Stable subgroups of graph products

Sahana H Balasubramanya, Marissa Chesser, Alice Kerr,

Johanna Mangahas, Marie Trin

November 17, 2025

Abstract

We extend the characterization of stable subgroups of right-angled Artin groups from [KMT17] to the case of graph products of infinite groups. Specifically, we show that the stable subgroups of such graph products are exactly the subgroups that quasi-isometrically embed in the associated contact graph. Equivalently, they are the subgroups that satisfy a condition arising from the defining graph: a stable subgroup is an almost join-free subgroup. In particular, we recover the equivalence between stable and purely loxodromic subgroups of [KMT17] in the case where all torsion subgroups of the vertex groups are finite.

Contents

1 Introduction

Given a simplicial graph Γ and a family of finitely generated groups $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ indexed by the vertices of Γ, the associated graph product $\Gamma\mathcal{G}$ is generated by the vertex groups G_v, with the additional relation that elements of adjacent vertex groups commute (as detailed in Section 2.3). Graph products were introduced by Green in her PhD thesis [Gre90] and have since been widely studied. They provide an interplay between free products and direct products of groups, and generalize both right-angled Artin groups (RAAGs) and right-angled Coxeter groups (RACGs). In this paper, we focus on the case where the defining graph Γ has no isolated vertices, and where the groups G_v are all finitely generated and infinite; this situation generalizes the one of RAAGs.

The goal of this paper is to characterize the *stable* subgroups of such graph products. A finitely generated subgroup H of a finitely generated group G is stable if its inclusion in G is a quasi-isometric embedding and quasi-geodesics from G with common endpoints in H fellow-travel (see Section 2.2 for details). In particular, they are hyperbolic subgroups, where the hyperbolicity is inherited from the ambient geometry of G.

This notion of stability was introduced in [DT15] to generalize both the quasi-convex subgroups of hyperbolic groups, and the convex cocompact subgroups of mapping class groups [FM02]. They have since been well studied for a broad collection of groups, see for example [Sis16, CD17, AMST19, Kim19, AM24]. In particular, [KMT17] proved a geometric characterization of the stable subgroups of RAAGs.

Theorem 1.1. [KMT17, Theo 1.1] *Let Γ be a finite connected graph and H a finitely generated subgroup of the right angled Artin group $A(\Gamma)$. The following are equivalent:*

1. *H is stable in $A(\Gamma)$.*
2. *The orbit maps of H into the extension graph are quasi-isometric embeddings.*
3. *H is purely loxodromic for the action on the extension graph.*

In this theorem stability is characterized in terms of action of the group on the extension graph. This space was introduced by Kim and Koberda in [KK13], and is a hyperbolic graph on which $A(\Gamma)$ acts by isometries [KK14]. In the language of [BCK+25], the equivalence of (1) and (2) in Theorem 1.1 implies that the extension graph is a *universal recognizing space* for the stable subgroups of $A(\Gamma)$. A universal recognizing space for a group G is a hyperbolic space on which G acts by isometries, and in which all the stable subgroups embed quasi-isometrically through their orbit maps for this action. This has also been called a *stability recognizing space*, with the additional (reverse) assumption that every subgroup quasi-isometrically embedded through its orbit maps is stable [CDEK25].

A universal recognizing space for a group can be expected to play a role analogous to the one of the curve complex for mapping class groups. In fact, Theorem 1.1 is also true for mapping class groups when the extension graph is replaced by the curve complex and with the additional assumption in condition (3) that the group H is undistorted [KL08, Ham05, BBKL20]. Conditions (1) and (2) are also more generally equivalent for hierarchically hyperbolic groups acting on the hyperbolic space associated to the maximal domain of some specific structure on the group [ABD21], as well as for CAT(0) groups acting on the curtain model [PSZ24], Morse-dichotomous groups acting on the contraction space [Zbi24, PZ24], and certain relatively hyperbolic groups acting on the coned-off Cayley graph [ADT17]. For graph products the analog of the curve graph that we are interested in is the *contact graph of the prism-complex*; one can refer to [Oya24] or [Gen17] for other possibilities.

Theorem 1.2. *Let Γ be a finite simple graph with no isolated vertices. Let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of finitely generated infinite groups, and $P(\Gamma\mathcal{G})$ the prism complex of the graph product $\Gamma\mathcal{G}$. Let H be a finitely generated subgroup of the graph product $\Gamma\mathcal{G}$. Then the following are equivalent.*

1. *H is stable in $\Gamma\mathcal{G}$.*
2. *The orbit maps of H into the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ are quasi-isometric embeddings.*
3. *H is almost join-free.*

The contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ is known to be a quasi-tree, and hence a hyperbolic space [Val21, Corollary C]. It is therefore a universal recognizing space for $\Gamma\mathcal{G}$ under the conditions of Theorem 1.2. It was also shown in [Val21, Corollary D] that, under these conditions, the action of $\Gamma\mathcal{G}$ on the contact graph is the largest acylindrical action. On the other hand, when isolated vertices are allowed, the authors showed in [BCK+25] that the largest acylindrical action of such a graph product does not necessarily provide a universal recognizing space for stable subgroups.

We note that the final equivalent condition in Theorem 1.2 is different from the one in Theorem 1.1. A subgroup of a graph product is called a *join subgroup* if it decomposes as a direct product of subgroups generated by disjoint subgraphs of the defining graph. The stable subgroups are then exactly the *almost join-free subgroups*, in other words those admitting only finite intersections with conjugates of join subgroups. (Note that our definition of (almost) join-free subgroups differs slightly from the one given by Tran in [Tra17]). In the case of RAAGs, the purely loxodromic condition is known to be equivalent to asking H to be (almost) join-free [KK14], and this equivalence is a key part of the proof of Theorem 1.1.

The reason this characterization cannot be extended to general graph products is the potential existence of infinite torsion subgroups. Under the added assumption that such subgroups do not exist, we are able to recover the final equivalent statement of Theorem 1.1; also, in this specific case our notion of almost join-free subgroups corresponds with the one of Tran [Tra17]. In the case where infinite torsion subgroups are allowed, stable implies purely loxodromic, but not the other way around (see Section 5).

Theorem 1.3. *Let Γ be a finite simple graph with no isolated vertices. Let $\mathcal{G} = \{G_v \,|\, v \in V(\Gamma)\}$ be a collection of infinite finitely generated groups with no infinite torsion subgroups, and $P(\Gamma\mathcal{G})$ the prism complex of the graph product $\Gamma\mathcal{G}$. Let H be a finitely generated subgroup of the graph product $\Gamma\mathcal{G}$. Then the following are equivalent.*

1. *H is stable in $\Gamma\mathcal{G}$.*
2. *The orbit maps of H into the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ are quasi-isometric embeddings.*
3. *H is almost join-free.*
4. *H is purely loxodromic with respect to the action on $\mathcal{C}P(\Gamma\mathcal{G})$.*

We remark that the hypotheses on the vertex groups and on the graph in Theorem 1.2 and Theorem 1.3 are necessary. Firstly, if Γ contains an isolated vertex then the associated vertex group will always be almost join-free, while its orbit in the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ will always have finite diameter, and the vertex group will be stable in $\Gamma\mathcal{G}$ if and only if it is hyperbolic. Secondly, if finite vertex groups were allowed, we could consider the case where Γ is a path with two vertices, and the vertex groups are $\mathbb{Z}$ and $\mathbb{Z}_2$. This would mean that that $\Gamma\mathcal{G}$ is an infinite join group (so not almost join-free), while being stable within itself, and with $\mathcal{C}P(\Gamma\mathcal{G})$ having finite diameter.

Note that this hypothesis on vertex groups excludes RACGs from our theorem; in that case an analogous version of the curve complex is given by the the graph of maximal product regions introduced by Oh [Oh22], and under some hypotheses Cashen, Dani, Edletzberger, and Karrer prove the equivalence between (1) and (2) for this graph [CDEK25, Cor 6.10]. As RACGs are hierarchically hyperbolic groups, the work in [ABD21] also applies. It is therefore natural to ask if a more general result exists for graph products.

Question 1.4. *Does a characterization of stable subgroups similar to Theorem 1.2 hold when the vertex groups are allowed to be finite?*

In [KMT17], as a consequence of the characterizations of stable subgroups in RAAGs, it is also shown that stable subgroups of RAAGs are free [KMT17, Prop 8.1, 8.2]. In the case of graph products, we get the following generalization as a direct consequence of Theorem 1.2, the contact graph being a quasi-tree, and the fact that groups which quasi-isometrically embed into quasi-trees are virtually free (see for example [But23, Corollary 10.7]).

Corollary 1.5. *Let Γ be a finite simple graph with no isolated vertices and $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ a collection of finitely generated infinite groups. The stable subgroups of $\Gamma\mathcal{G}$ are virtually free.*

In the non-hyperbolic cases of both RAAGs of connected graphs and mapping class groups, there is an additional characterization of stable subgroups: they are exactly the infinite-index Morse subgroups [Kim19, Tra19]. This is not the case for RACGs [Tra19, Theorem 1.11 and Corollary 1.12], however it could plausibly hold for graph products of infinite groups.

Question 1.6. *Let Γ be a finite connected simple graph and $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ a collection of finitely generated infinite groups. Do the stable subgroups of $\Gamma\mathcal{G}$ coincide with the infinite-index Morse subgroups?*

Remark 1.7. In private correspondence, Anthony Genevois suggested an alternate proof of Theorem 1.5 using Stallings-Dunwoody decompositions that would apply more generally to non-torsion infinite-index accessible Morse subgroups, further motivating Theorem 1.6.

Organization of the paper. Section 2 is dedicated to the introduction of the concepts and notation we will use throughout the paper. Section 3 is dedicated to the study of disk diagrams, which are the main technical tool of this paper. Their study serves in the proof of (3)$\Longrightarrow$(2) of Theorem 1.2. Section 4 is dedicated to proving Theorem 1.2; the diagram below describes how this theorem is proved for the convenience of the reader. Finally, in Section 5, we study the link between stable and purely loxodromic subgroups of graph products.

$$\begin{tikzcd}
\text{1)}H \text{ is stable in } \Gamma\mathcal{G} \arrow[d, Rightarrow, "Theorem\ 4.5"] & & \begin{array}{l}\text{2) The orbit of } H \text{ into the} \\ \text{contact graph } \mathcal{CP}(\Gamma\mathcal{G}) \text{ is a} \\ \text{quasi-isometric embedding}\end{array} \arrow[ll, Rightarrow, "Theorem\ 4.6"'] \\
\text{3)}H \text{ is almost join-free} \arrow[rr, Rightarrow, "Theorem\ 4.1", "Theorem\ 3.25"'] & & \begin{array}{l}(H, d_H) \text{ is q.i.e into } (\Gamma\mathcal{G}, d_p) \\ H \text{ is } N \text{ join-busting}\end{array} \arrow[u, Rightarrow, "Theorem\ 4.4"']
\end{tikzcd}$$

Acknowledgments. The authors are grateful to the organizers of the 2023 Women in Groups, Geometry, and Dynamics program, where our collaboration started. We would like to thank Anthony Genevois for his feedback, both at the beginning of this project and on a later draft, and for his enlightening comments on [Gen19a]. Our thanks also to Chris Cashen and Annette Karrer for their explanation of the case of right-angled Coxeter groups.

This work was supported by funding from the EPSRC grant EP/V027360/1 "Coarse geometry of groups and spaces" via the third author, the Simons Foundation (965204, JM) via the fourth author, and MPI MiS via the third and fifth authors.

2 Preliminaries

2.1 Coarse geometry

Two metric spaces (X, d_X) and (Y, d_Y) are said to be *quasi-isometric* if there is a map $f : X \to Y$ and two real numbers $\lambda \geqslant 1$ and $c \geqslant 0$ such that

- $\forall x, x' \in X \quad \frac{1}{\lambda} d_X(x, x') - c \leqslant d_Y(f(x), f(x')) \leqslant \lambda d_X(x, x') + c$
- $\forall y \in Y, \exists x \in X \quad d_Y(y, f(x)) \leqslant c.$

The map f is called a (λ, c)-*quasi-isometry*; the first condition alone makes the map f a (λ, c)–*quasi-isometric embedding*. A *quasi-isometry* (resp. *quasi-isometric embedding*) is a map which is a (λ, c)–quasi-isometry (resp. quasi-isometric embedding) for some pair (λ, c). A quasi-isometric embedding from an interval of $\mathbb{R}$ or $\mathbb{Z}$ is called a *quasi-geodesic*, which is a *geodesic* when $\lambda = 1$ and $c = 0$.

2.2 Stable subgroups

Let G be a finitely generated group. A finitely generated subgroup $H \leqslant G$ is said to be *stable* if both of the following conditions hold:

- H is undistorted in G, i.e. the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding.
- For each choice of constants $\lambda \geqslant 1, c \geqslant 0$, there is a constant $M > 0$ (depending on λ, c) such that any two (λ, c)–quasi-geodesics in G with the same end points in H are each contained in the M–neighborhood of the other.

Note that stability is independent of the choice of finite generating sets for H and G [DT15]. In particular, it is preserved under group automorphisms such as conjugation. We note this explicitly because much of the work in this paper is dedicated to the study of condition (3) in Theorem 1.2 which deals with conjugacy. Namely, we have:

Lemma 2.1. *Let H be a stable subgroup of a finitely generated group G. For any $g \in G$ the subgroup gHg^{-1} is also a stable subgroup of G.*

2.3 Graph products

Let $\Gamma = (V(\Gamma), E(\Gamma))$ be a graph, where the set $V(\Gamma)$ is the set of vertices of the graph and the set $E(\Gamma) \subseteq V(\Gamma) \times V(\Gamma)$ is the set of (non-oriented) edges of the graph. Let $\mathcal{G} = \{G_v | v \in V(\Gamma)\}$ be a family of groups called *vertex groups*. In this paper we will always assume that Γ is a finite simple graph and that the vertex groups are all non-trivial. The *graph product* $\Gamma\mathcal{G}$ associated to the pair $(\Gamma, \mathcal{G})$ is defined by

$$\Gamma\mathcal{G} := \left(\mathop{\ast}_{v\in V(\Gamma)} G_v \right) \Big/ \left\langle\!\left\langle [g,h] \;\mid\; g \in G_v, h \in G_u \text{ with } \{u,v\} \in E(\Gamma) \right\rangle\!\right\rangle.$$

Notice that if all the vertex groups are isomorphic to $\mathbb{Z}$ then $\Gamma\mathcal{G}$ is a right-angled Artin group, and if the vertex groups are all isomorphic $\mathbb{Z}_2$, then $\Gamma\mathcal{G}$ is a right-angled Coxeter group.

For every (induced) subgraph Γ' of Γ, the subgroup $G_{\Gamma'} = \langle G_v \mid v \in V(\Gamma')\rangle$ is known as a *parabolic subgroup* of $\Gamma\mathcal{G}$. For ease of notation, we will typically not distinguish between the vertex set $V(\Gamma')$ and the graph Γ'.

2.3.1 Join, star, and link subgroups

Recall that a graph Γ is said to be a *join* of the nonempty subgraphs Γ_1 and Γ_2 if $V(\Gamma) = V(\Gamma_1) \sqcup V(\Gamma_2)$, and for every $u \in V(\Gamma_1)$, $v \in V(\Gamma_2)$, we have that $\{u,v\} \in E(\Gamma)$. For any graph product associated to Γ, this implies that $\Gamma\mathcal{G} = G_{\Gamma_1} \times G_{\Gamma_2}$. We will refer to any subgraph Γ' of Γ that splits as a join as a *join subgraph* of Γ, and the corresponding parabolic subgroup $G_{\Gamma'}$ as a *join subgroup* of $\Gamma\mathcal{G}$. Note that when talking about subgraphs we will always mean induced subgraphs.

For a graph Γ, recall that the *link* of the vertex $v \in V(\Gamma)$, denoted $\mathsf{Link}(v)$, is the set of vertices adjacent to v, and the *star* of v is $v \cup \mathsf{Link}(v)$, which we denote by $\mathsf{Star}(v)$. We will refer to subgraphs induced by such sets of vertices as *link subgraphs* and *star subgraphs* of Γ respectively, and the corresponding parabolic subgroups $G_{\mathsf{Link}(v)}, G_{\mathsf{Star}(v)}$ as *link subgroups* and *star subgroups* of $\Gamma\mathcal{G}$.

We extend the notions of Link and Star to any subgraph Γ' of Γ, so that $\mathsf{Link}(\Gamma') = \bigcap_{v\in V(\Gamma')} \mathsf{Link}(v)$ (that is, the vertices adjacent to every vertex in Γ') and $\mathsf{Star}(\Gamma') = \Gamma' \cup \mathsf{Link}(\Gamma')$. These definitions imply the following lemma.

Lemma 2.2. *If Γ is a graph, and $\Lambda \subseteq \Lambda' \subseteq \Gamma$ where Λ' is a join, then $\mathsf{Star}(\Lambda)$ is a join.*

Proof. Let Λ' be the join of Γ_1 and Γ_2. Suppose $\Lambda \subseteq \Gamma_1$, then $\varnothing \neq \Gamma_2 \subseteq \mathsf{Link}(\Lambda)$, and $\mathsf{Star}(\Lambda)$ is the join of Λ with $\mathsf{Link}(\Lambda)$. The same holds if $\Lambda \subseteq \Gamma_2$. Otherwise, Λ is the join of $\Gamma_1 \cap \Lambda$ and $\Gamma_2 \cap \Lambda$, so $\mathsf{Star}(\Lambda)$ is a join whether or not $\mathsf{Link}(\Lambda)$ is nonempty. $\square$

Definition 2.3. A subgroup of a graph product is called *join-free* (respectively *star-free*) if it has trivial intersection with any conjugate of a join subgroup (respectively a star subgroup). A subgroup of a graph product is called *almost join-free* (respectively *almost star-free*) if it has finite intersection with any conjugate of a join subgroup (respectively a star subgroup).

Remark 2.4. In the case that the graph has no isolated vertices, every star subgroup is a join subgroup. Consequently, we also have that (almost) join-free implies (almost) star-free.

2.4 Associated graphs and complexes

Throughout this paper we consider several different Cayley graphs associated to $\Gamma\mathcal{G}$ and their corresponding metrics. For more details on these the reader can refer to [Gen17] or [Val21]. We recall here the basic definitions and facts we will need.

When necessary, we will assume that the vertex groups of $\Gamma\mathcal{G}$ are finitely generated, and fix finite generating sets $\{S_v\}_{v\in V(\Gamma)}$ for those vertex groups. In this case the graph product $\Gamma\mathcal{G}$ is also finitely generated, and comes with a natural finite generating set

$$S := \bigcup_{v\in V(\Gamma)} S_v.$$

When we talk about the *standard Cayley graph* of (a finitely generated) $\Gamma\mathcal{G}$, we refer to the Cayley graph $\mathrm{Cay}_S(\Gamma\mathcal{G})$, which is the Cayley graph with respect to this finite generating set S. For $g \in \Gamma\mathcal{G}$ we denote distance from the trivial element to g in $\mathrm{Cay}_S(\Gamma\mathcal{G})$ by $|g|_S$. By a *word in S* we mean the trivial word, or a word written in the letters of S or their inverses.

Remark 2.5. Although we need to assume that $\Gamma\mathcal{G}$ is finitely generated for Theorem 1.2, the majority of intermediate results in this paper do not require this assumption. In particular, it is not needed in the proof that (3) $\implies$ (2), which is what the majority of this paper works towards.

2.4.1 Prism Cayley graph and complex

Most of the work in this document uses what we call the *prism Cayley graph* of $\Gamma\mathcal{G}$, following Genevois' study of graph products and other groups in [Gen17]. This is the Cayley graph using the generating set S_p (the *prism generating set* for $\Gamma\mathcal{G}$) defined as follows:

$$S_p := \bigcup_{v\in V(\Gamma)} G_v\backslash\{\mathsf{id}_{G_v}\}.$$

The prism Cayley graph is denoted $\mathrm{Cay}_p(\Gamma\mathcal{G})$, geodesic words for the prism generating set are referred to as *prism geodesics* or *S_p–geodesics*, and $|\cdot|_p$ is the prism length for a group element. Genevois proved that the prism Cayley graph of a graph product is always a quasi-median graph [Gen17, Proposition 8.2].

To this prism Cayley graph is associated a cell complex: we say that the *prism complex* $P(\Gamma\mathcal{G})$ of $\Gamma\mathcal{G}$ is the 2–complex whose 1–skeleton is the prism Cayley graph (the 1–skeleton is labeled by the elements of the prism generating set) and whose 2–cells correspond to the following relations:

$$[g, h] = \mathsf{id} \text{ when } g \in G_v, h \in G_u \text{ with } \{u, v\} \in E(\Gamma) \qquad (\square)$$
$$g \cdot g' = (gg') \text{ when } g, g' \in G_v \backslash \{\mathsf{id}_{G_v}\}, g' \neq g^{-1} \qquad (\triangle)$$

This complex is therefore made out of squares and triangles as illustrated in Figure 1 (which motivates its name). Edges of a triangle are labeled by elements in the same vertex groups, as are elements on opposite edges of a square. In particular, labels on opposite edges of a square are labeled by the same element, but it is possible for labels on edges of a triangle to all be distinct.

Remark 2.6. In this paper we refer to such a cell complex as the prism complex of $\Gamma\mathcal{G}$ for convenience, however they are in fact a subclass of what Genevois calls prism complexes of quasi-median graphs [Gen17]. In fact, what we call the prism complex is the 2–skeleton of Genevois' quasi-median complex, which is obtained by filling every clique with a simplex, and filling every 1–skeleton of an n–cube with an n–cube. As we do not need to consider any higher dimensional cells, we omit them here.

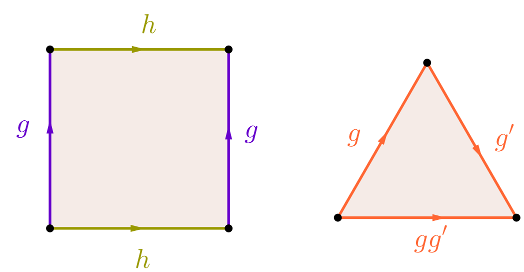

Figure 1: 2–cells of the prism complex

2.4.2 Hyperplanes and carriers

In the prism complex, a *hyperplane* represents an equivalence class of edges: two edges are equivalent if they belong to the same triangle or are opposite edges of the same square. All the edges in a given hyperplane are labeled by generators from the same vertex group; this defines the vertex (and vertex group) *associated to* that hyperplane. Each hyperplane corresponds to a dual *geometric hyperplane*: the union of the *midcubes* intersecting the edges of that hyperplane in $P(\Gamma\mathcal{G})$, where midcubes for the three kinds of cells are as illustrated in Figure 2.

A hyperplane may also be identified with its *carrier*: the union of closed cells that intersect the geometric hyperplane. From definitions one can verify that the 0–skeleton of a carrier is precisely some coset of $G_{\mathsf{Star}(v)}$, where v is the vertex associated to the hyperplane [Gen17, GM19]. We also make use of its 1–skeleton, the full subgraph

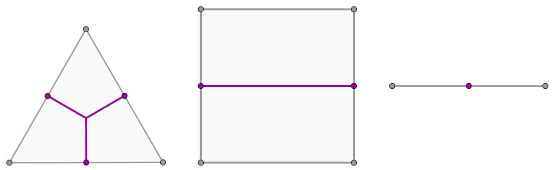

Figure 2: Midcubes for 1– and 2–cells of $P(\Gamma\mathcal{G})$

induced in $P(\Gamma\mathcal{G})$ by these vertices (following [Val21]). Note that the edges of the 1–skeleton that do not belong to the hyperplane itself are labeled by generators from $G_{\mathsf{Link}(v)}$.

Two distinct hyperplanes $\mathfrak{h}$ and $\mathfrak{h}'$ are said to be *transverse* if there exist a pair of equivalent edges in $\mathfrak{h}$ and a pair of equivalent edges in $\mathfrak{h}'$ which together form a square in the prism complex; equivalently, their corresponding geometric hyperplanes intersect. Note that this can only occur when the vertices corresponding to these two hyperplanes span an edge in Γ, so the vertex groups commute. Distinct hyperplanes are *tangent* if they are not transverse, but their carriers intersect.

2.4.3 Contact graph

The *contact graph* $\mathcal{C}P(\Gamma\mathcal{G})$ of the graph product $\Gamma\mathcal{G}$ is the graph whose vertices correspond to the hyperplanes of $P(\Gamma\mathcal{G})$, and whose edges connect pairs of vertices whose associated hyperplanes are either transverse or tangent; equivalently, there is an edge between two hyperplanes when their carriers intersect, an example is drawn in Figure 3. The natural action of the group $\Gamma\mathcal{G}$ on $P(\Gamma\mathcal{G})$ extends to an isometric action on $\mathcal{C}P(\Gamma\mathcal{G})$.

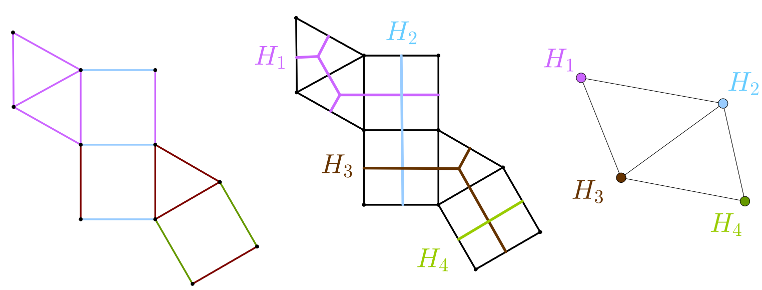

Figure 3: Hyperplanes, geometric hyperplanes, and contact graph. H_1 and H_2, H_2 and H_3, and H_3 and H_4 are transverse, while H_1 and H_3, and H_2 and H_4 are tangent.

The definition above is parallel to the construction of the contact graph of a CAT(0) cube complex, introduced by Hagen in [Hag14]. Several results from the cube complex case also hold for graph products, in particular the following is known.

Theorem 2.7. [Val21, Corollary C] *Let Γ be a bounded degree simplicial graph, and* *let $\Gamma\mathcal{G}$ be a graph product of non-trivial groups. Then $\mathcal{C}P(\Gamma\mathcal{G})$ is a quasi-tree (and therefore a hyperbolic space), and the action of $\Gamma\mathcal{G}$ on $\mathcal{C}P(\Gamma\mathcal{G})$ is acylindrical.*

2.4.4 Star Cayley graph

The last generating set we consider for $\Gamma\mathcal{G}$ comes from star subgroups:

$$S_\star := \bigcup_{v \in V(\Gamma)} G_{\mathsf{Star}(v)} \backslash \{\mathsf{id}_{G_{\mathsf{Star}(v)}}\}.$$

We call its associated Cayley graph $\mathrm{Cay}_\star(\Gamma\mathcal{G})$ the *star Cayley graph*, and apply analogous terminology and notations as for the previous Cayley graphs: geodesic words for this generating set are *star geodesics*, and $|\cdot|_\star$ denotes the star length of an element.

The following was observed by Valiunas in [Val21, Proof of Corollary D], and follows from [Gen17, Section 8.1] and [GM19, Theorem 2.10]. We include a short proof here for completeness.

Lemma 2.8. [Gen17, GM19] *For a graph product $\Gamma\mathcal{G}$, the contact graph of the prism complex is $\Gamma\mathcal{G}$-equivariantly quasi-isometric to the star Cayley graph.*

Proof. Let d be the metric on $\mathcal{C}P(\Gamma\mathcal{G})$. Identify hyperplanes in $P(\Gamma\mathcal{G})$ with the 0–skeletons of their carriers, which, as mentioned in Section 2.4.2, are of the form $gG_{\mathsf{Star}(v)}$ for some $g \in \Gamma\mathcal{G}$ and $v \in V(\Gamma)$. Recall that two hyperplanes are adjacent in $\mathcal{C}P(\Gamma\mathcal{G})$ if and only if their carriers intersect, and note that any such intersection implies the intersection of their 0–skeletons. Fix $v \in V(\Gamma)$, and define $f : \mathrm{Cay}_\star(\Gamma\mathcal{G}) \to \mathcal{C}P(\Gamma\mathcal{G})$ by $f(g) = gG_{\mathsf{Star}(v)}$. This map is clearly $\Gamma\mathcal{G}$-equivariant, so we just have to verify that it is a quasi-isometry.

Suppose $g \in \Gamma\mathcal{G}$ is such that $|g|_\star = 1$, so $g \in G_{\mathsf{Star}(w)}$ for some $w \in V(\Gamma)$. Then $G_{\mathsf{Star}(v)} \cap G_{\mathsf{Star}(w)}$ and $G_{\mathsf{Star}(w)} \cap gG_{\mathsf{Star}(v)}$ are both nonempty (containing id and g respectively), so $d(G_{\mathsf{Star}(v)}, gG_{\mathsf{Star}(v)}) \leqslant 2$. For general $g \in \Gamma\mathcal{G}$ the triangle inequality therefore tells us that $d(G_{\mathsf{Star}(v)}, gG_{\mathsf{Star}(v)}) \leqslant 2|g|_\star$.

For the opposite inequality, let $g \in G\backslash\{\mathsf{id}\}$, and let $k = d(G_{\mathsf{Star}(v)}, gG_{\mathsf{Star}(v)})$. This gives us a sequence of $k+1$ hyperplanes $G_{\mathsf{Star}(v)} = H_0, \ldots, H_k = gG_{\mathsf{Star}(v)}$ such that for every i the carriers of H_i and H_{i+1} intersect. By following paths along the 1–skeletons of these carriers, we can therefore write g as $h_0 \cdots h_k$, where each h_i belongs to the star subgroup corresponding to H_i. Therefore $|g|_\star \leqslant k + 1 \leqslant d(G_{\mathsf{Star}(v)}, gG_{\mathsf{Star}(v)}) + 1$.

Finally, let $g \in \Gamma\mathcal{G}$ and $w \in V(\Gamma)$. Then $gG_{\mathsf{Star}(w)} \cap gG_{\mathsf{Star}(v)} \neq \varnothing$, so $gG_{\mathsf{Star}(w)}$ is either adjacent to, or equal to, $f(g)$. □

Remark 2.9. This proof also shows that the contact graph is connected.

3 Disk diagrams

Many of the proofs in this paper are visualized using *disk diagrams*, a version of the *dual van Kampen diagrams* used for graph products by Valiunas [Val21, Section 6.2], and also by Genevois [Gen19b]. For finitely presented groups this notion is classical (cf. the *pictures* of [Fen83]), but our particular use of disk diagrams for graph products closely mirrors their more recent use for right-angled Artin groups ([CW04], [Kim08], [KK14]). In Section 3.1 we review the construction of such diagrams, and in Section 3.2 we note some basic lemmas that will be important to our applications. In Section 3.3, Section 3.4, and Section 3.5 we prove some more technical disk diagram results that will be key to the proof of the main theorem.

3.1 Constructing disk diagrams

By the nature of the relators for a presentation of $\Gamma\mathcal{G}$ using the prism generating set S_p, a van Kampen diagram $\mathcal{D}_{\mathsf{vK}}$ for an identity word $s = s_1 \ldots s_k$ will be a simply-connected complex made of squares, triangles, and edges embedded in $\mathbb{R}^2$, with edges labeled by the letters $s_i \in S_p$ around the boundary (see, for example, [LS77]). Our convention will be to read boundary words clockwise in van Kampen diagrams, their duals, and disk diagrams alike. Observe that (geometric) hyperplanes and their carriers can be defined for $\mathcal{D}_{\mathsf{vK}}$ in the same way as they are for the prism complex, using a cellular map from $\mathcal{D}_{\mathsf{vK}}$ to $P(\Gamma\mathcal{G})$ that respects edge labels. In particular, each geometric hyperplane and hyperplane carrier in a van Kampen diagram is a connected component of the preimage of a geometric hyperplane or hyperplane carrier in the prism complex, and two geometric hyperplanes can only intersect in a van Kampen diagram if their corresponding hyperplanes in the prism complex are transverse.

The dual of such a van Kampen diagram can be constructed as follows (and as illustrated in Figure 4): enlarge $\mathcal{D}_{\mathsf{vK}}$ to the embedded complex $\mathcal{D}'_{\mathsf{vK}}$ by adding one more vertex v_∞ to $\mathbb{R}^2 \backslash \mathcal{D}_{\mathsf{vK}}$, then adding edges between this vertex and each vertex encountered along the boundary of $\mathcal{D}_{\mathsf{vK}}$ (traversed cyclically), and finally adding 2–cells for each triangle bounded by a $\mathcal{D}_{\mathsf{vK}}$-boundary edge plus the two newly-added edges adjacent to it. The *dual van Kampen diagram* Δ is the dual of $\mathcal{D}'_{\mathsf{vK}}$ with the face corresponding to v_∞ deleted: it is a tessellated disk whose 1–skeleton consists of $\partial\Delta$ plus an interior subgraph, which may be disconnected.

Remark 3.1. Note that if one takes the van Kampen diagram together with its geometric hyperplanes, then Δ can be embedded in $\mathbb{R}^2$ such that it contains $\mathcal{D}_{\mathsf{vK}}$, as in Figure 4, and such that the intersection of the 1–skeleton of Δ together with $\mathcal{D}_{\mathsf{vK}}$ is exactly the union of the geometric hyperplanes of $\mathcal{D}_{\mathsf{vK}}$.

The boundary of a van Kampen diagram $\mathcal{D}_{\mathsf{vK}}$ is generally not an embedded circle, but $\partial\Delta$ is a topological circle (as it effectively comes from a neighborhood of $\mathcal{D}_{\mathsf{vK}}$), and furthermore we can identify Δ with the topological disk $\mathcal{D}_0$. We will define a disk diagram for w to be a certain labeling of the topological disk $\mathcal{D}_0$ along with an "almost cellular" map $\phi : \mathcal{D}_0 \to \mathcal{D}_{\mathsf{vK}}$ with certain properties (Proposition 3.2 below), with dual van Kampen diagrams serving as an intermediate step in the proof of their existence.

Proposition 3.2 (Disk diagram existence)**.** *Let s be an identity word in the prism generating set for a graph product $\Gamma\mathcal{G}$, and let $\mathcal{D}_{\mathsf{vK}}$ be a van Kampen diagram for s.*

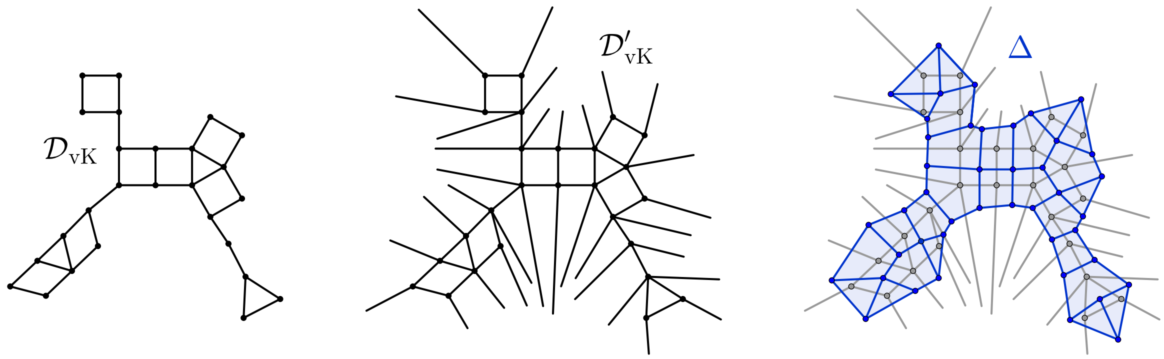

Figure 4: From van Kampen diagram to dual van Kampen diagram

There exists a map ϕ from the disk $\mathcal{D}_0$ to $\mathcal{D}_{\mathsf{vK}}$ such that:

(a) the image of $\partial\mathcal{D}_0$ is the circuit along the boundary of $\mathcal{D}_{\mathsf{vK}}$ following edges labeled in order by the letters of s;

(b) ϕ is a local homeomorphism on the pre-image of the interior of any 2–cell in $\mathcal{D}_{\mathsf{vK}}$;

(c) the pre-image of any geometric hyperplane in $\mathcal{D}_{\mathsf{vK}}$ is an embedded graph in $\mathcal{D}_0$.

Proof. Letting $\mathcal{D}_{\mathsf{vK}}$ be as in the statement, we start with the dual van Kampen diagram construction Δ described above, recalling that Δ can be seen as a neighborhood of $\mathcal{D}_{\mathsf{vK}}$, and in particular observing that Δ deformation retracts to $\mathcal{D}_{\mathsf{vK}}$. We also have that there is a natural homeomorphism between $\mathcal{D}_0$ and Δ.

As noted above (see Theorem 3.1), we can arrange the composition $\mathcal{D}_0 \to \Delta \to \mathcal{D}_{\mathsf{vK}}$ such that the geometric hyperplanes in $\mathcal{D}_{\mathsf{vK}}$ pull back to subgraphs of the inner 1–skeleton of Δ, and therefore pull back to embedded graphs in $\mathcal{D}_0$. We let $\phi : \mathcal{D}_0 \to \mathcal{D}_{\mathsf{vK}}$ come from the composition, which concludes the proof. $\square$

This enables the following definition.

Definition 3.3. Let $s = s_1 \cdots s_k$ be an identity word in S_p. A *disk diagram with boundary word s* is a disk $\mathcal{D}$ with a marking by a map ϕ as in Theorem 3.2 such that

- $\partial\mathcal{D}$ is segmented by the pre-images of the boundary edges of $\mathcal{D}_{\mathsf{vK}}$, and each segment is marked by the corresponding letter s_i of s,
- the interior of $\mathcal{D}$ is marked by the pre-images of the geometric hyperplanes.

The connected components of such pre-images of hyperplanes are called *dual graphs*, their dual edges in the boundary and their corresponding generator labels are called *roots* of the dual graph, and the dual graph is said to be *rooted* in its dual edges of s.

Remark 3.4. Note that properties of the geometric hyperplanes also apply to the dual graphs. In particular, each dual graph is associated to a given vertex of Γ, and all its roots belong to the corresponding vertex group. Moreover, distinct dual graphs may only intersect if their corresponding vertices have an edge between them in Γ.

Remark 3.5. The construction of dual van Kampen diagrams serves to prove the existence of disk diagrams, but the point the reader should keep in mind is the correspondence between van Kampen diagrams and disk diagrams. Most of the proofs in the rest of paper are based on disk diagrams. The pictures to keep in mind are illustrated in Figure 5. Here the colors can be thought of as corresponding to the vertex groups associated to the hyperplanes.

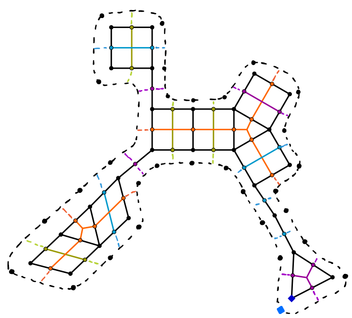

Figure 5: From (dual) van Kampen diagram to disk diagram

Convention. As already mentioned, we will work with disk diagrams oriented clockwise. Most of the time we will use disk diagrams to study identity words of the form $s = g\overline{w}$, where $w = w_1 \cdots w_m$ is a prism geodesic representative for g, and $\overline{w}$ is its inverse; then g will be read clockwise and w counter-clockwise as illustrated in Figure 6.

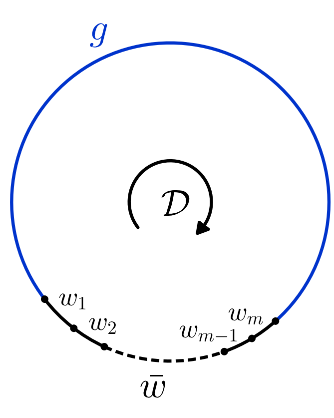

Figure 6: Convention for disk diagrams

We will be interested especially in paths in $\mathcal{D}$ that map to edge paths in $\mathcal{D}_{\mathsf{vK}}$. Already we may consider $\partial\mathcal{D}$ as having inherited a cellular decomposition into vertices and edges, corresponding to its image circuit in the van Kampen diagram. Between vertices on $\partial\mathcal{D}$ that are adjacent to edges that are roots of the same dual graph α, a *word along the carrier of the dual graph* means the product of a sequence of generators labeling an edge path in $\mathcal{D}_{\mathsf{vK}}$ connecting the corresponding vertices of $\partial\mathcal{D}_{\mathsf{vK}}$, which follows the 1–skeleton of the carrier of the hyperplane $\phi(\alpha)$. See examples in Figure 7. In particular, if α corresponds to $v \in V(\Gamma\mathcal{G})$, then words along its carrier spell elements of $G_{\mathsf{Star}(v)}$.

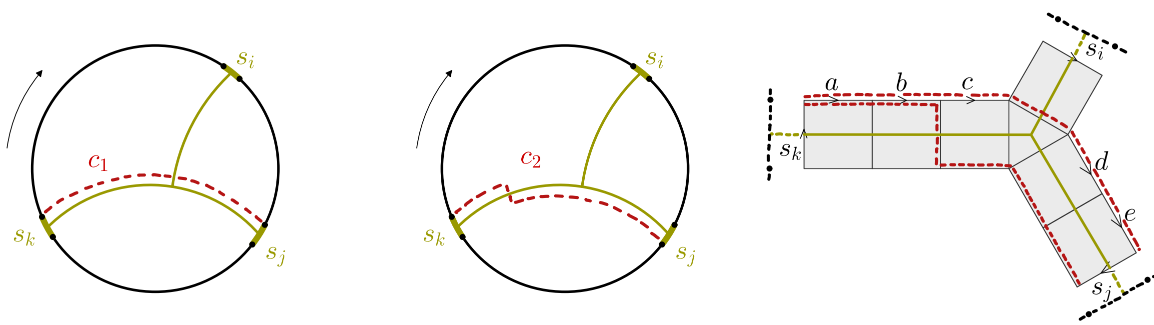

Figure 7: Words along a carrier: $c_1 = abcs_ide$, $c_2 = ab\bar{s}_kcde$

3.2 Disk diagram lemmas

Here we collect some fundamental properties of disk diagrams that will be used in this paper. First, we note that disk diagrams concatenate in the same way van Kampen diagrams do.

Lemma 3.6 (Concatenating disk diagrams)**.** *Let s_1 and s_2 be identity words in the prism generating set for $\Gamma\mathcal{G}$, with respective disk diagrams $\mathcal{D}_1$ and $\mathcal{D}_2$ given by Proposition 3.2. Suppose b is a subword of s_1 so that s_1 is the product a_1bc_1 of three subwords (with one of a_1 and c_1 possibly empty), and $\bar{b}$ a subword of s_2 with $s_2 = a_2\bar{b}c_2$. Then identifying $\mathcal{D}_1$ and $\mathcal{D}_2$ along the subsets of their boundaries labeled by the letters of b provides a disk diagram for $a_1c_2a_2c_1$.*

Proof. The disk diagram for Lemma 3.6 corresponds to the map ϕ defined as follows. The marking of $\mathcal{D}_1$ comes from some map $\phi_1 : \mathcal{D}_1 \to \mathcal{D}^{(1)}_{\mathsf{vK}}$, where $\mathcal{D}^{(1)}_{\mathsf{vK}}$ is a van Kampen diagram for s_1. Similarly, the marking of $\mathcal{D}_2$ comes from some map $\phi_2 : \mathcal{D}_2 \to \mathcal{D}^{(2)}_{\mathsf{vK}}$, where $\mathcal{D}^{(2)}_{\mathsf{vK}}$ is a van Kampen diagram for s_2. These van Kampen diagrams can be identified along the subsets of their boundaries labeled by b to obtain a van Kampen diagram $\mathcal{D}_{\mathsf{vK}}$ for $a_1c_2a_2c_1$. Then $\phi_1 \cup \phi_2 : \mathcal{D}_1 \cup \mathcal{D}_2 \to \mathcal{D}_{\mathsf{vK}}$ induces a well-defined map $\phi : \mathcal{D} \to \mathcal{D}_{\mathsf{vK}}$ where $\mathcal{D} = (\mathcal{D}_1 \cup \mathcal{D}_2)/\sim$ is the 2–to–1 identification along b of points in $\partial\mathcal{D}_1$ and $\partial\mathcal{D}_2$, which have the same image in $\mathcal{D}_{\mathsf{vK}}$ under ϕ_1 and ϕ_2 respectively. ∎

We highlight the next observation as we will use it repeatedly.

Fact 3.7. *If $\mathcal{D}$ is the concatenation of $\mathcal{D}_1$ and $\mathcal{D}_2$ then its dual graphs are concatenations of dual graphs from $\mathcal{D}_1$ and $\mathcal{D}_2$.*

Dual graphs in disk diagrams can be thought of as recording "cancellation patterns" for S_p–letters of identity words, as illustrated in the following lemma.

Lemma 3.8. *Let α be a dual graph in a disk diagram $\mathcal{D}$, and let $a_1, \dots, a_n$ be all the roots of α, in order around the boundary of $\mathcal{D}$. Then every a_i belongs to the same vertex group, and $a_1 \cdots a_n$ is an identity word. In particular, any dual graph has at least two roots.*

Proof. The first claim restates the fact that generators labeling edges in the same hyperplane must belong to the same vertex group, say G_v. To show that $a_1 \cdots a_n$ is the identity, consider the (potentially empty) boundary subwords c_i between a_i and a_{i+1}, along with c_n between a_n and a_1, so $a_1 c_1 a_2 \dots c_{n-1} a_n c_n = \mathsf{id}$. We can then consider a word c'_i along the carrier of α from the inner end of a_i to the inner end of a_{i+1}, so that $c'_i \overline{c_i}$ forms a closed circuit, and thus c_i and c'_i are group-equivalent. Therefore $a_1 c'_1 a_2 \dots c'_{n-1} a_n c'_n = \mathsf{id}$.

Note that as $\{a_1, \dots, a_n\}$ are all of the roots of α, the subword c_i does not contain any roots of α. We can therefore choose the path c'_i in the van Kampen diagram corresponding to $\mathcal{D}$ such that it does not cross the image of α, which means that every edge it traverses is labeled by an element of $G_{\mathsf{Link}(v)}$. As G_v and $G_{\mathsf{Link}(v)}$ form a direct product subgroup of $\Gamma\mathcal{G}$, we get that $a_1 \cdots a_n c'_1 \cdots c'_n = \mathsf{id}$, which implies that $a_1 \cdots a_n = \mathsf{id}$ due to the direct product structure. Moreover, because id is not a generator, $n \geqslant 2$. □

Our next lemma, about geodesic subwords of an identity word, recovers as a corollary that two vertices in $\mathrm{Cay}_p(\Gamma\mathcal{G})$ have only finitely many geodesics between them (whereas $\mathrm{Cay}_p(\Gamma\mathcal{G})$ is locally infinite). This fact is also a direct consequence of the normal form for graph products [Gre90].

Lemma 3.9 (Single-rooted graphs)**.** *Suppose a disk diagram has along its boundary an S_p–geodesic subword w. Then each dual graph roots in at most one edge of w.*

Proof. We prove the contrapositive. Let $\mathcal{D}$ be a disk diagram whose boundary word has a subword w. Suppose that α is a dual graph with roots in w labeled by generators l and r, appearing in that order in w, and let $v \in \Gamma$ be the vertex corresponding to α. Let c be the (possibly empty) subword of w between l and r on $\partial\mathcal{D}$. We can assume without loss of generality that c does not contain any roots of α. We can then consider a word c' along the carrier of α from the inner end of l to the inner end of r, so that $c'\overline{c}$ forms a closed circuit, and thus c and c' are group-equivalent. As in the proof of Theorem 3.8, we can choose the path c' in the van Kampen diagram so that it does not cross the image of α, which means that every edge it traverses is labeled by an element of $G_{\mathsf{Link}(v)}$. As $l, r \in G_v$, we see that c' commutes with both of these. Thus, as group elements, $lcr = lc'r = lrc' = kc' = kc$, where $k = lr$ is also in S_p. The word kc is strictly shorter than lcr, implying that w is not geodesic. $\square$

Corollary 3.10. [Gre90, Theorem 3.9] *The S_p–geodesic word for a given element in $\Gamma\mathcal{G}$ is unique up to ordering of its letters, so there are only finitely many prism geodesics between two vertices in $Cay_p(\Gamma\mathcal{G})$.*

Proof. Suppose u and v are prism geodesics between the same pair of vertices, so that $u\overline{v}$ labels the boundary word of a disk diagram. By Theorem 3.8 and Theorem 3.9, each dual graph of this diagram has exactly two roots, one in u and the other in $\overline{v}$. Theorem 3.8 ensures that the letters associated to those to roots are inverse to each other. Therefore u and v consist of the same letters and only differ in their ordering. $\square$

From these lemmas one may derive Genevois' characterization of prism-geodesic words; we give a proof for completeness of exposition.

Lemma 3.11. [Gen19b, Theorem 3.2] *Let $w = w_1 \cdots w_n$ be a non-trivial prism word of a graph product $\Gamma\mathcal{G}$. This word is geodesic if and only if for any $i < j$, if w_i and w_j belong to the same vertex group G_v, then there exists $i < k < j$ whose associated vertex in Γ does not belong to $\mathsf{Star}(v)$.*

Proof. Take a prism word $w = w_1 \cdots w_n$ admitting letters w_i and w_j which belong to the same vertex group G_v, and such that all the intermediate letters w_k with $k \in \{i+1, \ldots, j-1\}$ belong to vertex groups in $G_{\mathsf{Star}(v)}$. Hence, by construction w_i commutes with $w_{i+1} \cdots w_{j-1}$. Treating $w_i w_j$ as one generator, the prism word $w_1 \cdots w_{i-1} w_{i+1} \cdots w_{j-1}(w_i w_j) w_{j+1} \cdots w_n$ is a prism word of length at most $n-1$ representing the same group element as w, which is therefore not a prism-geodesic word.

Now suppose that w satisfies the latter condition of the lemma, and take u to be a prism-geodesic word representing the same group element. Fix $\mathcal{D}$ a disk diagram for $w\overline{u}$. By Theorem 3.9 we know that any graph dual to a letter in u has only one root in u. If the same is also true for w, then every dual graph in $\mathcal{D}$ has exactly one root in w and one in u, meaning that w and u have the same prism length, which would tell us that w is geodesic.

Assume for contradiction that there are two distinct letters a and b of w dual to the same dual graph α (in particular they belong to the same vertex group G_v); we can choose a and b such that the distance between a and b in w is the shortest distance relative to pairs of letters coming from dual graphs with at least two roots in w. Since a and b belong to the same vertex group, our assumption on w implies that there must be a letter c between a and b such that the vertex of Γ associated to c is not in $\mathsf{Star}(v)$. Theorem 3.8 ensures that the graph dual to c has a second root in the boundary of the diagram. By the minimality of a and b, this second root cannot be between a and b. This implies that the graph dual to c crosses α, meaning c belongs to a vertex group whose associated vertex is adjacent to v, contradicting that the associated vertex is not in $\mathsf{Star}(v)$. $\square$

3.3 Combing

Our goal in this section is to construct disk diagrams that have certain favorable properties, with respect to some geodesic subword w appearing backwards along their boundary. (Recall that we read w along the opposite orientation to fit our typical use of boundary words $g\overline{w}$ starting with some non-geodesic part g.) The ability to construct such disk diagrams will be critical to the proof of Theorem 3.25.

Let $g = g_1 \cdots g_n$ be a prism word and $w = w_1 \cdots w_m$ a prism-geodesic representative for g (so both g_j and w_i are in S_p). Given a disk diagram $\mathcal{D}$ for $g\overline{w}$, we define the *beginning function* (with respect to w) $b^{\mathcal{D}} : \{1, \ldots, m\} \to \{1, \ldots, n\}$ and the *ending function* (with respect to w) $e^{\mathcal{D}} : \{1, \ldots, m\} \to \{1, \ldots, n\}$ which associate to each index $i \in \{1, \ldots, m\}$ the index $j \in \{1, \ldots, n\}$ of the first, and respectively last, letters of g which are dual to the same dual graph as w_i in $\mathcal{D}$. In order to use beginning and ending functions flexibly, i.e. for any geodesic subword of a disk diagram boundary, we observe the following:

Fact 3.12. *A disk diagram $\mathcal{D}$ with boundary word s is also a disk diagram for the boundary word s' where s' is any cyclic permutation of s.*

Definition 3.13. Let $\overline{w}$ be a prism-geodesic subword of the boundary word for a disk diagram $\mathcal{D}$, so that after cyclic permutation this boundary word is $g\overline{w}$. Let $b^{\mathcal{D}}$ and $e^{\mathcal{D}}$ be the beginning and ending functions of $\mathcal{D}$ with respect to w. We say $\mathcal{D}$ is *w–left-combed* if, for any $1 \leqslant i < j \leqslant m$, $b^{\mathcal{D}}(i) < b^{\mathcal{D}}(j)$; likewise $\mathcal{D}$ is *w–right-combed* if for any $i < j$, $e^{\mathcal{D}}(i) < e^{\mathcal{D}}(j)$ as illustrated in Figure 8.

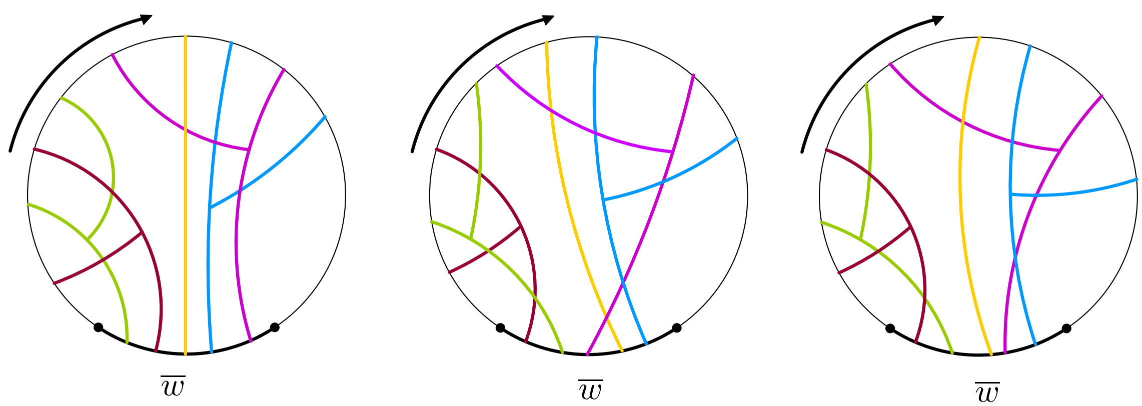

(a) Not combed (b) w-left-combed (c) w-right-combed

Figure 8: Example of combed disk diagrams.

For the proof of Theorem 3.25 we should notice that these definitions imply:

Fact 3.14. *If w' is a subword of a geodesic word w along the boundary of a w–left-combed (or w–right-combed) disk diagram $\mathcal{D}$, then $\mathcal{D}$ is also w'–left-combed (or w'–right-combed, respectively).*

Before establishing the existence of left/right-combed disk diagrams (Theorem 3.16 below), we record an elementary fact about disk diagrams and commutation. Given a prism word s that is equal to the identity, and a disk diagram $\mathcal{D}$ for s, if a and b are two successive letters in s which belong to two adjacent vertex groups, then we define *the commuting operation* on $\mathcal{D}$ as follow : If a and b belong to two adjacent vertex groups, then the square with sides labeled with $a, b, \bar{a}, \bar{b}$ is a Van Kampen diagram for $aba^{-1}b^{-1}$ that can directly be seen as a disk diagram. The commuting operation consists of gluing $\mathcal{D}$ to this square along ab, as in Theorem 3.6. This yields a disk diagram for the same word as $\mathcal{D}$ up to commutating a and b.

In a certain sense, the commuting operation does not change beginning and ending functions. To make this precise, recall that by Theorem 3.10, any two prism geodesics for g differ only by the order of their letters. Given such a prism geodesic $w = w_1 \cdots w_m$, and a permutation σ of $\{1, \ldots, m\}$ such that $w' = w_{\sigma(1)} \cdots w_{\sigma(m)}$ is also a prism-geodesic representative for g, we say that a disk-diagram $\mathcal{D}$ for $g\overline{w}$ and a disk-diagram $\mathcal{D}'$ for $g\overline{w}'$ have *the same* beginning function (resp. ending function) if for all i, $b^{\mathcal{D}'}(i) = b^{\mathcal{D}}(\sigma(i))$ (resp. $e^{\mathcal{D}'}(i) = e^{\mathcal{D}}(\sigma(i))$)

Lemma 3.15. *Let g be a prism word and w a prism-geodesic representative. Given a disk diagram $\mathcal{D}$ for $g\overline{w}$, any commuting operation applied to pairs of adjacent letters of w produces a disk-diagram with the same beginning and ending functions as $\mathcal{D}$.*

Proof. This is direct by construction. Indeed, as noted previously, when constructing a disk diagram by gluing, the new dual graphs are concatenations of dual graphs coming from the initial disk diagrams. □

Lemma 3.16 (Combing)**.** *Given a prism word g, a prism-geodesic representative w for g, and a disk-diagram $\mathcal{D}$ for $g\overline{w}$, there exists a reordering of the letters of w into a prism representative $w^{(l)}$ (resp. $w^{(r)}$) of g and a $w^{(l)}$-left-combed disk-diagram $\mathcal{D}_l$ for $g\overline{w}^{(l)}$ (resp. a $w^{(r)}$-right-combed disk-diagram $\mathcal{D}_r$ for $g\overline{w}^{(r)}$) with the same beginning and ending function as $\mathcal{D}$.*

Proof. Given g, we prove the lemma for the left-combed representative, as right-combing follows the same idea. Start with any prism geodesic $w = w_1 \cdots w_m$ for g, and fix a disk diagram $\mathcal{D}$ for the word $g\overline{w}$. For $1 \leqslant i \leqslant n-1$, if $b^{\mathcal{D}}(i) > b^{\mathcal{D}}(i+1)$ then the graphs dual to w_i and w_{i+1} cross in $\mathcal{D}$, meaning that w_i and w_{i+1} are in adjacent vertex groups, and therefore commute. Hence $w_1 \cdots w_{i-1} w_{i+1} w_i w_{i+2} \cdots w_m$ is still a geodesic word for g. Theorem 3.15 ensures that the commuting operation produces a disk-diagram for $g\overline{w_1 \cdots w_{i-1} w_{i+1} w_i w_{i+2} \cdots w_m}$ with the same ending and beginning functions as $\mathcal{D}$. The bubble sort algorithm can then be applied to conclude the proof. □

Remark 3.17. Although we will not need it, it is worth noticing that the combing process does more than leaving the beginning and ending functions unchanged. If $w^{(l)}$ differs from w by a permutation σ of the letters of w, then for any i, $w_i^{(l)}$ is dual in $\mathcal{D}^{(l)}$ to exactly the same letters of g as $w_{\sigma(i)}$ in $\mathcal{D}$.

We record here an application of combing to geodesic representatives of words in a subgroup H, which provides quantitative control which will be important for the proof of Theorem 3.25.

Lemma 3.18 (Reduction control)**.** *Let H be a finitely generated subgroup of $\Gamma\mathcal{G}$ with finite generating set S_H, and let $h = h_1 \cdots h_n$ be an S_H–geodesic. Write each h_i as an S_p–geodesic, so that h also appears as a prism word. Let s be an S_p–geodesic representative for h, and let $w = w_1 \cdots w_m$ be a subword of s. Consider a w-left-combed disk diagram $\mathcal{D}$ for the prism word $h\overline{s}$, and let $b^{\mathcal{D}}$ be the beginning function of $\mathcal{D}$ with respect to w. Given $1 \leqslant \ell < r \leqslant m$, let L and R be the indices of the subwords h_i of h containing the S_p–letters indexed by $b^{\mathcal{D}}(\ell)$ and $b^{\mathcal{D}}(r)$ respectively. Then*

$$R - L \geqslant \frac{r - l}{\mathsf{m}_H},$$

where m_H is the maximal S_p–length for elements in S_H.

Proof. First note that, by left-combing, for any $\ell < k < r$ we have that $b^{\mathcal{D}}(l) < b^{\mathcal{D}}(k) < b^{\mathcal{D}}(r)$. Each graph dual to w_k has its first root on some edge of a subword h_i for $L \leqslant i \leqslant R$, and each such subword has at most m_H edges; two of these $(R-L+1)\cdot\mathsf{m}_H$ total possibilities are already dual to w_ℓ and w_r. Therefore

$$r - \ell - 1 \leqslant (R - L + 1)\mathsf{m}_H - 2 < (R - L + 1)\mathsf{m}_H - 1.$$

□

3.4 Disk diagrams for almost star-free subgroups

We now apply disk diagrams to study finitely generated almost star-free subgroups of $\Gamma\mathcal{G}$. Given such a subgroup H together with a finite generating set S_H, we always consider elements of S_H as represented by S_p–geodesics. Hence if h is an element of H written as an S_H–geodesic $h_1 \cdots h_n$, then it is also represented by a prism word, and if w is a prism-geodesic representative for h, then a disk diagram for $h\overline{w}$ means a disk diagram for the associated prism word obtained by replacing the letters h_i of h by their S_p–geodesic representatives. Recall from Theorem 3.18 that we denote by m_H the maximum prism-length of an element in S_H.

The first of this series of lemmas gives a quantitative control on how S_H–geodesic words reduce to S_p–geodesic words. It is adapted from Lemma 4.1 in [KMT17], a right-angled Artin group counterpart bounding "cancellation diameter"; in RAAG disk diagrams, the analogue for dual graphs are arcs that represent the cancellation of a generator with its inverse, where these are letters from the (finite) standard generating set for the RAAG. In our setting of graph products, dual graphs represent the combining of prism generators from the same vertex group to yield the identity (as observed in Theorem 3.8).

Lemma 3.19 (Bounded combination diameter)**.** *Suppose H is a finitely generated almost star-free subgroup of a graph product $\Gamma\mathcal{G}$, and let S_H be a finite generating set of H. Then there exists a constant $D = D(S_H)$ with the following property: Suppose $h \in H$ is an S_H–geodesic word $h = h_1 \cdots h_n$ equivalent to the S_p–geodesic word w. Then in any disk diagram for the identity word $h\overline{w}$, if a dual graph has roots in h_i and h_j, then $|i - j| < D$.*

Proof. Let us suppose, by contradiction, that no such D exists. This means that there exist sequences of S_H–geodesic and S_p–geodesic words

$$h^{(k)} = h_1^{(k)} \cdots h_{n(k)}^{(k)} \qquad \text{where } h_i^{(k)} \in S_H \text{ and } n(k) \text{ is minimal, and}$$

$$w^{(k)} = h^{(k)} \qquad \text{where } w^{(k)} \text{ is an } S_p\text{–geodesic;}$$

along with disk diagrams $\mathcal{D}^{(k)}$ for the identity words $h^{(k)}\overline{w}^{(k)}$ (recall that each element of S_H comes with a prism geodesic representative, allowing us to see $h^{(k)}$ as a prism word); and indices $i(k) < j(k)$ such that some dual graph $\alpha^{(k)}$ of $\mathcal{D}^{(k)}$ roots in letters of $h_{i(k)}^{(k)}$ and $h_{j(k)}^{(k)}$, where $j(k) - i(k)$ grows arbitrarily large.

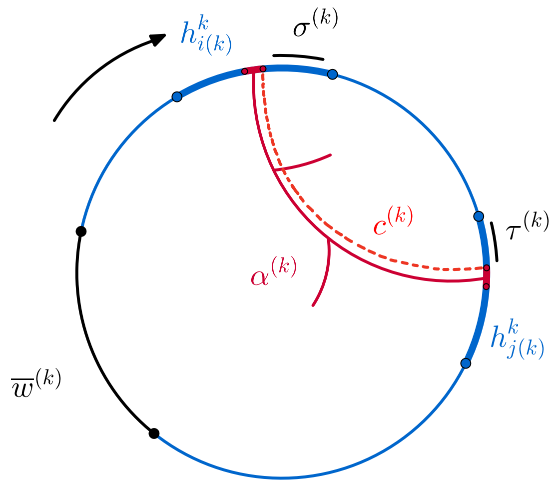

Figure 9: Decomposition of $h^{(k)}\overline{w}^{(k)}$

Let $v^{(k)}$ be the vertex corresponding to the dual graph $\alpha^{(k)}$. Taking words along the carriers of each $\alpha^{(k)}$, connecting letters in $h_{i(k)}^{(k)}$ and $h_{j(k)}^{(k)}$ respectively, this sequence of disk diagrams implies the existence of a sequence of $c^{(k)} \in G_{\mathsf{Star}(v^{(k)})}$ that are equal to subwords of $h^{(k)}$ of the form

$$c^{(k)} = \sigma^{(k)} h_{i(k)+1}^{(k)} \cdots h_{j(k)-1}^{(k)} \tau^{(k)}, \text{ where}$$

$$\sigma^{(k)} \text{ is a suffix of } h_{i(k)}^{(k)} \text{ and}$$

$$\tau^{(k)} \text{ is a prefix of } h_{j(k)}^{(k)}.$$

This construction is illustrated in Figure 9.

Because Γ is finite, we may pass to a subsequence so that $c^{(k)}$ are all in the same $G_{\mathsf{Star}(v)}$. Because Lemma 3.10 bounds the number of prism-geodesic ways to write the finitely many elements of S_H, we can pass to a further subsequence so that $\sigma^{(k)}$ and $\tau^{(k)}$ are also constant, equal to σ and τ respectively.

In this subsequence, $c^{(k)} = \sigma h'^{(k)}\tau$ where $h'^{(k)} \in H$, and $|h'^{(k)}|_{S_H} = j(k) - i(k) - 1$ grows arbitrarily large. We can therefore pass to a further subsequence such that $|h'^{(k)}|_{S_H} \neq |h'^{(l)}|_{S_H}$ for all $k \neq l$, and so $h'^{(k)} \neq h'^{(l)}$ for all $k \neq l$. As a result, we have that

$$h'^{(1)}(h'^{(k)})^{-1} = \sigma^{-1}c^{(1)}(c^{(k)})^{-1}\sigma \in \sigma^{-1}G_{\mathsf{Star}(v)}\sigma$$

is an infinite sequence of distinct elements of H, so H is not almost star-free. □

Like the previous lemma, the following lemma mirrors a right-angled Artin group version in [KMT17], this time of the same name. This concerns *vanishing subwords* for S_H–geodesic words h, defined as follows. Suppose $h = h_1 \cdots h_n$ where $h_i \in S_H$ are S_p–geodesic words and n is minimal. We say the subword $h_i \cdots h_j$ is *vanishing* if there exists an S_p–geodesic word w equal to h with a disk diagram for $h\overline{w}$ where no dual graph rooted in $\overline{w}$ intersects $h_i \cdots h_j$. In this sense, the subword $h_i \cdots h_j$ does not contribute to the letters in the prism-geodesic word w.

Lemma 3.20 (Bounded non-contribution)**.** *Suppose H is a finitely generated almost star-free subgroup of a graph product $\Gamma\mathcal{G}$, and let S_H be a finite generating set of H. Then there exists a constant $K = K(S_H)$ such that if $h_i \cdots h_j$ is a vanishing subword of $h \in H$ as in the definition above, then $j - i < K$.*

Proof. Let h be a S_H–geodesic with length n in S_H. Consider a vanishing subword $h' = h_i \cdots h_j$ of h in a disk diagram $\mathcal{D}$ for $h\overline{w}$ where w is an S_p–geodesic. Let $D = D(S_H)$ be as in Theorem 3.19. Let τ be the empty word if $i = 1$, and otherwise the subword $h_{i'} \cdots h_{i-1}$ of h, where $i' = \max\{1, i - D\}$. Likewise, let σ be the the empty word if $j = n$, and otherwise $h_{j+1} \cdots h_{j'}$ where $j' = \min\{n, j + D\}$. Note that, by Theorem 3.19, and since h' is vanishing in $\mathcal{D}$, any dual graph in $\mathcal{D}$ rooted in h' can only have endpoints in $\tau h'\sigma$, which is an S_H–geodesic subword of h.

We first want to show that h' is also a vanishing subword for $\tau h'\sigma$. Let w' be a prism-geodesic representative for $\tau h'\sigma$; we will find a disk diagram for $\tau h'\sigma\overline{w'}$ that evidences this vanishing of h'. Let $w'' = \overline{h_{i'-1} \cdots h_1} w \overline{h_n \cdots h_{j'+1}}$, without the first portion if either $i = 1$ or $i' = 1$, and without the last portion if either $j = n$ or $j' = n$. This means that w' is also a prism-geodesic representative for w''. Let $\mathcal{D}'$ be a disk diagram for $w''\overline{w}'$, and let $\mathcal{D}''$ be the concatenation of $\mathcal{D}$ and $\mathcal{D}'$ along w'', see Figure 10. By Theorem 3.6 this is a disk diagram for $\tau h'\sigma\overline{w}'$. Because the dual graphs in $\mathcal{D}''$ are concatenations of dual graphs in $\mathcal{D}$ and $\mathcal{D}'$, no graph dual to w' in $\mathcal{D}''$ roots in h'; thus h' is vanishing for $\tau h'\sigma$ as claimed.

To complete the proof, it suffices to show that there are only finitely many possibilities for $h' = \overline{\tau}w'\overline{\sigma}$, since this automatically bounds $|h'|_H = j - i + 1$. First note that

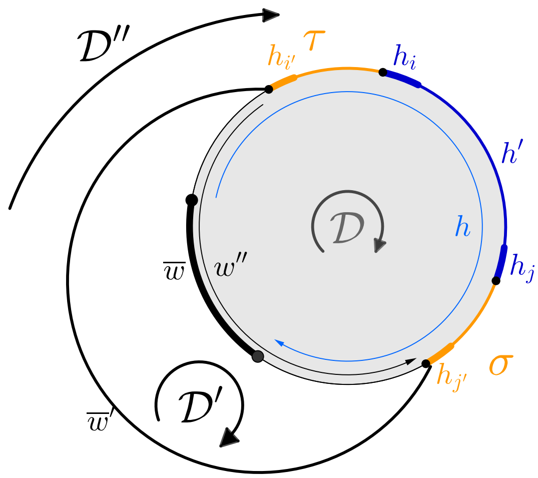

Figure 10: Creation of a disk diagram $\mathcal{D}''$ where h' vanishes

the S_H–lengths of τ and σ are bounded by D, so as S_H is finite, there are only finitely many possibilities for $\overline{\tau}$ and $\overline{\sigma}$. By Theorem 3.8, we have that each edge labeled a in w' satisfies $a = a_1 \dots a_m$, where the a_i are the roots in either τ or σ of the graph dual to a, in order along the boundary. As the S_H–lengths of τ and σ are bounded by D, and Theorem 3.10 bounds the number of prism-geodesic ways to write the finitely many elements of S_H, this gives only finitely many possibilities for each prism letter a of w'. As the S_p–lengths of τ and σ are bounded by $D\mathsf{m}_H$, the S_p–length of w' is bounded by $2D\mathsf{m}_H$. Therefore there are only finitely many possibilities for w', and likewise only finitely many possibilities for h'. ∎

In contrast to the previous two results, Lemma 3.21 addresses new behavior arising in disk diagrams for graph products. Because disk diagram concatenation also concatenates dual graphs, we must verify we do not lose the quantitative control from Theorem 3.19.

Lemma 3.21 (Bounded post-concatenation combination)**.** *Suppose H is a finitely generated almost star-free subgroup of a graph product $\Gamma\mathcal{G}$, and let S_H be a finite generating set of H. Then there exists a constant $C = C(S_H)$ with the following property. Suppose $h \in H$ is an S_H–geodesic word $h = h_1 \cdots h_n$ equivalent to the S_p–geodesic word w, and that the sub-word $h' = h_i \cdots h_j$ of h is equivalent to the S_p–geodesic word u. Let $\mathcal{D}$ be a disk diagram for the identity word $h\overline{w}$, let $\mathcal{D}'$ be a disk diagram for the identity word $\bar{h}'u$, and let $\mathcal{D}''$ be the concatenation of $\mathcal{D}$ and $\mathcal{D}'$ along h'. If a dual graph in $\mathcal{D}''$ has roots in u, $h_1 \cdots h_{i-1}$, and $h_{j+1} \cdots h_n$, then $|i - j| < C$.*

Proof. Consider the restriction of the relevant dual graph in $\mathcal{D}''$ to $\mathcal{D}$. This gives us a dual graph in $\mathcal{D}$ that is rooted in both $h_1 \cdots h_{i-1}$ and some $h_{i'}$ in h', and a dual graph in $\mathcal{D}$ that is rooted in both $h_{j+1} \cdots h_n$ and some $h_{j'}$ in h' (possibly the same graph).

We can assume without loss of generality that $i' \leqslant j'$. By Theorem 3.19 applied to $\mathcal{D}$, we therefore have that $|i - i'| < D$ and $|j - j'| < D$, where D is the constant from Theorem 3.19.

Let v be the vertex associated to the relevant dual graph in $\mathcal{D}''$. By taking a word W along the carrier of this dual graph, we can see that h' can be written as $\sigma W \tau$, where σ is an S_p–subword of $h_i \cdots h_{i'}$, τ is an S_p–subword of $h_{j'} \cdots h_j$, and $W \in G_{\mathsf{Star}(v)}$, as illustrated in Figure 11.

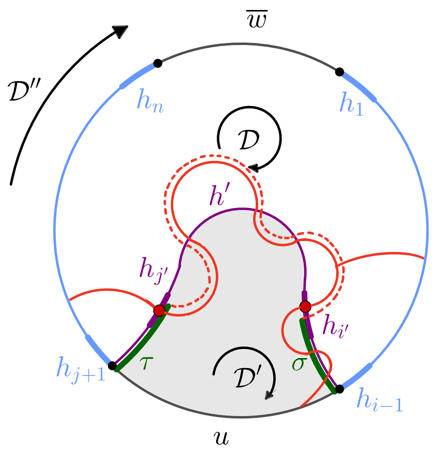

Figure 11: The relevant dual graph in $\mathcal{D} \cup \mathcal{D}'$.

Now suppose, by contradiction, that no bound on $|i - j|$ exists. We can therefore pick a sequence of $h^{(k)}$ and $h'^{(k)}$ in H where the $h'^{(k)}$ have strictly increasing S_H–length, and such that disk diagrams and dual graphs exist for each which satisfy the hypotheses of this lemma. In particular, we can write $h'^{(k)}$ as $\sigma^{(k)} W^{(k)} \tau^{(k)}$, where $\sigma^{(k)}$ is an S_p-subword of $h^{(k)}_{i(k)} \cdots h^{(k)}_{i'(k)}$, $\tau^{(k)}$ is an S_p–subword of $h^{(k)}_{j'(k)} \cdots h^{(k)}_{j(k)}$, and $W^{(k)} \in G_{\mathsf{Star}(v^{(k)})}$. As there are only finitely many possibilities for $v^{(k)}$, and both $|i(k) - i'(k)| < D$ and $|j(k) - j'(k)| < D$, in the same way as in previous lemmas in this section, we can pass to a subsequence such that $h'^{(k)} = \sigma W^{(k)} \tau$ for some constant words σ and τ, and such that each $W^{(k)} \in G_{\mathsf{Star}(v)}$ for the same vertex v.

Recall that the S_H–length of the $h'^{(k)}$ strictly increase as k increases, so in particular no two are equal as elements of H. We therefore have that

$$h'^{(1)}(h'^{(k)})^{-1} = \sigma W^{(1)}(W^{(k)})^{-1}\sigma^{-1} \in \sigma G_{\mathsf{Star}(v)}\sigma^{-1}$$

is an infinite sequence of distinct elements of H. This contradicts H being almost star-free. $\square$

3.5 Join-busting

In this section we prove a key property of almost join-free subgroups of $\Gamma\mathcal{G}$; this will make use of the disk diagram tools established thus far.

Definition 3.22. A *join (sub)word* is a (sub)word consisting of letters from S_p such that all these letters belong to the same join subgroup of $\Gamma\mathcal{G}$. A subgroup H of a graph product $\Gamma\mathcal{G}$ is said to be *N–join-busting* for some $N > 0$ if for any prism geodesic w representing an element $h \in H$, the prism length of any join subword of w is bounded by N.

We have one more technical result, Theorem 3.24, to record before we proceed to the main theorem of this section.

Definition 3.23. We call an S_p–word *M–homogeneous with respect to* Λ for some subgraph Λ of Γ, if all of its letters are contained in G_Λ, and all of its subwords of prism length at least M contain at least one letter from each vertex group of Λ.

Lemma 3.24. *If $(u_k)_{k\in\mathbb{N}}$ is a sequence of S_p–geodesic join words, such that $|u_k|_p$ tends to infinity with k, then there exists a sequence of S_p–geodesic words $(w_n)_{n\in\mathbb{N}}$ such that:*

- *the w_n are subwords of a subsequence of u_k,*
- *the lengths $|w_n|_p$ tend to infinity with n,*
- *there exists $M > 0$ and a join Λ such that the w_n are all M–homogeneous with respect to the same subgraph of Λ.*

Proof. This proof works in the same way as [KMT17, Lemma 5.1]. □

Theorem 3.25 (Join-busting)**.** *Let Γ be a graph with no isolated vertices. Suppose H is a finitely generated almost join-free subgroup of a graph product $\Gamma\mathcal{G}$, and let S_H be a finite generating set of H. Then there exists a constant $N = N(S_H)$ such that H is N–join-busting.*

Proof. We begin by noticing that as Γ has no isolated vertices, the assumption that H is almost join-free implies that H is also almost star-free. This allows us to apply the results from Section 3.4.

Towards contradiction, assume that H is not N–join-busting for any N. Then we can find a sequence $(h^{(k)})_{k\in\mathbb{N}}$ of elements in H such that some prism-geodesic representatives can be written $a^{(k)}w^{(k)}b^{(k)}$, where the $w^{(k)}$ are join words with $\lim_{k\to\infty} |w^{(k)}|_p = \infty$. For each k, we fix an S_H–geodesic representative

$$h^{(k)} = h_1^{(k)} \cdots h_{n(k)}^{(k)},$$

for $h^{(k)}$, meaning $n(k)$ is the length of $h^{(k)}$ in S_H and each $h_i^{(k)}$ is an element of S_H. We consider $h^{(k)}$ as a word in S_p by representing each $h_i^{(k)}$ by a prism-geodesic word in S_p.

By Theorem 3.16, we can assume that we have a $w^{(k)}$–left-combed disk diagram $\mathcal{D}^{(k)}$ for each identity word $h^{(k)}\overline{a^{(k)}w^{(k)}b^{(k)}}$, after re-ordering the letters of $w^{(k)}$. Applying Fact 3.14 and Lemma 3.24, we may also assume that there exists $M > 0$ such that the $w^{(k)}$ are all M–homogeneous with respect to the same subgraph Λ of Γ, where Λ is contained in a join. That is, we pass to a subsequence and replace $w^{(k)}$ by subwords, effectively enlarging the parts of the prism-geodesic representative for $h^{(k)}$ labeled by $a^{(k)}$ and $b^{(k)}$.

By Lemma 3.9, each dual graph in $\mathcal{D}^{(k)}$ rooted in $w^{(k)}$ has its remaining roots only in $h^{(k)}$. Our goal now is to find an S_H–subword of $h^{(k)}$ whose prism geodesic representative is of the form $\sigma^{(k)}W^{(k)}\tau^{(k)}$, such that $W^{(k)}$ is a join word whose S_p–length tends to infinity with k, whereas $\sigma^{(k)}$ and $\tau^{(k)}$ each range over finitely many possibilities.

Towards this, first let us reduce notation by fixing k. We set $h = h^{(k)} = h_1 \cdots h_n$, $a = a^{(k)}$, $b = b^{(k)}$, $\mathcal{D} = \mathcal{D}^{(k)}$, and $w = w^{(k)} = w_1 \cdots w_m$. Assume without loss of generality that $m > 2M$. Since $\mathcal{D}$ is w–left-combed, we may apply Theorem 3.18 for $\ell = M$ and $r = m - M$. Thus, where h_L and h_R contain the first edges dual to the same graphs as w_ℓ and w_r respectively, we have by Theorem 3.18

$$R - L \geqslant \frac{m - 2M}{\mathsf{m}_H}.$$

In particular, longer S_p–length of w implies longer S_H–length of $h_{L+1} \cdots h_{R-1}$; denote the latter word by h'. Henceforth we assume w is long enough so that $R - L > C + 2D$, where D and C are the constants taken from Theorem 3.19 and Theorem 3.21. We assume without loss of generality that D is an integer, and that $D \geqslant 2$.

Let x be a prism letter of h' whose dual graph in $\mathcal{D}$ has a root in awb. Either it is rooted in w, in which case x belongs to G_Λ, or it is rooted in one of a or b. Since $\mathcal{D}$ is w–left-combed, all the dual graphs rooted in w_i for $1 \leqslant i \leqslant M$ have roots in $h_1 \cdots h_R$, so a graph dual to x rooted in a will cross all of them. Similarly, a graph dual to x rooted in b will cross all the dual graphs rooted in w_i for $m - M \leqslant i \leqslant m$. In either case, because w is M–homogeneous for Λ, the graph dual to x crosses a graph rooted in an edge corresponding to each vertex group of Λ; therefore x belongs to $G_{\mathsf{Star}(\Lambda)}$. Thus we have:

Fact 3.26. *If x is a letter of h' with dual graph rooted in awb, then $x \in G_{\mathsf{Star}(\Lambda)}$.*

Now we consider the subword $h'' = h_{L+D} \cdots h_{R-D}$, and an S_p–geodesic representative $u = u_1 \cdots u_t$, where $u_i \in S_p$ and $t = |h''|_p$. By assumption on $R - L$, the S_H–length of h'' is positive. Note also that, since S_H is finite, we have finitely many possibilities for the subwords $h_{L+1} \cdots h_{L+D-1}$ and $h_{R-D+1} \cdots h_{R-1}$. We also note that, by Theorem 3.19, any graph dual to a letter of h'' can only have roots in h' or awb.

Let $\mathcal{D}''$ be a disk-diagram for $\overline{h''}u$, and concatenate it with $\mathcal{D}$ to obtain a disk diagram $\mathcal{P}$ for the word $uh_{R-D+1} \cdots h_n\overline{awb}h_1 \cdots h_{L+D-1}$. This is the situation illustrated in Figure 12a. By construction, and by Theorem 3.7 and Theorem 3.9, a graph dual to a letter of u in $\mathcal{P}$ has its remaining roots in some combination of $h_{L+1} \cdots h_{L+D-1}$, $h_{R-D+1} \cdots h_{R-1}$, and awb. In particular, by Theorem 3.21, and our assumption that $R - L > C + 2D$, any graph dual to a letter in a u can only have roots in exactly one of $h_{L+1} \cdots h_{L+D-1}$ or $h_{R-D+1} \cdots h_{R-1}$. As a consequence, if a graph dual in $\mathcal{P}$ to a letter of u has a root in awb (it cannot have more than one by Theorem 3.9), then it should be its beginning or ending one (in the sense of beginning and ending functions with respect to u, as defined in Section 3.3). Moreover, by Theorem 3.7 and Theorem 3.26, any letter of u whose dual graph has a root in awb must belong to $x \in G_{\mathsf{Star}(\Lambda)}$.

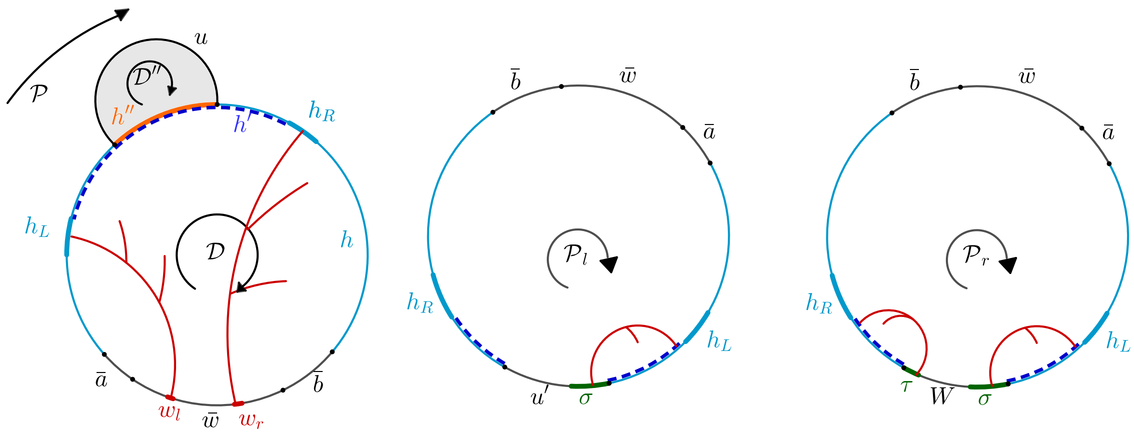

(a) Concatenation (b) Left-combing (c) Right-combing

Figure 12: Illustration of the different steps of the proof

Applying Theorem 3.16 to $\mathcal{P}$ produces a prism-geodesic $u^{(l)}$ for the word h'' and a disk-diagram $\mathcal{P}_l$ for $u^{(l)}h_{R-D+1}\cdots h_n\overline{awb}h_1\cdots h_{L+D-1}$ which is $\overline{u}^{(l)}$–left-combed and has the same beginning and ending functions with respect to $u^{(l)}$ as $\mathcal{P}$ with respect to u. In particular, $u^{(l)}$ can be written $\sigma u'$, where the subword σ is dual to graphs that only root in $h_{L+1}\cdots h_{L+D-1}$ (and no letter of u' has this property) and thus has prism length at most $D\mathsf{m}_H$. Applying Theorem 3.16 to $\mathcal{P}_l$ allows us to obtain a $\overline{u}^{(r)}$–right-combed disk diagram $\mathcal{P}_r$ for $\sigma u^{(r)}h_{R-D+1}\cdots h_n\overline{awb}h_1\cdots h_{L+D-1}$, where $u^{(r)}$ is a prism-geodesic representative for u' that is written in the form $W\tau$, where the graphs dual to τ only root in $h_{R-D+1}\cdots h_{R-1}$ (and no letter of W has this property) and is also of prism-length at most $D\mathsf{m}_H$. Therefore we have a prism-geodesic representative for h'' of the form $\sigma W\tau$, where the word W contains only u_i with dual graphs that have either their beginning or ending root in awb. As the beginning and ending roots are not affected by combing, we can conclude that $u_i \in G_{\mathsf{Star}(\Lambda)}$.

Moreover, there are only finitely many possibilities for σ and τ. Indeed, consider a letter u_j of σ, by Theorem 3.8 we have that $u_j = s_1\cdots s_p$ where the s_i are the letters on which the graph dual to u_j roots in $\mathcal{P}$ and $p \leqslant D\mathsf{m}_H$. Then the s_i are letters in the prism-geodesic expression of the elements of S_H. Since S_H is finite, and we fixed prism-geodesic representatives for its elements at the beginning, we have finitely many possibilities for each u_j. Hence, as σ is of bounded prism-length, there are also finitely many possibilities for σ. The same argument applies to τ.

Applying the above process in all the $\mathcal{P} = \mathcal{P}^{(k)}$, we get a sequence

$$\sigma^{(k)}W^{(k)}\tau^{(k)} = (h'')^{(k)} \in H$$

such that $W^{(k)} \in G_{\mathsf{Star}(\Lambda)}$. Because $\sigma^{(k)}$ and $\tau^{(k)}$ each range over finitely many possibilities, we pass to a subsequence $(h'')^{(k)} = \sigma W^{(k)}\tau$. Note that the length of $(h'')^{(k)}$ grows with the the length of $(h')^{(k)}$, which itself grows with the length of $w^{(k)}$. Therefore there is $j_0 \in \mathbb{N}$ and infinitely many $j \neq j_0$ such that the

$$(h'')^{(j_0)}((h'')^{(j)})^{-1} = \sigma W^{(1)}(W^{(j)})^{-1}\sigma^{-1}$$

are distinct and non-trivial elements of H conjugate by the same element of $\Gamma\mathcal{G}$ into $G_{\mathsf{Star}(\Lambda)}$. As Λ is a subgraph of a join, by Theorem 2.2 we have that $\mathsf{Star}(\Lambda)$ is a join, so this contradicts H being almost join-free. □

Remark 3.27. Note that if Λ were a subgraph of a star, then $\mathsf{Star}(\Lambda)$ would not necessarily be a subgraph of a star. This is the key reason why we need to use almost join-free and join-busting in the statement of Theorem 3.25, rather than the star equivalents.

4 Proof of Theorem 1.2

This section is dedicated to proving the main theorem of this paper. Beforehand, let us recall the statement.

Theorem 1.2. *Let Γ be a finite simple graph with no isolated vertices. Let $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of finitely generated infinite groups, and $P(\Gamma\mathcal{G})$ the prism complex of the graph product $\Gamma\mathcal{G}$. Let H be a finitely generated subgroup of the graph product $\Gamma\mathcal{G}$. Then the following are equivalent.*

1. *H is stable in $\Gamma\mathcal{G}$.*

2. *The orbit of H into the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ is a quasi-isometric embedding.*

3. *H is almost join-free.*

We will prove (3) $\implies$ (2) in Section 4.1, (1) $\implies$ (3) in Section 4.2, and (2) $\implies$ (1) in Section 4.3.

4.1 Almost join-free subgroups embed quasi-isometrically into the contact graph

We start the proof of Theorem 1.2 by proving (3) $\implies$ (2). Relying on Theorem 2.8, it is enough to show that H quasi-isometrically embeds into $\Gamma\mathcal{G}$ for the star metric. To do so, we will use Theorem 3.25 together with the next two results.

First of all, recall that given a finitely generated subgroup H of $\Gamma\mathcal{G}$ together with a finite generating set S_H, we implicitly consider all the elements in S_H to be written as S_p–geodesic words, and we denote the maximum S_p–length of such elements by m_H.

Proposition 4.1. *If H is a finitely generated almost star-free subgroup of a graph product $\Gamma\mathcal{G}$, then H is quasi-isometrically embedded in the prism Cayley graph of $\Gamma\mathcal{G}$.*

Proof. The proof follows [KMT17, Proposition 4.4]. Recall that $|\cdot|_H$ is the S_H–length of an element in H, and $|\cdot|_p$ the prism-length of an element in $\Gamma\mathcal{G}$. By definition of m_H, for every $h \in H$ we have that $|h|_H \geqslant |h|_p/\mathsf{m}_H$. Now consider a prism-geodesic word w for h, and a disk diagram for $h\overline{w}$. By Theorem 3.9, any dual graph with a root in $\overline{w}$ has its remaining roots in h. By Theorem 3.19 each graph dual to a letter of $\overline{w}$ has its extremal roots in h at S_H–distance at most D, and Theorem 3.20 ensures there is no vanishing subword of h of S_H–length greater than K. We can see that $|h|_H$ is maximal relative to $|h|_p = |w|_p$ when none of the dual graphs rooted in both h and w cross each other, all of the roots in h of each of these dual graphs are as far apart as possible, and the roots of different dual graphs are separated by maximally large vanishing subwords, as illustrated in Figure 13. This implies that $|h|_H \leqslant (K + D)|h|_p + K$. □

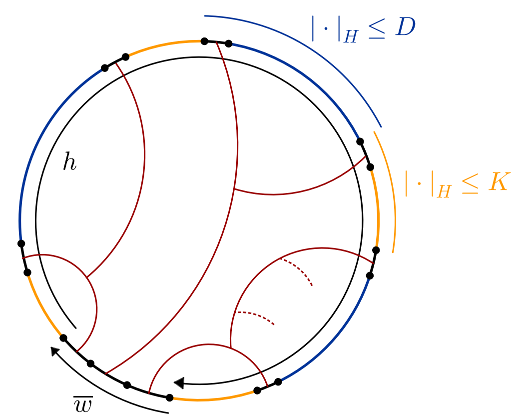

Figure 13: Hyperplane distribution in a disk diagram for $h\overline{w}$ with greatest difference between $|h|_H$ and $|h|_p$.

Lemma 4.2. *Let g be a non-trivial element of a graph product $\Gamma\mathcal{G}$. There is a star-geodesic word $s_1 \cdots s_n$ for g whose letters admit prism-geodesic representatives $s_i = w_{i,1} \cdots w_{i,j(i)}$ such that their concatenation $w_{1,1} \cdots w_{n,j(n)}$ is a prism-geodesic representative for g.*

Proof. Fix a star-geodesic word $s_1 \cdots s_n$ for g, and prism-geodesic words $w_{i,1} \cdots w_{i,j(i)}$ for each of the s_i. For all i let $v_i \in V(\Gamma)$ be such that $s_i \in G_{\mathsf{Star}(v_i)}$, and for $1 \leqslant \ell \leqslant j(i)$ let $v_{i,\ell} \in V(\Gamma)$ be such that $w_{i,\ell} \in G_{v_{i,\ell}}$, noting that $v_{i,\ell} \in \mathsf{Star}(v_i)$. If the prism word

$$(w_{1,1} \cdots w_{1,i(1)}) \cdots (w_{k,1} \cdots w_{k,j(k)}) \cdots (w_{n,1} \cdots w_{n,j(n)})$$

is not geodesic, then by Theorem 3.11 we can find $k < k'$ and p, p' such that $v_{k,p} = v_{k',p'}$, and such that each letter $w_{i,\ell}$ between $w_{k,p}$ and $w_{k',p'}$ has corresponding vertex $v_{i,\ell} \in \mathsf{Star}(v_{k,p})$. In particular $w_{k',p'}$ commutes with all these letters. Let $x = w_{k,p}w_{k',p'}$, represented by the empty word in the case $w_{k,p} = (w_{k',p'})^{-1}$. Using the prism geodesic criteria from Theorem 3.11, and the fact that $x, w_{k,p}$, and $w_{k',p'}$ are all in the same vertex group, one can verify that g has star-geodesic representative

$$s_1 \cdots s_{k-1} s'_k s_{k+1} \cdots s_{k'-1} s'_{k'} s_{k'+1} \cdots s_n,$$

where $s'_k \in G_{\mathsf{Star}(v_k)}$ and $s'_{k'} \in G_{\mathsf{Star}(v_{k'})}$ have respective prism-geodesic representatives

$$\begin{aligned} s'_k &= w_{k,1} \cdots w_{k,p-1} x w_{k,p+1} \cdots w_{k,j(k)} \qquad \text{and} \\ s'_{k'} &= w_{k',1} \cdots w_{k',p'-1} w_{k',p'+1} \cdots w_{k,'j(k')}. \end{aligned}$$

Thus we obtain a star geodesic for g whose letters have prism geodesic representatives such that their total prism word length is strictly lower than that of the word we started with. The proof ends by applying this process finitely many times. □

Theorem 4.3. *If H is a finitely generated almost join-free subgroup of a graph product $\Gamma\mathcal{G}$, where Γ has no isolated vertices, then H is quasi-isometrically embedded in the star Cayley graph of $\Gamma\mathcal{G}$.*

Proof. Recall that, by Theorem 4.1, H quasi-isometrically embeds into the prism Cayley graph of $\Gamma\mathcal{G}$. It therefore suffices to show that the star and prism metrics are quasi-isometric on H.

We first observe that $S_p \subset S_\star$ implies $|\cdot|_\star \leqslant |\cdot|_p$. For the other bound, take $h \in H$, and let $s_1 \cdots s_n$ be a star-geodesic representative obtained from Theorem 4.2, so that $|h|_p = |s_1|_p + \cdots + |s_n|_p$. As H is almost join-free, Theorem 3.25 ensures that H is N–join-busting for some $N > 0$. As Γ has no isolated vertices, every star is a join, so the prism length of each s_i is at most N. We therefore have that $|h|_p \leqslant N|h|_\star$. □

Using Theorem 2.8, (3) $\implies$ (2) is now a direct corollary of Theorem 4.3.

Corollary 4.4. *Suppose H is a finitely generated almost join-free subgroup of a graph product $\Gamma\mathcal{G}$, where Γ has no isolated vertices. Then H is quasi-isometrically embedded in the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ via its orbit maps.*

4.2 Stable subgroups are almost join-free

In this section we prove (1) $\implies$ (3) in Theorem 1.2. This also follows from the more general well-known fact that any infinite intersection with a quasi-isometrically embedded direct product of infinite groups is an obstruction to a subgroup being stable. As we have not yet found a direct reference for this, we include the below for completeness.

Proposition 4.5. *Let $\Gamma\mathcal{G}$ be a graph product of finitely generated infinite groups. If $H \leqslant \Gamma\mathcal{G}$ is stable, then H is almost join-free.*

Proof. Suppose H is not almost join-free, so there exists infinite $A \subset H$ such that A is conjugate into a join subgroup G_Λ. As stability is preserved under conjugation (Theorem 2.1), we may assume A itself is contained in G_Λ.

Consider a standard generating set S for $\Gamma\mathcal{G}$ (*ie.* the generating set obtained by taking a union of finite generating sets for each vertex group in $\mathcal{G}$). By the normal form for graph products [Gre90], under this generating set every parabolic subgroup can be considered as being isometrically embedded in $\Gamma\mathcal{G}$. As Λ is a join of two subgraphs Γ_1 and Γ_2 of Γ, recall that G_Λ is isomorphic to $G_{\Gamma_1} \times G_{\Gamma_2}$, and each one of G_Λ, G_{Γ_1}, and G_{Γ_2} are isometrically embedded in $\Gamma\mathcal{G}$ under the generating set S.

As $A = H \cap G_\Lambda$ is infinite, the projection of A to at least one of G_{Γ_1} and G_{Γ_2} must also be infinite. Call these projections X and Y respectively, and without loss of generality assume that X is infinite. For every $M > 0$ we can therefore pick $x_1, x_2 \in X$ and $y_1, y_2 \in Y$ such that $d_S(x_1, x_2) > M$ and $x_1y_1, x_2y_2 \in A$. It follows from [DT15, Lemma 6.2] that there exist (3,0)-quasi-geodesics in G_Λ, and therefore in $\Gamma\mathcal{G}$, between x_1y_1 and x_2y_2 such that at least one of these is not in the M-neighborhood of the other. This means that H cannot be stable. □

4.3 Quasi-isometrically embedded subgroups are stable

The following proves (2) $\implies$ (1) in Theorem 1.2. This follows as a simple consequence of Theorem 2.7, and the work of Abbott and Manning in [AM24].

Proposition 4.6. *Let H be a finitely generated subgroup of a finitely generated graph product $\Gamma\mathcal{G}$. If the orbit maps of H into the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ are quasi-isometric embeddings, then H is stable.*

Proof. If H is finite, then it is automatically stable, so assume that H is infinite. Recall from Theorem 2.7 that $\mathcal{C}P(\Gamma\mathcal{G})$ is hyperbolic, and that the action of $\Gamma\mathcal{G}$ on $\mathcal{C}P(\Gamma\mathcal{G})$ is acylindrical. Because the orbit of H into $\mathcal{C}P(\Gamma\mathcal{G})$ is a quasi-isometric embedding by assumption, we can conclude that $(\Gamma\mathcal{G}, \mathcal{C}P(\Gamma\mathcal{G}), H)$ is an A/QI triple, as defined in [AM24, Definition 1.2]. We also have that $\Gamma\mathcal{G}$ is finitely generated by assumption. Thus, [AM24, Theorem 1.5] implies that H is stable. □

5 Purely loxodromic subgroups

In [KMT17, Theorem 1.1], it is shown that the stable subgroups of right-angled Artin groups are exactly the finitely generated purely loxodromic subgroups. We recall that an isometry g of a hyperbolic space X is called *loxodromic* if $\langle g \rangle$ has unbounded orbits, and fixes exactly two points on the Gromov boundary ∂X. If G is a group acting by isometries on X, a subgroup $H \leqslant G$ is *purely loxodromic* if every infinite order $h \in H$ acts loxodromically on X.

We can show that every stable subgroup of a graph product is purely loxodromic via the following standard observation, combined with Theorem 1.2 and the fact that the action of $\Gamma\mathcal{G}$ on $\mathcal{C}P(\Gamma\mathcal{G})$ is acylindrical (see Theorem 2.7). Note that in an acylindrical action on a hyperbolic space, every element with unbounded orbits is loxodromic [Bow08, Lemma 2.2].

Lemma 5.1. *Let G be a group acting acylindrically on a hyperbolic space X. Let H be a finitely generated subgroup of G. If the orbit of H into X is a quasi-isometric embedding, then H is purely loxodromic with respect to that action.*

Proof. Let $h \in H$ be an infinite order element, so $\langle h \rangle$ has infinite diameter in H. By assumption, the orbit of H into X is a quasi-isometric embedding, so $\langle h \rangle$ has unbounded orbits in X. As the action of G on X is acylindrical, h must be loxodromic. □

On the other hand, if a subgroup of a graph product is purely loxodromic, then it does not in general have to be stable. If Γ is connected, and some vertex group G_v is an infinite finitely generated torsion group, then G_v is (vacuously) purely loxodromic. On the other hand, G_v has infinite intersection with a join subgroup, so is not almost join-free, and therefore is not stable by Theorem 4.5.

However, if we add the assumption that $\Gamma\mathcal{G}$ cannot contain infinite torsion subgroups (for example if $\Gamma\mathcal{G}$ is virtually torsion-free, as in the right-angled Artin group case), we are able to obtain equivalence. The following proposition (which will help us establish the equivalence) is known, see for example [FK24, Proposition 2.10], however we include a short proof here for completeness.

Proposition 5.2. *Let $\Gamma\mathcal{G}$ be a graph product, and $P(\Gamma\mathcal{G})$ the prism complex of the graph product $\Gamma\mathcal{G}$. Then every element in $\Gamma\mathcal{G}$ that is conjugate into a join subgroup has bounded orbits in the action of $\Gamma\mathcal{G}$ on the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$.*

Proof. Let $G_\Lambda \leqslant \Gamma\mathcal{G}$ be a join subgroup, and let Γ_1 and Γ_2 be subgraphs of Γ such that $G_\Lambda = G_{\Gamma_1} \times G_{\Gamma_2}$. Let $v_1 \in \Gamma_1$ and $v_2 \in \Gamma_2$. Then $\Gamma_1 \subset \mathsf{Star}(v_2)$ and $\Gamma_2 \subset \mathsf{Star}(v_1)$. Therefore every $h \in G_\Lambda$ can be written as $h = h_1 h_2$, with $h_1 \in G_{\mathsf{Star}(v_2)}$ and $h_2 \in G_{\mathsf{Star}(v_1)}$.

This implies that G_Λ lies in a ball of radius 2 around the identity in the star Cayley graph of $\Gamma\mathcal{G}$. As a consequence, any $g \in G_\Lambda$ must have bounded orbits in the star Cayley graph, and so any conjugate must also have bounded orbits. Recall that the star Cayley graph is $\Gamma\mathcal{G}$–equivariantly quasi-isometric to $\mathcal{C}P(\Gamma\mathcal{G})$, which concludes the proof. □

Corollary 5.3. *Let $\Gamma\mathcal{G}$ be a graph product with no infinite torsion subgroups. If $H \leqslant \Gamma\mathcal{G}$ is purely loxodromic with respect to the action on $\mathcal{C}P(\Gamma\mathcal{G})$, then H is almost join-free.*

Proof. Suppose H is purely loxodromic but has infinite intersection with some conjugate of a join subgroup. By our assumption on $\Gamma\mathcal{G}$, this intersection contains an element of infinite order, which as H is purely loxodromic must have unbounded orbits in the action on $\mathcal{C}P(\Gamma\mathcal{G})$. This contradicts Theorem 5.2. □

As a result, in this case we obtain a fourth condition equivalent to stability.

Theorem 5.4. *Let Γ be a finite simple graph with no isolated vertices, $\mathcal{G} = \{G_v \mid v \in V(\Gamma)\}$ be a collection of infinite finitely generated groups with no infinite torsion subgroups, and $P(\Gamma\mathcal{G})$ the prism complex of the graph product $\Gamma\mathcal{G}$. Let H be a finitely generated subgroup of the graph product $\Gamma\mathcal{G}$. Then the following are equivalent.*

1. *H is stable in $\Gamma\mathcal{G}$.*

2. *The orbit maps of H into the contact graph $\mathcal{C}P(\Gamma\mathcal{G})$ are quasi-isometric embeddings.*

3. *H is almost join-free.*

4. *H is purely loxodromic with respect to the action on $\mathcal{C}P(\Gamma\mathcal{G})$.*

References

[ABD21] Carolyn Abbott, Jason Behrstock, and Matthew Gentry Durham, *Largest acylindrical actions and stability in hierarchically hyperbolic groups*, Trans. Amer. Math. Soc. Ser. B **8** (2021), 66–104, With an appendix by Daniel Berlyne and Jacob Russell.

[ADT17] Tarik Aougab, Matthew Gentry Durham, and Samuel J. Taylor, *Pulling back stability with applications to* $\mathrm{Out}(F_n)$ *and relatively hyperbolic groups*, J. Lond. Math. Soc. (2) **96** (2017), no. 3, 565–583.

[AM24] Carolyn R. Abbott and Jason F. Manning, *Acylindrically hyperbolic groups and their quasi-isometrically embedded subgroups*, Michigan Mathematical Journal **72** (2024), 357–402.

[AMST19] Yago Antolín, Mahan Mj, Alessandro Sisto, and Samuel J. Taylor, *Intersection properties of stable subgroups and bounded cohomology*, Indiana University Mathematics Journal **68** (2019), no. 1, 179–199.

[BBKL20] Mladen Bestvina, Kenneth Bromberg, Autumn E. Kent, and Christopher J. Leininger, *Undistorted purely pseudo-Anosov groups*, J. Reine Angew. Math. **760** (2020), 213–227.

[BCK+25] Sahana Balasubramania, Marissa Chesser, Alice Kerr, Johanna Mangahas, and Marie Trin, *(Non-)recognizing spaces for stable subgroups*, Proc. Amer. Math. Soc. **153** (2025), 3197–3207.

[Bow08] Brian H. Bowditch, *Tight geodesics in the curve complex*, Inventiones mathematicae **171** (2008), 281–300.

[But23] J. O. Button, *Quasi-actions whose quasi-orbits are quasi-isometric to trees*, Preprint, arXiv:2305.13105 (2023).

[CD17] Matthew Cordes and Matthew Gentry Durham, *Boundary convex cocompactness and stability of subgroups of finitely generated groups*, International Mathematics Research Notices **2019** (2017), no. 6, 1699–1724.

[CDEK25] Christopher H. Cashen, Pallavi Dani, Alexandra Edletzberger, and Annette Karrer, *RAAGedy right-angled Coxeter groups*, Preprint, arXiv:2506.16789 (2025).

[CW04] John Crisp and Bert Wiest, *Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups*, Algebr. Geom. Topol. **4** (2004), 439–472.

[DT15] Matthew Durham and Samuel J Taylor, *Convex cocompactness and stability in mapping class groups*, Algebraic & Geometric Topology **15** (2015), no. 5, 2839–2859.

[Fen83] Roger A. Fenn, *Techniques of geometric topology*, London Mathematical Society Lecture Note Series, vol. 57, Cambridge University Press, Cambridge, 1983. MR 787801

[FK24] Elia Fioravanti and Alice Kerr, *Short positive loxodromics in graph products*, Preprint, arXiv:2410.06751 (2024).

[FM02] Benson Farb and Lee Mosher, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. **6** (2002), 91–152 (English).

[Gen17] Anthony Genevois, *Cubical-like geometry of quasi-median graphs and applications to geometric group theory*, Preprint, arXiv:1712.01618 (2017).

[Gen19a] ———, *Hyperbolicities in* CAT(0) *cube complexes*, Enseign. Math. **65** (2019), no. 1-2, 33–100.

[Gen19b] ———, *On the geometry of van kampen diagrams of graph products of groups*, Preprint, arXiv:1901.04538 (2019).

[GM19] Anthony Genevois and Alexandre Martin, *Automorphisms of graph products of groups from a geometric perspective*, Proceedings of the London Mathematical Society **119** (2019), 1745–1779.

[Gre90] Elisabeth Ruth Green, *Graph products of groups*, Ph.D. thesis, University of Leeds, 1990.

[Hag14] Mark F. Hagen, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, Journal of Topology **7** (2014), 385–418.

[Ham05] Ursula Hamenstaedt, *Word hyperbolic extensions of surface groups*, Preprint, arXiv:math/0505244 (2005).

[Kim08] Sang-hyun Kim, *Co-contractions of graphs and right-angled Artin groups*, Algebr. Geom. Topol. **8** (2008), no. 2, 849–868.

[Kim19] Heejoung Kim, *Stable subgroups and morse subgroups in mapping class groups*, International Journal of Algebra and Computation **29** (2019), no. 05, 893–903.

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

[KK14] Sang-Hyun Kim and Thomas Koberda, *The geometry of the curve graph of a right-angled Artin group*, Internat. J. Algebra Comput. **24** (2014), no. 2, 121–169.

[KL08] Autumn E. Kent and Christopher J. Leininger, *Shadows of mapping class groups: Capturing convex cocompactness*, GAFA Geom. funct. anal. **18** (2008), 1270–1325.

[KMT17] Thomas Koberda, Johanna Mangahas, and Samuel J. Taylor, *The geometry of purely loxodromic subgroups of right-angled Artin groups*, Trans. Amer. Math. Soc. **369** (2017), no. 11, 8179–8208.

[LS77] Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete [Results in Mathematics and Related Areas], vol. Band 89, Springer-Verlag, Berlin-New York, 1977.

[Oh22] Sangrok Oh, *Quasi-isometry invariants of weakly special square complexes*, Topology Appl. **307** (2022), Paper No. 107945, 46.

[Oya24] Koichi Oyakawa, *Infinite graph product of groups I: geometry of the extension graph*, Preprint, arXiv:2409.09527 (2024).

[PSZ24] Harry Petyt, Davide Spriano, and Abdul Zalloum, *Hyperbolic models for CAT(0) spaces*, Adv. Math. **450** (2024).

[PZ24] Harry Petyt and Abdul Zalloum, *Constructing metric spaces from systems of walls*, Preprint, arXiv:2404.12057 (2024).

[Sis16] Alessandro Sisto, *Quasi-convexity of hyperbolically embedded subgroups*, Math. Z. **283** (2016), no. 3-4, 649–658.

[Tra17] Hung Kong Tran, *Malnormality and join-free subgroups in right-angled coxeter group*, Preprint, arXiv:1703.09032 (2017).

[Tra19] Hung Cong Tran, *On strongly quasiconvex subgroups*, Geometry & Topology **23** (2019), 1173–1235.

[Val21] Motiejus Valiunas, *Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products*, Groups Geom. Dyn. **15** (2021), no. 1, 143–195.

[Zbi24] Stefanie Zbinden, *Hyperbolic spaces that detect all strongly-contracting directions*, Preprint, arXiv:2404.12162 (2024).