_j) = 1$, and so $d_{\Delta X/G}(H_1^G, H_j^G) = 1$. As by assumption H_1^G has at most D neighbours in $\Delta X/G$, it follows that $k \leq D + 1$, and so $|\mathcal{H}| \leq D + 1$, as required. \square

Theorem 4.6. *Suppose a group G acts specially on a quasi-median graph X , and suppose there exists some $D \in \mathbb{N}$ such that $|\text{Stab}_G(w)| \leq D$ for any $w \in V(X)$ and any vertex of $\Delta X/G$ has at most D neighbours. Then the induced action $G \curvearrowright \mathcal{C}X$ is acylindrical, and the acylindricity constants D_ε and N_ε can be expressed as functions of ε and D only.*

Proof. Let $\varepsilon \in \mathbb{N}$. We claim that the acylindricity condition in Definition 1.2 is satisfied for $D_\varepsilon = 2\varepsilon + 6$ and $N_\varepsilon = N^{2(\varepsilon+3)}D/(N-1)^2$, where $N = (D+1)^{2^{D+1}}$.

Indeed, let $h, k \in \Delta X$ be such that $d_{\mathcal{C}X}(h, k) \geq D_\varepsilon$. Let $H, K \in V(\mathcal{C}X)$ be hyperplanes such that $d_{\mathcal{C}X}(H, h) \leq 1/2$ and $d_{\mathcal{C}X}(K, k) \leq 1/2$, and note that we have $d_{\mathcal{C}X}(H, K) \geq D_\varepsilon - 1 = 2\varepsilon + 5$. Let $\mathcal{G}_\varepsilon(h, k) = \{g \in G \mid d_{\mathcal{C}X}(h, h^g) \leq \varepsilon, d_{\mathcal{C}X}(k, k^g) \leq \varepsilon\}$, and note that we have $\mathcal{G}_\varepsilon(h, k) \subseteq \mathcal{G}_{\varepsilon+1}(H, K)$. We thus aim to show that $|\mathcal{G}_{\varepsilon+1}(H, K)| \leq N_\varepsilon$.

Pick vertices $v \in \mathcal{N}(K)$ and $p \in \mathcal{N}(H)$, and an element $g \in \mathcal{G}_{\varepsilon+1}(H, K)$. Let $\mathfrak{S} = (H_0, \dots, H_n, p_0, \dots, p_{n+1})$ be a v -minimal contact sequence for the tuple (H, H^g, p, p^g) with v -gate (g_0, \dots, g_n) and (v, G) -orbit sequence $(\mathcal{H}_0, \dots, \mathcal{H}_n, \mathcal{H}'_0, \dots, \mathcal{H}'_n)$; as $g \in \mathcal{G}_{\varepsilon+1}(H, K)$, we have $n \leq \varepsilon + 1$. For $0 \leq i \leq n$, set $\mathfrak{g}_i = \mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_i)$ and $\mathfrak{g}'_i = \mathfrak{g}(v; p^g, \mathcal{H}'_0, \dots, \mathcal{H}'_i)$; let also $\mathfrak{g}_{-1} = p$ and $\mathfrak{g}'_{n+1} = p^g$.

We first claim that there exist hyperplanes $A_0, \dots, A_n \in V(\mathcal{C}X)$ such that $\mathfrak{g}_{i-1}, \mathfrak{g}_i \in \mathcal{N}(A_i)$ for each i . Indeed, this is clear if $\mathfrak{g}_i = p_{i+1}$ for each i , as in that case we may simply take $A_i = H_i$ for each i . Otherwise, let $k \in \{0, \dots, n\}$ be minimal such that $\mathfrak{g}_k \neq p_{k+1}$, and let A be a hyperplane separating \mathfrak{g}_k and p_{k+1} such that $\mathfrak{g}_k \in \mathcal{N}(A)$. For $0 \leq i \leq k-1$ we may take $A_i = H_i$, while for $k \leq i \leq n$ we can show (by induction on i , say) that $\mathfrak{g}_i \in \mathcal{N}(A)$. Indeed, the base case ($i = k$) is clear by construction; and if $\mathfrak{g}_{i-1} \in \mathcal{N}(A)$ for some $i > k$ and $\mathfrak{g}_{i-1} = q_0, \dots, q_m = \mathfrak{g}_i$ is a geodesic in X , then A cannot osculate with the hyperplane separating q_{j-1} and q_j by Proposition 4.3 (ii) and (iii), and so $q_j \in \mathcal{N}(A)$ by induction on j . Thus we may take $A_i = A$ for $k \leq i \leq n$, so that $\mathfrak{g}_{i-1}, \mathfrak{g}_i \in \mathcal{N}(A_i)$ for each i , as claimed. A symmetric argument shows that there exist hyperplanes $B_n, \dots, B_0 \in V(\mathcal{C}X)$ such that $\mathfrak{g}'_{i+1}, \mathfrak{g}'_i \in \mathcal{N}(B_i)$ for each i .

Now, we pass the sequence $(\mathfrak{g}_{-1}, \dots, \mathfrak{g}_n)$ to a subsequence $(\mathfrak{g}_{k_0}, \dots, \mathfrak{g}_{k_a})$ by removing those \mathfrak{g}_i for which $\mathfrak{g}_{i-1} = \mathfrak{g}_i$. It then follows that $\mathfrak{g}_{k_{i-1}} \neq \mathfrak{g}_{k_i}$ and that $\mathfrak{g}_{k_{i-1}}, \mathfrak{g}_{k_i} \in \mathcal{N}(A_{k_i})$ for $1 \leq i \leq a$, where $a \leq n+1 \leq \varepsilon+2$. Similarly, we may pass the sequence $(\mathfrak{g}'_{n+1}, \dots, \mathfrak{g}'_0)$ to a subsequence $(\mathfrak{g}'_{k'_0}, \dots, \mathfrak{g}'_{k'_b})$ such that $\mathfrak{g}'_{k'_{i-1}} \neq \mathfrak{g}'_{k'_i}$ and that $\mathfrak{g}'_{k'_{i-1}}, \mathfrak{g}'_{k'_i} \in \mathcal{N}(B_{k'_i})$ for $1 \leq i \leq b$, where $b \leq n+1 \leq \varepsilon+2$.

Now as $d_{\mathcal{C}X}(H, H^g) + d_{\mathcal{C}X}(K, K^g) + 3 \leq 2(\varepsilon+1) + 3 = 2\varepsilon + 5 \leq d_{\mathcal{C}X}(H, K)$, it follows from Corollary 4.4 that \mathfrak{S} is also a v^g -minimal contact sequence and that $\mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_n) = \mathfrak{g}(v^g; p, \mathcal{H}_0, \dots, \mathcal{H}_n)$. Therefore, by Proposition 4.3 (i) and the discussion above,

$$\begin{aligned} \mathfrak{g}(v; p, \mathcal{H}_{k_1}, \dots, \mathcal{H}_{k_a}) &= \mathfrak{g}(v; p, \mathcal{H}_0, \dots, \mathcal{H}_n) = \mathfrak{g}(v^g; p, \mathcal{H}_0, \dots, \mathcal{H}_n) \\ &= \mathfrak{g}(v^g; p^g, \mathcal{H}'_n, \dots, \mathcal{H}'_0) = \mathfrak{g}(v; p, \mathcal{H}'_n, \dots, \mathcal{H}'_0)^g \quad (3) \\ &= \mathfrak{g}(v; p, \mathcal{H}'_{k'_1}, \dots, \mathcal{H}'_{k'_b})^g. \end{aligned}$$

As the stabiliser of $\mathfrak{g}(v; p, \mathcal{H}_{k_1}, \dots, \mathcal{H}_{k_a})$ has cardinality $\leq D$, it follows that, given any subsets $\mathcal{H}_{k_1}, \dots, \mathcal{H}_{k_a}, \mathcal{H}'_{k'_1}, \dots, \mathcal{H}'_{k'_b} \subseteq V(\mathcal{C}X/G)$, there are at most D elements $g \in G$ satisfying (3). But as $\mathfrak{g}_{k_{i-1}} \neq \mathfrak{g}_{k_i}$, as \mathfrak{g}_{k_i} lies on a geodesic between $\mathfrak{g}_{k_{i-1}}$ and v , and as $\mathfrak{g}_{k_{i-1}}, \mathfrak{g}_{k_i} \in \mathcal{N}(A_i)$, it follows from Lemma 4.5 that there are at most $(D+1)^2$ possible choices for A_{k_i} (for $1 \leq i \leq a$). Moreover, given a choice

of A_{k_i} , as $\mathcal{H}_{k_i} \subseteq \text{star}_{\Delta X/G}(A_{k_i}^G)$ and by assumption $|\text{star}_{\Delta X/G}(A_{k_i}^G)| \leq D + 1$, there exist at most 2^{D+1} choices for \mathcal{H}_{k_i} . It follows that there exist at most N^a choices for the subsets $\mathcal{H}_{k_1}, \dots, \mathcal{H}_{k_a} \subseteq V(\mathcal{C}X/G)$, where $N = (D + 1)2^{D+1}$; similarly, there exist at most N^b choices for the subsets $\mathcal{H}'_{k'_1}, \dots, \mathcal{H}'_{k'_b} \subseteq V(\mathcal{C}X/G)$.

In particular,

$$|\mathcal{G}_{\varepsilon+1}(H, K)| \leq D \left(\sum_{a=0}^{\varepsilon+2} N^a \right) \left(\sum_{b=0}^{\varepsilon+2} N^b \right) < D \left(\frac{N^{\varepsilon+3}}{N-1} \right)^2 = N_\varepsilon,$$

as required. □

5. Application to graph products

We use this section to deduce results about graph products from Theorems A and B: namely, we show Corollary C in Section 5.1 and Corollary D in Section 5.2. Throughout this section, let Γ be a simplicial graph and let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of non-trivial groups. Let X be the quasi-median graph associated to $\Gamma\mathcal{G}$, as given by Theorem 1.3: namely, X is the Cayley graph of $\Gamma\mathcal{G}$ with respect to the (not necessarily finite) generating set $S = \bigsqcup_{v \in V(\Gamma)} (G_v \setminus \{1\})$.

For the rest of the paper, given a subset $A \subseteq V(\Gamma)$, we write Γ_A for the full subgraph of Γ spanned by A and denote $\{G_v \mid v \in A\}$ by \mathcal{G}_A , and we view the graph product $\Gamma_A\mathcal{G}_A$ as a subgroup of $\Gamma\mathcal{G}$. We will use the following result.

Theorem 5.1 (Genevois [11, Section 8.1]; Genevois and Martin [15, Theorem 2.10]). *For $v \in V(\Gamma)$, let H_v be the hyperplane dual to the clique $G_v \subseteq X$. Then any hyperplane H in X is of the form H_v^g for some $v \in V(\Gamma)$ and $g \in \Gamma\mathcal{G}$. Moreover, the vertices in $\mathcal{N}(H)$ are precisely $\Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}g \subseteq V(X)$.*

Remark 5.2. Due to our convention to consider only right actions, the Cayley graph $X = \text{Cay}(\Gamma\mathcal{G}, S)$ defined in Theorem 1.3 is the *left* Cayley graph: for $s \in S$ and $g \in \Gamma\mathcal{G}$, an edge labelled s joins $g \in V(X)$ to $sg \in V(X)$. Therefore, contrary to the convention in [11, 15], the vertices in the carrier of a hyperplane will form a *right* coset of $\Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}$ for some $v \in V(\Gamma)$.

5.1. Acylindrical hyperbolicity. Here we prove Corollary C. It is clear from Theorem 1.3 and Lemma 1.4 that we may apply Theorems A and B to the quasi-median graph X associated to a graph product $\Gamma\mathcal{G}$. In particular, it follows that the contact graph $\mathcal{C}X$ is a quasi-tree and $\Gamma\mathcal{G}$ acts on it acylindrically. We thus only need to show that, given that $|V(\Gamma)| \geq 2$ and the complement Γ^C of Γ is connected, either the action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ is non-elementary or $\Gamma\mathcal{G} \cong C_2 * C_2$.

Lemma 5.3. *The following are facts equivalent:*

- (i) $\mathcal{C}X$ is unbounded;
- (ii) $\Gamma^{\mathcal{C}}$ is connected and $|V(\Gamma)| \geq 2$.

Proof. We first show (i) \implies (ii). Indeed, if Γ is a single vertex v , then X is a single clique and so $\mathcal{C}X$ is a single vertex. On the other hand, if $\Gamma^{\mathcal{C}}$ is disconnected, then we have a partition $V(\Gamma) = A \sqcup B$ where A and B are adjacent and non-empty. In particular, $\Gamma\mathcal{G} = \Gamma_A\mathcal{G}_A \times \Gamma_B\mathcal{G}_B$, and so any vertex $g \in \Gamma\mathcal{G}$ of X can be expressed as $g = g_A g_B$ for some $g_A \in \Gamma_A\mathcal{G}_A$ and $g_B \in \Gamma_B\mathcal{G}_B$. Thus, if $H \in V(\mathcal{C}X)$ then by Theorem 5.1 $\mathcal{N}(H) = \Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}g_A g_B$ for some $g_A \in \Gamma_A\mathcal{G}_A$, $g_B \in \Gamma_B\mathcal{G}_B$ and $v \in V(\Gamma)$: without loss of generality, suppose $v \in A$. Then $g_B \in \Gamma_B\mathcal{G}_B \leq \Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}$ and $g_A \in \Gamma_A\mathcal{G}_A \leq \Gamma_{\text{star}(u)}\mathcal{G}_{\text{star}(u)}$ for any $u \in B$, and so $g_A \in \mathcal{N}(H) \cap \mathcal{N}(H_u)$; therefore, $d_{\mathcal{C}X}(H, H_u) \leq 1$. Since $1 \in \mathcal{N}(H_u) \cap \mathcal{N}(H_v)$ and so $d_{\mathcal{C}X}(H_u, H_v) \leq 1$ for any $u, v \in V(\Gamma)$, it follows that $d_{\mathcal{C}X}(H, H') \leq 3$ for any $H, H' \in \mathcal{C}X$ and so $\mathcal{C}X$ is bounded, as required.

To show (ii) \implies (i), suppose that Γ is a graph with at least 2 vertices and connected complement. Thus, there exists a closed walk $(v_0, v_1, \dots, v_\ell)$ on the complement of Γ that visits every vertex—in particular, we have $v_i \in V(\Gamma)$ with $v_\ell = v_0$ and $v_{i-1} \neq v_i$, $(v_{i-1}, v_i) \notin E(\Gamma)$ for $1 \leq i \leq \ell$. Pick arbitrary non-identity elements $g_i \in G_{v_i}$ for $i = 1, \dots, \ell$, and consider the element $g = g_1 \cdots g_\ell \in \Gamma\mathcal{G}$.

Now let $n \in \mathbb{N}$, and let $A, B \in V(\mathcal{C}X)$ be such that $1 \in \mathcal{N}(A)$ and $g^n \in \mathcal{N}(B)$. Let $A = A_0, \dots, A_m = B$ be the geodesic in $\mathcal{C}X$ and let $1 = p_0, \dots, p_{m+1} = g^n$ be the vertices in X given by Proposition 3.1. It follows from the normal form theorem for graph products [16, Theorem 3.9] that

$$\overbrace{(g_1 \cdots g_\ell) \cdots (g_1 \cdots g_\ell)}^n$$

is the *unique* normal form for the element g^n . In particular, as geodesics in X are precisely the words spelling out normal forms of elements of $\Gamma\mathcal{G}$, we have $p_i = g^{\ell n - c_i + 1} g^{\ell n - c_i + 2} \cdots g^{\ell n}$, where $0 = c_0 \leq c_1 \leq \dots \leq c_{m+1} = \ell n$ and indices are taken modulo ℓ .

We now claim that $c_{i+1} - c_i < \ell$ for each i . Indeed, suppose $c_{i+1} - c_i \geq \ell$ for some i . Note that, as $p_i, p_{i+1} \in \mathcal{N}(A_i)$, it follows from Theorem 5.1 that $\Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}p_i = V(\mathcal{N}(A_i)) = \Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}p_{i+1}$ for some $v \in V(\Gamma)$, and therefore we have $p_{i+1}p_i^{-1} \in \Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}$. But as $g^{\ell n - c_{i+1}} g^{\ell n - c_{i+1} + 1} \cdots g^{\ell n - c_i - 1}$ is a normal form for $p_{i+1}p_i^{-1}$ (where indices are taken modulo ℓ), it follows that $v_j \in \text{star}(v)$ for $\ell n - c_{i+1} \leq j < \ell n - c_i$ (with indices again modulo ℓ). But as by assumption $c_{i+1} - c_i \geq \ell$ and as $\{v_1, \dots, v_\ell\} = V(\Gamma)$, this implies that $\text{star}(v) = V(\Gamma)$, and so v is an isolated vertex of $\Gamma^{\mathcal{C}}$. This contradicts the fact that $\Gamma^{\mathcal{C}}$ is connected; thus $c_{i+1} - c_i < \ell$ for each i , as claimed.

In particular, we get $\ell n = \sum_{i=0}^m (c_{i+1} - c_i) < (m + 1)\ell$, and so $m + 1 > n$. Thus $d_{\mathcal{C}X}(A, B) = m \geq n$ and so $\mathcal{C}X$ is unbounded, as required. \square

It is now easy to deduce when the action of $\Gamma\mathcal{G}$ on $\mathcal{C}X$ is non-elementary acylindrical.

Proof of Corollary C. By the argument above, we only need to show the last part. Thus, suppose that Γ is a graph with at least 2 vertices and connected complement. Then, by Lemma 5.3, the graph $\mathcal{C}X$ is unbounded. In particular, given any $H \in V(\mathcal{C}X)$ and $n \in \mathbb{N}$, we may pick $H' \in V(\mathcal{C}X)$ such that $d_{\mathcal{C}X}(H, H') \geq n + 1$. Since the action $\Gamma\mathcal{G} \curvearrowright X$ is transitive on vertices, it follows that given any vertex $p \in \mathcal{N}(H)$ there exists $g \in \Gamma\mathcal{G}$ such that $p^g \in \mathcal{N}(H')$, and in particular $d_{\mathcal{C}X}(H^g, H') \leq 1$. Thus $d_{\mathcal{C}X}(H, H^g) \geq n$, and so the action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ has unbounded orbits.

It is left to show that $\Gamma\mathcal{G}$ is either isomorphic to $C_2 * C_2$ or not virtually cyclic. Indeed, if Γ has only two vertices (v and w , say) and $G_v \cong G_w \cong C_2$, then Γ is a discrete graph with two vertices and $\Gamma\mathcal{G} \cong G_v * G_w \cong C_2 * C_2$. Otherwise, we claim that there exists a subset $A \subseteq V(\Gamma)$ such that $\Gamma_A\mathcal{G}_A$ is a free product of two non-trivial groups, not both C_2 . If $|V(\Gamma)| = 2$ (and so $V(\Gamma) = \{v, w\}$) then $\Gamma\mathcal{G} \cong G_v * G_w$ and so we may take $A = V(\Gamma)$. Otherwise, since $|V(\Gamma)| \geq 3$ and Γ^C is connected, Γ^C contains a path of length 2, and so there exist vertices $v_1, v_2, w \in \Gamma$ such that $v_1 \sim w \sim v_2$. Let $A = \{v_1, v_2, w\}$ and $H = \Gamma_{\{v_1, v_2\}}\mathcal{G}_{\{v_1, v_2\}}$ (so either $H \cong G_{v_1} \times G_{v_2}$ or $H \cong G_{v_1} * G_{v_2}$). Since the groups G_v are non-trivial for each $v \in V(\Gamma)$, we have $|H| \geq 4 > 2$ and so $\Gamma_A\mathcal{G}_A \cong G_w * H$ satisfies our conditions, which proves the claim.

Thus, in either case $\Gamma_A\mathcal{G}_A$ has infinitely many ends, and so is not virtually cyclic. In particular, since the subgroup $\Gamma_A\mathcal{G}_A \leq \Gamma\mathcal{G}$ is not virtually cyclic, neither is $\Gamma\mathcal{G}$, as required. \square

Remark 5.4. After appearance of the first version of this paper in a form of a preprint, it has been brought to the author’s attention that most of the results stated in Corollary C have already been proved by Genevois. In [12, Theorem 4.11], Genevois shows that ΔX is quasi-isometric to a tree whenever it is connected and Γ is finite, so in particular, by Theorem B (i), $\mathcal{C}X$ is a quasi-tree as well. Moreover, methods used by Genevois to prove [13, Theorem 4.1] can be adapted to show that the action of $\Gamma\mathcal{G}$ on $\mathcal{C}X$ is non-uniformly acylindrical; here, the *non-uniform acylindricity* of an action $G \curvearrowright X$ is a weaker version of acylindricity, defined by replacing the phrase “is bounded above by N_ε ” by “is finite” in Definition 1.2. Corollary C strengthens this statement.

5.2. \mathcal{AH} -accessibility. Here we study \mathcal{AH} -accessibility, introduced in [1] by Abbott, Balasubramanya and Osin, of graph products. In particular, we show that if Γ is connected, non-trivial, and the groups in \mathcal{G} are infinite, then the action of $\Gamma\mathcal{G}$ on $\mathcal{C}X$ is the “largest” acylindrical action of $\Gamma\mathcal{G}$ on a hyperbolic metric space. Hence we prove Corollary D.

We briefly recall the terminology of [1]. Given two isometric actions $G \curvearrowright X$ and $G \curvearrowright Y$ of a group G , we say $G \curvearrowright X$ *dominates* $G \curvearrowright Y$, denoted $G \curvearrowright Y \preceq G \curvearrowright X$, if there exist $x \in X, y \in Y$ and a constant C such that

$$d_Y(y, y^g) \leq C d_X(x, x^g) + C$$

for all $g \in G$. The actions $G \curvearrowright X$ and $G \curvearrowright Y$ are said to be *weakly equivalent* if $G \curvearrowright X \preceq G \curvearrowright Y$ and $G \curvearrowright Y \preceq G \curvearrowright X$. This partitions all such actions into equivalence classes.

It is easy to see that \preceq defines a preorder on the set of all isometric actions of G on metric spaces. Therefore, \preceq defines a partial order on the set of equivalence classes of all such actions. We may restrict this to a partial order on the set $\mathcal{AH}(G)$ of equivalence classes of *acylindrical* actions of G on a *hyperbolic* space. We then say the group G is *\mathcal{AH} -accessible* if the partial order $\mathcal{AH}(G)$ has a largest element (which, if exists, must necessarily be unique), and we say G is *strongly \mathcal{AH} -accessible* if a representative of this largest element is a Cayley graph of G .

Recall that for an action $G \curvearrowright X$ by isometries with X hyperbolic, an element $g \in G$ is said to be *loxodromic* if, for some (or any) $x \in X$, the map $\mathbb{Z} \rightarrow X$ given by $n \mapsto x^{g^n}$ is a quasi-isometric embedding. It is clear from the definitions that the “largest” action $G \curvearrowright X$ will also be *universal*, in the sense that every element of G , that is loxodromic with respect to *some* acylindrical action of G on a hyperbolic space, will be loxodromic with respect to $G \curvearrowright X$.

In [1, Theorem 2.18 (c)], it is shown that the all right-angled Artin groups are \mathcal{AH} -accessible (and more generally, so are all hierarchically hyperbolic groups—in particular, groups acting properly and cocompactly on a CAT(0) cube complex possessing a factor system [2, Theorem A]). Here we generalise this result to “most” graph products of infinite groups. The proof is very similar to that of [1, Lemma 7.16].

Proof of Corollary D. It is easy to show—for instance, by Theorem 5.1—that $\mathcal{C}X$ is (G -equivariantly) quasi-isometric to the Cayley graph of $\Gamma\mathcal{G}$ with respect to $\bigcup_{v \in V(\Gamma)} \Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)}$.

We prove the statement by induction on $|V(\Gamma)|$. If $|V(\Gamma)| = 1$ ($V(\Gamma) = \{v\}$, say), then v is an isolated vertex of Γ and so, by the assumption, $\Gamma\mathcal{G} \cong G_v$ is strongly \mathcal{AH} -accessible.

Suppose now that $|V(\Gamma)| \geq 2$. If Γ has an isolated vertex ($\Gamma = \Gamma_A \sqcup \{v\}$ for some partition $V(\Gamma) = A \sqcup \{v\}$, say), then $\Gamma\mathcal{G} \cong \Gamma_A\mathcal{G}_A * G_v$ is hyperbolic relative to $\{\Gamma_A\mathcal{G}_A, G_v\}$. By the induction hypothesis, both $\Gamma_A\mathcal{G}_A$ and G_v are strongly \mathcal{AH} -accessible, and hence, by [1, Theorem 7.9], so is $\Gamma\mathcal{G}$. If, on the other hand, the complement Γ^C of Γ is disconnected ($\Gamma^C = \Gamma_A^C \sqcup \Gamma_B^C$ for some partition $V(\Gamma) = A \sqcup B$, say), then $\Gamma\mathcal{G} \cong \Gamma_A\mathcal{G}_A \times \Gamma_B\mathcal{G}_B$ is not acylindrically hyperbolic by [25, Corollary 7.2], as both $\Gamma_A\mathcal{G}_A$ and $\Gamma_B\mathcal{G}_B$ are infinite. It then follows from [1, Example 7.8] that $\Gamma\mathcal{G}$ is strongly \mathcal{AH} -accessible; it also follows that *any* acylindrical

action of $\Gamma\mathcal{G}$ on a hyperbolic metric space $(\Gamma\mathcal{G} \curvearrowright \mathcal{C}X, \text{ say})$ represents the largest element of $\mathcal{A}\mathcal{H}(\Gamma\mathcal{G})$.

Hence, we may without loss of generality assume that Γ is a graph with no isolated vertices and connected complement. It then follows that $|V(\Gamma)| \geq 4$, and so by Corollary C, $\mathcal{C}X$ is a hyperbolic metric space and $\Gamma\mathcal{G}$ acts on it non-elementarily acylindrically. It is easy to see from Theorem 5.1 that, given two hyperplanes $H, H' \in V(\mathcal{C}X)$, they are adjacent in $\mathcal{C}X$ if and only if there exist distinct $u, v \in V(\Gamma)$ and $g \in \Gamma\mathcal{G}$ such that $H = H_u^g$ and $H' = H_v^g$. It follows that the quotient space $\mathcal{C}X/\Gamma\mathcal{G}$ is the complete graph on $|V(\Gamma)|$ vertices, and in particular, the action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ is cocompact.

Moreover, it follows from Theorem 5.1 that the stabiliser of an arbitrary vertex H_v^g of $\mathcal{C}X$ is precisely $G \cong (\Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)})^g \cong G_v^g \times (\Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)})^g$. Since Γ has no isolated vertices, $\text{link}(v) \neq \emptyset$, and so, as all groups in \mathcal{G} are infinite, both G_v^g and $(\Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)})^g$ are infinite groups. Thus, G is a direct product of two infinite groups, and so—by [25, Corollary 7.2], say— G does not possess a non-elementary acylindrical action on a hyperbolic space. Since G is not virtually cyclic, for every acylindrical action of $\Gamma\mathcal{G}$ on a hyperbolic space Y , the induced action of G on Y has bounded orbits. It then follows from [1, Proposition 4.13] that $\Gamma\mathcal{G}$ is strongly $\mathcal{A}\mathcal{H}$ -accessible—and in particular, $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ represents the largest element of $\mathcal{A}\mathcal{H}(\Gamma\mathcal{G})$. \square

Remark 5.5. Corollary D gives some explicit descriptions for the class of hierarchically hyperbolic groups, introduced by Behrstock, Hagen and Sisto in [5]. In particular, a result by Berlai and Robbio [6, Theorem C] says that if all vertex groups G_v are hierarchically hyperbolic with the intersection property and clean containers, then the same can be said about $\Gamma\mathcal{G}$. Moreover, Abbott, Behrstock and Durham show in [2, Theorem A] that all hierarchically hyperbolic groups are $\mathcal{A}\mathcal{H}$ -accessible, which implies Corollary D in the case when the vertex groups G_v are hierarchically hyperbolic with the intersection property and clean containers.

More precisely, every hierarchically hyperbolic group G comes with an action on a space \mathcal{X} , such that there exist projections $\pi_Y: \mathcal{X} \rightarrow 2^{\hat{\mathcal{C}}Y}$ to some collection of δ -hyperbolic spaces $\{\hat{\mathcal{C}}Y \mid Y \in \mathfrak{S}\}$, where \mathfrak{S} is a partial order that contains a (unique) largest element, $S \in \mathfrak{S}$, say. Moreover, the action of G on \mathcal{X} induces an action of G on (a space quasi-isometric to) $\mathcal{U} = \bigcup_{x \in \mathcal{X}} \pi_S(x) \subseteq \hat{\mathcal{C}}S$, and in [5, Theorem 14.3] it is shown that this action is acylindrical. In [2], this construction is modified so that the action $G \curvearrowright \mathcal{U}$ represents the largest element of $\mathcal{A}\mathcal{H}(G)$. If Γ is connected, non-trivial, and the groups G_v are infinite and hierarchically hyperbolic (with the intersection property and clean containers), then the proof of Corollary D gives this action $\Gamma\mathcal{G} \curvearrowright \mathcal{U}$ explicitly. This is potentially useful for studying hierarchical hyperbolicity of graph products.

Remark 5.6. Note that the condition on the G_v being infinite is necessary for the proof to work. Indeed, suppose $\Gamma = \overset{a}{\bullet} \text{---} \overset{b}{\bullet} \text{---} \overset{c}{\bullet} \text{---} \overset{d}{\bullet}$ is a path of length 3, and $G_v = \langle g_v \rangle \cong C_2$ for each $v \in V(\Gamma)$, so that $\Gamma\mathcal{G}$ is the right-angled Coxeter group over Γ . Notice that $\Gamma\mathcal{G} \cong A *_C B$, where $A = G_b \times (G_a * G_c)$, $B = G_c \times G_d$ and $C = G_c$. In particular, since C is finite, $\Gamma\mathcal{G}$ is hyperbolic relative to $\{A, B\}$. Hence the Cayley graph $\text{Cay}(\Gamma\mathcal{G}, A \cup B)$ is hyperbolic and the usual action of $\Gamma\mathcal{G}$ on it is acylindrical.

It is easy to verify from the normal form theorem for amalgamated free products that the element $g_b g_d$ will be loxodromic with respect to $\Gamma\mathcal{G} \curvearrowright \text{Cay}(\Gamma\mathcal{G}, A \cup B)$. However, as $g_b g_d \in \Gamma_{\text{star}(c)} \mathcal{G}_{\text{star}(c)}$, we know that $g_b g_d$ stabilises the hyperplane dual to $G_c \subseteq V(X)$ under the action of $\Gamma\mathcal{G}$ on $\mathcal{C}X$, and so $g_b g_d$ is not loxodromic with respect to $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$. In particular, the equivalence class of $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ cannot be the largest element of $\mathcal{A}\mathcal{H}(\Gamma\mathcal{G})$. It is straightforward to generalise this argument to show that if $c \in V(\Gamma)$ is a separating vertex of a connected finite simplicial graph Γ , then for *any* graph product $\Gamma\mathcal{G}$ with G_c finite, the action $\Gamma\mathcal{G} \curvearrowright \mathcal{C}X$ will not be the “largest” one.

On the other hand, note that this particular group $\Gamma\mathcal{G}$ (and indeed any right-angled Coxeter group) will be $\mathcal{A}\mathcal{H}$ -accessible: see [2, Theorem A (4)].

6. Equational noetherianity of graph products

In this section we prove Theorem E. To do this, we use the methods that Groves and Hull exhibited in [17]. Here we briefly recall their terminology.

The approach to equationally noetherian groups used in [17] is through sequences of homomorphisms. In particular, let G be any group, let F be a finitely generated group and let $\varphi_i: F \rightarrow G$ be a sequence of homomorphisms ($i \in \mathbb{N}$). Let $\omega: \mathcal{P}(\mathbb{N}) \rightarrow \{0, 1\}$ be a non-principal ultrafilter. We say a sequence of properties $(P_i)_{i \in \mathbb{N}}$ holds ω -almost surely if $\omega(\{i \in \mathbb{N} \mid P_i \text{ holds}\}) = 1$. We define the ω -kernel of F with respect to (φ_i) to be

$$F_{\omega, (\varphi_i)} = \{f \in F \mid \varphi_i(f) = 1 \text{ } \omega\text{-almost surely}\};$$

we write F_ω for $F_{\omega, (\varphi_i)}$ if the sequence (φ_i) is clear. It is easy to check that F_ω is a normal subgroup of F . We say φ_i factors through F_ω ω -almost surely if $F_\omega \subseteq \ker(\varphi_i)$ ω -almost surely.

The idea behind all these definitions is the following result.

Theorem 6.1 (Groves and Hull [17, Theorem 3.5]). *Let ω be a non-principal ultrafilter. Then the following are equivalent for any group G :*

- (i) G is equationally noetherian;
- (ii) for any finitely generated group F and any sequence of homomorphisms $(\varphi_i: F \rightarrow G)$, φ_i factors through F_ω ω -almost surely.

Remark 6.2. Note that Definition 1.6 differs from the usual definition of equationally noetherian groups, as we do not allow “coefficients” in our equations: that is, we restrict to subsets $S \subseteq F_n$ instead of $S \subseteq G * F_n$. However, the two concepts coincide when G is finitely generated—see [4, §2.2, Proposition 3]. We use this (weaker) definition of equationally noetherian groups as it is more suitable for our methods. In particular, we use an equivalent characterisation of equationally noetherian groups given by Theorem 6.1.

The structure of this section is as follows. In Section 6.1, we introduce “admissible” graphs and show that being equationally noetherian is preserved under taking graph products over connected admissible graphs. In Section 6.2, we introduce (polygonal) dual van Kampen diagrams for graph products—a tool that we use in our proof of Theorem E. In Section 6.3, we introduce minimal polygonal representations for words in a free group, closely related to polygonal van Kampen diagrams defined in the preceding subsection. In Section 6.4, we give our main technical results in order to show that all connected graphs of girth ≥ 6 are admissible. In Section 6.5, we combine these results to prove Theorem E.

6.1. Reduction to sequences of linking homomorphisms. Suppose now that the group G acts by isometries on a metric space (Y, d) . As before, let F be a finitely generated group, ω a non-principal ultrafilter, and $(\varphi_i: F \rightarrow G)_{i=1}^\infty$ a sequence of homomorphisms. Pick a finite generating set S for F . We say that the sequence of homomorphisms (φ_i) is *non-divergent* if

$$\liminf_{\omega} \max_{y \in Y} \max_{s \in S} d(y, y^{\varphi_i(s)}) < \infty.$$

We say that (φ_i) is *divergent* otherwise. It is easy to see that this does not depend on the choice of a generating set S for F .

The main technical result of [17] states that in case Y is hyperbolic and the action of G on Y is non-elementary acylindrical, it is enough to consider non-divergent sequences of homomorphisms (cf Theorem 6.1).

Theorem 6.3 (Groves and Hull [17, Theorem B]). *Let Y be a hyperbolic metric space and G a group acting non-elementarily acylindrically on Y . Suppose that for any finitely generated group F and any non-divergent sequence of homomorphisms $(\varphi_i: F \rightarrow G)$, φ_i factors through F_ω ω -almost surely. Then G is equationally noetherian.*

We now consider the particular case when G is a graph product and $Y = \mathcal{C}X$ is the crossing graph of the quasi-median graph X associated to G . Thus, as before, let Γ be a finite simplicial graph and let $\mathfrak{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of non-trivial groups. It turns out that in this case we may reduce any non-divergent sequence of homomorphisms to one of the following form: see the proof of Theorem 6.6.

Definition 6.4. Let F be a finitely generated free group, let $S \subset F$ be a free basis for F , and let $\varphi: F \rightarrow \Gamma\mathcal{G}$ be a homomorphism. We say φ is *linking* if for each $s \in S$, there exists $v = v(s) \in V(\Gamma)$ such that $\varphi(s) \in \Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)}$; although in general this depends on the choice of the free basis S , we will usually fix S and omit references to it. We say the graph Γ is *admissible* if for every collection of non-trivial equationally noetherian groups $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ and every sequence of linking homomorphisms $(\varphi_i: F \rightarrow \Gamma\mathcal{G})_{i=1}^\infty$, φ_i factors through F_ω ω -almost surely.

The proof of Theorem 6.6 uses the following result.

Lemma 6.5. *Full subgraphs of admissible graphs are admissible.*

Proof. Let Γ be a admissible graph, let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of non-trivial equationally noetherian groups, and let F be a finitely generated free group. Let $A \subseteq V(\Gamma)$, so that Γ_A is a full subgraph of Γ , and let $(\varphi_i^A: F \rightarrow \Gamma_A\mathcal{G}_A)_{i=1}^\infty$ be a sequence of linking homomorphisms. Let ω be a non-principal ultrafilter. We aim to show that φ_i^A factors through $F_{\omega,(\varphi_i^A)}$ ω -almost surely.

Note that we have a canonical inclusion of subgroup $\iota_A: \Gamma_A\mathcal{G}_A \rightarrow \Gamma\mathcal{G}$. For each i , let $\varphi_i = \iota_A \circ \varphi_i^A: F \rightarrow \Gamma\mathcal{G}$. It is easy to see that the φ_i are linking homomorphisms. In particular, since Γ is admissible, we have $F_{\omega,(\varphi_i)} \subseteq \ker \varphi_i$ ω -almost surely. Moreover, since ι_A is injective, we have $\ker \varphi_i^A = \ker \varphi_i$ for each i and $F_{\omega,(\varphi_i^A)} = F_{\omega,(\varphi_i)}$. Thus $F_{\omega,(\varphi_i^A)} \subseteq \ker \varphi_i^A$ ω -almost surely, and so Γ_A is admissible, as required. \square

Theorem 6.6. *For any connected admissible graph Γ and any collection $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ of equationally noetherian groups, the graph product $\Gamma\mathcal{G}$ is equationally noetherian.*

Proof. We proceed by induction on $|V(\Gamma)|$. If $|V(\Gamma)| = 1$ ($V(\Gamma) = \{v\}$, say) then $\Gamma\mathcal{G} \cong G_v$, and so the result is clear. Thus, assume that $|V(\Gamma)| \geq 2$.

If the complement of Γ is disconnected, then we have a partition $V(\Gamma) = A \sqcup B$ such that $\Gamma\mathcal{G} \cong \Gamma_A\mathcal{G}_A \times \Gamma_B\mathcal{G}_B$. By Lemma 6.5, both Γ_A and Γ_B are admissible, and so $\Gamma_A\mathcal{G}_A$ and $\Gamma_B\mathcal{G}_B$ are equationally noetherian by the induction hypothesis. It is clear from the definition that a direct product $G \times H$ of equationally noetherian groups G and H is equationally noetherian: indeed, this follows from the cartesian product decomposition $V_{G \times H}(S) = V_G(S) \times V_H(S)$, for any $S \subseteq F_n$. Thus $\Gamma\mathcal{G}$ is equationally noetherian.

Therefore, we may without loss of generality assume that Γ is a connected graph with a connected complement and with $|V(\Gamma)| \geq 2$ (and, therefore, $|V(\Gamma)| \geq 4$). In this case, Corollary C shows that $\mathcal{C}X$ is a hyperbolic metric space and the action of $\Gamma\mathcal{G}$ on it is non-elementary acylindrical. We thus may use Theorem 6.3 to show that $\Gamma\mathcal{G}$ is equationally noetherian.

In particular, let F be a finitely generated group and let $(\varphi_i: F \rightarrow \Gamma\mathcal{G})_{i=1}^\infty$ be a non-divergent sequence of homomorphisms. By Theorem 6.3, it is enough to show that φ_i factors through F_ω ω -almost surely.

We proceed as in the proof of [17, Theorem D]. Let S be a finite generating set for F . Note that, by Theorem 5.1, we may conjugate each φ_i (if necessary) to assume that the minimum (over all hyperplanes H in X) of $\max_{s \in S} d_{eX}(H, H^{\varphi_i(s)})$ is attained for $H = H_u$ for some $u \in V(\Gamma)$. Moreover, it is easy to see from Theorem 5.1 that $\|g\|_* - d_{eX}(H_u, H_u^g) \leq 1$ for any $g \in \Gamma\mathcal{G}$ and $u \in V(\Gamma)$, where we write $\|g\|_*$ for the minimal integer $\ell \in \mathbb{N}$ such that $g = g_1 \cdots g_\ell$ and $g_i \in \Gamma_{\text{star}(v_i)}\mathcal{G}_{\text{star}(v_i)}$ for some $v_i \in V(\Gamma)$. In particular, since the sequence (φ_i) is non-divergent, it follows that

$$\lim_{\omega} \max_{s \in S} \|\varphi_i(s)\|_* < \infty.$$

It follows that for each $s \in S$, there exists $\hat{n}_s \in \mathbb{N}$ such that $\|\varphi_i(s)\|_* = \hat{n}_s$ ω -almost surely. Moreover, for each $s \in S$, there exist $\hat{v}_{s,1}, \dots, \hat{v}_{s,\hat{n}_s} \in V(\Gamma)$ such that we have

$$\varphi_i(s) = \hat{g}_{i,s,1} \cdots \hat{g}_{i,s,\hat{n}_s}$$

with $\hat{g}_{i,s,j} \in \Gamma_{\text{star}(\hat{v}_{s,j})}\mathcal{G}_{\text{star}(\hat{v}_{s,j})}$ ω -almost surely. But since $\Gamma_{\text{star}(v)}\mathcal{G}_{\text{star}(v)} = G_v \times \Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)}$ for each $v \in V(\Gamma)$, we can write $\hat{g}_{i,s,j} = g_{i,s,2j-1}g_{i,s,2j}$, where $g_{i,s,2j-1} \in G_{\hat{v}_{s,j}} \leq \Gamma_{\text{link}(v_{s,2j-1})}\mathcal{G}_{\text{link}(v_{s,2j-1})}$ with any choice of vertex $v_{s,2j-1} \in \text{link}(\hat{v}_{s,j})$ (which exists since Γ is connected and $|V(\Gamma)| \geq 2$), and $g_{i,s,2j} \in \Gamma_{\text{link}(v_{s,2j})}\mathcal{G}_{\text{link}(v_{s,2j})}$ with $v_{s,2j} = \hat{v}_{s,j}$. It follows that, after setting $n_s = 2\hat{n}_s$, we ω -almost surely may write

$$\varphi_i(s) = g_{i,s,1} \cdots g_{i,s,n_s}$$

with $g_{i,s,j} \in \Gamma_{\text{link}(v_{s,j})}\mathcal{G}_{\text{link}(v_{s,j})}$.

Now for each $s \in S$, define abstract letters $h_{s,1}, \dots, h_{s,n_s}$. For each $v \in V(\Gamma)$, let

$$S_v = \{h_{s,j} \mid v_{s,j} = v\},$$

let $\bar{S} = \sqcup_{v \in V(\Gamma)} S_v$, and let $\bar{F} = F(\bar{S})$, the free group on \bar{S} . We can define a map from S to F by sending $s \in S$ to $h_{s,1} \cdots h_{s,n_s}$. Let N be the normal subgroup of \bar{F} generated by images of all the relators of F under this map. This gives a group homomorphism $\rho: F \rightarrow \bar{F}/N$.

The map $\hat{\varphi}_i: \bar{F}/N \rightarrow \Gamma\mathcal{G}$, obtained by sending $h_{s,j}N$ to $g_{i,s,j}$, is ω -almost surely a well-defined homomorphism. Indeed, all the relators in \bar{F}/N are of the form $\phi(\{h_{s,1} \cdots h_{s,n_s} \mid s \in S\})$, where $\phi(S)$ is a relator in F , and so ω -almost surely map to the identity under $\hat{\varphi}_i$. It is also clear that $\varphi_i = \hat{\varphi}_i \circ \rho$ ω -almost surely.

Now let $\pi: \bar{F} \rightarrow \bar{F}/N$ be the quotient map. Then, by construction, the homomorphisms $\varphi'_i = \hat{\varphi}_i \circ \pi: \bar{F} \rightarrow \Gamma\mathcal{G}$ are linking (when they are well defined). Since Γ is admissible and the groups G_v are equationally noetherian, it follows

that φ'_i factors through \bar{F}_ω ω -almost surely. Since π is surjective, this implies that $(\bar{F}/N)_\omega \subseteq \ker \hat{\varphi}_i$ ω -almost surely. Thus $\varphi_i = \hat{\varphi}_i \circ \rho$ factors through $F_\omega = \rho^{-1}((\bar{F}/N)_\omega)$ ω -almost surely, as required. \square

We expect that the class of equationally noetherian groups is closed under taking arbitrary graph products. Although we are not able to show this in full generality, in the next subsections we show that any connected graph Γ that is triangle-free, square-free and pentagon-free, is admissible, and therefore (by Theorem 6.6) the class of equationally noetherian groups is closed under taking graph products over such graphs Γ .

6.2. Dual van Kampen diagrams. Before embarking on a proof of Theorem E, let us define the following notion. Following methods of [10] and [21], we consider *dual van Kampen diagrams* for words representing the identity in $\Gamma\mathcal{G}$; recently, dual van Kampen diagrams for graph products have been independently introduced by Genevois in [14]. Here we explain their construction and properties.

We consider van Kampen diagrams in the quasi-median graph X given by Theorem 1.3, viewed as a Cayley graph. In particular, note that we have a presentation

$$\Gamma\mathcal{G} = \langle S \mid R_\Delta \sqcup R_\square \rangle \tag{4}$$

with generators

$$S = \bigsqcup_{v \in V\Gamma} (G_v \setminus \{1\})$$

and relators of two types: the “triangular” relators

$$R_\Delta = \bigsqcup_{v \in V(\Gamma)} \{ghk^{-1} \mid g, h, k \in G_v \setminus \{1\}, gh = k \text{ in } G_v\}$$

and the “rectangular” relators

$$R_\square = \bigsqcup_{(v,w) \in E(\Gamma)} \{[g_v, g_w] \mid g_v \in G_v \setminus \{1\}, g_w \in G_w \setminus \{1\}\}.$$

We now dualise the notion of van Kampen diagrams with respect to the presentation (4). Let $D \subseteq \mathbb{R}^2$ be a van Kampen diagram with boundary label w , for some word $w \in S^*$ representing the identity in $\Gamma\mathcal{G}$, with respect to the presentation (4). It is convenient to pick a colouring $V(\Gamma) \rightarrow \mathbb{N}$ and to colour edges of D according to their labels. Suppose that $w = g_1 \cdots g_n$ for some syllables g_i , and let e_1, \dots, e_n be the corresponding edges on the boundary of D . We add a “vertex at infinity” ∞ somewhere on $\mathbb{R}^2 \setminus D$, and for each $i = 1, \dots, n$, we attach to D a triangular “boundary” face whose vertices are the endpoints of e_i and ∞ . We get the dual van Kampen diagram Δ corresponding to D by taking the dual of

D as a polyhedral complex and removing the face corresponding to ∞ : thus, Δ is a tessellation of a disk. See Figure 7.




Figure 7. Van Kampen diagram (D , left) and its dual (Δ , right) with the word $a_1 b_1 c_1 a_2 b_2 c_2 a_3^{-1} c_2^{-1} a_3 b_2^{-1} c_3^{-1} a_4^{-1} c_4^{-1} b_1^{-1}$ as its boundary label, where $a_i \in G_a$ with $a_1 a_2 = a_4$, $b_i \in G_b$, $c_i \in G_c$ with $c_4 c_3 = c_1$, and $b \sim a \sim c$ in Γ . The black edges on D represent the boundary faces attached: the non-visible endpoint of each black edge is the point ∞ . The dual van Kampen diagram Δ contains 6 components in total: 2 components corresponding to each of the vertices a , b and c .

We lift the colouring of edges in D to a colouring of edges of Δ : this gives a corresponding vertex $v \in V(\Gamma)$ for each internal edge of Δ . We say a 1-subcomplex (a subgraph) of Δ is a v -component (or just a component) for some $v \in V(\Gamma)$ if it is a maximal connected subgraph each of whose edges correspond to the vertex v . We call a vertex of Δ an *intersection point* (respectively *branch point*, *boundary point*) if it comes from a triangular (respectively rectangular, boundary) face in D . It is easy to see that boundary, intersection and branch points lying on a component C will be precisely the vertices of C of degree 1, 2, and 3, respectively.

We now equip dual van Kampen diagrams with some additional structure. Some of the concepts introduced in the following Definition are displayed in Figure 8.

Definition 6.7. Let \mathbf{W} be a cyclic word over $\bigsqcup_{v \in V(\Gamma)} (G_v \setminus \{1\})$ such that $\mathbf{W} = 1$ in $\Gamma \mathcal{G}$. A *polygonal van Kampen diagram* for \mathbf{W} is a dual van Kampen diagram with boundary label \mathbf{W} together with the following information:

- a subdivision $\mathbf{W} = g_1 \cdots g_k$ of \mathbf{W} into geodesic subwords g_ℓ , called *van Kampen pieces*, such that for each $\ell \in \{1, \dots, k\}$ we have $\text{supp } g_\ell \subseteq \text{link}(v)$ for some $v \in V(\Gamma)$;

- an assignment of a *label* from $V(\Gamma) \times \{', \wedge\}$ for each van Kampen piece, such that $\text{supp } g_\ell \subseteq \{v\}$ if g_ℓ has label $(v, ')$ and $\text{supp } g_\ell \subseteq \text{link}(v)$ if g_ℓ has label (v, \wedge) .

For any $v \in V(\Gamma)$, we write \dot{v} for $(v, ')$, and call each van Kampen piece with label \dot{v} a \dot{v} -piece; we write \hat{v} for (v, \wedge) , and call each van Kampen piece with label \hat{v} a \hat{v} -piece.

We say a van Kampen piece g is *trivial* if $|g| = 0$, *small* if $|g| = 1$, and *large* otherwise. With g_1, \dots, g_k the van Kampen pieces as above, we say g_{j_1}, \dots, g_{j_m} are *consecutive* van Kampen pieces if $2 \leq m \leq k$ and $j_{i+1} \equiv j_i + 1 \pmod{k}$ for $1 \leq i \leq m - 1$. We say van Kampen pieces g and h are *adjacent* if either g, h or h, g are consecutive.

We adopt the convention that \mathbf{W} is read out by following the boundary of Δ clockwise. For a van Kampen piece g of Δ , we write $\mathcal{J}_g(\Delta)$ for the subinterval of the boundary $\partial\Delta$ of Δ corresponding to the subword g . Given two distinct points $P, Q \in \partial\Delta$, we write $[P, Q]$ for the closed subinterval of $\partial\Delta$ going clockwise from P to Q . We say a component C of Δ is *supported* at a van Kampen piece g if it has a boundary point on $\mathcal{J}_g(\Delta)$. A *supporting interval* of C is an interval of the form $[P, Q]$, where $P, Q \in \partial\Delta$ are distinct boundary points of C such that all boundary points of C are contained in $[P, Q]$.

An *oriented component* of Δ is a pair $(C, [P, Q])$, where C is a component of Δ and $[P, Q]$ is a supporting interval of C . An oriented component $(C, [P, Q])$ is said to *enclose* an oriented component $(C', [P', Q'])$ (or we may simply say that $(C, [P, Q])$ *encloses* C') if $[P', Q'] \subsetneq [P, Q]$. We say that $(C, [P, Q])$ is *trivial* if $P \in \mathcal{J}_g(\Delta)$, $Q \in \mathcal{J}_h(\Delta)$ and g, h are consecutive van Kampen pieces; $(C, [P, Q])$ is said to be *minimal* if it is not trivial and does not enclose any non-trivial components, and *almost minimal* if it is neither trivial nor minimal and does not enclose any non-trivial non-minimal components. We say $(C, [P, Q])$ is an oriented component *starting* (respectively *ending*) at a van Kampen piece g if $P \in \mathcal{J}_g(\Delta)$ (respectively $Q \in \mathcal{J}_g(\Delta)$).

We order the components supported at a given van Kampen piece g by looking at their boundary points on $\mathcal{J}_g(\Delta)$. In particular, for $i \geq 1$ and a van Kampen piece g , we say a component C supported at g is the *i -th component supported at g* if the boundary point of C in $\mathcal{J}_g(\Delta)$ is the i -th boundary point on $\mathcal{J}_g(\Delta)$ when following $\partial\Delta$ clockwise. We similarly define the *last* (or *penultimate*) *component supported at g* .

The following two Lemmas will be used in Section 6.4.

Lemma 6.8. *Let Δ be a polygonal van Kampen diagram, and suppose that Γ is triangle-free. Then for any van Kampen piece g of Δ , no two components supported at g intersect.*

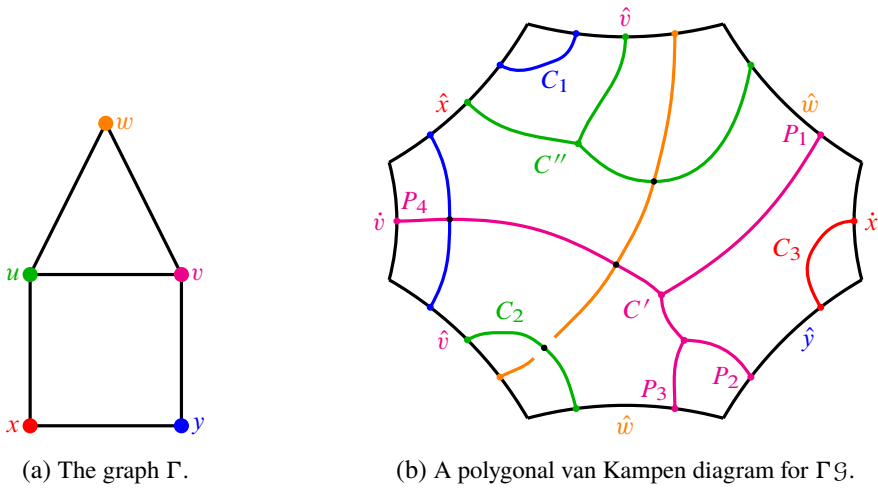


Figure 8. An example of a polygonal van Kampen diagram Δ . The interval $\mathcal{J}_g(\Delta)$ corresponding to a van Kampen piece g is drawn as a smooth curve, with the label of g written next to it. The trivial components are C_1, C_2 and C_3 (with the obvious choices of supporting intervals). The oriented component $(C', [P_1, P_4])$ is minimal, but for $1 \leq i \leq 3$, the oriented component $(C', [P_{i+1}, P_i])$ is not minimal as it encloses C'' . The components C_3 and C' are the first (penultimate) and second (last) components, respectively, supported at the unique \hat{y} -piece of Δ .

Proof. Let C be a v -component and C' be a w -component (for some $v, w \in V(\Gamma)$), both supported at a van Kampen piece g . Suppose for contradiction that C and C' intersect: then $w \in \text{link}(v)$. But we have $\{v, w\} \subseteq \text{supp } g \subseteq \text{link}(u)$ for some $u \in V(\Gamma)$. Hence u, v, w span a triangle in Γ , contradicting the fact that Γ is triangle-free. Thus C and C' cannot intersect, as required. \square

Lemma 6.9. *Let Δ be a polygonal van Kampen diagram, and suppose that Γ has girth ≥ 5 . Suppose C_1 and C_2 are distinct components of Δ supported at a van Kampen piece g of Δ , and C is a v -component (for some $v \in V(\Gamma)$) that either intersects or coincides with each of C_1 and C_2 . Then g is a \hat{v} -piece, and C intersects both C_1 and C_2 .*

Proof. Let C_1 and C_2 be the i -th and the j -th pieces (respectively) supported at g , and suppose, without loss of generality, that $i < j$.

Suppose first that $j \geq i + 2$. Let $P_1, P_2 \in \mathcal{J}_p(\Delta)$ be the boundary points of C_1, C_2 , respectively. Note that, by the assumption, any continuous path in Δ between a point in $[P_1, P_2] \setminus \{P_1, P_2\}$ and a point in $[P_2, P_1] \setminus \{P_1, P_2\}$ must intersect either C or C_1 or C_2 . Thus, if C' is the $(i + 1)$ -st component supported at p , then C' must either intersect or coincide with one of C, C_1 and C_2 . But C' cannot intersect C_1 or C_2 by Lemma 6.8, and it cannot coincide with either C_1

or C_2 since g is a geodesic. Thus C' must either intersect or coincide with C ; by replacing C_2 with C' if necessary, we may therefore assume, without loss of generality, that $j = i + 1$.

Now let $v_1, v_2 \in V(\Gamma)$ be such that C_1 and C_2 are a v_1 -component and a v_2 -component, respectively. As g is a geodesic and $j = i + 1$, it follows that $v_1 \neq v_2$. This implies that we cannot have $C_1 = C = C_2$; on the other hand, if one of the C_i intersected C and the other one was equal to C , then C_1 and C_2 would intersect, contradicting Lemma 6.8. Thus both C_1 and C_2 must intersect C , and so $v_1, v_2 \in \text{link}(v)$. As Γ is square-free and $v_1 \neq v_2$, this implies that g must be a \hat{v} -piece, as required. \square

6.3. Minimal polygonal representations. Throughout this and the next subsections, let Γ be a connected finite simple graph with at least two vertices, and let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of groups. Let $F = \langle S_0 \mid \rangle$ be a finitely generated free group (so that $|S_0| < \infty$), let $S := S_0 \sqcup S_0^{-1} \subset F$, and let $\varphi_i: F \rightarrow \Gamma\mathcal{G}$ be a sequence of homomorphisms. For each $v \in V(\Gamma)$, let $S_v \subseteq S$ be the set of all $s \in S$ such that $\text{supp } \varphi_i(s) \subseteq \text{link}(v)$ ω -almost surely. Suppose that the homomorphisms φ_i are linking, and in particular, $S = \bigcup_{v \in V(\Gamma)} S_v$.

Now for each $v \in V(\Gamma)$, define abstract letters $T_v^{(0)} = \{t_v(s) \mid s \in S_0\}$. Let \bar{F} be a free group with free basis $\bar{S}_0 := S_0 \sqcup (\bigsqcup_{v \in V(\Gamma)} T_v^{(0)})$, so that F is a free factor of \bar{F} . We may extend the homomorphisms $\varphi_i: F \rightarrow \Gamma\mathcal{G}$ to homomorphisms $\bar{\varphi}_i: \bar{F} \rightarrow \Gamma\mathcal{G}$ by setting $\bar{\varphi}_i(t_v(s)) = \pi_v(\varphi_i(s))$, where $\pi_v: \Gamma\mathcal{G} \rightarrow G_v$ is the canonical projection. Let $T_v := T_v^{(0)} \sqcup (T_v^{(0)})^{-1} \subset \bar{F}$ for $v \in V(\Gamma)$, and let $\bar{S} := \bar{S}_0 \sqcup \bar{S}_0^{-1} \subset \bar{F}$.

For each $v \in V(\Gamma)$, let \dot{F}_v and \hat{F}_v be the subgroups of \bar{F} generated by T_v and $\hat{S}_v := S_v \sqcup (\bigsqcup_{w \in \text{link}(v)} T_w)$, respectively. Note that $\bar{\varphi}_i(\dot{F}_v) \subseteq G_v$ and $\bar{\varphi}_i(\hat{F}_v) \subseteq \Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)}$ ω -almost surely. For convenience, for any $v, w \in V(\Gamma)$ and $s \in S_0$ we also set

$$t_v(t_w(s)) = \begin{cases} t_w(s) & \text{if } v = w, \\ 1 & \text{otherwise,} \end{cases}$$

so that we have well-defined functions $t_v: \bar{S}_0 \rightarrow T_v^{(0)} \cup \{1\}$ extending to homomorphisms $t_v: \bar{F} \rightarrow \dot{F}_v$.

The aim of this and the next subsections is to introduce several technical lemmas we will use in a proof of the following result.

Proposition 6.10. *If Γ has girth ≥ 6 , then $\bar{F}_{\omega, (\bar{\varphi}_i)}$ is the normal closure K of*

$$\mathcal{K} = \left(\bigcup_{v \in V(\Gamma)} (\bar{F}_{\omega, (\bar{\varphi}_i)} \cap \dot{F}_v) \right) \cup \left(\bigcup_{v \in V(\Gamma)} [\dot{F}_v, \hat{F}_v] \right)$$

in \bar{F} .

By construction, $[\dot{F}_v, \hat{F}_v] \subseteq \bar{F}_\omega$ ω -almost surely for all $v \in V(\Gamma)$, and so $\mathcal{K} \subseteq \bar{F}_\omega$. Thus, we only need to show that K contains every $f \in \bar{F}$ such that $\bar{\varphi}_i(f) = 1$ ω -almost surely. We will use the remainder of this subsection to prove this. In order to do that, we work with “polygonal representations,” which are defined below and closely related to polygonal van Kampen diagrams introduced in Section 6.2.

Definition 6.11. Let $W = s_1 \cdots s_n$ be a cyclic word over the alphabet \bar{S} . A *polyg-
onal representation* for W consists of the following data.

- A subdivision $W = |s_{j_0} \cdots s_{j_1-1}| |s_{j_1} \cdots s_{j_2-1}| \cdots |s_{j_{k-1}} \cdots s_{j_k-1}|$ of W , where $j_k = j_0$, into subwords $p_\ell = s_{j_{\ell-1}} \cdots s_{j_\ell-1}$ called *pieces*, such that for each $\ell \in \{1, \dots, k\}$ we have $s_{j_{\ell-1}}, \dots, s_{j_\ell-1} \in \hat{S}_v$ for some $v \in V(\Gamma)$.
- An assignment of a *label* from $V(\Gamma) \times \{', \hat{\cdot}\}$ for each piece. For any $v \in V(\Gamma)$, we write \dot{v} for $(v, ')$, and call each piece with label \dot{v} a \dot{v} -piece; we write \hat{v} for $(v, \hat{\cdot})$, and call each piece with label \hat{v} a \hat{v} -piece. We require the assignment to be such that for any piece $p_\ell = s_{j_{\ell-1}} \cdots s_{j_\ell-1}$, we have $s_{j_{\ell-1}}, \dots, s_{j_\ell-1} \in T_v$ if p_ℓ is a \dot{v} -piece, and $s_{j_{\ell-1}}, \dots, s_{j_\ell-1} \in \hat{S}_v$ if p_ℓ is a \hat{v} -piece.

The *length* of a piece p_ℓ is the number $\lim^\omega |\bar{\varphi}_i(p_\ell)| \in \mathbb{N} \cup \{\infty\}$, where $|g|$ is the word length of $g \in \Gamma\mathcal{G}$ over the alphabet $\bigsqcup_{v \in V(\Gamma)} (G_v \setminus \{1\})$. Clearly, any \dot{v} -piece has length ≤ 1 for any $v \in V(\Gamma)$. If a piece p_ℓ has length $L < \infty$, then (as Γ is finite) there exist vertices $v_1, \dots, v_L \in V(\Gamma)$ such that we ω -almost surely have $\bar{\varphi}_i(p_\ell) = g_{i,1} \cdots g_{i,L}$ for some $g_{i,m} \in G_{v_m}$; if in addition we have $v_m \neq v_{m'}$ for $1 \leq m < m' \leq L$, then p_ℓ is said to be a *small* piece. A piece that is not small is called *large*.

Given a finite sequence p_{j_1}, \dots, p_{j_m} of pieces, we say p_{j_1}, \dots, p_{j_m} are *consec-
utive* if $2 \leq m \leq k$ and $j_{i+1} \equiv j_i + 1 \pmod{k}$ for $1 \leq i \leq m - 1$. We say a piece p is *adjacent* to a piece q if either p, q or q, p are consecutive.

It is clear that every word has at least one polygonal representation: for instance, the one where pieces are single letters.

We say two words U, U' over \bar{S} are *equivalent* if UK and $U'K$ are conjugate in \bar{F}/K , and we say two cyclic words W, W' over \bar{S} are *equivalent* if some (equivalently, any) representative words U of W and U' of W' are equivalent. It is clear that this defines an equivalence relation on the set of cyclic words over \bar{S} .

Given a cyclic word W' with a polygonal representation \mathcal{D}' , we denote by $\dot{P}(W', \mathcal{D}')$ and $\hat{P}(W', \mathcal{D}')$ the number of \dot{v} -pieces (for $v \in V(\Gamma)$) and the number of \hat{v} -pieces (for $v \in V(\Gamma)$), respectively, in \mathcal{D}' . A polygonal representation \mathcal{D} for a word W is said to be *minimal* if, for all cyclic words W' equivalent to W and all polygonal representations \mathcal{D}' for W' , we have $(\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D})) \leq (\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}'))$ in the lexicographical order.

Remark 6.12. There is an a priori noticeable discrepancy between the definition of a “small/large piece” in Definition 6.11 and the definition of a “trivial/small/large van Kampen piece” in Definition 6.7. However, in view of parts (i) and (ii) of Lemma 6.13 below, the two definitions are related in the following sense.

Let W be a cyclic word over \bar{S} , let \mathcal{D} be a polygonal representation for W with pieces p_1, \dots, p_k , and suppose that \mathcal{D} is a minimal polygonal representation. Suppose moreover that $W \in \bar{F}_\omega$: in particular, ω -almost surely there exists a polygonal van Kampen diagram Δ_i with van Kampen pieces $g_{i,1}, \dots, g_{i,k}$, where $g_{i,\ell}$ is a geodesic word representing $\bar{\varphi}_i(p_\ell)$ for each ℓ , and where the label of the van Kampen piece $g_{i,\ell}$ coincides with the label of the piece p_ℓ . Then p_ℓ is a small (respectively large) piece if and only if $g_{i,\ell}$ is ω -almost surely a small (respectively large) van Kampen piece.

Since in the subsequent argument we do not consider polygonal representations that are not minimal, we chose our terminology the way we did in Definitions 6.7 and 6.11.

Lemma 6.13. *A minimal polygonal representation satisfies the following:*

- (i) *there are no pieces of length 0;*
- (ii) *every small piece is a \hat{v} -piece for some $v \in V(\Gamma)$ (in particular, any small piece has length 1, and there are no pieces of length 2);*
- (iii) *if p_1, \dots, p_m are consecutive pieces, p_1 and p_m are a \hat{w}_1 -piece and a \hat{v} -piece (respectively) for some $w_1, v \in V(\Gamma)$, and, for each $j \in \{2, \dots, m\}$, p_j is a \hat{w}_j -piece for some $w_j \in V(\Gamma)$, then we have $\{w_1, w_2, \dots, w_{m-1}\} \not\subseteq \text{link}(v)$;*
- (iv) *if p_1, \dots, p_m are consecutive pieces, p_1 and p_m are \hat{v} -pieces for some $v \in V(\Gamma)$, and, for each $j \in \{2, \dots, m-1\}$, p_j is a \hat{w}_j -piece for some $w_j \in V(\Gamma)$, then $\{w_2, \dots, w_{m-1}\} \not\subseteq \text{link}(v)$;*
- (v) *no two pieces with the same label are adjacent;*
- (vi) *for any $v \in V(\Gamma)$ and $w \in \text{link}(v)$, no \hat{w} -piece is adjacent to a \hat{v} -piece;*
- (vii) *for any $v \in V(\Gamma)$ and $w \in \text{link}(v)$, no \hat{v} -piece is adjacent to both a \hat{w} -piece and a \hat{v} -piece;*
- (viii) *for any $v \in V(\Gamma)$, no \hat{v} -piece is adjacent to two \hat{v} -pieces;*
- (ix) *for any $v \in V(\Gamma)$, no \hat{v} -piece is adjacent to two \hat{v} -pieces.*

Proof. Let \mathcal{D} be a minimal polygonal representation for a word W . We claim that \mathcal{D} satisfies the conclusions of the Lemma. We check the parts (i)–(ix) in order. Some of these parts are visualised in Figure 9.

- (i) Note that if p is a piece of length 0, then $\bar{\varphi}_i(p) = 1$ ω -almost surely, and so $p \in \bar{F}_\omega \cap \hat{F}_u$, where $u = v$ if p is a \hat{v} -piece and u is any vertex in $\text{link}(v)$ if p is a \hat{v} -piece (such a vertex u exists since Γ is connected).

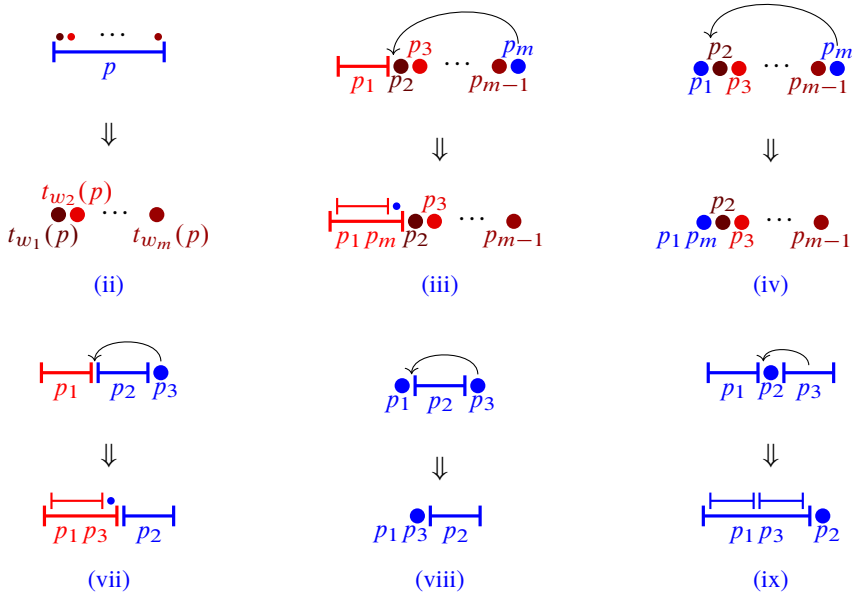


Figure 9. Proof of Lemma 6.13. The colours represent different vertices, v being blue and w_i reddish. The \hat{u} -pieces and \hat{u} -pieces (for some $u \in V(\Gamma)$) are denoted by \bullet and --- , respectively.

In particular, $p \in \mathcal{K} \subseteq K$, and so deleting the subword p of W and the piece p from \mathcal{D} results in a new word W' , equivalent to W , and a polygonal representation \mathcal{D}' for W' . But then we either have $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D}) - 1)$ or $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}) - 1, \dot{P}(W, \mathcal{D}))$. This contradicts minimality of \mathcal{D} .

- (ii) Suppose that p is a small \hat{v} -piece for some $v \in V(\Gamma)$. Then there exist distinct vertices $w_1, \dots, w_m \in \text{link}(v)$ such that we have $\bar{\varphi}_i(p) = g_{i,1} \cdots g_{i,m}$ ω -almost surely, where $g_{i,\ell} \in G_{w_\ell}$. In particular, we have $(t_{w_1}(p) \cdots t_{w_m}(p))^{-1} p \in \bar{F}_\omega \cap \hat{F}_v \subseteq K$. Thus, replacing the subword p of W with a subword $t_{w_1}(p) \cdots t_{w_m}(p)$ results in a new word W' , equivalent to W , and replacing the \hat{v} -piece p in \mathcal{D} with m pieces $t_{w_1}(p), \dots, t_{w_m}(p)$ (with labels $\hat{w}_1, \dots, \hat{w}_m$, respectively) we obtain a polygonal representation \mathcal{D}' for W' . However, then \mathcal{D}' satisfies $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}) - 1, \dot{P}(W, \mathcal{D}) + \ell)$, contradicting minimality of \mathcal{D} .
- (iii) Suppose for contradiction that this is not the case. We then have $[p_m, p_2 \cdots p_{m-1}] \in [\hat{F}_v, \bar{F}_v] \subseteq K$, and so multiplying (a conjugate of) the word $W = p_1 \cdots p_m W_0$ by $[p_m, p_2 \cdots p_{m-1}]$ we get $W' = p_1 p_m p_2 \cdots p_{m-1} W_0$ equivalent to W . We define a polygonal representation \mathcal{D}' for W' by letting the pieces be $p_1 p_m, p_2, \dots, p_{m-1}$ together with the pieces of W_0 ,

where $p_1 p_m$ is a \hat{w}_1 -piece and all the other pieces have the same labels as the corresponding pieces in \mathcal{D} . But then $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D}) - 1)$, again contradicting minimality of \mathcal{D} .

- (iv) Suppose for contradiction that this is not the case. We construct a word W' equivalent to W and a polygonal representation \mathcal{D}' for W the same way as in part (iii), except for letting $p_1 p_m$ be a \hat{v} -piece this time. Then again $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D}) - 1)$, contradicting minimality of \mathcal{D} .
- (v) If two adjacent pieces p and q have the same label, then we can concatenate them (resulting in a single piece with the same label) to obtain a new polygonal representation \mathcal{D}' of W . But then we have $(\hat{P}(W, \mathcal{D}'), \dot{P}(W, \mathcal{D}')) = (\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D}) - 1)$ if p and q are \hat{v} -pieces, and $(\hat{P}(W, \mathcal{D}'), \dot{P}(W, \mathcal{D}')) = (\hat{P}(W, \mathcal{D}) - 1, \dot{P}(W, \mathcal{D}))$ if p and q are \hat{v} -pieces (for some $v \in V(\Gamma)$). This contradicts minimality of \mathcal{D} .
- (vi) The case when p, q are consecutive pieces such that p is a \hat{w} -piece and q is a \hat{w} -piece (for some $w \in \text{link}(v)$) is exactly the case $m = 2$ in part (iii). If instead q, p are consecutive, a symmetric argument works.
- (vii) Suppose that p_1, p_2, p_3 are consecutive pieces such that p_1, p_2 and p_3 have labels \hat{w}, \hat{v} and \hat{v} , respectively; if instead p_1, p_2 and p_3 have labels \hat{v}, \hat{v} , and \hat{w} , respectively, then a symmetric argument works. We then have $[p_3, p_2] \in [\hat{F}_v, \hat{F}_v] \subseteq K$, and so multiplying (a conjugate of) $W = p_1 p_2 p_3 W_0$ by $[p_3, p_2]$ we get a word $W' = p_1 p_3 p_2 W_0$ equivalent to W . We define a polygonal representation \mathcal{D}' for W' by letting the pieces be $p_1 p_3$ and p_2 together with the pieces of W_0 , where $p_1 p_3$ is a \hat{w} -piece and all the other pieces have the same labels as the corresponding pieces in \mathcal{D} . But then $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}), \dot{P}(W, \mathcal{D}) - 1)$, again contradicting minimality of \mathcal{D} .
- (viii) This follows exactly as in part (vii), except that the labels of the piece p_1 of \mathcal{D} and the piece $p_1 p_3$ of \mathcal{D}' are now both \hat{v} instead of \hat{w} .
- (ix) Suppose that p_1, p_2, p_3 are consecutive pieces with labels \hat{v}, \hat{v} and \hat{w} , respectively. We then have $[p_2, p_3] \in [\hat{F}_v, \hat{F}_v] \subseteq K$, and so multiplying a conjugate of $W = p_1 p_2 p_3 W_0$ by $[p_2, p_3]^{-1}$ we get a word $W' = p_1 p_3 p_2 W_0$ equivalent to W . We define a polygonal representation \mathcal{D}' for W' by letting the pieces be $p_1 p_3$ and p_2 together with the pieces of W_0 , where $p_1 p_3$ is a \hat{v} -piece and all the other pieces have the same labels as the corresponding pieces in \mathcal{D} . But then $(\hat{P}(W', \mathcal{D}'), \dot{P}(W', \mathcal{D}')) = (\hat{P}(W, \mathcal{D}) - 1, \dot{P}(W, \mathcal{D}))$, again contradicting minimality of \mathcal{D} . \square

6.4. Graphs of large girth. The idea of our proof of Proposition 6.10 is to use properties of polygonal van Kampen diagrams analogous to the ones enjoyed by minimal polygonal representations. In particular, we define the following.

Definition 6.14. A polygonal van Kampen diagram is said to be *minimal* if Δ satisfies the conditions (i)–(ix) in Lemma 6.13, with each appearance of the word “piece” replaced with “van Kampen piece.”

In particular, we will refer to Lemma 6.13 when talking about minimal polygonal van Kampen diagrams. Therefore, in the following discussion we will use the fact that minimal polygonal van Kampen diagrams have no trivial pieces (by Lemma 6.13 (i)), and so all their van Kampen pieces are either small or large. The following results (Lemma 6.15, Lemma 6.16 and Corollary 6.17) are the main auxiliary results required in our proof of Proposition 6.10.

Lemma 6.15. *Let Δ be a minimal polygonal van Kampen diagram, and suppose that Γ has girth ≥ 6 . If $(C, [P, Q])$ is a component of Δ that is either trivial or minimal, then the following hold:*

- (i) $(C, [P, Q])$ does not enclose any other components;
- (ii) if C is supported at a van Kampen piece g , then g is large;
- (iii) if C is supported at a van Kampen piece g , then either $P \in \mathcal{J}_g(\Delta)$ and C is the last component supported at g , or $Q \in \mathcal{J}_g(\Delta)$ and C is the first component supported at g .

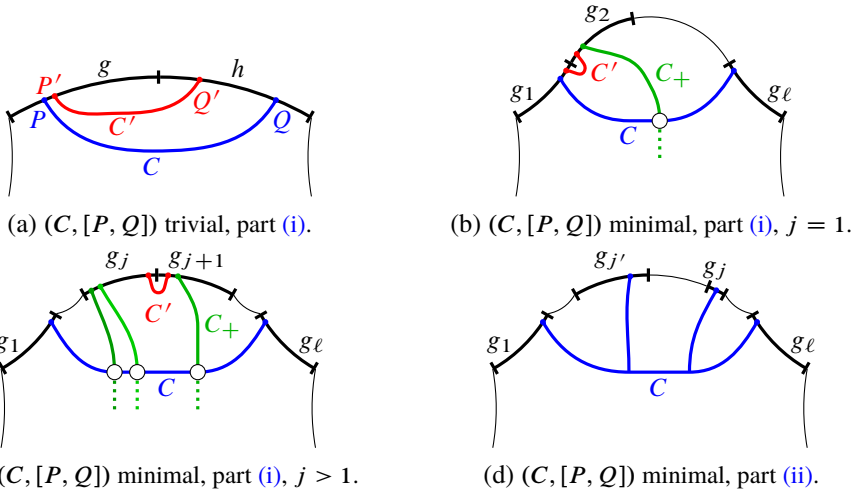


Figure 10. Proof of Lemma 6.15. $\begin{matrix} C'_0 \\ C_0 \end{matrix} \begin{matrix} \diagup \\ \diagdown \end{matrix}$ denotes that a component C'_0 either intersects or coincides with a component C_0 .

Proof. Let $(C, [P, Q])$ be an oriented v -component. We will show properties (i)–(iii) when $(C, [P, Q])$ is trivial (see (T.i)–(T.iii) below) and when it is minimal (see (M.i)–(M.iii) below) separately.

($C, [P, Q]$) is trivial. There are two consecutive van Kampen pieces g, h such that $P \in \mathcal{J}_g(\Delta)$ and $Q \in \mathcal{J}_h(\Delta)$.

(T.i) Let $i \geq 1$ be such that C is the i -th component supported at g . Suppose that $(C, [P, Q])$ encloses a component $(C', [P', Q'])$. Then we must have $P' \in \mathcal{J}_g(\Delta)$, $Q' \in \mathcal{J}_h(\Delta)$, and C' must be the j -th component supported at g for some $j > i$; see Figure 10a. But then, whenever C'' is the j' -th component supported at g , where $i + 1 \leq j' \leq j$, we know by Lemma 6.8 that C and C'' do not intersect, and so $(C, [P, Q])$ encloses C'' . In particular, we may assume, without loss of generality, that $j = i + 1$.

Now let $v' \in V(\Gamma)$ be the label of C' . As g is a geodesic, it follows that $v' \neq v$. As Γ is square-free, it follows that $|\text{link}(v) \cap \text{link}(v')| \leq 1$. Therefore, as both C and C' are supported at both g and h , it follows that $\text{link}(v) \cap \text{link}(v') = \{w\}$ and that g and h are both \hat{w} -pieces. But this contradicts Lemma 6.13 (v). Thus $(C, [P, Q])$ cannot enclose any other component, as required.

(T.ii) If g (or h) is small then it is a \hat{v} -component by Lemma 6.13 (ii), and if it is large then it is a \hat{w} -component for some $w \in \text{link}(v)$. Therefore, if g and h were both small, then they would be two adjacent \hat{v} -pieces, contradicting Lemma 6.13 (v); if one of them was small and the other one large, then they would be a \hat{v} -piece adjacent to a \hat{w} -piece for some $w \in \text{link}(v)$, contradicting Lemma 6.13 (vi). Thus both g and h must be large, as required.

(T.iii) As g and h are adjacent, it is clear that C is not supported at van Kampen pieces other than g and h . If C is not the first component supported at h , then let C' be the first component supported at h . By Lemma 6.8, C and C' do not intersect, and so $(C, [P, Q])$ must enclose C' —but this contradicts part (T.i). Thus C must be the first component supported at h ; similarly, C must be the last component supported at g .

($C, [P, Q]$) is minimal. Let g_1, \dots, g_ℓ be consecutive van Kampen pieces of Δ such that $P \in \mathcal{J}_{g_1}(\Delta)$ and $Q \in \mathcal{J}_{g_\ell}(\Delta)$.

(M.i) Suppose that $(C, [P, Q])$ encloses a v' -component $(C', [P', Q'])$ for some $v' \in V(\Gamma)$. As $(C, [P, Q])$ is minimal, it follows that $(C', [P', Q'])$ is trivial, and so $P' \in \mathcal{J}_{g_j}(\Delta)$ and $Q' \in \mathcal{J}_{g_{j+1}}(\Delta)$ for some $j \in \{1, \dots, \ell-1\}$. Without loss of generality, assume that C' is chosen in such a way that j is as small as possible. By part (T.ii) applied to $(C', [P', Q'])$, g_{j+1} is a large van Kampen piece, and so there are at least 3 components supported at g_{j+1} . Therefore, the second component C_+ supported at g_{j+1} —say a v_+ -component for some $v_+ \in V(\Gamma)$ —is not trivial by part (T.iii), and so, as $(C, [P, Q])$ is minimal, C_+ must either intersect C or coincide with C . As g_{j+1} is a geodesic and as C', C_+ are the first and the second

components supported at g_{j+1} (respectively), we also have $v_+ \neq v'$. Moreover, as g_{j+1} is large, it must be a \hat{u} -piece for some $u \in V(\Gamma)$ by Lemma 6.13 (ii).

Suppose first that $j = 1$; see Figure 10b. Then the components C and C' are supported at g_1 , and so g_1 is a large van Kampen piece; therefore, by Lemma 6.13 (ii), g_1 is a \hat{w} -piece for some $w \in \text{link}(v) \cap \text{link}(v')$. Now if we had $C = C_+$, then we would have $v = v_+ \neq v'$, and so, as C and C' are supported at both g_1 and g_2 , both g_1 and g_2 would have to be \hat{u} -pieces for the *unique* (as Γ is square-free) vertex $u \in \text{link}(v) \cap \text{link}(v')$, contradicting Lemma 6.13 (v). On the other hand, if C_+ intersected C , then v, w, v', u, v_+, v would be a closed walk on Γ of length 5, contradicting the fact that Γ is triangle-free and pentagon-free. Thus we arrive at a contradiction in either case, and so we cannot have $j = 1$.

Suppose now that $j > 1$: then $\mathcal{J}_{g_j}(\Delta) \subseteq [P, Q]$, and so, by minimality of $(C, [P, Q])$, C must either intersect or coincide with every non-trivial component supported at g_j . Moreover, g_j must be a large piece by part (T.ii), and so there must be at least 3 components supported at g_j by Lemma 6.13 (ii). In particular, by part (T.iii) and the minimality of j , the first and the second components supported at g_j are non-trivial, and so they must each either intersect or coincide with C ; see Figure 10c. It follows from Lemma 6.9 that g_j is a \hat{v} -piece, and in particular $v' \in \text{link}(v)$. But now if we had $C_+ = C$, then v, v', u would span a triangle in Γ , contradicting the fact that Γ is triangle-free. On the other hand, if C_+ intersected C , then—as $v_+ \neq v'$ and C', C_+ are both supported at g_{j+1} , and as Γ is square-free—it would follow that g_{j+1} must be a \hat{v} -piece, as well as g_j , contradicting Lemma 6.13 (v). Thus we cannot have $j \neq 1$ either.

(M.ii) Suppose that C is supported at a small van Kampen piece g_j , and so g_j is a \hat{v} -piece. Then either $j > 1$ or $j < \ell$; without loss of generality, assume that $j > 1$. Let $j' < j$ be largest such that C is supported at j' , so that $j' \geq 1$; see Figure 10d. It follows from part (M.i) that $(C, [P, Q])$ does not enclose any components. In particular, for each $m \in \{j'+1, \dots, j-1\}$, C intersects every component supported at m , and so g_m is either a \hat{v} -piece or a \hat{w}_m -piece for some $w_m \in \text{link}(v)$.

Now if g_m is a \hat{v} -piece for some $m \in \{j'+1, \dots, j-1\}$, then we must have $j = j' + 2$ and $m = j' + 1$: indeed, otherwise one of g_{m-1} and g_{m+1} is either a \hat{v} -piece or a \hat{w} -piece for some $w \in \text{link}(v)$, contradicting either Lemma 6.13 (v) or Lemma 6.13 (vi). But then $g_{j'+1}$ is a \hat{v} -piece adjacent to a \hat{v} -piece $g_{j'+2}$ and to $g_{j'}$, which is either a \hat{v} -piece or a \hat{w} -piece for some $w \in \text{link}(v)$ (as C is supported at j'); this contradicts either Lemma 6.13 (vii) or Lemma 6.13 (viii). On the other hand, if for

all $m \in \{j' + 1, \dots, j - 1\}$, g_m is a \hat{w}_m -piece for some $w_m \in \text{link}(v)$, then, as before, $g_{j'}$ is either a \hat{v} -piece or a \hat{w} -piece for some $w \in \text{link}(v)$, contradicting either Lemma 6.13 (iii) or Lemma 6.13 (iv). Thus in either case, g_j cannot be small, as required.

(M.iii) Suppose that C is supported at g_j for some $j \in \{1, \dots, \ell\}$. If C is not the first component supported at g_j , then the first component C' supported at g_j does not intersect C by Lemma 6.8; but $(C, [P, Q])$ does not enclose C' by part (M.i), implying that $j = 1$. Similarly, if C is not the last component supported at g_j , then $j = \ell$. But by part (M.ii), the van Kampen piece g_j is large, and so C cannot be both the first and the last component supported at g_j . As $1 \neq \ell$, it follows that either $j = 1$ and C is the last component supported at g_j , or $j = \ell$ and C is the first component supported at g_j , as required. \square

Lemma 6.16. *Let Δ be a minimal polygonal van Kampen diagram, and suppose that Γ has girth ≥ 6 . Then Δ has no almost minimal components.*

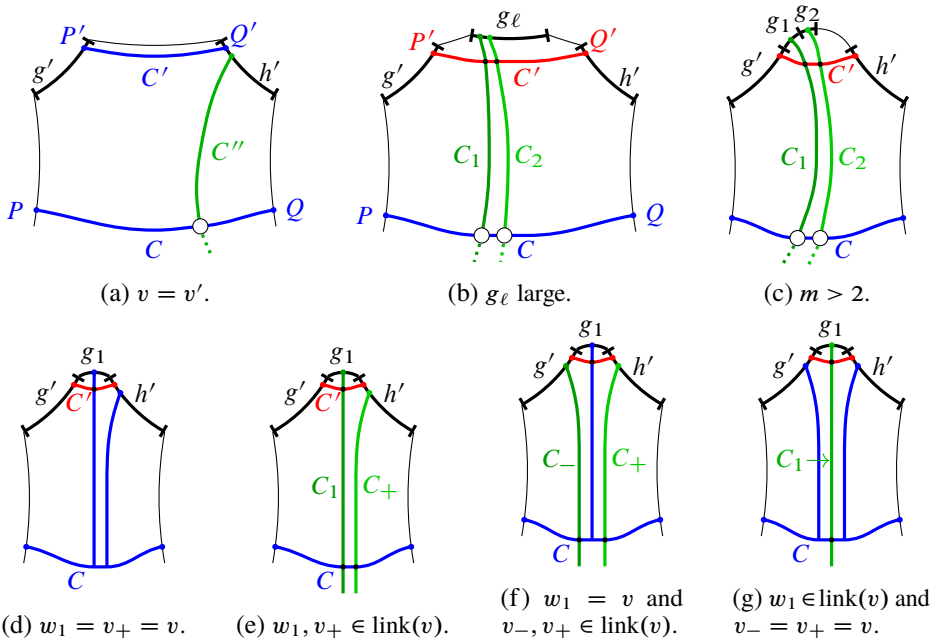


Figure 11. Ruling out impossible scenarios in the proof of Lemma 6.16.

Proof. Suppose that $(C, [P, Q])$ is an almost minimal v -component of Δ . Let g and h be the van Kampen pieces of Δ such that $P \in \mathcal{J}_g(\Delta)$ and $Q \in \mathcal{J}_h(\Delta)$. Then $(C, [P, Q])$ must enclose a minimal v' -component $(C', [P', Q'])$ (for some

$v' \in V(\Gamma)$); let g' and h' be the van Kampen pieces of Δ such that $P' \in \mathcal{J}_{g'}(\Delta)$ and $Q' \in \mathcal{J}_{h'}(\Delta)$. Suppose, without loss of generality, that $(C', [P', Q'])$ is chosen in such a way that $[P, P']$ is as small as possible: that is, if $(C'_0, [P'_0, Q'_0])$ is any minimal component enclosed by $(C, [P, Q])$, then $[P, P'] \subseteq [P, P'_0]$.

We first claim that $v \neq v'$. Indeed, as $(C', [P', Q'])$ is minimal and C' is supported at h' , it follows from Lemma 6.15 (ii) that h' is a large van Kampen piece, and so, by Lemma 6.13 (ii), there are at least 3 components supported at h' . Therefore, the second component supported at h' —a v'' -component C'' , say—is neither trivial nor minimal by Lemma 6.15 (iii), and so it must either intersect or coincide with C ; see Figure 11a. As h' is a geodesic and C' and C'' are the first (by Lemma 6.15 (iii)) and the second components (respectively) supported at h' , it follows that $v' \neq v''$. If $C = C''$, then this immediately implies $v \neq v'$. On the other hand, as C' and C'' are supported at h' , it follows that h' must be a \hat{u} -piece for some $u \in \text{link}(v') \cap \text{link}(v'')$. Thus, if C and C'' intersect, then $v'' \in \text{link}(v)$, and so the fact that Γ is triangle-free implies that $v \neq v'$. Thus we have $v \neq v'$ in either case, as claimed.

Now let $g' = g_0, g_1, \dots, g_m = h'$ be consecutive van Kampen pieces in Δ . We claim that the pieces g_1, \dots, g_{m-1} are small. Indeed, suppose for contradiction that g_ℓ is large for some $\ell \in \{1, \dots, m-1\}$, and so, by Lemma 6.13 (ii), there are at least three components supported at g_ℓ . Let C_1 and C_2 be the first and the second components supported at g_ℓ , respectively. As $(C', [P', Q'])$ is minimal and not supported at g_ℓ , it follows from Lemma 6.15 (i) that C' intersects both C_1 and C_2 , and that neither C_1 nor C_2 is trivial; see Figure 11b. As C_2 is neither the first nor the last component supported at g_ℓ , it follows from Lemma 6.15 (iii) that C_2 is not minimal, and so must either intersect or coincide with C . On the other hand, if C_1 is not minimal, then it must likewise either intersect or coincide with C ; whereas if $(C_1, [P_1, Q_1])$ is minimal for some $[P_1, Q_1]$, then we cannot have $P_1 \in [P', Q']$ by minimality of $(C', [P', Q'])$, and we cannot have $P_1 \in [P, P']$ by minimality of $[P, P']$, so $P_1 \notin [P, Q]$ and C must intersect C_1 . Thus C_1 and C_2 are two components supported at g_ℓ and each either intersecting or coinciding with C and C' , so by Lemma 6.9 g_ℓ must be both a \hat{v} -piece and a \hat{v}' -piece, which is impossible as $v \neq v'$.

Thus the pieces g_1, \dots, g_{m-1} are all small, as claimed. It follows that, for each $\ell \in \{1, \dots, m-1\}$, g_ℓ is a \hat{w}_ℓ -piece for some $w_\ell \in V(\Gamma)$, and there is a unique w_ℓ -component C_ℓ supported at g_ℓ . Thus, by Lemma 6.15 (i), each C_ℓ intersects C' and so $w_\ell \in \text{link}(v')$ for each ℓ . Moreover, it also follows from Lemma 6.15 (ii) that each C_ℓ is neither trivial nor minimal, and so must either intersect or coincide with C ; therefore, $w_\ell \in \text{link}(v) \cup \{v\}$.

We now claim that $m = 2$. Indeed, as $(C', [P', Q'])$ is not trivial, we cannot have $m < 2$. On the other hand, if $m > 2$, then consider $w_1, w_2 \in \text{link}(v')$; see Figure 11c. By Lemma 6.13 (v), we have $w_1 \neq w_2$, and so we cannot have $w_1 = v$ and $w_2 = v$. We also cannot have $w_1 = v$ and $w_2 \in \text{link}(v)$: indeed, otherwise v, v', w_2 span a triangle in Γ , contradicting the fact that Γ is triangle-free; similarly, $w_1 \in \text{link}(v)$

and $w_2 = v$ is impossible. Finally, we have $|\text{link}(v) \cap \text{link}(v')| \leq 1$ (as Γ is square free and $v \neq v'$), so as $w_1 \neq w_2$, we cannot have $w_1, w_2 \in \text{link}(v)$. Thus we must have $m = 2$, as claimed. In particular, the van Kampen pieces g', g_1, h' are consecutive, and g_1 is a \hat{w}_1 -piece for some $w_1 \in \text{link}(v') \cap (\text{link}(v) \cup \{v\})$.

Now by Lemma 6.15 (ii) and Lemma 6.15 (iii), C' is the last component supported at g' and the first component supported at h' , and both g', h' are large. In particular, by Lemma 6.13 (ii), there are at least 3 components supported at each g' and h' . Moreover, by Lemma 6.15 (iii), the penultimate component supported at g' and the second component supported at h' —say a v_- -component C_- and a v_+ -component C_+ , respectively—are both not trivial and not minimal, and so they must each either intersect or coincide with C ; in particular, $v_-, v_+ \in \text{link}(v) \cup \{v\}$.

Now if $w_1 = v_+ = v$ (see Figure 11d), then $C_+ = C = C_1$ and in particular C_+ and C' intersect, contradicting Lemma 6.8 (as both C_+ and C' are supported at h'); similarly, we cannot have $w_1 = v_- = v$. If $w_1, v_+ \in \text{link}(v)$ (see Figure 11e) then, as h' is large, we know that h' is a \hat{u} -piece for some $u \in \text{link}(v') \cap \text{link}(v_+)$, and v, v_+, u, v', w_1, v is a closed walk in Γ of length 5, contradicting the fact that Γ is triangle-free and pentagon-free; similarly, we cannot have $w_1, v_- \in \text{link}(v)$.

It follows that either $w_1 = v$ and $v_-, v_+ \in \text{link}(v)$ (see Figure 11f), or $w_1 \in \text{link}(v)$ and $v_- = v_+ = v$ (see Figure 11g). In either case, C' and C_- are two pieces supported at g' and intersecting C_1 , and so g' is a \hat{w}_1 -piece by Lemma 6.9; similarly, h' is a \hat{w}_1 -piece. But then g_1 is a \hat{w}_1 -piece adjacent to two \hat{w}_1 -pieces, contradicting Lemma 6.13 (ix). This concludes the proof. \square

Corollary 6.17. *Let Δ be a minimal polygonal van Kampen diagram, and suppose that Γ has girth ≥ 6 . Then Δ has 0 van Kampen pieces.*

Proof. We claim first that all components of Δ are either minimal or trivial. Indeed, otherwise Δ has a non-minimal non-trivial component $(C, [P, Q])$ with $[P, Q]$ minimal (with respect to inclusion) among all such components. Thus $(C, [P, Q])$ cannot enclose any non-minimal non-trivial component by minimality of $[P, Q]$, and so $(C, [P, Q])$ is almost minimal. This contradicts Lemma 6.16.

Thus, as all components of Δ are either minimal or trivial, it follows by Lemma 6.15 (ii) that Δ has no small van Kampen pieces. Suppose Δ has at least one van Kampen piece. Then any piece p of Δ has length ≥ 3 by Lemma 6.13 (ii). Now if C is the second component of Δ supported at a piece p , then C is a component that is either minimal or trivial, but C is neither the first nor the last component supported at p . This contradicts Lemma 6.15 (iii).

Thus Δ has no van Kampen pieces, as required. \square

6.5. Proofs of Proposition 6.10 and Theorem E. In this subsection, we deduce Proposition 6.10 from the results in Section 6.4. We then use Proposition 6.10 and Theorem 6.6 to prove Theorem E.

Proof of Proposition 6.10. As noted after the statement of the Proposition, it is clear that $K \subseteq \bar{F}_\omega$, and so we only need to show that $\bar{F}_\omega \subseteq K$. This is equivalent to saying that every (cyclic) word $W \in \bar{F}_\omega$ is equivalent to the trivial word $1 \in \bar{F}_\omega$ of length 0. As the trivial word has a polygonal representation \mathcal{D}_0 with 0 pieces, which is clearly minimal, it is enough to show that \mathcal{D}_0 is the *unique* minimal polygonal representation for any word $W \in \bar{F}_\omega$.

Thus, let \mathcal{D} be a minimal polygonal representation for a word $W \in \bar{F}_\omega$, and let p_1, \dots, p_k be the pieces of \mathcal{D} , so that $W = p_1 \cdots p_k$. Since $\bar{\varphi}_i(W) = 1$ ω -almost surely, we may ω -almost surely define a polygonal van Kampen diagram Δ_i for $W_i = g_{i,1} \cdots g_{i,k}$ whose van Kampen pieces are $g_{i,1}, \dots, g_{i,k}$, where $g_{i,\ell}$ is a geodesic word over $\bigsqcup_{v \in V(\Gamma)} (G_v \setminus \{1\})$ representing $\bar{\varphi}_i(p_\ell)$, such that the label of a piece $g_{i,\ell}$ is v whenever $g_{i,\ell}$ represents an element in $G_v \setminus \{1\}$ for some $v \in V(\Gamma)$. Let $\mathcal{A} \subseteq \mathbb{N}$ for the set of all $i \in \mathbb{N}$ such that

- $\bar{\varphi}_i(W) = 1$;
- for each $\ell \in \{1, \dots, k\}$ and each $j \in \{0, 1\}$, the piece p_ℓ has length j if and only if the van Kampen piece $g_{i,\ell}$ has length j ;
- for each $\ell \in \{1, \dots, k\}$, the label of the piece p_ℓ of \mathcal{D} coincides with the label of the van Kampen piece $g_{i,\ell}$ of Δ_i .

It follows from Lemma 6.13 (i), Lemma 6.13 (ii) and the construction that \mathcal{A} is an ω -large subset of \mathbb{N} . But note that as \mathcal{D} is minimal, the polygonal van Kampen diagram Δ_i is minimal for each $i \in \mathcal{A}$. It then follows from Corollary 6.17 that Δ_i has 0 van Kampen pieces for each $i \in \mathcal{A}$ (and so ω -almost surely). Thus \mathcal{D} has 0 pieces, and so $\mathcal{D} = \mathcal{D}_0$, as required. \square

Proof of Theorem E. If $|V(\Gamma)| = 1$, then we have $\Gamma\mathcal{G} \cong G_v$ for the unique vertex v of Γ , and so $\Gamma\mathcal{G}$ is equationally noetherian by the assumption. Thus, we may assume, without loss of generality, that $|V(\Gamma)| \geq 2$.

Suppose first that Γ is connected. We claim that Γ is admissible. Thus, let F be a finitely generated free group, and let $(\varphi_i : F \rightarrow \Gamma\mathcal{G})_{i=1}^\infty$ be a sequence of linking homomorphisms. As in the beginning of Section 6.3, define a free group \bar{F} containing F and subgroups $\hat{F}_v, \tilde{F}_v \leq \bar{F}$ for each $v \in V(\Gamma)$, and extend the homomorphisms φ_i to $\bar{\varphi}_i : \bar{F} \rightarrow \Gamma\mathcal{G}$. It follows from the construction that $\ker(\varphi_i) = \ker(\bar{\varphi}_i) \cap F$ for each i , and so $F_{\omega,(\varphi_i)} = \bar{F}_{\omega,(\bar{\varphi}_i)} \cap F$.

Now for any $v \in V(\Gamma)$, consider homomorphisms $\varphi_i^{(v)} : \hat{F}_v \rightarrow \Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)}$ given by $\varphi_i^{(v)}(g) = \bar{\varphi}_i(g)$. As Γ is square-free, we have $\Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)} \cong *_{w \in \text{link}(v)} G_w$; in particular, it follows from Theorem 1.7 (and $|V(\Gamma)| < \infty$) that $\Gamma_{\text{link}(v)}\mathcal{G}_{\text{link}(v)}$ is equationally noetherian. It follows that $(\hat{F}_v)_{\omega,(\varphi_i^{(v)})} \subseteq \ker(\varphi_i^{(v)})$ ω -almost surely. But it follows from the construction that $\ker(\varphi_i^{(v)}) = \ker(\bar{\varphi}_i) \cap \hat{F}_v$ for each i , and in particular $(\hat{F}_v)_{\omega,(\varphi_i^{(v)})} = \bar{F}_{\omega,(\bar{\varphi}_i)} \cap \hat{F}_v$. Thus $\bar{F}_{\omega,(\bar{\varphi}_i)} \cap \hat{F}_v \subseteq \ker(\bar{\varphi}_i)$ ω -almost surely.

Note that, by our construction, we have $[\hat{F}_v, \hat{F}_v] \subseteq \ker(\bar{\varphi}_i)$ for each i and each $v \in V(\Gamma)$. Together with the fact that $\bar{F}_{\omega, (\bar{\varphi}_i)} \cap \hat{F}_v \subseteq \ker(\bar{\varphi}_i)$ ω -almost surely for each $v \in V(\Gamma)$, this implies (by Proposition 6.10) that $\bar{F}_{\omega, (\bar{\varphi}_i)} \subseteq \ker(\bar{\varphi}_i)$ ω -almost surely. In particular, it follows that $F_{\omega, (\varphi_i)} = \bar{F}_{\omega, (\bar{\varphi}_i)} \cap F \subseteq \ker(\bar{\varphi}_i) \cap F = \ker(\varphi_i)$ ω -almost surely. Thus Γ is admissible, as claimed. In particular, it follows from Theorem 6.6 that $\Gamma\mathcal{G}$ is equationally noetherian.

Finally, suppose that the graph Γ is not connected. Then we have a partition $V(\Gamma) = A_1 \sqcup \cdots \sqcup A_m$ into non-empty subsets such that $\Gamma = \Gamma_{A_1} \sqcup \cdots \sqcup \Gamma_{A_m}$ is a disjoint union of connected subgraphs Γ_{A_i} , and so $\Gamma\mathcal{G} \cong \Gamma_{A_1}\mathcal{G}_{A_1} * \cdots * \Gamma_{A_m}\mathcal{G}_{A_m}$. By the argument above, it follows that $\Gamma_{A_i}\mathcal{G}_{A_i}$ is equationally noetherian for each i . Therefore, by Theorem 1.7, $\Gamma\mathcal{G}$ is equationally noetherian as well, as required. \square

References

- [1] C. Abbott, S. Balasubramanya, and D. V. Osin, Hyperbolic structures on groups. *Algebr. Geom. Topol.* **19** (2019), no. 4, 1747–1835. [Zbl 07121514](#) [MR 3995018](#)
- [2] C. Abbott, J. Behrstock, and M. G. Durham, Largest acylindrical actions and stability in hierarchically hyperbolic groups. *Trans. Amer. Math. Soc. Ser. B* **8** (2021), 66–104. [Zbl 07319368](#) [MR 4215647](#)
- [3] H.-J. Bandelt, H. M. Mulder, and E. Wilkeit, Quasi-median graphs and algebras. *J. Graph Theory* **18** (1994), no. 7, 681–703. [Zbl 0810.05057](#) [MR 1297190](#)
- [4] G. Baumslag, A. Myasnikov, and V. Remeslennikov, Algebraic geometry over groups. I. Algebraic sets and ideal theory. *J. Algebra* **219** (1999), no. 1, 16–79. [Zbl 0938.20020](#) [MR 1707663](#)
- [5] J. Behrstock, M. F. Hagen, and A. Sisto, Hierarchically hyperbolic spaces. I. Curve complexes for cubical groups. *Geom. Topol.* **21** (2017), no. 3, 1731–1804. [Zbl 1439.20043](#) [MR 3650081](#)
- [6] F. Berlai and B. Robbio, A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.* **14** (2020), no. 4, 1127–1203. [MR 4186470](#)
- [7] B. Brešar, J. Chalopin, V. Chepoi, T. Gologranc, and D. Osajda, Bucolic complexes. *Adv. Math.* **243** (2013), 127–167. [Zbl 1314.05238](#) [MR 3062742](#)
- [8] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften, 319. Springer-Verlag, Berlin, 1999. [Zbl 0988.53001](#) [MR 1744486](#)
- [9] M. Burger and S. Mozes, Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.* **324** (1997), no. 7, 747–752. [Zbl 0966.20013](#) [MR 1446574](#)
- [10] J. Crisp and B. Wiest, Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.* **4** (2004), 439–472. [Zbl 1057.20028](#) [MR 2077673](#)

- [11] A. Genevois, *Cubical-like geometry of quasi-median graphs and applications to geometric group theory*. Ph.D. thesis. Aix-Marseille Université, Marseille, 2017.
- [12] A. Genevois, Automorphisms of graph products of groups and acylindrical hyperbolicity. Preprint, 2018. [arXiv:1807.00622v3](https://arxiv.org/abs/1807.00622v3) [math.GR]
- [13] A. Genevois, A cylindrical hyperbolicity from actions on CAT(0) cube complexes: a few criteria. *New York J. Math.* **25** (2019), 1214–1239. [Zbl 07130524](https://zbmath.org/journal/10.1007/s00036-019-01214-1) [MR 4028832](https://doi.org/10.1007/s00036-019-01214-1)
- [14] A. Genevois, On the geometry of van Kampen diagrams of graph products of groups. Preprint, 2019. [arXiv:1901.04538v2](https://arxiv.org/abs/1901.04538v2) [math.GR]
- [15] A. Genevois and A. Martin, *Automorphisms of graph products of groups from a geometric perspective*, P. Lond. Math. Soc. **119** (2019), no. 6, 1745–1779.
- [16] E. R. Green, *Graph products of groups*. Ph.D. thesis. University of Leeds, Leeds, 1990.
- [17] D. Groves and M. Hull, Homomorphisms to acylindrically hyperbolic groups. I. Equationally Noetherian groups and families. *Trans. Amer. Math. Soc.* **372** (2019), no. 10, 7141–7190. [Zbl 07124920](https://zbmath.org/journal/10.1090/tran/7141) [MR 4024550](https://doi.org/10.1090/tran/7141)
- [18] M. F. Hagen, Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.* **7** (2014), no. 2, 385–418. [Zbl 1348.20044](https://zbmath.org/journal/10.1017/S1446788713000044) [MR 3217625](https://doi.org/10.1017/S1446788713000044)
- [19] F. Haglund and D. T. Wise, Special cube complexes. *Geom. Funct. Anal.* **17** (2008), no. 5, 1551–1620. [Zbl 1155.53025](https://zbmath.org/journal/10.1007/s00036-008-0155-1) [MR 2377497](https://doi.org/10.1007/s00036-008-0155-1)
- [20] S. Kim and T. Koberda, Embedability between right-angled Artin groups. *Geom. Topol.* **17** (2013), no. 1, 493–530. [Zbl 1278.20049](https://zbmath.org/journal/10.1017/S1446788712000499) [MR 3039768](https://doi.org/10.1017/S1446788712000499)
- [21] S. Kim and T. Koberda, The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.* **24** (2014), no. 2, 121–169. [Zbl 1342.20042](https://zbmath.org/journal/10.1112/S1074549913000042) [MR 3192368](https://doi.org/10.1112/S1074549913000042)
- [22] P. H. Kropholler and A. Martino, Graph-wreath products and finiteness conditions. *J. Pure Appl. Algebra* **220** (2016), no. 1, 422–434. [Zbl 1328.18018](https://zbmath.org/journal/10.1016/j.jpaa.2015.12.011) [MR 3393469](https://doi.org/10.1016/j.jpaa.2015.12.011)
- [23] J. F. Manning, Geometry of pseudocharacters. *Geom. Topol.* **9** (2005), 1147–1185. [Zbl 1083.20038](https://zbmath.org/journal/10.1017/S1446788705000388) [MR 2174263](https://doi.org/10.1017/S1446788705000388)
- [24] A. Minasyan and D. V. Osin, Acylindrical hyperbolicity of groups acting on trees. *Math. Ann.* **362** (2015), no. 3–4, 1055–1105. [Zbl 1360.20038](https://zbmath.org/journal/10.1007/s00208-015-1105-5) [MR 3368093](https://doi.org/10.1007/s00208-015-1105-5)
- [25] D. V. Osin, Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.* **368** (2016), no. 2, 851–888. [Zbl 1380.20048](https://zbmath.org/journal/10.1090/tran/368) [MR 3430352](https://doi.org/10.1090/tran/368)
- [26] M. Sageev, Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3) **71** (1995), no. 3, 585–617. [Zbl 0861.20041](https://zbmath.org/journal/10.1017/S0025571895000041) [MR 1347406](https://doi.org/10.1017/S0025571895000041)
- [27] Z. Sela, Diophantine geometry over groups X: The elementary theory of free products. Preprint, 2010. [arXiv:1012.0044v1](https://arxiv.org/abs/1012.0044v1) [math.GR]

Received April 11, 2019

Motiejus Valiunas, Instytut Matematyczny, Uniwersytet Wrocławski,
plac Grunwaldzki 2/4, 50-384 Wrocław, Poland

e-mail: valiunas@math.uni.wroc.pl