Groups Geom. Dyn. 16 (2022), 523–580
DOI 10.4171/GGD/652

Hierarchical hyperbolicity of graph products

Daniel Berlyne and Jacob Russell

Abstract. We show that any graph product of finitely generated groups is hierarchically hyperbolic relative to its vertex groups. We apply this result to answer two questions of Behrstock, Hagen, and Sisto: we show that the syllable metric on any graph product forms a hierarchically hyperbolic space, and that graph products of hierarchically hyperbolic groups are themselves hierarchically hyperbolic groups. This last result is a strengthening of a result of Berlai and Robbio by removing the need for extra hypotheses on the vertex groups. We also answer two questions of Genevois about the geometry of the electrification of a graph product of finite groups.

Contents

1. Introduction

There have been many attempts to generalise the notion of hyperbolicity of a group since it was first introduced by Gromov [19]. One of these, *hierarchical hyperbolicity*, was developed by Behrstock, Hagen, and Sisto [3, 5] as a way of describing hyperbolic behaviour in quasi-geodesic metric spaces via hierarchy machinery akin to that constructed for mapping class groups by Masur and Minsky [22, 23]. The work of Behrstock, Hagen, and Sisto originally focused on developing such machinery for right-angled Artin groups, but also encompasses a wide variety of groups and spaces, such as virtually cocompact special groups [3], 3-manifold groups with no Nil or Sol components [3], Teichmüller space with either the Teichmüller or Weil–Petersson metric [5, 7, 9, 10, 13, 22, 26], and graph products of hyperbolic groups [8]. Hierarchical hyperbolicity has deep geometric consequences for a space, including a Masur and Minsky style distance formula [3], a

2020 Mathematics Subject Classification. Primary 20F65; Secondary 20F67, 20F55, 20E08.
Keywords. Hierarchically hyperbolic, graph product.

quadratic isoperimetric inequality [3], rank rigidity and Tits alternative theorems [11, 12], control over top-dimensional quasi-flats [6], and bounds on the asymptotic dimension [4].

A hierarchically hyperbolic structure on a quasi-geodesic space $\mathcal{X}$ is a collection of uniformly hyperbolic spaces $C(W)$ indexed by the elements W of an index set $\mathfrak{S}$. For each $W \in \mathfrak{S}$, there is a projection map from $\mathcal{X}$ onto the hyperbolic space $C(W)$, and every pair of elements of $\mathfrak{S}$ is related by one of three mutually exclusive relations: orthogonality, nesting, and transversality. This data then satisfies a collection of axioms that allow for the coarse geometry of the entire space to be recovered from the projections to the hyperbolic spaces $C(W)$.

In the present paper, we construct an explicit hierarchy structure for any graph product, using right-angled Artin groups as our motivating example. Given a finite simplicial graph Γ with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$, we define the *right-angled Artin group* A_Γ by

$$A_\Gamma = \big\langle V(\Gamma) \mid [v, w] = e \ \forall \{v, w\} \in E(\Gamma) \big\rangle.$$

More generally, if we associate to each vertex v of Γ a finitely-generated group G_v, then we define the *graph product* G_Γ by

$$G_\Gamma = \Big(\mathop{\ast}_{v \in V(\Gamma)} G_v \Big) \Big/ \big\langle\!\big\langle [g_v, g_w] \mid g_v \in G_v,\ g_w \in G_w,\ \{v, w\} \in E(\Gamma) \big\rangle\!\big\rangle,$$

so that A_Γ is obtained as the special case where the vertex groups are $G_v = \mathbb{Z}$ for all $v \in V(\Gamma)$.

For right-angled Artin groups A_Γ, a hierarchically hyperbolic structure was constructed by Behrstock, Hagen, and Sisto by considering the collection of induced subgraphs of the defining graph Γ [5]. Each induced subgraph Λ of Γ generates a new right-angled Artin group A_Λ, which is realised as a subgroup of A_Γ. The Cayley graph of A_Γ is the 1-skeleton of a CAT(0) cube complex X, which comes equipped with a projection to a hyperbolic space $C(X)$ called the *contact graph*. Since each induced subgraph Λ of Γ generates its own right-angled Artin group with associated cube complex $Y \subseteq X$, the subgroup A_Λ has its own associated contact graph $C(Y)$. Since edges of Γ correspond to commuting relations in A_Γ, join subgraphs of Γ (that is, subgraphs of the form $\Lambda_1 \sqcup \Lambda_2$ where every vertex of Λ_1 is joined by an edge to every vertex of Λ_2) generate direct product subgroups of A_Γ. This provides us with an intuitive notion of *orthogonality* within our hierarchy. Set containment of subgraphs of Γ provides a natural partial order in the hierarchy, which we call *nesting*, and any subgraphs that are not orthogonal or nested are considered *transverse*. Collectively, the hyperbolic spaces $C(Y)$ allow us to recover the entire geometry of A_Γ, via projections to the subcomplexes $Y \subseteq X$ and through the nesting, orthogonality, and transversality relations defined above.

Since the nesting and orthogonality relations for a right-angled Artin group are intrinsic to the defining graph Γ, it is sensible to attempt to generalise this hierarchy structure to arbitrary graph products. It is important to note, however, that arbitrary graph products may *not* be hierarchically hyperbolic, since we have no control over the vertex groups. For example, the vertex groups could be copies of $\mathrm{Out}(F_3)$, which is known not to be hierarchically hyperbolic [3]. However, this is the only roadblock. Specifically, we show that graph products are *relatively hierarchically hyperbolic*; that is, graph products admit a structure satisfying all of the axioms of hierarchical hyperbolicity with the exception that the spaces associated to the nesting-minimal sets (the vertex groups) are not necessarily hyperbolic.

Theorem A. *Let Γ be a finite simplicial graph, with each vertex v labelled by a finitely-generated group G_v. The graph product G_Γ is a hierarchically hyperbolic group relative to the vertex groups.*

The notion of relative hierarchical hyperbolicity was originally developed by Behrstock, Hagen, and Sisto in [3] and is explored further in [4]. Despite the lack of hyperbolicity in the nesting-minimal sets, many of the consequences of hierarchical hyperbolicity are preserved in the relatively hierarchically hyperbolic setting. In particular, Theorem A implies that the graph product G_Γ has a Masur and Minsky style distance formula and an acylindrical action on the nesting-maximal hyperbolic space; see Corollaries 4.23 and 4.24.

Another way of asserting control over the vertex groups is by replacing the word metric on G_Γ with the *syllable metric*, which measures the length of an element $g \in G_\Gamma$ by counting the minimal number of elements needed to express g as a product of vertex group elements. This has the effect of making all vertex groups diameter 1, and therefore hyperbolic. The syllable metric on a right-angled Artin group was studied by Kim and Koberda as an analogue of the Weil–Petersson metric on Teichmüller space (the Weil–Petersson metric is quasi-isometric to the space obtained from the mapping class group by coning off all cyclic subgroups generated by Dehn twists) [20]. Kim and Koberda produce several hierarchy-like results for the syllable metric on a right-angled Artin group with triangle- and square-free defining graph, including a Masur and Minsky style distance formula and an acylindrical action on a hyperbolic space. This inspired Behrstock, Hagen, and Sisto to ask if the syllable metric on a right-angled Artin group is a hierarchically hyperbolic space [3]. We give a positive answer to this question, not just for right-angled Artin groups but for all graph products.

Corollary B. *Let Γ be a finite simplicial graph, with each vertex v labelled by a group G_v. The graph product G_Γ endowed with the syllable metric is a hierarchically hyperbolic space.*

To prove Theorem A and Corollary B, we utilise techniques developed by Genevois and Martin in [14, 17], which exploit the cubical-like geometry of a graph product when endowed with the syllable metric. This allows us to adapt proofs from the right-angled Artin group case, which rely heavily on geometric properties of cube complexes. While the syllable metric does not appear in the statement of Theorem A, it is an integral part of the proof, acting as a middle ground where geometric computations are performed before projecting to the associated hyperbolic spaces. This also allows Theorem A and Corollary B to be proved essentially simultaneously.

Our primary application of Theorem A is showing that a graph product of hierarchically hyperbolic groups is itself hierarchically hyperbolic. This gives a positive answer to another question of Behrstock, Hagen, and Sisto [3, Question D].

Theorem C. *Let Γ be a finite simplicial graph, with each vertex v labelled by a group G_v. If each G_v is a hierarchically hyperbolic group, then the graph product G_Γ is a hierarchically hyperbolic group.*

Berlai and Robbio have established a combination theorem for graphs of hierarchically hyperbolic groups that they use to prove Theorem C when the vertex groups satisfy some natural, but non-trivial, additional hypotheses [8]. For the specific case of graph products, Theorem C improves upon Berlai and Robbio's result by removing the need for these additional hypotheses, as well as providing an explicit description of the hierarchically hyperbolic structure in terms of the defining graph.

We also use our relatively hierarchically hyperbolic structure for graph products to answer two questions of Genevois about a new quasi-isometry invariant for graph products of finite groups called the *electrification* of G_Γ. Graph products of finite groups form a particularly interesting class, as they include right-angled Coxeter groups and are the only cases where the syllable metric and word metric are quasi-isometric. Genevois defines the electrification $\mathbb{E}(\Gamma)$ of a graph product of finite groups by taking the syllable metric on G_Γ and adding edges between elements g, h whenever $g^{-1}h \in G_\Lambda \leq G_\Gamma$ and Λ is a *minsquare* subgraph of Γ, that is, a minimal subgraph that contains opposite vertices of a square if and only if it contains the whole square. Motivated by an analogy with relatively hyperbolic groups, Genevois proved that any quasi-isometry between graph products of finite groups induces a quasi-isometry between their electrifications, and used this invariant to distinguish several quasi-isometry classes of right-angled Coxeter groups [16]. Geometrically, the electrification sits between the syllable metric on G_Γ and the nesting-maximal hyperbolic space in our hierarchically hyperbolic structure on G_Γ. We exploit this situation to classify when the electrification has bounded diameter and when it is a quasi-line, answering Questions 8.3 and 8.4 of [16].

Theorem D. *Let G_Γ be a graph product of finite groups and let $\mathbb{E}(\Gamma)$ be its electrification.*

(1) $\mathbb{E}(\Gamma)$ *has bounded diameter if and only if Γ is either a complete graph, a minsquare graph, or the join of minsquare graph and a complete graph.*

(2) $\mathbb{E}(\Gamma)$ *is a quasi-line if and only if G_Γ is virtually cyclic.*

As a final application of Theorem A, we give a new proof of Meier's classification of hyperbolicity of graph products [24].

Outline of the paper. We begin by introducing the necessary tools from the geometry of graph products in Section 2.1 and reviewing the definition of a relative hierarchically hyperbolic group (HHG) in Section 2.2. In Section 3, we set up our proof of the relative hierarchical hyperbolicity of graph products by defining the necessary spaces, projections, and relations. In Section 4, we show that the spaces, projections, and relations defined in

Section 3 satisfy the axioms of a relative HHG (or non-relative hierarchically hyperbolic space in the case of the syllable metric). This completes the proofs of Theorem A and Corollary B. Section 5 is devoted to applications. We start by proving that graph products of HHGs are HHGs (Theorem C) in Section 5.1, which requires a technical result that can be found in the appendix of [1]. In Section 5.2, we record our proof of Meier's hyperbolicity criteria, and in Section 5.3, we classify when Genevois' electrification has infinite diameter and when it is a quasi-line, proving Theorem D.

2. Background

2.1. Graph products

Definition 2.1 (Graph product). Let Γ be a finite simplicial graph with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$, and with each vertex $v \in V(\Gamma)$ labelled by a group G_v. The *graph product* G_Γ is the group

$$G_\Gamma = \Big(\mathop{\ast}_{v \in V(\Gamma)} G_v \Big) \Big/ \langle\langle [g_v, g_w] \mid g_v \in G_v,\ g_w \in G_w,\ \{v, w\} \in E(\Gamma) \rangle\rangle.$$

We call the G_v the *vertex groups* of the graph product G_Γ.

By deleting any vertices labelled by the trivial group, every graph product is isomorphic to a graph product where each vertex group is non-trivial. We will therefore operate under the standing assumption that the vertex groups of our graph products are always non-trivial.

Note that if all vertex groups of G_Γ are copies of $\mathbb{Z}$, then G_Γ is the right-angled Artin group with defining graph Γ, and if all vertex groups are copies of $\mathbb{Z}/2\mathbb{Z}$, then G_Γ is the corresponding right-angled Coxeter group.

We wish to study the geometry of G_Γ by adapting the cubical geometry of right-angled Artin groups. To this end, we will first need to eliminate any badly behaved geometry occurring within vertex groups. We do this by replacing the usual word metric with the *syllable metric*.

Definition 2.2 (Syllable metric on a graph product). Let G_Γ be a graph product. The graph $S(\Gamma)$ is the metric graph whose vertices are elements of G_Γ and where $g, h \in G_\Gamma$ are joined by an edge of length 1 labelled by $g^{-1}h$ if there exists a vertex v of Γ such that $g^{-1}h \in G_v$. We denote the distance in $S(\Gamma)$ by $\mathsf{d}_{\mathrm{syl}}(\cdot,\cdot)$ and say $\mathsf{d}_{\mathrm{syl}}(g,h)$ is the *syllable distance* between g and h. When convenient, we will use $|g|_{\mathrm{syl}}$ to denote $\mathsf{d}_{\mathrm{syl}}(e,g)$ and call it the *syllable length* of g.

Notice that all cosets of vertex groups have diameter 1 under the syllable metric, thus trivialising their geometry. Therefore, when working with $S(\Gamma)$, instead of expressing an element $g \in G_\Gamma$ as a word in the generators of G_Γ, it is more geometrically meaningful to express g as a product of *any* elements of vertex groups.

Definition 2.3 (Syllable expressions). Let G_Γ be a graph product and $g \in G_\Gamma$. If $g = s_1 \dots s_n$ where each $s_i \in G_{v_i}$ for some $v_i \in V(\Gamma)$, then we say that $s_1 \dots s_n$ is a *syllable expression* for g. If $s_1 \dots s_n$ is a syllable expression for g and $n = \mathrm{d}_{\mathrm{syl}}(e, g)$, then we say that $s_1 \dots s_n$ is a *reduced syllable expression* for g. In this case, n is the smallest number of terms possible for any syllable expression of g.

A foundational fact about graph products is that any syllable expression can be reduced by applying a sequence of canonical moves.

Theorem 2.4 (Reduction algorithm for graph products; [18, Theorem 3.9]). *Let G_Γ be a graph product and $g \in G_\Gamma$. If $s_1 \dots s_n$ is a reduced syllable expression for g and $t_1 \dots t_m$ is a syllable expression for g, then $t_1 \dots t_m$ can be transformed into $s_1 \dots s_n$ by applying a sequence of the following three moves:*

- *remove a term t_i if $t_i = e$;*
- *replace consecutive terms t_i and t_{i+1} belonging to the same vertex group G_v with the single term $t_i t_{i+1} \in G_v$;*
- *exchange the position of consecutive terms t_i and t_{i+1} when $t_i \in G_v$ and $t_{i+1} \in G_w$ with v joined to w by an edge in Γ.*

The next corollary shows that when each of the vertex groups of the graph product is finitely generated, Theorem 2.4 implies that the word length of any $g \in G_\Gamma$ will be the sum of the word lengths of the terms in any reduced syllable expression for g.

Corollary 2.5 (Reduced syllable expressions minimise word length). *Let G_Γ be a graph product of finitely generated groups. For each $v \in V(\Gamma)$, let S_v be a finite generating set for the vertex group G_v, and let $|g|$ be the word length of $g \in G_\Gamma$ with respect to the finite generating set $S = \bigcup_{v \in V(\Gamma)} S_v$. For all $g \in G_\Gamma$, if $s_1 \dots s_n$ is a reduced syllable expression for g, then*

$$|g| = \sum_{i=1}^{n} |s_i|.$$

Proof. Let $s_1 \dots s_n$ be a reduced syllable expression for $g \in G_\Gamma$. There exist $w_1, \dots, w_m \in S$ such that $|g| = m$ and $g = w_1 \dots w_m$. Since every element of S is an element of one of the vertex groups of G_Γ, the product $w_1 \dots w_m$ is also a syllable expression for g. Thus, by applying a finite number of the moves from Theorem 2.4, we can transform $w_1 \dots w_m$ into $s_1 \dots s_n$. We can therefore write each s_i as a product $s_i = w_{\sigma_i(1)} \dots w_{\sigma_i(m_i)}$, where $m_i \le m$ and σ_i is a permutation of $\{1, \dots, m\}$. Further, if $i \ne k$, then $\{\sigma_i(1), \dots, \sigma_i(m_i)\} \cap \{\sigma_k(1), \dots, \sigma_k(m_k)\} = \emptyset$. Thus,

$$\sum_{i=1}^{n} |s_i| \le \sum_{i=1}^{n} m_i \le m.$$

However, $m = |g| \le \sum_{i=1}^{n} |s_i|$ by definition, so $|g| = \sum_{i=1}^{n} |s_i|$. ∎

Another critical consequence of Theorem 2.4 is that the terms in a reduced syllable expression for an element of a graph product are well defined up to applying the commutation relation. This ensures that the following notions are well defined for an element of a graph product.

Definition 2.6 (Syllables and support of an element). Let G_Γ be a graph product and let $g \in G_\Gamma$. If $s_1 \dots s_n$ is a reduced syllable expression for g, then we call the s_i the *syllables* of g and use $\operatorname{supp}(g)$ to denote the maximal subgraph of Γ with vertex set $\{v_1, \dots, v_n\}$, where $s_i \in G_{v_i}$. We call $\operatorname{supp}(g)$ the *support* of g.

Convention 2.7. Whenever we consider a subgraph $\Lambda \subseteq \Gamma$, we will assume that Λ is both non-empty and an induced subgraph of Γ. That is, whenever v, w are vertices of Λ that are joined by an edge of Γ, then v and w are joined by an edge of Λ as well.

Another hallmark feature of graph products is their rich collection of subgroups corresponding to subgraphs of the defining graph.

Definition 2.8 (Graphical subgroups). Let G_Γ be a graph product with vertex groups $\{G_v : v \in V(\Gamma)\}$ and let $\Lambda \subseteq \Gamma$ be a subgraph. We use $\langle\Lambda\rangle$ to denote the subgroup of G_Γ generated by $\{G_v : v \in V(\Lambda)\}$. We call such subgroups the *graphical subgroups* of G_Γ. Note, each subgroup $\langle\Lambda\rangle$ is isomorphic to the graph product G_Λ.

Since the graphical subgroups are themselves graph products, we can also define the syllable metric on them and their cosets.

Definition 2.9 (Syllable metric on graphical subgroups). Let G_Γ be a graph product, $g \in G_\Gamma$, and $\Lambda \subseteq \Gamma$. Let $S(\Lambda)$ be the metric graph defined in Definition 2.2 for the graph product $\langle\Lambda\rangle$, and let $S(g\Lambda)$ denote the metric graph whose vertices are elements of the coset $g\langle\Lambda\rangle$ and where gx and gy are joined by an edge of length 1 if x and y are joined by an edge in $S(\Lambda)$.

Remark 2.10 (Graphical subgroups are convex in $S(\Gamma)$). Geodesics in $S(\Gamma)$ between two elements k and h are labelled by the reduced syllable forms of $k^{-1}h$. The induced subgraph of $S(\Gamma)$ with vertex set $g\langle\Lambda\rangle$ is therefore convex and graphically isomorphic to $S(g\Lambda)$ via the identity map. In particular, the distance between two vertices k, h of $S(g\Lambda)$ is $\mathsf{d}_{\mathrm{syl}}(k, h)$.

In order to analyse how the graphical subgroups of G_Γ interact, we make extensive use of the following definitions from graph theory.

Definition 2.11 (Star, link, and join). Let Γ be a finite simplicial graph and Λ a subgraph of Γ. The *link* of Λ, denoted by $\operatorname{lk}(\Lambda)$, is the subgraph spanned by the vertices of $\Gamma \smallsetminus \Lambda$ that are connected to every vertex of Λ. The *star* of Λ, denoted by $\operatorname{st}(\Lambda)$, is $\Lambda \cup \operatorname{lk}(\Lambda)$. We say that Λ is a *join* if the vertices of Λ can be expressed as $V(\Lambda) = V(\Lambda_1) \cup V(\Lambda_2)$, where Λ_1 and Λ_2 are disjoint subgraphs of Γ and every vertex of Λ_1 is connected to every vertex of Λ_2. We denote the join of Λ_1 and Λ_2 by $\Lambda_1 \bowtie \Lambda_2$. In particular, $\operatorname{st}(\Lambda)$ is the join $\Lambda \bowtie \operatorname{lk}(\Lambda)$.

Remark 2.12. The star, link, and join have important algebraic significance. A join subgraph of Γ generates a subgroup of G_Γ which splits as a direct product, while $\langle \mathrm{st}(\Lambda)\rangle$ is the largest subgroup of G_Γ which splits as a direct product with $\langle\Lambda\rangle$ as one of the factors: $\langle \mathrm{st}(\Lambda)\rangle = \langle\Lambda\rangle \times \langle \mathrm{lk}(\Lambda)\rangle$. Moreover, since every element of $\langle\Lambda\rangle$ commutes with every element of $\langle \mathrm{lk}(\Lambda)\rangle$, the reduced syllable form tells us that we can always write an element $g \in \langle \mathrm{st}(\Lambda)\rangle$ in the form $g = \lambda l$, where $\lambda \in \langle\Lambda\rangle$ and $l \in \langle \mathrm{lk}(\Lambda)\rangle$.

Genevois observed that the graph $S(\Gamma)$ is almost a cube complex, with the only non-cubical behaviour arising from the vertex groups. More precisely, he showed the following result.

Proposition 2.13 ([14, Lemmas 8.5 and 8.8]). *Two adjacent edges of $S(\Gamma)$ are edges of a triangle if and only if they are labelled by elements of the same vertex group. Two adjacent edges of $S(\Gamma)$ are edges of an induced square if and only if they are labelled by elements of adjacent vertex groups. In this case, opposite edges of the square are labelled by the same vertex groups.*

The above proposition means that while $S(\Gamma)$ is not a cube complex, it is the 1-skeleton of a complex built from *prisms* glued isometrically along subprisms. Henceforth, we will interchangeably refer to $S(\Gamma)$ and the canonical cell complex of which it is the 1-skeleton.

Definition 2.14 (Prism). A *prism* P of $S(\Gamma)$ is a subcomplex which can be written as a product of simplices $P = T_1 \times \cdots \times T_m$.

Since a cube is a product of 1-simplices, prisms generalise the cubes in a cube complex. Genevois used the prisms in $S(\Gamma)$ to build hyperplanes with very similar properties to those in CAT(0) cube complexes. We present a slightly different, but equivalent, construction of these hyperplanes in $S(\Gamma)$.

In a cube complex, hyperplanes are built from mid-cubes. If we view each cube in a cube complex as a product $[-\frac{1}{2}, \frac{1}{2}]^n$, we obtain a *mid-cube* by restricting one of the intervals $[-\frac{1}{2}, \frac{1}{2}]$ to 0. In much the same way, we obtain a *mid-prism* from a prism by performing a modified barycentric subdivision on one of its simplices. If this simplex is a 1-simplex, this just gives us the midpoint of the edge.

Definition 2.15 (Mid-prism). Given an n-simplex T in $S(\Gamma)$, perform a modified barycentric subdivision as follows. First, add a vertex at the barycentre of each sub-simplex of T. Then, for each $2 \le k \le n$, add edges connecting the barycentre of each k-simplex in T to the barycentres of each of its $(k-1)$-sub-simplices; see Figure 1. The complex we have added through this procedure is then the 1-skeleton of a canonical simply connected cell complex, which we denote by $K(T)$. We call $K(T)$ the *mid-prism* of T. More generally, we define a *mid-prism* K_i of a prism $P = T_1 \times \cdots \times T_m$ to be the product $K_i = T_1 \times \cdots \times T_{i-1} \times K(T_i) \times T_{i+1} \times \cdots \times T_m$.

Note that the simplices in $S(\Gamma)$ that arise from infinite vertex groups have infinitely many vertices. A simplex with infinitely many vertices may still be assigned a mid-prism,

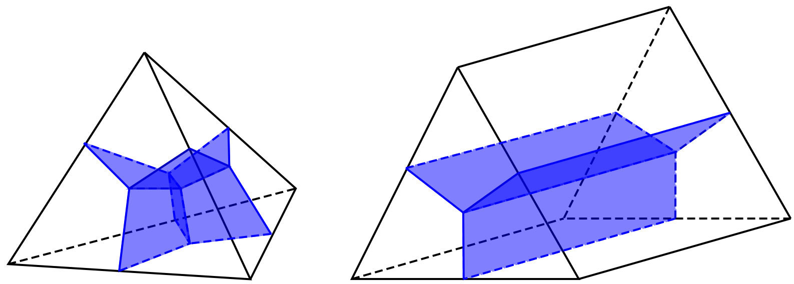

Figure 1. The mid-prism of a 3-simplex and a mid-prism of the product of a 2-simplex and a 1-simplex.

by constructing mid-prisms for each of its finite sub-simplices. The inductivity of the barycentric subdivision procedure ensures that these mid-prisms all agree with each other.

A *hyperplane* of a cube complex is defined to be a maximal connected union of mid-cubes. In the same way, we can construct hyperplanes in $S(\Gamma)$ by taking maximal connected unions of mid-prisms.

Definition 2.16 (Hyperplane, carrier). Construct an equivalence relation $\sim$ on the edges of $S(\Gamma)$ by defining $E_1 \sim E_2$ if E_1 and E_2 are either opposite sides of a square or two sides of a triangle, and then extending transitively. We say that the *hyperplane* dual to the equivalence class $[E]$ is the union of all mid-prisms that intersect edges of $[E]$; see Figure 2. The *carrier* of the hyperplane dual to $[E]$ is the union of all prisms that contain edges of $[E]$.

If a geodesic γ or a coset $g\langle\Lambda\rangle$ contains an edge that is dual to a hyperplane H, then we say that H *crosses* γ or $g\langle\Lambda\rangle$. We say that a hyperplane H *separates* two subsets X and Y of $S(\Gamma)$ if X and Y are each entirely contained in different connected components of $S(\Gamma) \smallsetminus H$.

Each hyperplane of a cube complex comes with two corresponding *combinatorial hyperplanes*, obtained by restricting intervals to $-\frac{1}{2}$ or $\frac{1}{2}$ instead of 0 when constructing mid-cubes. The advantage of these combinatorial hyperplanes is that they form subcomplexes of the cube complex. In $S(\Gamma)$, we obtain combinatorial hyperplanes by restricting a simplex to a vertex instead of performing barycentric subdivision when constructing mid-prisms.

Definition 2.17 (Combinatorial hyperplane). Let $P = T_1 \times \cdots \times T_m$ be a prism, where each T_i is an n_i-simplex. Each mid-prism K_i splits P into $n_i + 1$ sectors, each containing a subcomplex $T_1 \times \cdots \times \{v_k\} \times \cdots \times T_m$, where v_k is a vertex of T_i. Given a hyperplane H of $S(\Gamma)$, consider the union of all such subcomplexes obtained from the mid-prisms of H. We call each connected component of this union a *combinatorial hyperplane* associated to H; see Figure 2.

Remark 2.18 (Labelling hyperplanes). Proposition 2.13 tells us that if two edges E_1 and E_2 of $S(\Gamma)$ are sides of a common triangle or opposite sides of a square, then they are

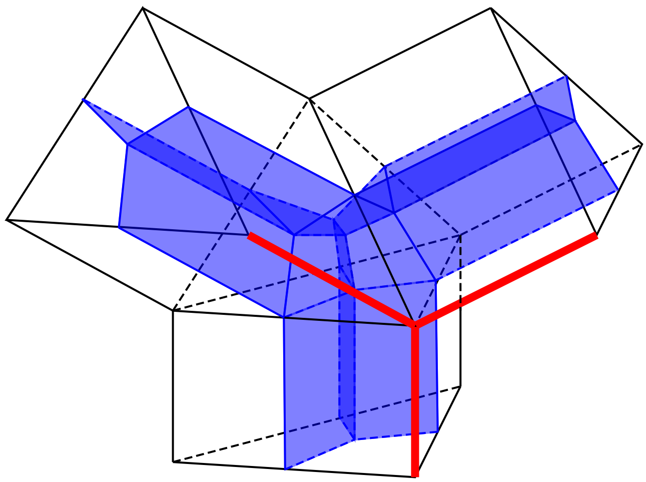

Figure 2. A hyperplane (blue) inside its carrier, and an associated combinatorial hyperplane (red).

labelled by elements of the same vertex group. It follows that all edges that a hyperplane H intersects are labelled by elements of the same vertex group G_v. We therefore label H with the vertex group G_v. Moreover, the edges of the associated combinatorial hyperplanes will then be labelled by elements of $\langle \mathrm{lk}(v) \rangle$. This fact will be exploited repeatedly in our proofs.

Genevois established that the hyperplanes of $S(\Gamma)$ maintain many of the fundamental properties from the cubical setting.

Proposition 2.19 (Properties of hyperplanes; [14, Section 2]). (1) *Every hyperplane of $S(\Gamma)$ separates $S(\Gamma)$ into at least two connected components.*

(2) *If H is a hyperplane of $S(\Gamma)$, then any combinatorial hyperplane for H is convex in $S(\Gamma)$.*

(3) *If H is a hyperplane of $S(\Gamma)$, then any connected component of $S(\Gamma) \smallsetminus H$ is convex in $S(\Gamma)$.*

(4) *A continuous path γ in $S(\Gamma)$ is a geodesic if and only if γ intersects each hyperplane at most once.*

(5) *If two hyperplanes cross, then they are labelled by adjacent vertex groups.*

Remark 2.20. Item (4) implies that a hyperplane H of $S(\Gamma)$ crosses a geodesic connecting a pair of points x, y if and only if H separates x and y. Thus, if $\gamma_1, \dots, \gamma_n$ is a collection of geodesics in $S(\Gamma)$ such that $\gamma_1 \cup \cdots \cup \gamma_n$ forms a loop and H is a hyperplane that crosses γ_i, then H must also cross γ_j for some $j \neq i$.

It is important to note that while we still use the terms “hyperplane” and “combinatorial hyperplane” here, they differ from those of cube complexes in a critical way: the complement of a hyperplane H in $S(\Gamma)$ may have more than two connected components, and thus H may have more than two associated combinatorial hyperplanes.

Genevois and Martin use the convexity of the cosets $g\langle\Lambda\rangle$ to construct a nearest point projection onto $g\langle\Lambda\rangle$, which we call a *gate map*. The map and its properties are given below, and will be essential tools throughout this paper.

Proposition 2.21 (Gate onto graphical subgroups; [17, Section 2]). *Let G_Γ be a graph product. For all $\Lambda \subseteq \Gamma$ and $g \in G_\Gamma$, there exists a map $\mathfrak{g}_{g\Lambda}\colon G_\Gamma \to g\langle\Lambda\rangle$ satisfying the following properties.*

(1) *For all $k, h \in G_\Gamma$, $\mathsf{d}_{\mathrm{syl}}(\mathfrak{g}_{g\Lambda}(h), \mathfrak{g}_{g\Lambda}(k)) \leq \mathsf{d}_{\mathrm{syl}}(h, k)$.*

(2) *For all $x, h \in G_\Gamma$, $h \cdot \mathfrak{g}_{g\Lambda}(x) = \mathfrak{g}_{hg\Lambda}(hx)$. In particular, $\mathfrak{g}_{g\Lambda}(x) = g \cdot \mathfrak{g}_{\Lambda}(g^{-1}x)$.*

(3) *For all $x \in G_\Gamma$, $\mathfrak{g}_{g\Lambda}(x)$ is the unique element of $g\langle\Lambda\rangle$ such that $\mathsf{d}_{\mathrm{syl}}(x, \mathfrak{g}_{g\Lambda}(x)) = \mathsf{d}_{\mathrm{syl}}(x, g\langle\Lambda\rangle)$.*

(4) *Any hyperplane in $S(\Gamma)$ that separates x from $\mathfrak{g}_{g\Lambda}(x)$ separates x from $g\langle\Lambda\rangle$.*

(5) *If $x, y \in G_\Gamma$ and H is a hyperplane in $S(\Gamma)$ separating $\mathfrak{g}_{g\Lambda}(x)$ and $\mathfrak{g}_{g\Lambda}(y)$, then H separates x and y, so that x and $\mathfrak{g}_{g\Lambda}(x)$ (resp. y and $\mathfrak{g}_{g\Lambda}(y)$) are contained in the same connected component of $S(\Gamma) \smallsetminus H$.*

We also obtain a convenient algebraic formulation for the gate map of an element g onto a graphical subgroup $\langle\Lambda\rangle$ by considering the collection of all possible initial subwords of g that are contained in $\langle\Lambda\rangle$.

Definition 2.22 (Prefixes and suffixes). Let $g \in G_\Gamma$. If there exist $p, s \in G_\Gamma$ so that $g = ps$ and $|g|_{\mathrm{syl}} = |p|_{\mathrm{syl}} + |s|_{\mathrm{syl}}$, we call p a *prefix* of g and s a *suffix* of g. We shall use $\mathrm{prefix}(g)$ and $\mathrm{suffix}(g)$ to respectively denote the collections of all prefixes and suffixes of g.

Lemma 2.23 (Algebraic description of the gate map). *For all $\Lambda \subseteq \Gamma$ and $g \in G_\Gamma$, there exists $p \in \mathrm{prefix}(g) \cap \langle\Lambda\rangle$ so that $\mathfrak{g}_\Lambda(g) = p$. Further, p is the element of $\mathrm{prefix}(g) \cap \langle\Lambda\rangle$ with the largest syllable length.*

Proof. Since $\mathrm{prefix}(g) \cap \langle\Lambda\rangle$ is a finite set, there exists $p \in \mathrm{prefix}(g) \cap \langle\Lambda\rangle$ so that $|p'|_{\mathrm{syl}} \leq |p|_{\mathrm{syl}}$ for all $p' \in \mathrm{prefix}(g) \cap \langle\Lambda\rangle$. Let $x = \mathfrak{g}_\Lambda(g)$ and let s be the suffix of g corresponding to p. If there exists a non-identity element $y \in \mathrm{prefix}(s) \cap \langle\Lambda\rangle$, then py would be an element of $\mathrm{prefix}(g) \cap \langle\Lambda\rangle$ with syllable length strictly larger than p. Since this is impossible by choice of p, we have $\mathrm{prefix}(s) \cap \langle\Lambda\rangle = \{e\}$. This implies $|x^{-1}ps|_{\mathrm{syl}} \geq |s|_{\mathrm{syl}}$ since $x^{-1}p \in \langle\Lambda\rangle$, and we have the following calculation:

$$\mathsf{d}_{\mathrm{syl}}(x, g) = |x^{-1}ps|_{\mathrm{syl}} \geq |s|_{\mathrm{syl}} = |p^{-1}g|_{\mathrm{syl}} = \mathsf{d}_{\mathrm{syl}}(p, g).$$

Since $p \in \langle\Lambda\rangle$, this implies $x = p$, as x is the unique element of $\langle\Lambda\rangle$ that minimises the syllable distance of g to $\langle\Lambda\rangle$ (Proposition 2.21 (3)). ■

Definition 2.24. Denote the element p of $\mathrm{prefix}(g) \cap \langle\Lambda\rangle$ with largest syllable length by $\mathrm{prefix}_\Lambda(g)$, and define $\mathrm{suffix}_\Lambda(g) = (\mathrm{prefix}_\Lambda(g^{-1}))^{-1}$.

2.2. Relative hierarchically hyperbolic groups

We break the definition of a relative hierarchically hyperbolic group (HHG) given by Behrstock, Hagen, and Sisto in [3] into three parts in order to more clearly organise the structure of our arguments. First we define what we call the *proto-hierarchy structure*, which sets up the defining information (relations and projections) for the HHG structure. We then give the more advanced geometric properties that we need to impose for the group to be a relative hierarchically hyperbolic space (HHS). We then define a *relative HHG* to be a group whose Cayley graph is a relative HHS in such a way that the relative HHS structure agrees with the group structure.

Definition 2.25 (Proto-hierarchy structure). Let $\mathcal{X}$ be a quasi-geodesic space and $E > 0$. An *E-proto-hierarchy structure* on $\mathcal{X}$ is an index set $\mathfrak{S}$ and a set $\{C(W) : W \in \mathfrak{S}\}$ of geodesic spaces $(C(W), \mathrm{d}_W)$ such that the following axioms are satisfied.

(1) (*Projections*) For each $W \in \mathfrak{S}$, there exists a *projection* $\pi_W\colon \mathcal{X} \to 2^{C(W)}$ such that for all $x \in \mathcal{X}$, $\pi_W(x) \neq \emptyset$ and $\operatorname{diam}(\pi_W(x)) \leq E$. Moreover, each π_W is (E, E)-coarsely Lipschitz and $C(W) \subseteq \mathcal{N}_E(\pi_W(\mathcal{X}))$ for all $W \in \mathfrak{S}$.

(2) (*Nesting*) If $\mathfrak{S} \neq \emptyset$, then $\mathfrak{S}$ is equipped with a partial order $\sqsubseteq$ and contains a unique $\sqsubseteq$-maximal element. When $V \sqsubseteq W$, we say that V is *nested* in W. For each $W \in \mathfrak{S}$, we denote by $\mathfrak{S}_W$ the set of all $V \in \mathfrak{S}$ with $V \sqsubseteq W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \sqsubsetneq W$ there is a specified non-empty subset $\rho^V_W \subseteq C(W)$ with $\operatorname{diam}(\rho^V_W) \leq E$.

(3) (*Orthogonality*) $\mathfrak{S}$ has a symmetric relation called *orthogonality*. If V and W are orthogonal, we write $V \perp W$ and require that V and W are not $\sqsubseteq$-comparable. Further, whenever $V \sqsubseteq W$ and $W \perp U$, we require that $V \perp U$. We denote by $\mathfrak{S}^{\perp}_W$ the set of all $V \in \mathfrak{S}$ with $V \perp W$.

(4) (*Transversality*) If $V, W \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say that V, W are *transverse*, denoted by $V \pitchfork W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \pitchfork W$, there are non-empty sets $\rho^V_W \subseteq C(W)$ and $\rho^W_V \subseteq C(V)$ each of diameter at most E.

We use $\mathfrak{S}$ to denote the entire proto-hierarchy structure, including the index set $\mathfrak{S}$, spaces $\{C(W) : W \in \mathfrak{S}\}$, projections $\{\pi_W : W \in \mathfrak{S}\}$, and relations $\sqsubseteq, \perp, \pitchfork$. We call the elements of $\mathfrak{S}$ the *domains* of $\mathfrak{S}$ and call the set ρ^V_W the *relative projection* from V to W. The number E is called the *hierarchy constant* for $\mathfrak{S}$.

Definition 2.26 (Relative HHS). An E-proto-hierarchy structure $\mathfrak{S}$ on a quasi-geodesic space $\mathcal{X}$ is a *relative E-hierarchically hyperbolic space structure* (relative E-HHS structure) on $\mathcal{X}$ if it satisfies the following additional axioms.

(1) (*Hyperbolicity*) For each $W \in \mathfrak{S}$, either W is $\sqsubseteq$-minimal or $C(W)$ is E-hyperbolic.

(2) (*Finite complexity*) Any set of pairwise $\sqsubseteq$-comparable elements has cardinality at most E.

(3) (*Containers*) For each $W \in \mathfrak{S}$ and $U \in \mathfrak{S}_W$ with $\mathfrak{S}_W \cap \mathfrak{S}_U^{\perp} \neq \emptyset$, there exists $Q \in \mathfrak{S}_W$ such that $V \sqsubseteq Q$ whenever $V \in \mathfrak{S}_W \cap \mathfrak{S}_U^{\perp}$. We call Q a *container of U in W*.

(4) (*Uniqueness*) There exists a function $\theta\colon [0,\infty) \to [0,\infty)$ so that for all $r \geq 0$, if $x, y \in \mathcal{X}$ and $\mathsf{d}_{\mathcal{X}}(x,y) \geq \theta(r)$, then there exists $W \in \mathfrak{S}$ such that

$$\mathsf{d}_W\big(\pi_W(x), \pi_W(y)\big) \geq r.$$

We call θ the *uniqueness function of* $\mathfrak{S}$.

(5) (*Bounded geodesic image*) For all $x, y \in \mathcal{X}$ and $V, W \in \mathfrak{S}$ with $V \sqsubsetneq W$, if $\mathsf{d}_V(\pi_V(x), \pi_V(y)) \geq E$, then every $C(W)$-geodesic from $\pi_W(x)$ to $\pi_W(y)$ must intersect the E-neighbourhood of ρ_W^V.

(6) (*Large links*) For all $W \in \mathfrak{S}$ and $x, y \in \mathcal{X}$, there exists $\{V_1, \dots, V_m\} \subseteq \mathfrak{S}_W \smallsetminus \{W\}$ such that $m \leq E\mathsf{d}_W(\pi_W(x), \pi_W(y)) + E$, and for all $U \in \mathfrak{S}_W \smallsetminus \{W\}$, either $U \in \mathfrak{S}_{V_i}$, for some i, or $\mathsf{d}_U(\pi_U(x), \pi_U(y)) \leq E$.

(7) (*Consistency*) If $V \pitchfork W$, then

$$\min\big\{\mathsf{d}_W\big(\pi_W(x), \rho_W^V\big), \mathsf{d}_V\big(\pi_V(x), \rho_V^W\big)\big\} \leq E$$

for all $x \in \mathcal{X}$. Further, if $U \sqsubseteq V$ and either $V \sqsubsetneq W$ or $V \pitchfork W$ and $W \not\perp U$, then $\mathsf{d}_W(\rho_W^U, \rho_W^V) \leq E$.

(8) (*Partial realisation*) If $\{V_i\}$ is a finite collection of pairwise orthogonal elements of $\mathfrak{S}$ and $p_i \in C(V_i)$ for each i, then there exists $x \in \mathcal{X}$ so that

- $\mathsf{d}_{V_i}(\pi_{V_i}(x), p_i) \leq E$ for all i,
- for each i and each $W \in \mathfrak{S}$, if $V_i \sqsubsetneq W$ or $W \pitchfork V_i$, we have $\mathsf{d}_W(\pi_W(x), \rho_W^{V_i}) \leq E$.

If $C(W)$ is E-hyperbolic for all $W \in \mathfrak{S}$, then $\mathfrak{S}$ is an *E-HHS structure* on $\mathcal{X}$. We call a quasi-geodesic space $\mathcal{X}$ a *(relative) E-HHS* if there exists a (relative) E-HHS structure on $\mathcal{X}$. We use the pair $(\mathcal{X}, \mathfrak{S})$ to denote a (relative) HHS equipped with the specific (relative) HHS structure $\mathfrak{S}$.

Definition 2.27 (Relative HHG). Let G be a finitely generated group and let X be the Cayley graph of G with respect to some finite generating set. We say that G is a (relative) E-HHG if the following hold.

(1) The space X admits a (relative) E-HHS structure $\mathfrak{S}$.

(2) There is a $\sqsubseteq$-, $\perp$-, and $\pitchfork$-preserving action of G on $\mathfrak{S}$ by bijections such that $\mathfrak{S}$ contains finitely many G-orbits.

(3) For each $W \in \mathfrak{S}$ and $g \in G$, there exists an isometry $g_W\colon C(W) \to C(gW)$ satisfying the following for all $V, W \in \mathfrak{S}$ and $g, h \in G$:

- the map $(gh)_W\colon C(W) \to C(ghW)$ is equal to the map $g_{hW} \circ h_W\colon C(W) \to C(ghW)$;
- for each $x \in X$, $g_W(\pi_W(x))$ and $\pi_{gW}(g \cdot x)$ are at most E-far apart in $C(gW)$;
- if $V \pitchfork W$ or $V \sqsubsetneq W$, then $g_W(\rho_W^V)$ and ρ_{gW}^{gV} are at most E-far apart in $C(gW)$.

The structure $\mathfrak{S}$ satisfying (1)–(3) is called a *(relative) E-HHG* structure on G. We use $(G, \mathfrak{S})$ to denote a group G equipped with a specific (relative) HHG structure $\mathfrak{S}$.

We build the proto-hierarchy structure for a graph product of finitely generated groups in Section 3 and spend Section 4 verifying that this structure satisfies the axioms of a relative HHS and respects the group structure.

3. The proto-hierarchy structure on a graph product

For this section, G_Γ will be a graph product of finitely generated groups. For each vertex group G_v, let S_v be a finite generating set for G_v, then define S to be $\bigcup_{v \in V(\Gamma)} S_v$. Throughout this section, d will denote the word metric on G_Γ with respect to S. We now begin to explicitly construct the HHS structure on G_Γ. We first define the index set, associated spaces, and projection maps in Section 3.1 and then define the relations and relative projections in Section 3.2.

3.1. The index set, associated spaces, and projections

The index set for our relative HHS structure on G_Γ is the set of parallelism classes of graphical subgroups. This mirrors the case of right-angled Artin groups studied in [5].

Definition 3.1 (Parallelism and the index set for a graph product). Let G_Γ be a graph product. For a subgraph $\Lambda \subseteq \Gamma$, we shall use $g\Lambda$ to denote the coset $g\langle\Lambda\rangle$ for ease of notation. We say that $g\Lambda$ and $h\Lambda$ are *parallel* if $g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$ and write $g\Lambda \| h\Lambda$. Let $[g\Lambda]$ denote the equivalence class of $g\Lambda$ under the parallelism relation $\|$. Define the index set $\mathfrak{S}_\Gamma = \{[g\Lambda] : g \in G_\Gamma,\ \Lambda \subseteq \Gamma\}$.

The geometric intuition for the definition of parallelism comes from the fact that if two cosets $g\langle\Lambda\rangle$ and $h\langle\Lambda\rangle$ satisfy $g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$, then they are each crossed by precisely the same set of hyperplanes of $S(\Gamma)$. Again, it is important to note that these hyperplanes, introduced in Definition 2.16, are generalisations of those in cube complexes.

Proposition 3.2 (Parallel cosets have the same hyperplanes). *Let $\Lambda \subseteq \Gamma$ and $g, h \in G_\Gamma$. If $g\langle\Lambda\rangle \| h\langle\Lambda\rangle$, then every hyperplane of $S(\Gamma)$ crossing $g\langle\Lambda\rangle$ must also cross $h\langle\Lambda\rangle$.*

Proof. Since $g\langle\Lambda\rangle \| h\langle\Lambda\rangle$, we have $g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$ and there exist $\lambda \in \langle\Lambda\rangle$ and $l \in \langle \mathrm{lk}(\Lambda)\rangle$ such that $g^{-1}h = \lambda l$ (Remark 2.12). Since λ and l commute, $g^{-1}h\langle\Lambda\rangle = l\langle\Lambda\rangle$.

Let H be a hyperplane in $S(\Gamma)$ crossing $g\langle\Lambda\rangle$. In particular, H separates two adjacent points ga and gb in $g\langle\Lambda\rangle$. Translating by g^{-1}, we have that $g^{-1}H$ separates a and b in $\langle\Lambda\rangle$. Let $s_1 \dots s_n$ be a reduced syllable expression for l. Thus, there is a geodesic from a to la and a geodesic from b to lb each labelled by $s_1 \dots s_n$, where each $s_i \in \langle \mathrm{lk}(\Lambda)\rangle$. Since $b^{-1}a$ labels an edge of $\langle\Lambda\rangle$, $b^{-1}a$ and s_i span a square for each $i \in \{1, \dots, n\}$. Thus, we have a strip of squares joining the edge between a and b to the edge between la and lb with the hyperplane $g^{-1}H$ running through the middle. Hence $g^{-1}H$ crosses $l\langle\Lambda\rangle = g^{-1}h\langle\Lambda\rangle$ and by translating by g, H crosses $h\langle\Lambda\rangle$. ∎

The hierarchy structure on a graph product on n vertices can be thought of as being built up in n levels, with level k consisting of the subgraphs with k vertices. Whenever we build up to the next level in the hierarchy, we need to record precisely the geometry we have just added; any less will violate the uniqueness axiom, while any more may violate hyperbolicity. When defining our spaces $C(g\Lambda)$, we therefore do not want to record any distance travelled in strict subgraphs of Λ. This leads us to the *subgraph metric*:

Definition 3.3 (Subgraph metric on a graph product). Let G_Γ be a graph product. Define $C(\Gamma)$ to be the graph whose vertices are elements of G_Γ and where $g, h \in G_\Gamma$ are joined by an edge if there exists a proper subgraph $\Lambda \subsetneq \Gamma$ such that $g^{-1}h \in \langle\Lambda\rangle$, or if $g^{-1}h$ is an element of the generating set S defined at the beginning of the section. We denote the distance in $C(\Gamma)$ by $\mathsf{d}_\Gamma(\cdot,\cdot)$ and say that $\mathsf{d}_\Gamma(g,h)$ is the *subgraph distance* between g and h. When Γ is a single vertex v, $C(\Gamma) = C(v)$ is the Cayley graph of the vertex group G_v with respect to the finite generating set S. Otherwise, $\mathsf{d}_\Gamma(e,g)$ is equal to the smallest n such that $g = \lambda_1 \dots \lambda_n$ with $\operatorname{supp}(\lambda_i)$ a proper subgraph of Γ for each $i \in \{1,\dots,n\}$.

If $g = \lambda_1 \dots \lambda_n$, where $\operatorname{supp}(\lambda_i)$ is a proper subgraph of Γ for each $i \in \{1,\dots,n\}$, then we call $\lambda_1 \dots \lambda_n$ a *subgraph expression* for g. If $n = \mathsf{d}_\Gamma(e,g)$, then $\lambda_1 \dots \lambda_n$ is a *reduced subgraph expression* for g. Note that when Γ is a single vertex, there are no subgraph expressions.

Remark 3.4. When Γ has at least two vertices, $S(\Gamma)$ is obtained from $\operatorname{Cay}(G_\Gamma, S)$ by adding extra edges, where S is the generating set defined at the beginning of the section. Likewise, $C(\Gamma)$ is then obtained from $S(\Gamma)$ by adding even more edges. It therefore follows that $\mathsf{d}_\Gamma \leq \mathsf{d}_{\mathrm{syl}} \leq \mathsf{d}$, where d is the word metric on G_Γ induced by S.

In a reduced subgraph expression $g = \lambda_1 \dots \lambda_n$, we may assume

$$\operatorname{suffix}_{\Lambda_{i+1}}(\lambda_1 \dots \lambda_i) = e$$

for each $i \in \{1,\dots,n-1\}$ by removing any non-trivial suffix from the end of $\lambda_1 \dots \lambda_i$ and attaching it to the beginning of λ_{i+1}. By repeating this procedure for each i in ascending order and then writing reduced syllable expressions for each λ_i, we then obtain a reduced syllable expression for g.

Lemma 3.5. *If Γ contains at least two vertices, then for each $g \in G_\Gamma$, there exist $\lambda_1,\dots,\lambda_n \in G_\Gamma$ with* $\operatorname{supp}(\lambda_i) = \Lambda_i \subsetneq \Gamma$ *such that the following hold.*

(1) *$\lambda_1 \dots \lambda_n$ is a reduced subgraph expression for g.*

(2) *For each $i \in \{1,\dots,n-1\}$,* $\operatorname{suffix}_{\Lambda_{i+1}}(\lambda_1 \dots \lambda_i) = e$.

(3) $|g|_{\mathrm{syl}} = |\lambda_1 \dots \lambda_n|_{\mathrm{syl}} = \sum_{j=1}^{n} |\lambda_j|_{\mathrm{syl}}$.

In particular, for each $x, y \in G_\Gamma$, there exists an $S(\Gamma)$-geodesic γ connecting x and y such that if $\lambda_1 \dots \lambda_n$ is the above reduced subgraph expression for $x^{-1}y$, then the element $x\lambda_1 \dots \lambda_i$ is a vertex of γ for each $i \in \{1,\dots,n\}$.

Proof. We begin by noting how the final conclusion of the lemma follows from the main conclusion. Let $\lambda_1 \dots \lambda_n$ be a reduced subgraph expression for $x^{-1}y$ that satisfies (3).

For each $i \in \{1, \dots, n\}$, let $s_1^i \dots s_{m_i}^i$ be a reduced syllable expression for λ_i. Since $|x^{-1}y|_{\mathrm{syl}} = |\lambda_1 \dots \lambda_n|_{\mathrm{syl}} = \sum_{j=1}^{n} |\lambda_j|_{\mathrm{syl}}$, it follows that $(s_1^1 \dots s_{m_1}^1) \dots (s_1^n \dots s_{m_n}^n)$ is a reduced syllable expression for $x^{-1}y$. Hence there exists an $S(\Gamma)$-geodesic η from e to $x^{-1}y$ whose edges are labelled by $(s_1^1 \dots s_{m_1}^1) \dots (s_1^n \dots s_{m_n}^n)$, and this implies that the element $\lambda_1 \dots \lambda_i$ appears as a vertex of η for each $i \in \{1, \dots, n\}$. Translating by x gives $\gamma = x\eta$ as the desired geodesic.

We now prove that we can find a reduced subgraph expression satisfying (2) and (3) for any element of G_Γ. Our proof proceeds by induction on $n = \mathsf{d}_\Gamma(e, g)$. If $n = 1$, then $\mathrm{supp}(g)$ is a proper subgraph of Γ and the conclusion is trivially true.

Assume that the lemma holds for all $h \in G_\Gamma$ with $\mathsf{d}_\Gamma(e, h) \leq n - 1$ and let $g \in G_\Gamma$ with $\mathsf{d}_\Gamma(e, g) = n$. Let $\omega_1 \dots \omega_n$ be a reduced subgraph expression for g. Let $\Omega_i = \mathrm{supp}(\omega_i)$ for each $i \in \{1, \dots, n\}$. By the induction hypothesis, we can assume that $g_0 = \omega_1 \dots \omega_{n-1}$ satisfies the conclusion of the lemma. Hence $|\omega_1 \dots \omega_{n-1}|_{\mathrm{syl}} = \sum_{j=1}^{n-1} |\omega_j|_{\mathrm{syl}}$ and $\mathrm{suffix}_{\Omega_{i+1}}(\omega_1 \dots \omega_i) = e$ for $i \in \{1, \dots, n-2\}$.

Let $\sigma = \mathrm{suffix}_{\Omega_n}(\omega_1 \dots \omega_{n-1})$. For each $i \in \{1, \dots, n-1\}$, let $s_1^i \dots s_{m_i}^i$ be a reduced syllable expression for ω_i. Now $(s_1^1 \dots s_{m_1}^1) \dots (s_1^{n-1} \dots s_{m_{n-1}}^{n-1})$ is a reduced syllable expression for $\omega_1 \dots \omega_{n-1}$ as $|\omega_1 \dots \omega_{n-1}|_{\mathrm{syl}} = \sum_{j=1}^{n-1} |\omega_j|_{\mathrm{syl}}$. Thus, each syllable of σ is a syllable of one of $\omega_1, \dots, \omega_{n-1}$. For each $i \in \{1, \dots, n-1\}$, let $j_1 < \dots < j_i$ be the elements of $\{1, \dots, m_i\}$ such that $s_{j_1}^i, \dots, s_{j_i}^i$ are the syllables of ω_i that are not syllables of σ. For $i \in \{1, \dots, n-1\}$, let $\omega_i' = s_{j_1}^i \dots s_{j_i}^i$. Thus, we have $\omega_1 \dots \omega_{n-1} = \omega_1' \dots \omega_{n-1}' \sigma$ where $\mathrm{suffix}_{\Omega_n}(\omega_1' \dots \omega_{n-1}') = e$.

Let $\omega_n' = \sigma\omega_n$. Then $\omega_1' \dots \omega_{n-1}' \omega_n'$ is a reduced subgraph expression for g with $\mathrm{supp}(\omega_n') = \Omega_n$ and $\mathrm{suffix}_{\Omega_n}(\omega_1' \dots \omega_{n-1}') = e$. Let $g' = \omega_1' \dots \omega_{n-1}'$. Since $\omega_1' \dots \omega_n'$ is a reduced subgraph expression for g, then $\omega_1' \dots \omega_{n-1}'$ is a reduced subgraph expression for g'. Hence $\mathsf{d}_\Gamma(e, g') = n - 1$ and the induction hypothesis says that there exists a reduced subgraph expression $\lambda_1 \dots \lambda_{n-1}$ for g' such that $\mathrm{suffix}_{\mathrm{supp}(\lambda_{i+1})}(\lambda_1 \dots \lambda_i) = e$ for $i \in \{1, \dots, n-2\}$ and $|\lambda_1 \dots \lambda_{n-1}|_{\mathrm{syl}} = \sum_{j=1}^{n-1} |\lambda_j|_{\mathrm{syl}}$. Further, $\mathrm{suffix}_{\Omega_n}(\lambda_1 \dots \lambda_{n-1}) = e$ as $\lambda_1 \dots \lambda_{n-1} = g' = \omega_1' \dots \omega_{n-1}'$.

Now let $\lambda_n = \omega_n'$ and $\Lambda_i = \mathrm{supp}(\lambda_i)$ for each $i \in \{1, \dots, n\}$. We verify that $\lambda_1, \dots, \lambda_n$ satisfies the conclusion of the lemma for g.

(1) $\lambda_1 \dots \lambda_n$ is a reduced subgraph expression for g as each $\Lambda_i = \mathrm{supp}(\lambda_i)$ is a proper subgraph of Γ and $\mathsf{d}_\Gamma(e, g) = n$.

(2) For each $i \in \{1, \dots, n-1\}$, the above shows that $\mathrm{suffix}_{\Lambda_{i+1}}(\lambda_1 \dots \lambda_i) = e$.

(3) We prove that writing each λ_i in a reduced syllable form produces a reduced syllable form for the product $\lambda_1 \dots \lambda_n$. For each $i \in \{1, \dots, n\}$, let $t_1^i \dots t_{k_i}^i$ be a reduced syllable expression for λ_i. Since $|\lambda_1 \dots \lambda_{n-1}|_{\mathrm{syl}} = \sum_{j=1}^{n-1} |\lambda_j|_{\mathrm{syl}}$, we know that $(t_1^1 \dots t_{k_1}^1) \dots (t_1^{n-1} \dots t_{k_{n-1}}^{n-1})$ is a reduced syllable expression for $\lambda_1 \dots \lambda_{n-1}$. Therefore, if

$$(t_1^1 \dots t_{k_1}^1) \dots (t_1^n \dots t_{k_n}^n)$$

is not a reduced syllable expression for $\lambda_1 \dots \lambda_n$, then Theorem 2.4 implies that there must exist syllables t_j^i of $\lambda_1 \dots \lambda_{n-1}$ and t_ℓ^n of λ_n such that

$$\operatorname{supp}(t_j^i) = \operatorname{supp}(t_\ell^n)$$

and t_j^i can be moved to be adjacent to t_ℓ^n using a number of commutation relations. However, this implies that t_j^i is a suffix for $\lambda_1 \dots \lambda_{n-1}$ with support in Λ_n. This is impossible as $\operatorname{suffix}_{\Lambda_n}(\lambda_1 \dots \lambda_{n-1}) = e$. Therefore, $(t_1^1 \dots t_{k_1}^1) \dots (t_1^n \dots t_{k_n}^n)$ must be a reduced syllable expression for $\lambda_1 \dots \lambda_n$ and hence

$$|\lambda_1 \dots \lambda_n|_{\mathrm{syl}} = |\lambda_1|_{\mathrm{syl}} + \cdots + |\lambda_n|_{\mathrm{syl}}$$

as desired. ∎

We can now define the geodesic spaces associated to elements of the index set. In the next section, we will show that they are hyperbolic.

Definition 3.6. Let G_Γ be a graph product. For each $g \in G_\Gamma$ and $\Lambda \subseteq \Gamma$, let $C(g\Lambda)$ denote the graph whose vertices are elements of the coset $g\langle\Lambda\rangle$ and where gx and gy are joined by an edge if x and y are joined by an edge in $C(\Lambda)$. The metric on $C(g\Lambda)$ is denoted by $\mathsf{d}_{g\Lambda}(\cdot,\cdot)$.

Remark 3.7. If $\Lambda \subseteq \Gamma$ is a join $\Lambda = \Lambda_1 \bowtie \Lambda_2$, then every element $\lambda \in \langle\Lambda\rangle$ can be written as $\lambda = \lambda_1\lambda_2$, where $\lambda_1 \in \langle\Lambda_1\rangle$ and $\lambda_2 \in \langle\Lambda_2\rangle$. Since Λ_1 and Λ_2 are proper subgraphs of Λ, this implies that $C(\Lambda)$, and therefore $C(g\Lambda)$, has diameter at most 2 whenever Λ splits as a join.

We now wish to use our gate map from Proposition 2.21 to define projections for our hierarchy structure. Since $\mathfrak{S}_\Gamma$ is the set of parallelism classes of cosets of graphical subgroups, we must verify that the gate map is well behaved under parallelism.

Lemma 3.8 (Gates to parallelism classes are well defined). *If $g\Lambda \| h\Lambda$, then for all $x \in G_\Gamma$, $\mathfrak{g}_{h\Lambda}(x) = \mathfrak{g}_{h\Lambda} \circ \mathfrak{g}_{g\Lambda}(x)$. In particular, if $g\Lambda \| h\Lambda$, then $\mathfrak{g}_{h\Lambda}|_{g\langle\Lambda\rangle}\colon g\langle\Lambda\rangle \to h\langle\Lambda\rangle$ agrees with the isometry of $S(\Gamma)$ induced by the element hpg^{-1}, where $p = \operatorname{prefix}_\Lambda(h^{-1}g)$.*

Proof. Suppose that $\mathfrak{g}_{h\Lambda}(x) \neq \mathfrak{g}_{h\Lambda}(\mathfrak{g}_{g\Lambda}(x))$. There must then exist a hyperplane H separating $\mathfrak{g}_{h\Lambda}(x)$ and $\mathfrak{g}_{h\Lambda}(\mathfrak{g}_{g\Lambda}(x))$ in $S(\Gamma)$. By (4) and (5) of Proposition 2.21, H separates x and $\mathfrak{g}_{g\Lambda}(x)$ and thus cannot cross $g\langle\Lambda\rangle$. However, H crosses $h\langle\Lambda\rangle$, and so must cross $g\langle\Lambda\rangle$ by Proposition 3.2. As this is a contradiction, we must have that $\mathfrak{g}_{h\Lambda}(x) = \mathfrak{g}_{h\Lambda}(\mathfrak{g}_{g\Lambda}(x))$.

Note that, if $g\lambda \in g\langle\Lambda\rangle$, then the equivariance (Proposition 2.21 (2)) plus the prefix description of the gate map (Lemma 2.23) imply that

$$\mathfrak{g}_{h\Lambda}(g\lambda) = h \cdot \mathfrak{g}_\Lambda(h^{-1}g\lambda) = h \cdot \operatorname{prefix}_\Lambda(h^{-1}g\lambda).$$

Since $h^{-1}g \in \langle \operatorname{st}(\Lambda)\rangle$, we can write $h^{-1}g = pl$, where $p \in \langle\Lambda\rangle$ and $l \in \langle \operatorname{lk}(\Lambda)\rangle$. Therefore, $\mathfrak{g}_{h\Lambda}(g\lambda) = h \cdot \operatorname{prefix}_\Lambda(pl\lambda) = hp\lambda$, that is, $\mathfrak{g}_{h\Lambda}|_{g\langle\Lambda\rangle}$ agrees with the isometry induced by hpg^{-1}. ∎

Since $\mathrm{Cay}(G_\Gamma, S)$, $S(\Gamma)$, and $C(\Gamma)$ differ only in that the latter two have extra edges, we can easily promote our gate map to a projection map.

Definition 3.9. For all $\Lambda \subseteq \Gamma$ and $g \in G_\Gamma$, define $\pi_{g\Lambda}\colon G_\Gamma \to C(g\Lambda)$ by $i_{g\Lambda} \circ \mathfrak{g}_{g\Lambda}$, where $i_{g\Lambda}$ is the inclusion map from $g\langle\Lambda\rangle$ into $C(g\Lambda)$.

Remark 3.10. Combining the prefix description of the gate map (Lemma 2.23) with equivariance (Proposition 2.21 (2)), we have that $\mathfrak{g}_{g\Lambda}(x) = g \cdot \mathrm{prefix}_\Lambda(g^{-1}x)$ for all $x \in G_\Gamma$. Since the only difference between $\pi_{g\Lambda}$ and $\mathfrak{g}_{g\Lambda}$ is the metric on the image, this means that $\pi_{g\Lambda}(x) = g \cdot \mathrm{prefix}_\Lambda(g^{-1}x)$ as well.

Note that any coset of $\langle\Lambda\rangle$ can be expressed in the form $g\langle\Lambda\rangle$, where $\mathrm{suffix}_\Lambda(g) = e$ (and thus $\mathrm{prefix}_\Lambda(g^{-1}) = e$). Indeed, let $h\langle\Lambda\rangle$ be a coset of $\langle\Lambda\rangle$, and suppose that $\mathrm{suffix}_\Lambda(h) = \lambda$. Then we can write $h = g\lambda$, where $\mathrm{suffix}_\Lambda(g) = e$. It therefore follows that $h\langle\Lambda\rangle = g\lambda\langle\Lambda\rangle = g\langle\Lambda\rangle$. The next proposition shows that choosing the representative of $g\langle\Lambda\rangle$ in this way ensures that $\mathrm{prefix}_\Lambda(g^{-1}x)$ contains only syllables of x. This is particularly helpful when considering the prefix description of $\pi_{g\Lambda}(x)$.

Proposition 3.11. *Let $\Lambda \subseteq \Gamma$ and let $g \in G_\Gamma$. Then, for all $x, y \in G_\Gamma$, every syllable of $(\mathfrak{g}_{g\Lambda}(x))^{-1} \cdot \mathfrak{g}_{g\Lambda}(y)$ is a syllable of $x^{-1}y$. In particular, if g is the representative of $g\langle\Lambda\rangle$ with* $\mathrm{suffix}_\Lambda(g) = e$ *and $h \in G_\Gamma$, then every syllable of* $\mathrm{prefix}_\Lambda(g^{-1}h) = \mathfrak{g}_\Lambda(g^{-1}h)$ *is a syllable of h.*

Proof. Let $x, y \in G_\Gamma$, then let $p_x = \mathfrak{g}_{g\Lambda}(x)$ and $p_y = \mathfrak{g}_{g\Lambda}(y)$. Let η be an $S(\Gamma)$-geodesic connecting p_x and p_y and let γ be an $S(\Gamma)$-geodesic connecting x and y. Let $s_1, \dots, s_n$ be the elements of the vertex groups of G_Γ that label the edges of η. This means that $s_1, \dots, s_n$ are the syllables of $p_x^{-1}p_y$. For each $i \in \{1, \dots, n\}$, let H_i be the hyperplane dual to the edge of η that is labelled by s_i and let v_i be the vertex of Γ such that $s_i \in G_{v_i}$.

Since each H_i separates $\mathfrak{g}_{g\Lambda}(x)$ and $\mathfrak{g}_{g\Lambda}(y)$, each H_i must also cross γ by Proposition 2.19 (4) and Proposition 2.21 (5). For $i \in \{1, \dots, n\}$, let E_i be the edge of γ dual to H_i. Note that every edge dual to H_i is labelled by an element of the vertex group G_{v_i}, but not necessarily by the same element of G_{v_i}.

If E_i is not labelled by $s_i \in G_{v_i}$, then the hyperplane H_i must encounter a triangle of $S(\Gamma)$ between η and γ. This creates a branch of the hyperplane H_i that cannot cross either η or γ by Proposition 2.19 (4). Thus, this branch must cross either an $S(\Gamma)$-geodesic connecting x and p_x or an $S(\Gamma)$-geodesic connecting y and p_y; see Figure 3. Without loss of generality, assume that H_i crosses an $S(\Gamma)$-geodesic connecting x and $p_x = \mathfrak{g}_{g\Lambda}(x)$. This means that H_i separates x from $\mathfrak{g}_{g\Lambda}(x)$, and thus H_i must separate x from all of $g\langle\Lambda\rangle$ (Proposition 2.19 (4)). However, this is impossible as H_i crosses $g\langle\Lambda\rangle$. Therefore, H_i cannot encounter a triangle between η and γ, and E_i must therefore be labelled by the element s_i. Since the elements labelling the edges of γ are the syllables of $x^{-1}y$, this implies that every syllable of $p_x^{-1}p_y$ is also a syllable of $x^{-1}y$.

For the final clause of the proposition, note that $\mathrm{suffix}_\Lambda(g) = e$ implies that $\mathfrak{g}_\Lambda(g^{-1}) = \mathrm{prefix}_\Lambda(g^{-1}) = e$. Thus, we can apply the above with $x = g^{-1}$ and $y = g^{-1}h$ to con-

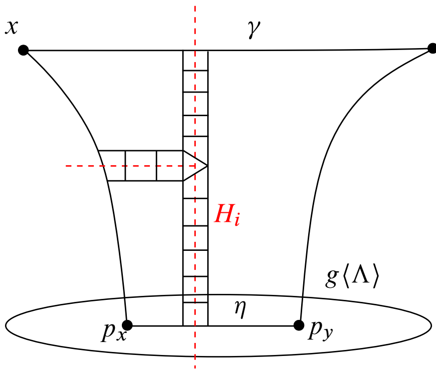

Figure 3. If the the hyperplane H_i encounters a triangle of $S(\Gamma)$ between η and γ, then a branch of H_i must cross an $S(\Gamma)$-geodesic from x to p_x (shown) or from y to p_y.

clude that every syllable of $(\mathfrak{g}_\Lambda(g^{-1}))^{-1}\mathfrak{g}_\Lambda(g^{-1}h) = \mathfrak{g}_\Lambda(g^{-1}h)$ is also a syllable of $(g^{-1})^{-1}g^{-1}h = h$. ∎

Given that $h, k \in G_\Gamma$, we shall employ a common abuse of notation by using $\mathsf{d}_{g\Lambda}(h,k)$ to denote $\mathsf{d}_{g\Lambda}(\pi_{g\Lambda}(h), \pi_{g\Lambda}(k))$. We can now prove our first HHS axiom.

Lemma 3.12 (Projections). *For each $g \in G_\Gamma$ and $\Lambda \subseteq \Gamma$, the projection $\pi_{g\Lambda}$ is $(1,0)$-coarsely Lipschitz.*

Proof. We want to show that $\mathsf{d}_{g\Lambda}(x,y) \le \mathsf{d}(x,y)$ for all $x, y \in G_\Gamma$. First assume that Λ consists of a single vertex v. Let p_x and p_y be $\mathfrak{g}_{g\Lambda}(x) = \pi_{g\Lambda}(x)$ and $\mathfrak{g}_{g\Lambda}(y) = \pi_{g\Lambda}(y)$, respectively. Since Λ is the single vertex v, $C(\Lambda)$ is the Cayley graph of G_v with respect to our fixed finite generating set, and $C(g\Lambda)$ is a coset of $C(\Lambda)$. Thus, it suffices to prove that $|p_x^{-1}p_y|$ is bounded above by $|x^{-1}y|$, where $|\cdot|$ is the word length on G_Γ with respect to the generating set S defined at the beginning of the section.

Let $s = p_x^{-1}p_y \in G_v$. By Proposition 3.11, s must be a syllable of $x^{-1}y$, that is, s appears in a reduced syllable expression for $x^{-1}y$. Recall that if $s_1 \dots s_n$ is a reduced syllable expression for $x^{-1}y$, then $|x^{-1}y| = \sum_{i=1}^{n} |s_i|$ (Corollary 2.5). Thus, $|x^{-1}y| \ge |s| = |p_x^{-1}p_y|$.

Now assume that Λ contains at least two vertices. By Proposition 2.21 (1), we have

$$\mathsf{d}_{\mathrm{syl}}(\mathfrak{g}_{g\Lambda}(x), \mathfrak{g}_{g\Lambda}(y)) \le \mathsf{d}_{\mathrm{syl}}(x,y) \le \mathsf{d}(x,y).$$

Furthermore, $C(g\Lambda)$ is obtained from $S(g\Lambda)$ by adding edges as Λ contains at least two vertices. Thus, we have

$$\mathsf{d}_{g\Lambda}(x,y) \le \mathsf{d}_{\mathrm{syl}}\big(\mathfrak{g}_{g\Lambda}(x), \mathfrak{g}_{g\Lambda}(y)\big) \le \mathsf{d}_{\mathrm{syl}}(x,y) \le \mathsf{d}(x,y).$$ ∎

Given an $S(\Gamma)$-geodesic γ, there is a natural order on its vertices which arises from orienting γ. The distances between the vertices of γ under the projection $\pi_{g\Lambda}$ then satisfy the following monotonicity property with respect to this order.

Lemma 3.13 (Subgraph distance along $S(\Gamma)$-geodesics). *Let γ be an $S(\Gamma)$-geodesic connecting two elements $x, y \in G_\Gamma$. For each vertex q of γ, each element $g \in G_\Gamma$, and each subgraph $\Lambda \subseteq \Gamma$, we have*

$$\mathsf{d}_{g\Lambda}(x, q) \leq \mathsf{d}_{g\Lambda}(x, y) \quad \textit{and} \quad \mathsf{d}_{g\Lambda}(q, y) \leq \mathsf{d}_{g\Lambda}(x, y).$$

Proof. Fix $g \in G_\Gamma$ and a subgraph $\Lambda \subseteq \Gamma$. Let $p_x = \mathfrak{g}_{g\Lambda}(x)$, $p_y = \mathfrak{g}_{g\Lambda}(y)$, and $p_q = \mathfrak{g}_{g\Lambda}(q)$.

First suppose that Λ consists of a single vertex of Γ. Then the $S(\Gamma)$-diameter of $g\langle\Lambda\rangle$ is 1 and there exists a single hyperplane H so that every edge of $g\langle\Lambda\rangle$ is dual to H. If $p_q \neq p_x$ and $p_q \neq p_y$, then H must separate p_q from both p_x and p_y. Therefore, H must cross γ between x and q and again between q and y by Proposition 2.21 (5). However, this is impossible as H cannot cross γ twice (Proposition 2.19 (4)). Thus, we must have either $p_q = p_x$ or $p_q = p_y$. The conclusion of the lemma then automatically holds as $\pi_{g\Lambda}(q) = \pi_{g\Lambda}(x)$ or $\pi_{g\Lambda}(q) = \pi_{g\Lambda}(y)$.

Now assume that Λ has at least two vertices and $p_q \neq p_x$ and $p_q \neq p_y$. Let $\lambda_1 \dots \lambda_m$ be a reduced subgraph expression for $p_x^{-1} p_y$ of the form provided by Lemma 3.5. In particular, there exists an $S(\Gamma)$-geodesic η connecting p_x and p_y whose vertices include $p_x \lambda_1 \dots \lambda_i$ for each $i \in \{1, \dots, m\}$.

Let α and β be $S(\Gamma)$-geodesics connecting p_x to p_q and p_q to p_y, respectively. Any hyperplane that crosses α must also cross γ and separate x and q by Proposition 2.21 (5). Similarly, any hyperplane that crosses β must also cross γ and separate y and q. Thus, a hyperplane that crosses both α and β would cross the $S(\Gamma)$-geodesic γ twice. Since no hyperplane of $S(\Gamma)$ can cross the same geodesic twice (Proposition 2.19 (4)), it follows that any hyperplane that crosses α (resp. β) cannot cross β (resp. α). By Remark 2.20, any hyperplane that crosses either α or β must therefore cross η as $\alpha \cup \beta \cup \eta$ forms a loop in $S(\Gamma)$.

We now prove that $\mathsf{d}_{g\Lambda}(x, q) \leq \mathsf{d}_{g\Lambda}(x, y)$. The proof for $\mathsf{d}_{g\Lambda}(q, y) \leq \mathsf{d}_{g\Lambda}(x, y)$ is nearly identical with β replacing α. Let $E_1, \dots, E_k$ be the edges of α and let H_j be the hyperplane that crosses E_j for $j \in \{1, \dots, k\}$. We say that two hyperplanes H_j and H_ℓ *cross between α and η* if there exists a vertex a of α such that for each vertex b of η, either H_j or H_ℓ separates a from b; see Figure 4.

Claim 3.14. There exists an $S(\Gamma)$-geodesic α' that connects p_x and p_q such that no two of $H_1, \dots, H_k$ cross between α' and η.

Proof. Let $\alpha_1 = \alpha$ and let K_i be the number of times two of $H_1, \dots, H_k$ cross between α_i and η. Note that $K_1 \leq \frac{k(k-1)}{2}$. If $K_1 = 0$, we are done. Otherwise, there exists $j \in \{1, \dots, k\}$ such that H_j is the first hyperplane, where H_{j-1} and H_j cross between α_1 and η. Since H_{j-1} and H_j cross, Proposition 2.19 (5) tells us that the edges E_{j-1} and E_j are labelled by elements of adjacent vertex groups. By Proposition 2.13, E_{j-1} and E_j are two sides of a square S of $S(\Gamma)$ inside which H_{j-1} and H_j cross. Let α_2 be the $S(\Gamma)$-geodesic obtained from α_1 by replacing the edges E_{j-1} and E_j with the other two sides of the square S; see Figure 5.

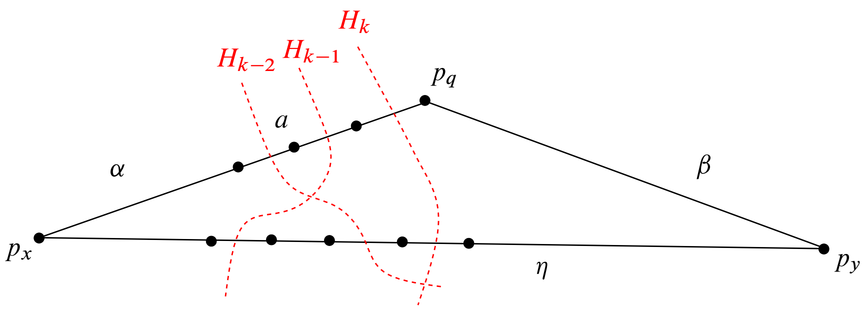

Figure 4. The hyperplanes H_{k-2} and H_{k-1} cross between α and η because the vertex a is separated from every vertex of η by either H_{k-2} or H_{k-1}. Even though H_{k-2} and H_k cross, they do not cross *between* α and η.

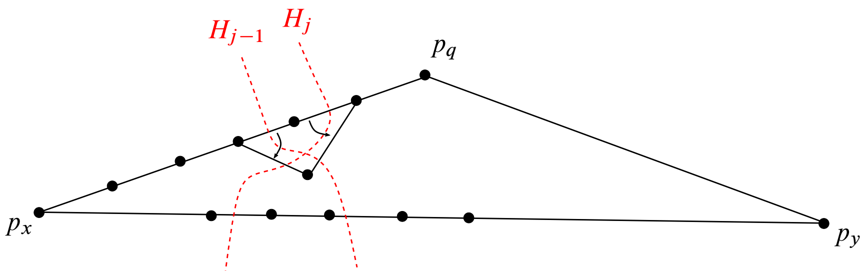

Figure 5. The edges E_{j-1} and E_j can be replaced with the other two edges of the square S to obtain a new $S(\Gamma)$-geodesic with $K_2 = K_1 - 1$.

Since H_{j-1} and H_j crossed between α_1 and η, we now have $K_2 = K_1 - 1$; that is, that the number of times two of $H_1, \dots, H_k$ cross between α_2 and η is one less than the number of times two of $H_1, \dots, H_k$ crossed between α_1 and η. Reindex $H_1, \dots, H_k$ such that H_j crosses the jth edge of α_2.

If $K_2 = 0$, we are done, with $\alpha' = \alpha_2$. Otherwise, we can repeat this argument at most $\frac{k(k-1)}{2}$ times to construct a sequence of geodesics $\alpha_1, \alpha_2, \dots, \alpha_r$, where $K_{i+1} = K_i - 1$ and $K_r = 0$. Then $\alpha' = \alpha_r$. ■

Let α' be as in Claim 3.14 and reindex $H_1, \dots, H_k$ so that H_j crosses the jth edge of α' for each $j \in \{1, \dots, k\}$. Since H_j crosses η for each $j \in \{1, \dots, k\}$, the labels for the edges of α' are a subset of the labels of η. Further, since no two of $H_1, \dots, H_k$ cross between α' and η, the order in which the labels of edges appear along α' is the same as the order in which they appear along η. Since the vertices of η include $p_x\lambda_1 \dots \lambda_i$ for each $i \in \{1, \dots, m\}$, this implies that we can write $p_x^{-1} p_q = \lambda'_1 \dots \lambda'_m$, where $\operatorname{supp}(\lambda'_i) \subseteq \operatorname{supp}(\lambda_i)$ for each $i \in \{1, \dots, m\}$. It therefore follows that the $C(g\Lambda)$-distance between p_x and p_q is bounded above by the $C(g\Lambda)$-distance between p_x and p_y, and so we have $\mathsf{d}_{g\Lambda}(x, q) \leq \mathsf{d}_{g\Lambda}(x, y)$. ■

3.2. The relations

Here we define the nesting, orthogonality, and transversality relations in the proto-hierarchy structure, and prove that they have the desired properties. We tackle the nesting relation first.

Definition 3.15 (Nesting). Let G_Γ be a graph product and let $\mathfrak{S}_\Gamma$ be the index set of parallelism classes of cosets of graphical subgroups described in Definition 3.1. We say that $[g\Lambda] \sqsubseteq [h\Omega]$ if $\Lambda \subseteq \Omega$ and there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\Omega] = [h\Omega]$.

Lemma 3.16. *The relation $\sqsubseteq$ is a partial order.*

Proof. The only property that requires checking is transitivity; that is, if $[g_1\Lambda_1] \sqsubseteq [g_2\Lambda_2] \sqsubseteq [g_3\Lambda_3]$, then $[g_1\Lambda_1] \sqsubseteq [g_3\Lambda_3]$.

Since $\subseteq$ is transitive, we have $\Lambda_1 \subseteq \Lambda_3$. Furthermore, there exist $a, b \in G_\Gamma$ such that $[g_1\Lambda_1] = [a\Lambda_1]$, $[a\Lambda_2] = [g_2\Lambda_2] = [b\Lambda_2]$, $[g_3\Lambda_3] = [b\Lambda_3]$; that is, $g_1^{-1}a \in \langle \mathrm{st}(\Lambda_1)\rangle$, $g_2^{-1}a, g_2^{-1}b \in \langle \mathrm{st}(\Lambda_2)\rangle$, and $g_3^{-1}b \in \langle \mathrm{st}(\Lambda_3)\rangle$. Thus,

$$g_1^{-1}a = l_1\lambda_1, \quad g_2^{-1}a = l_2\lambda_2, \quad g_2^{-1}b = l_2'\lambda_2', \quad g_3^{-1}b = l_3\lambda_3,$$

where $\lambda_i, \lambda_i' \in \langle \Lambda_i \rangle$ and $l_i, l_i' \in \langle \mathrm{lk}(\Lambda_i)\rangle$ for each i. Let $c = b(\lambda_2')^{-1}\lambda_2$. Then $g_3^{-1}c = g_3^{-1}b(\lambda_2')^{-1}\lambda_2 \in \langle \mathrm{st}(\Lambda_3)\rangle$ since $\Lambda_2 \subseteq \Lambda_3$. Moreover, since $\mathrm{lk}(\Lambda_2) \subseteq \mathrm{lk}(\Lambda_1)$,

$$g_1^{-1}c = g_1^{-1}aa^{-1}g_2g_2^{-1}bb^{-1}c = l_1\lambda_1\lambda_2^{-1}l_2^{-1}l_2'\lambda_2'(\lambda_2')^{-1}\lambda_2 = l_1l_2^{-1}l_2'\lambda_1 \in \big\langle \mathrm{st}(\Lambda_1)\big\rangle.$$

Thus, $[g_1\Lambda_1] = [c\Lambda_1]$ and $[g_3\Lambda_3] = [c\Lambda_3]$, verifying that $[g_1\Lambda_1] \sqsubseteq [g_3\Lambda_3]$. ∎

Definition 3.17 (Upwards relative projection). If $[g\Lambda] \sqsubsetneq [h\Omega]$, for any choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$, define $\rho^{g\Lambda}_{h\Omega} \subseteq C(h\Omega)$ to be

$$\rho^{g\Lambda}_{h\Omega} = \bigcup_{k\Lambda \| g\Lambda} \pi_{h\Omega}\big(k\langle\Lambda\rangle\big) = \pi_{h\Omega}\big(g\big\langle \mathrm{st}(\Lambda)\big\rangle\big).$$

The equality between $\bigcup_{k\Lambda\|h\Lambda} \pi_{h\Omega}(k\langle\Lambda\rangle)$ and $\pi_{h\Omega}(g\langle\mathrm{st}(\Lambda)\rangle)$ is a consequence of the definition that $k\Lambda \| g\Lambda$ if and only if $g^{-1}k \in \langle\mathrm{st}(\Lambda)\rangle$. Indeed, $g\langle\mathrm{st}(\Lambda)\rangle = gg^{-1}k\langle\mathrm{st}(\Lambda)\rangle = k\langle\mathrm{st}(\Lambda)\rangle \supseteq k\langle\Lambda\rangle$ for all $k\Lambda \| g\Lambda$. Conversely, each element of $g\langle\mathrm{st}(\Lambda)\rangle$ can be written as $gl\lambda$ where $l \in \langle\mathrm{lk}(\Lambda)\rangle$ and $\lambda \in \langle\Lambda\rangle$, so that $gl\lambda \in gl\langle\Lambda\rangle$ where $g^{-1}gl = l \in \langle\mathrm{st}(\Lambda)\rangle$ and hence $g\Lambda \| gl\Lambda$.

Lemma 3.18 (Upwards relative projections have bounded diameter). *If $[g\Lambda] \sqsubsetneq [h\Omega]$, then for any choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$, we have* $\mathrm{diam}(\rho^{g\Lambda}_{h\Omega}) \leq 2$.

Proof. Let $g\Lambda$ and $h\Omega$ be fixed representatives of $[g\Lambda]$ and $[h\Omega]$, respectively. Suppose first that Ω splits as a join. Then $\mathrm{diam}(C(h\Omega)) = 2$ by Remark 3.7, and hence $\mathrm{diam}(\rho^{g\Lambda}_{h\Omega}) \leq 2$. For the remainder of the proof, we will therefore assume that Ω does not split as a join. Note that this implies that $\mathrm{st}(\Lambda) \cap \Omega \subsetneq \Omega$. Indeed, suppose that

$\mathrm{st}(\Lambda) \cap \Omega = \Omega$. Then $\Omega \subseteq \mathrm{st}(\Lambda)$, so either $\Omega \subseteq \Lambda$, $\Omega \subseteq \mathrm{lk}(\Lambda)$, or Ω splits as a join. The first two cases are impossible as $\Lambda \subsetneq \Omega$, and the last case is ruled out by assumption.

Let $a \in G_\Gamma$ be such that $[a\Lambda] = [g\Lambda]$ and $[a\Omega] = [h\Omega]$. Since $[a\Lambda] = [g\Lambda]$, we have $g^{-1}a \in \langle \mathrm{st}(\Lambda)\rangle$, so $g\langle \mathrm{st}(\Lambda)\rangle = gg^{-1}a\langle \mathrm{st}(\Lambda)\rangle = a\langle \mathrm{st}(\Lambda)\rangle$. Thus, $\rho^{g\Lambda}_{h\Omega} = \pi_{h\Omega}(g\langle \mathrm{st}(\Lambda)\rangle) = \pi_{h\Omega}(a\langle \mathrm{st}(\Lambda)\rangle)$. Note that any element of $a\langle \mathrm{st}(\Lambda)\rangle$ can be expressed in the form $a\lambda l$, where $\lambda \in \langle \Lambda \rangle$ and $l \in \langle \mathrm{lk}(\Lambda)\rangle$. Using the equivariance (Proposition 2.21 (2)) and the prefix description of the gate map (Lemma 2.23), we have

$$\mathfrak{g}_{a\Omega}(a\lambda l) = a \cdot \mathfrak{g}_\Omega(a^{-1}a\lambda l) = a \cdot \mathrm{prefix}_\Omega(\lambda l) = a\lambda \cdot \mathrm{prefix}_\Omega(l).$$

This implies that $\mathfrak{g}_{a\Omega}(a\lambda l) = a\lambda l_0$, where $l_0 = \mathrm{prefix}_\Omega(l) \in \langle \mathrm{lk}(\Lambda) \cap \Omega\rangle$ and so

$$\mathrm{supp}(\lambda l_0) \subseteq \Lambda \cup \big(\mathrm{lk}(\Lambda) \cap \Omega\big) = \mathrm{st}(\Lambda) \cap \Omega \subsetneq \Omega.$$

Moreover, by Lemma 3.8, $\mathfrak{g}_{h\Omega}(a\lambda l) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{a\Omega}(a\lambda l)) = \mathfrak{g}_{h\Omega}(a\lambda l_0)$.

Since $a\Omega \| h\Omega$, the gate map from $a\langle\Omega\rangle$ to $h\langle\Omega\rangle$ agrees with the isometry of $S(\Gamma)$ induced by the element hpa^{-1}, where $p = \mathrm{prefix}_\Omega(h^{-1}a)$ (Lemma 3.8). Since $\mathrm{supp}(\lambda l_0) \subsetneq \Omega$, this implies that

$$\mathfrak{g}_{h\Omega}(a\lambda l_0) = hpa^{-1} \cdot a\lambda l_0 = hp\lambda l_0.$$

Therefore, given two arbitrary elements $a\lambda l, a\lambda' l' \in a\langle \mathrm{st}(\Lambda)\rangle$, we have

$$\big(\mathfrak{g}_{h\Omega}(a\lambda l)\big)^{-1} \mathfrak{g}_{h\Omega}(a\lambda' l') = l_0^{-1}\lambda^{-1}\lambda' l_0',$$

where

$$\mathrm{supp}(l_0^{-1}\lambda^{-1}\lambda' l_0') \subseteq \mathrm{st}(\Lambda) \cap \Omega \subsetneq \Omega.$$

This implies that the $C(h\Omega)$-diameter of $\pi_{h\Omega}(g\langle \mathrm{st}(\Lambda)\rangle) = \rho^{g\Lambda}_{h\Omega}$ is at most 1 in this case. ∎

Next we deal with the orthogonality relation.

Definition 3.19 (Orthogonality). Let G_Γ be a graph product and let $\mathfrak{S}_\Gamma$ be the index set of parallelism classes of cosets of graphical subgroups described in Definition 3.1. We say that $[g\Lambda] \perp [h\Omega]$ if $\Lambda \subseteq \mathrm{lk}(\Omega)$ and there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\Omega] = [h\Omega]$.

Lemma 3.20 (Orthogonality axiom). *The relation $\perp$ has the following properties:*

(1) *$\perp$ is symmetric;*

(2) *if $[g\Lambda] \perp [h\Omega]$, then $[g\Lambda]$ and $[h\Omega]$ are not $\sqsubseteq$-comparable;*

(3) *if $[g\Lambda] \sqsubseteq [h\Omega]$ and $[h\Omega] \perp [k\Pi]$, then $[g\Lambda] \perp [k\Pi]$.*

Proof. (1) If $\Lambda \subseteq \mathrm{lk}(\Omega)$, then all vertices of Λ are connected to all vertices of Ω, hence $\Omega \subseteq \mathrm{lk}(\Lambda)$ too. Thus, the relation $\perp$ is symmetric.

(2) Any graph is disjoint from its own link, hence if $[g\Lambda] \perp [h\Omega]$, then $[g\Lambda]$ and $[h\Omega]$ cannot be $\sqsubseteq$-comparable.

(3) Suppose that $[g\Lambda] \sqsubseteq [h\Omega]$ and $[h\Omega] \perp [k\Pi]$. Then $\Lambda \subseteq \Omega \subseteq \mathrm{lk}(\Pi)$, and there exist $a, b \in G_\Gamma$ such that

$$[a\Lambda] = [g\Lambda], \quad [a\Omega] = [h\Omega] = [b\Omega], \quad \text{and} \quad [b\Pi] = [k\Pi].$$

In particular, this means that $b^{-1}a \in \langle \mathrm{st}(\Omega) \rangle$, hence we can write $b^{-1}a = \omega l$, where $\omega \in \langle \Omega \rangle$ and $l \in \langle \mathrm{lk}(\Omega) \rangle$. Then $\omega^{-1}b^{-1}a = l \in \langle \mathrm{lk}(\Omega) \rangle \subseteq \langle \mathrm{lk}(\Lambda) \rangle \subseteq \langle \mathrm{st}(\Lambda) \rangle$, and so $[a\Lambda] = [b\omega\Lambda]$. On the other hand, $\omega^{-1}b^{-1}b = \omega^{-1} \in \langle \Omega \rangle \subseteq \langle \mathrm{lk}(\Pi) \rangle \subseteq \langle \mathrm{st}(\Pi) \rangle$, and so $[b\Pi] = [b\omega\Pi]$. Therefore, $[g\Lambda] \perp [k\Pi]$, because $\Lambda \subseteq \mathrm{lk}(\Pi)$ and $[g\Lambda] = [b\omega\Lambda]$, $[k\Pi] = [b\omega\Pi]$. ∎

Our final relation is transversality, which is a little more nuanced, since our $[g\Lambda]$ and $[h\Omega]$ need not have a common representative k in this case.

Definition 3.21 (Transversality and lateral relative projections). If $[g\Lambda], [h\Omega] \in \mathfrak{S}_\Gamma$ are not orthogonal and neither is nested in the other, then we say $[g\Lambda]$ and $[h\Omega]$ are transverse, denoted by $[g\Lambda] \pitchfork [h\Omega]$. When $[g\Lambda] \pitchfork [h\Omega]$, for each choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$, define $\rho^{h\Omega}_{g\Lambda} \subseteq C(g\Lambda)$ by

$$\rho^{h\Omega}_{g\Lambda} = \bigcup_{k\Omega \parallel h\Omega} \pi_{g\Lambda}\big(k\langle \Omega \rangle\big) = \pi_{g\Lambda}\big(h\big\langle \mathrm{st}(\Omega) \big\rangle\big).$$

The next lemma verifies that $\rho^{h\Omega}_{g\Lambda}$ has diameter at most 2.

Lemma 3.22. *If $[g\Lambda] \pitchfork [h\Omega]$, then for any choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$, we have* $\mathrm{diam}(\pi_{g\Lambda}(h\langle \mathrm{st}(\Omega) \rangle)) \le 2$ *and* $\mathrm{diam}(\pi_{h\Omega}(g\langle \mathrm{st}(\Lambda) \rangle)) \le 2$.

Proof. We provide the proof for $\mathrm{diam}(\pi_{g\Lambda}(h\langle \mathrm{st}(\Omega) \rangle)) \le 2$. The other case is identical.

Let $x, y \in h\langle \mathrm{st}(\Omega) \rangle$. Define $p_x = \pi_{g\Lambda}(x) = \mathfrak{g}_{g\Lambda}(x)$ and $p_y = \pi_{g\Lambda}(y) = \mathfrak{g}_{g\Lambda}(y)$. If Λ splits as a join $\Lambda_1 \bowtie \Lambda_2$, then $\mathsf{d}_{g\Lambda}(p_x, p_y) \le \mathrm{diam}(C(g\Lambda)) \le 2$ by Remark 3.7.

Now suppose that Λ does not split as a join. Since $p_x, p_y \in g\langle \Lambda \rangle$, we have

$$\mathrm{supp}(p_x^{-1} p_y) \subseteq \Lambda.$$

If $\mathrm{supp}(p_x^{-1} p_y)$ is a proper subgraph of Λ, then the $C(g\Lambda)$-distance between p_x and p_y will be at most 1. Thus, it suffices to prove that $\mathrm{supp}(p_x^{-1} p_y) \ne \Lambda$.

Since $[g\Lambda] \pitchfork [h\Omega]$, we have that $[g\Lambda] \not\perp [h\Omega]$, $[g\Lambda] \not\sqsubseteq [h\Omega]$, and $[h\Omega] \not\sqsubseteq [g\Lambda]$. This can occur in two different ways: either $\Lambda \not\subseteq \mathrm{lk}(\Omega)$, $\Omega \not\subseteq \Lambda$ and $\Lambda \not\subseteq \Omega$, or there does not exist $k \in G_\Gamma$ so that $[g\Lambda] = [k\Lambda]$ and $[h\Omega] = [k\Omega]$.

First assume that $\Lambda \not\subseteq \mathrm{lk}(\Omega)$ and $\Lambda \not\subseteq \Omega$. Then $\Lambda \not\subseteq \mathrm{st}(\Omega)$, as Λ also does not split as a join. This implies that $\mathrm{st}(\Omega) \cap \Lambda \ne \Lambda$. By Proposition 3.11, every syllable of $p_x^{-1} p_y$ is a syllable of $x^{-1}y$. Since $x^{-1}y \in \langle \mathrm{st}(\Omega) \rangle$, this implies that $\mathrm{supp}(p_x^{-1} p_y) \subseteq \mathrm{st}(\Omega) \cap \Lambda \ne \Lambda$ as desired.

Now assume that $\Lambda \subseteq \mathrm{lk}(\Omega)$ or $\Lambda \subseteq \Omega$. Thus, there does not exist $k \in G_\Gamma$ so that $[g\Lambda] = [k\Lambda]$ and $[h\Omega] = [k\Omega]$. For the purposes of contradiction, suppose that $\mathrm{supp}(p_x^{-1} p_y) = \Lambda$.

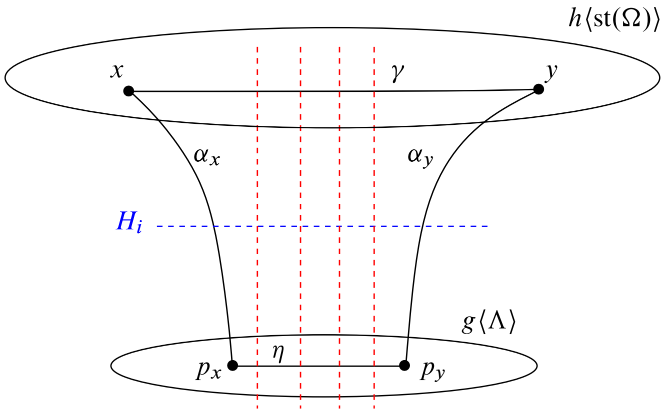

Figure 6. Any hyperplane that crosses α_x and α_y must cross all of the hyperplanes separating p_x and p_y.

Let s_x and s_y be the suffixes of x and y, respectively, such that $x = p_x s_x$ and $y = p_y s_y$. Select the following $S(\Gamma)$-geodesics: α_x connecting x and p_x, α_y connecting y and p_y, η connecting p_x and p_y, γ connecting x and y; see Figure 6.

Let $t_1 \dots t_n$ be the reduced syllable expression for s_x corresponding to the geodesic α_x. For each $i \in \{1, \dots, n\}$, let H_i be the hyperplane crossing the edge of α_x labelled by t_i. Recall that a hyperplane in $S(\Gamma)$ crosses a geodesic segment if and only if it separates the end points of the segment (Proposition 2.19 (4)). Each H_i therefore separates x and $p_x = \mathfrak{g}_{g\Lambda}(x)$, so each H_i must separate x from all of $g\langle\Lambda\rangle$ by Proposition 2.21 (4). In particular, no H_i crosses η. Thus, by Remark 2.20, each H_i must cross either γ or α_y. If H_i crosses γ, then $t_i \in \langle \mathrm{st}(\Omega)\rangle$. On the other hand, if H_i crosses α_y, then H_i must cross every hyperplane that separates p_x and p_y; see Figure 6. Because $\mathrm{supp}(p_x^{-1} p_y) = \Lambda$, it follows that for every vertex v of Λ there exists a hyperplane that separates p_x and p_y and is labelled by v. Hence, if H_i crosses α_y, then H_i crosses at least one hyperplane that is labelled by each vertex of Λ. By Proposition 2.19 (5), if two hyperplanes cross then they are labelled by adjacent vertices in Γ. Thus, the vertex labelling H_i must be in the link of Λ. In particular, $t_i \in \langle \mathrm{lk}(\Lambda)\rangle$.

The above shows that $t_i \in \langle \mathrm{st}(\Omega)\rangle$ or $t_i \in \langle \mathrm{lk}(\Lambda)\rangle$ for each $i \in \{1, \dots, n\}$. Further, $t_i \in \langle \mathrm{st}(\Omega)\rangle$ if H_i crosses γ and $t_i \in \langle \mathrm{lk}(\Lambda)\rangle$ if H_i crosses α_y. Now suppose that $i < j$ and that H_i crosses γ, but H_j crosses α_y. As shown in Figure 7, this forces H_i to cross H_j, which implies that t_i and t_j commute by Proposition 2.19 (5). Thus, by commuting the syllables of s_x, we have $s_x = l_x \omega_x$ where $\omega_x \in \langle \mathrm{st}(\Omega)\rangle$ and $l_x \in \langle \mathrm{lk}(\Lambda)\rangle$.

Now, since $x \in h\langle \mathrm{st}(\Omega)\rangle$, we have $h^{-1}x \in \langle \mathrm{st}(\Omega)\rangle$, which implies that $[h\Omega] = [x\Omega]$. Since $x = p_x s_x = p_x l_x \omega_x$, we have $[x\Omega] = [p_x l_x \omega_x \Omega] = [p_x l_x \Omega]$. Similarly, $p_x \in g\langle\Lambda\rangle$, so $g^{-1}p_x \in \langle\Lambda\rangle$, which implies that $[g\Lambda] = [p_x\Lambda]$. Now $[p_x\Lambda] = [p_x l_x \Lambda]$ as $p_x^{-1}(p_x l_x) = l_x \in \langle \mathrm{lk}(\Lambda)\rangle \subseteq \langle \mathrm{st}(\Lambda)\rangle$. Thus, we have

$$[h\Omega] = [p_x l_x \Omega] \quad \text{and} \quad [g\Lambda] = [p_x l_x \Lambda].$$

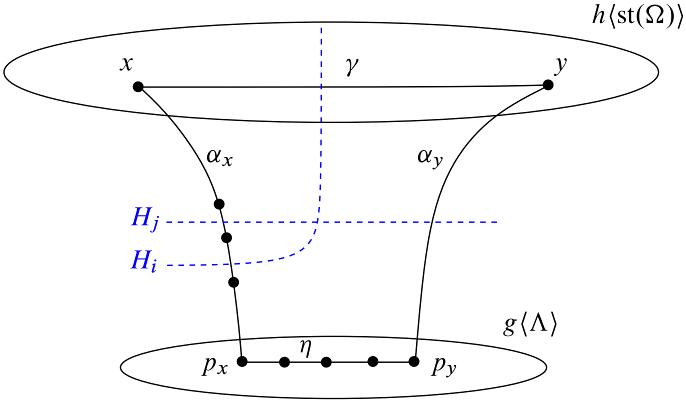

Figure 7. The hyperplane H_i crosses α_x and γ while H_j crosses α_x and α_y. Since H_i appears before H_j along α_x, H_i must cross H_j.

However, this contradicts our assumption that there is no $k \in G_\Gamma$ such that $[h\Omega] = [k\Omega]$ and $[g\Lambda] = [k\Lambda]$, proving that we must have $\mathrm{supp}(p_x^{-1} p_y) \neq \Lambda$ as desired. ∎

3.3. The proto-hierarchy structure

We now combine the work in this section to give a proto-hierarchy structure for G_Γ.

Theorem 3.23. *Let G_Γ be a graph product of finitely generated groups. For each parallelism class $[g\Lambda] \in \mathfrak{S}_\Gamma$, fix a representative $g\Lambda \in [g\Lambda]$. The following is a 2-proto-hierarchy structure for (G_Γ, d).*

- *The index set is the set of parallelism classes $\mathfrak{S}_\Gamma$ defined in Definition 3.1.*
- *The space $C([g\Lambda])$ associated to $[g\Lambda]$ is the space $C(g\Lambda)$ from Definition 3.3, where $g\Lambda$ is the fixed representative of $[g\Lambda]$.*
- *The projection map $\pi_{[g\Lambda]} \colon G_\Gamma \to C([g\Lambda])$ is the map $\pi_{g\Lambda} \colon G_\Gamma \to C(g\Lambda)$ from Definition 3.9 for the fixed representative $g\Lambda \in [g\Lambda]$.*
- *$[g\Lambda] \sqsubseteq [h\Omega]$ if $\Lambda \subseteq \Omega$ and there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\Omega] = [h\Omega]$.*
- *The upwards relative projection $\rho^{[g\Lambda]}_{[h\Omega]}$ when $[g\Lambda] \sqsubsetneq [h\Omega]$ is the set $\rho^{g\Lambda}_{h\Omega}$ from Definition 3.17, where $g\Lambda$ and $h\Omega$ are the fixed representatives for $[h\Omega]$ and $[g\Lambda]$.*
- *$[g\Lambda] \perp [h\Omega]$ if $\Lambda \subseteq \mathrm{lk}(\Omega)$ and there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\Omega] = [h\Omega]$.*
- *$[g\Lambda] \pitchfork [h\Omega]$ whenever $[g\Lambda]$ and $[h\Omega]$ are not orthogonal and neither is nested into the other.*
- *The lateral relative projection $\rho^{[g\Lambda]}_{[h\Omega]}$ when $[g\Lambda] \pitchfork [h\Omega]$ is the set $\rho^{g\Lambda}_{h\Omega}$ from Definition 3.21, where $g\Lambda$ and $h\Omega$ are the fixed representatives for $[h\Omega]$ and $[g\Lambda]$.*

Proof. The projection map $\pi_{[g\Lambda]}$ is shown to be $(1,0)$-coarsely Lipschitz in Lemma 3.12. Nesting is shown to be a partial order in Lemma 3.16. The upward relative projection has diameter at most 2 by Lemma 3.18. Lemma 3.20 shows that orthogonality is symmetric and mutually exclusive of nesting, and that nested domains inherit orthogonality. The lateral relative projections have diameter at most 2 by Lemma 3.22. ∎

4. Graph products are relative HHGs

In this section, we complete our proof that graph products of finitely generated groups are relative HHGs (Theorem 4.22) by proving the eight remaining HHS axioms and showing that the group structure is compatible with our hierarchy structure. In Section 4.1, we prove hyperbolicity of $C(g\Lambda)$ whenever Λ contains at least two vertices. Section 4.2 is devoted to proving the finite complexity and containers axioms. Section 4.3 deals with the uniqueness axiom, and in Section 4.4, we the prove the bounded geodesic image and large links axioms. In Section 4.5, we verify partial realisation, and Section 4.6 deals with the consistency axiom. Finally, in Section 4.7, compatibility of the relative HHS structure with the group structure is checked.

We also obtain some auxiliary results along the way: in Section 4.1, we show that not only are the spaces $C(g\Lambda)$ hyperbolic whenever Λ contains at least two vertices, but they are also quasi-trees; and in Section 4.3, we use uniqueness to give a classification of when $C(g\Lambda)$ has infinite diameter.

We conclude the section by remarking that the syllable metric on G_Γ is an HHS. This is true even when the vertex groups are not finitely generated. However, until then we will continue to assume that G_Γ is a graph product of finitely generated groups and that d is the word metric on G_Γ, where the generating set for G_Γ is given by taking a union of finite generating sets for each vertex group.

4.1. Hyperbolicity

Lemma 4.1 (Hyperbolicity). *For each $[g\Lambda] \in \mathfrak{S}_\Gamma$, either $[g\Lambda]$ is $\sqsubseteq$-minimal or $C(g\Lambda)$ is $\frac{7}{2}$-hyperbolic.*

Remark 4.2. The hyperbolicity of $C(g\Lambda)$ can also be deduced from [15, Proposition 6.4]. The proof presented below uses a different argument that produces the explicit hyperbolicity constant of $\frac{7}{2}$.

Proof. Take $[g\Lambda] \in \mathfrak{S}_\Gamma$ and suppose it is not $\sqsubseteq$-minimal; i.e., Λ contains at least two vertices. Let $x, y, z \in C(g\Lambda)$ be three distinct points and let $\gamma_1, \gamma_2, \gamma_3$ be three $C(g\Lambda)$-geodesics connecting the pairs $\{y, z\}$, $\{z, x\}$, $\{x, y\}$, respectively. We wish to show that this triangle is $\frac{7}{2}$-slim, that is, we will show that γ_1 is contained in the $\frac{7}{2}$-neighbourhood of $\gamma_2 \cup \gamma_3$. Since $C(g\Lambda)$ is a metric graph whose edges have length 1, it suffices to show that any vertex of γ_1 is at distance at most 3 from $\gamma_2 \cup \gamma_3$.

Let $p^i_1, \dots, p^i_{m_i}$ be the vertices of γ_i, and let γ'_i be the path in $S(g\Lambda)$ obtained by connecting each pair of consecutive vertices p^i_j and p^i_{j+1} with an $S(g\Lambda)$-geodesic α^i_j. Since α^i_j is labelled by vertices of $\operatorname{supp}((p^i_j)^{-1} p^i_{j+1})$, which is a proper subgraph of Λ,

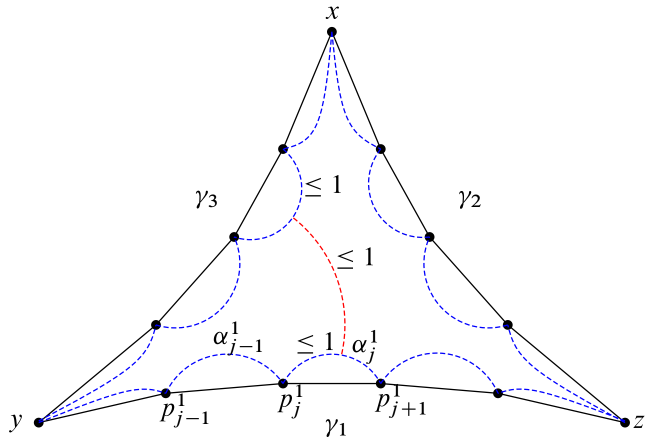

Figure 8. For each edge of the $C(g\Lambda)$-geodesic triangle, we construct an $S(g\Lambda)$-geodesic segment α^i_j between its endpoints (shown in blue). To show that the triangle is $\frac{7}{2}$-slim, it then suffices to show that for each j, $\alpha^1_{j-1} \cup \alpha^1_j$ is $C(g\Lambda)$-distance 1 from some α^i_t with $i \neq 1$.

the $C(g\Lambda)$-distance between any vertex of α^i_j and p^i_j or p^i_{j+1} is at most 1. It therefore suffices to show that given any vertex p^1_j of γ_1, either α^1_{j-1} or α^1_j is $C(g\Lambda)$-distance 1 from some α^i_t with $i = 2$ or 3; see Figure 8.

If Λ has no edges, then $\langle\Lambda\rangle$ is the free product of the vertex groups, hence $S(g\Lambda)$ is a tree of simplices, that is, any cycle in $S(g\Lambda)$ is contained in a single simplex (a coset of a vertex group). Therefore, any two paths in $S(g\Lambda)$ with the same endpoints are contained in the 1-neighbourhood of each other, and in particular γ'_1 is contained in the 1-neighbourhood of $\gamma'_2 \cup \gamma'_3$. Thus, any vertex of γ_1 is at distance at most 3 from $\gamma_2 \cup \gamma_3$ in $C(g\Lambda)$.

Now suppose that Λ has at least one edge, so that it has a vertex w with non-empty link. We may also assume that Λ does not split as a join; otherwise, $C(g\Lambda)$ has diameter 2 by Remark 3.7 and hence is clearly $\frac{7}{2}$-hyperbolic. Take a vertex p^1_j of γ_1. If p^1_j is one of the first or last four vertices of γ_1, then it is at distance at most 3 from γ_2 or γ_3 in $C(g\Lambda)$. Otherwise, p^1_j is an endpoint of two consecutive edges L_{j-1} and L_j of γ_1 labelled by strict subgraphs Λ_{j-1} and Λ_j of Λ. We must have $\Lambda_{j-1} \cup \Lambda_j = \Lambda$, as otherwise we could replace these two edges with a single edge labelled by $\Lambda_{j-1} \cup \Lambda_j$, contradicting γ_1 being a $C(g\Lambda)$-geodesic. It follows that all vertices of Λ appear as labels on the edges of the geodesic segments α^1_{j-1} and α^1_j of γ'_1 corresponding to L_{j-1} and L_j. Consider the collection $\mathcal{E}_w$ of edges of $\alpha^1_{j-1} \cup \alpha^1_j$ labelled by the fixed vertex w with $\mathrm{lk}(w) \cap \Lambda \neq \emptyset$, and consider the collection $\mathcal{H}_w$ of hyperplanes in $S(g\Lambda)$ dual to the edges in $\mathcal{E}_w$. We proceed to construct an $S(g\Lambda)$-path from an edge of $\mathcal{E}_w$ to some α^i_t with $i = 2$ or 3, either by travelling through the carrier of a single hyperplane, labelled by $\mathrm{st}(w) \cap \Lambda \subsetneq \Lambda$, or by following a sequence of combinatorial hyperplanes labelled by $\mathrm{lk}(w) \cap \Lambda \subsetneq \Lambda$. Since this path will be labelled by a proper subgraph of Λ, the $C(g\Lambda)$-distance between its endpoints will be 1.

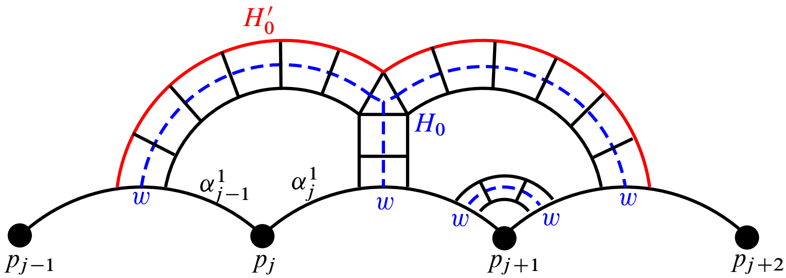

Figure 9. The outermost hyperplane H_0 of $\mathcal{H}_w$ and its outermost combinatorial hyperplane H_0'.

Suppose some hyperplane $H \in \mathcal{H}_w$ also crosses a geodesic segment α^i_t of $\gamma_2' \cup \gamma_3'$. Since the carrier of H is labelled by vertices of $\mathrm{st}(w) \cap \Lambda$, and $\mathrm{st}(w) \cap \Lambda$ is a strict subgraph of Λ because Λ does not split as a join, it follows that p^1_j is at most $C(g\Lambda)$-distance 3 from either γ_2 or γ_3, as desired.

Suppose therefore that no hyperplane of $\mathcal{H}_w$ crosses $\gamma_2' \cup \gamma_3'$. This means that each $H \in \mathcal{H}_w$ must cross γ_1' a second time (Remark 2.20). Further, Proposition 2.19 (5) tells us that no two hyperplanes labelled by the same vertex may cross each other. It follows that there exists an outermost hyperplane H_0 of $\mathcal{H}_w$; that is, no hyperplane of $\mathcal{H}_w$ crosses edges of γ_1' both earlier and later than H_0 does. Moreover, H_0 has an outermost combinatorial hyperplane H_0'; see Figure 9. Note that since this combinatorial hyperplane is labelled by vertices of $\mathrm{lk}(w) \cap \Lambda \subsetneq \Lambda$, the $C(g\Lambda)$-distance between any two points on H_0' is 1. In particular, since γ_1 is a $C(g\Lambda)$-geodesic, it follows that the segments α^1_r and α^1_k that H_0' intersects must satisfy $|k-r| \leq 2$. As we know that H_0 crosses $\alpha^1_{j-1} \cup \alpha^1_j$, this implies that H_0' must intersect $\alpha^1_{j-1} \cup \alpha^1_j$ too. Recalling that a hyperplane may not cross the same geodesic twice (Proposition 2.19 (4)), we may therefore suppose without loss of generality that $r = j$ and $j < k \leq j+2$ (the cases where $j-2 \leq k < j$ or $r = j-1$ proceed similarly).

Let E_0 be the edge of $\mathcal{E}_w$ on α^1_j that H_0 crosses, and let e_1 and e_2 denote its endpoints. Let F_0 be the edge of α^1_k labelled by w that H_0 crosses, and denote its endpoints by f_1 and f_2. Then there is a path η connecting e_1 and f_2 that is contained in the combinatorial hyperplane H_0' labelled by vertices of $\mathrm{lk}(w) \cap \Lambda \subsetneq \Lambda$. Furthermore, if w does not appear as a label of an earlier edge of α^1_j or a later edge of α^1_k, then $\mathrm{d}_{g\Lambda}(p^1_j, p^1_{k+1}) = 1$ as the path obtained by travelling from p^1_j to e_1 along α^1_j, then from e_1 to f_2 along η, then from f_2 to p^1_{k+1} along α^1_k is labelled by the proper subgraph $\Lambda \smallsetminus w$. This contradicts the assumption that γ_1 is a $C(g\Lambda)$-geodesic. On the other hand, if w appears as a label of an earlier edge E_{-1} of α^1_j (take the closest one to E_0) but not a later edge of α^1_k, then the corresponding hyperplane H_{-1} must cross a segment α^1_l with $l < j$ (since H_0 is outermost), and there exists an $S(g\Lambda)$-path ξ labelled by $\Lambda \smallsetminus w$ connecting e_1 and α^1_l. Then the $C(g\Lambda)$-distance between the endpoints of the path $\xi \cup \eta$ is 1 and so we obtain $\mathrm{d}_{g\Lambda}(p^1_l, p^1_{k+1}) \leq 2$, a contradiction. There therefore exists some edge labelled by w which appears after F_0 on α^1_k. Let E_1 be the closest such edge to H_0, and consider the hyperplane H_1 dual to E_1.

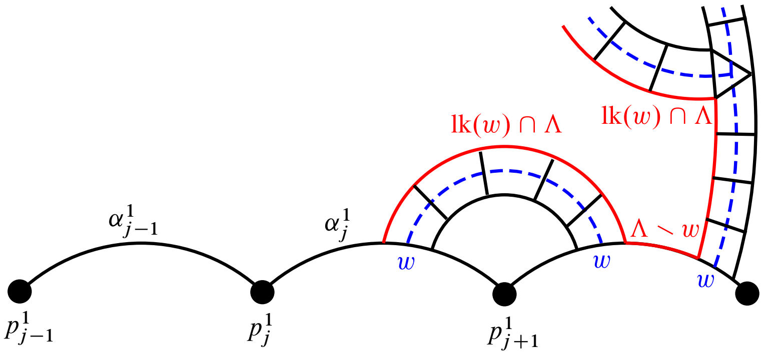

Figure 10. By following a sequence of combinatorial hyperplanes, we obtain a path labelled by $\Lambda \smallsetminus w$ (shown in red) that must eventually leave γ'_1 and cross $\gamma'_2 \cup \gamma'_3$.

If H_1 crosses α^1_s with $|s - j| \geq 3$, then we obtain a contradiction since we have a path in $C(g\Lambda)$ from p^1_j to p^1_{s+1} (or p^1_{j+1} to p^1_s if $s < j$) of length at most 3. If H_1 crosses α^1_s with $|s - k| \geq 3$, then similarly we obtain a contradiction. Assume therefore that $|s - j| \leq 2$ and $|s - k| \leq 2$. Note that since H_0 and H_1 cannot cross, we must have $s < j$ or $s > k$.

If $s < j$, then we must have $k = j + 1$ and $s = j - 1$. In this case, H_1 crosses α^1_{j-1}, which contradicts our assumption that H_0 is an outermost hyperplane of $\mathcal{H}_w$. Thus, H_1 cannot cross any α^1_s with $s < k$. This implies that if H_1 crosses a segment α^i_s with $i = 2$ or 3, then we can conclude that p^1_j is at most $C(g\Lambda)$-distance 3 from either γ_2 or γ_3, by following a sequence of geodesics labelled by vertices of $\mathrm{lk}(w) \cap \Lambda$ and contained in combinatorial hyperplanes associated to H_0 and H_1; see Figure 10.

On the other hand, if H_1 crosses α^1_s with $s > k$, then $k = j + 1$ and $s = j + 2$. Repeating the same process, there must exist a later edge of α^1_s labelled by w. Let H_2 be the hyperplane dual to the closest such edge to H_1. If H_2 also crosses α^1_t where $t \neq s$, then we must have $t < j = s - 2$ or $t > s = j + 2$, as H_2 cannot cross the previous hyperplanes. However, the first case results in $|t - s| \geq 3$, and the second case gives $|t - j| \geq 3$, both of which give a contradiction. Therefore, H_2 must cross α^i_t where $i = 2$ or 3. Following the sequence of geodesics labelled by vertices of $\Lambda \smallsetminus w$, we again see that p^1_j is at most $C(g\Lambda)$-distance 3 from either γ_2 or γ_3. ■

A similar technique can moreover show that the spaces $C(g\Lambda)$ are quasi-trees, by applying Manning's bottleneck criterion.

Theorem 4.3 (Bottleneck criterion [21]). *Let Y be a geodesic metric space. The following are equivalent:*

(1) *Y is quasi-isometric to some simplicial tree T;*

(2) *there is some $\Delta > 0$ so that for all $y, z \in Y$ there is a midpoint $m = m(y,z)$ with $\mathsf{d}(y,m) = \mathsf{d}(z,m) = \frac{1}{2}\mathsf{d}(y,z)$ and the property that any path from y to z must pass within a distance Δ of m.*

Theorem 4.4. *For each $[g\Lambda] \in \mathfrak{S}_\Gamma$, either $[g\Lambda]$ is $\sqsubseteq$-minimal or $C(g\Lambda)$ is a quasi-tree.*

The proof of Theorem 4.4 proceeds similarly to the proof of Lemma 4.1, with the role of γ_1 being played by a geodesic from y to z containing the midpoint $m(y,z)$, and replacing $\gamma_2 \cup \gamma_3$ with an arbitrary path from y to z.

Proof. Suppose that $[g\Lambda]$ is not $\sqsubseteq$-minimal. Let $x, y \in C(g\Lambda)$, let γ be a $C(g\Lambda)$-geodesic connecting x and y, and let β be another $C(g\Lambda)$-path from x to y. From γ and β, we may obtain paths γ' and β' in $S(g\Lambda)$ by replacing each edge with a geodesic segment in $S(g\Lambda)$. Note that any point on such a segment is $C(g\Lambda)$-distance 1 from the endpoints of the segment. Let m be the midpoint of γ, so that m is either a vertex of γ or a midpoint of an edge.

If Λ has no edges, then $S(g\Lambda)$ is a tree of simplices in the same manner as in the previous proof, and in particular any two paths in $S(g\Lambda)$ between x and y are contained in the 1-neighbourhood of each other. Applying this to γ' and β' shows that m is at distance at most $\Delta = \frac{7}{2}$ from β.

Now suppose that Λ has at least one edge, and let L_1 and L_2 be two edges of γ adjacent to m (if m is the midpoint of an edge L, pick L and one edge adjacent to it). Then L_1 and L_2 are labelled by strict subgraphs Λ_1 and Λ_2 of Λ such that $\Lambda_1 \cup \Lambda_2 = \Lambda$. Thus, either Λ_1 or Λ_2 contains a vertex w with non-empty link, and w therefore appears as a label of a hyperplane crossing an edge of the corresponding geodesic segments α_1 and α_2 of γ'.

We can now repeat the argument in the proof of Lemma 4.1 to find a path connecting $\alpha_1 \cup \alpha_2$ to β' that is labelled by a proper subgraph of Λ. It follows that m is at most $C(g\Lambda)$-distance $\Delta = \frac{7}{2}$ from β. ■

4.2. Finite complexity and containers

Lemma 4.5 (Finite complexity). *Any set of pairwise $\sqsubseteq$-comparable elements has cardinality at most $|V(\Gamma)|$.*

Proof. If $[g\Lambda] \sqsubseteq [h\Omega]$ and Λ and Ω have the same number of vertices, then we must have $\Lambda = \Omega$ and $[g\Lambda] = [k\Lambda] = [k\Omega] = [h\Omega]$ for some $k \in G_\Gamma$. Therefore, any two distinct $\sqsubseteq$-comparable elements must have different numbers of vertices. Thus, any set of pairwise $\sqsubseteq$-comparable elements has cardinality at most $|V(\Gamma)|$. ■

Lemma 4.6 (Containers). *Let $[h\Omega] \sqsubsetneq [g\Lambda]$ be elements of $\mathfrak{S}_\Gamma$. If there exists $[k\Pi] \in \mathfrak{S}_\Gamma$ such that $[k\Pi] \sqsubseteq [g\Lambda]$ and $[k\Pi] \perp [h\Omega]$, then $[k\Pi] \sqsubseteq [a(\mathrm{lk}(\Omega) \cap \Lambda)] \sqsubsetneq [a\Lambda]$, where $a \in G_\Gamma$ satisfies $[a\Lambda] = [g\Lambda]$ and $[a\Omega] = [h\Omega]$.*

Proof. First, since $[k\Pi] \sqsubseteq [g\Lambda]$ and $[k\Pi] \perp [h\Omega]$, we have $\Pi \subseteq \Lambda$ and $\Pi \subseteq \mathrm{lk}(\Omega)$, hence $\Pi \subseteq \mathrm{lk}(\Omega) \cap \Lambda \subsetneq \Lambda$. Next, let $b \in G_\Gamma$ be such that $[b\Pi] = [k\Pi]$ and $[b\Omega] = [h\Omega]$, and let $c \in G_\Gamma$ be such that $[c\Pi] = [k\Pi]$ and $[c\Lambda] = [g\Lambda]$. We claim that there exists $d \in G_\Gamma$ such that $[k\Pi] = [d\Pi]$ and $[a(\mathrm{lk}(\Omega) \cap \Lambda)] = [d(\mathrm{lk}(\Omega) \cap \Lambda)]$, which would complete our proof.

Indeed, $k^{-1}a = k^{-1}bb^{-1}a = k^{-1}cc^{-1}a$, and we know that

$$\mathrm{supp}(k^{-1}b) \subseteq \mathrm{st}(\Pi), \quad \mathrm{supp}(b^{-1}a) \subseteq \mathrm{st}(\Omega)$$

and

$$\mathrm{supp}(k^{-1}c) \subseteq \mathrm{st}(\Pi), \quad \mathrm{supp}(c^{-1}a) \subseteq \mathrm{st}(\Lambda).$$

Writing $p = \mathrm{prefix}_{\mathrm{st}(\Pi)}(k^{-1}a)$, we have $p^{-1}k^{-1}a = s$, where s satisfies $\mathrm{prefix}_{\mathrm{st}(\Pi)}(s) = e$. That is, $\mathrm{prefix}_{\mathrm{st}(\Pi)}(p^{-1}k^{-1}bb^{-1}a) = e$. Since $p^{-1}k^{-1}b \in \langle \mathrm{st}(\Pi)\rangle$ and $b^{-1}a \in \langle \mathrm{st}(\Omega)\rangle$, this implies that $p^{-1}k^{-1}a \in \langle \mathrm{st}(\Omega)\rangle$. Similarly, writing $k^{-1}a = k^{-1}cc^{-1}a$ shows us that $p^{-1}k^{-1}a \in \langle \mathrm{st}(\Lambda)\rangle$.

That is to say, we can write $k^{-1}a = ps$, where $p \in \langle \mathrm{st}(\Pi)\rangle$ and $s \in \langle \mathrm{st}(\Omega) \cap \mathrm{st}(\Lambda)\rangle$. But $\Omega \subseteq \Lambda$ and $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Omega)$, hence

$$\mathrm{st}(\Omega) \cap \mathrm{st}(\Lambda) = \Omega \cup \mathrm{lk}(\Lambda) \cup \big(\mathrm{lk}(\Omega) \cap \Lambda\big).$$

Moreover, $\Omega \cup \mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\mathrm{lk}(\Omega) \cap \Lambda)$, hence $s \in \langle \mathrm{st}(\mathrm{lk}(\Omega) \cap \Lambda)\rangle$. Thus, $k^{-1}as^{-1} = p \in \langle \mathrm{st}(\Pi)\rangle$ and $a^{-1}as^{-1} \in \langle \mathrm{st}(\mathrm{lk}(\Omega) \cap \Lambda)\rangle$. Letting $d = as^{-1}$, we have $[k\Pi] = [d\Pi]$ and $[a(\mathrm{lk}(\Omega) \cap \Lambda)] = [d(\mathrm{lk}(\Omega) \cap \Lambda)]$ as desired. ∎

4.3. Uniqueness

Here we prove the uniqueness axiom, which tells us that all geometry of G_Γ is witnessed by some associated space $C(g\Lambda)$. This means we do not lose any geometric information through our projections. We also use this axiom to classify boundedness of the hyperbolic spaces $C(g\Lambda)$. In what follows, $|\cdot|_{G_\Gamma}$ denotes the word length on G_Γ with respect to the generating set S defined at the beginning of Section 3.

Lemma 4.7 (Uniqueness). *Let G_Γ be a graph product of finitely generated groups. For all $g \in G_\Gamma$, if $\mathsf{d}_{h\Lambda}(e,g) \leq r$ for all $h \in G_\Gamma$ and subgraphs $\Lambda \subseteq \Gamma$, then*

$$|g|_{G_\Gamma} \leq (2^{|V(\Gamma)|}r + 2)^{|V(\Gamma)|}.$$

Proof. Let $r \geq 0$. If Γ is a single vertex, then the conclusion is immediate as the only subgraph is Γ and $d_\Gamma(e,g) = |g|_{G_\Gamma} = r$. Suppose that Γ contains $n+1$ vertices and assume that the lemma holds for any graph product of finitely generated groups whose defining graph contains at most n vertices. Suppose that $g \in G_\Gamma$ with $\mathsf{d}_{h\Lambda}(e,g) \leq r$ for all $h \in G_\Gamma$ and subgraphs $\Lambda \subseteq \Gamma$.

Since $\mathsf{d}_\Gamma(e,g) \leq r$, there exist proper subgraphs $\Lambda_i \subsetneq \Gamma$ and elements λ_i with $\mathrm{supp}(\lambda_i) = \Lambda_i$ so that $g = \lambda_1 \ldots \lambda_m$ and $\mathsf{d}_\Gamma(e,g) = m \leq r$. We shall see that $\mathsf{d}_{h\Omega}(e,g) \leq r$ implies that $\mathsf{d}_{h\Omega}(e,\lambda_i)$ is uniformly bounded for each $\Omega \subseteq \Lambda_i$ and $h \in \langle \Lambda_i \rangle$. Since each $\langle \Lambda_i \rangle$ is a graph product on at most n vertices, induction will imply that the word length of each λ_i is bounded, which in turn will bound the word length of g.

If Γ splits as a join $\Gamma = \Lambda_1 \bowtie \Lambda_2$, then any element $g \in G_\Gamma$ can be written in the form $g = \lambda_1\lambda_2$, where $\lambda_i \in \langle \Lambda_i \rangle$ for $i = 1,2$ and $|g|_{G_\Gamma} = |\lambda_1|_{G_\Gamma} + |\lambda_2|_{G_\Gamma}$. Moreover, if

$h \in \langle \Lambda_i \rangle$ and $\Omega \subseteq \Lambda_i$, then $\mathfrak{g}_{h\Omega}(g) = h \cdot \operatorname{prefix}_\Omega(h^{-1}g) = h \cdot \operatorname{prefix}_\Omega(h^{-1}\lambda_i) = \mathfrak{g}_{h\Omega}(\lambda_i)$. Therefore, $\mathsf{d}_{h\Omega}(e, \lambda_i) = \mathsf{d}_{h\Omega}(e, g) \le r$ and by induction $|\lambda_i|_{G_\Gamma} \le (2^n r + 2)^n$ for $i = 1, 2$. Thus, $|g|_{G_\Gamma} \le 2(2^n r + 2)^n \le (2^{n+1} r + 2)^{n+1}$.

Suppose that Γ does not split as a join, and define $p_0 = e$ and $p_i = \lambda_1 \ldots \lambda_i$ for $i \in \{1, \ldots, m\}$. Note that the p_i are the vertices of the $C(\Gamma)$-geodesic connecting e and g with edges labelled by the λ_i. By Lemma 3.5, we can assume that $\operatorname{suffix}_{\Lambda_i}(p_{i-1}) = e$ for each $i \in \{2, \ldots, m\}$ and that there exists an $S(\Gamma)$-geodesic connecting e to g that contains each p_i as a vertex. Fix $i \in \{1, \ldots, m\}$, $h \in \langle \Lambda_i \rangle$, and $\Omega \subseteq \Lambda_i$.

As stated above, we wish to show that $\mathsf{d}_{h\Omega}(e, \lambda_i)$ is bounded uniformly in terms of r so that we can apply the induction hypothesis. Since $\mathsf{d}_{h\Omega}(e, \lambda_i)$ is independent of the choice of representative of the coset $h\langle \Omega \rangle$, we can assume that $\operatorname{suffix}_\Omega(h) = e$. To achieve the bound on $\mathsf{d}_{h\Omega}(e, \lambda_i)$, we use the following two claims plus the assumption that $\mathsf{d}_{h\Omega}(e, g) \le r$.

Claim 4.8. $\pi_{p_{i-1}h\Omega}(p_{i-1}) = \pi_{p_{i-1}h\Omega}(e)$.

Proof. By equivariance and the prefix description of the gate map (Lemma 2.23),

$$\mathfrak{g}_{p_{i-1}h\Omega}(p_{i-1}) = p_{i-1}h \cdot \operatorname{prefix}_\Omega(h^{-1})$$

and

$$\mathfrak{g}_{p_{i-1}h\Omega}(e) = p_{i-1}h \cdot \operatorname{prefix}_\Omega(h^{-1} p_{i-1}^{-1}).$$

Since $\operatorname{prefix}_{\Lambda_i}(p_{i-1}^{-1}) = e$, we have $\operatorname{prefix}_\Omega(p_{i-1}^{-1}) = e$ too. Moreover, since $h \in \langle \Lambda_i \rangle$ and $\operatorname{prefix}_\Omega(p_{i-1}^{-1}) = e$, we have $\operatorname{prefix}_\Omega(h^{-1} p_{i-1}^{-1}) = \operatorname{prefix}_\Omega(h^{-1})$ and so $\mathfrak{g}_{p_{i-1}h\Omega}(p_{i-1}) = \mathfrak{g}_{p_{i-1}h\Omega}(e)$. This implies that $\pi_{p_{i-1}h\Omega}(p_{i-1}) = \pi_{p_{i-1}h\Omega}(e)$. ∎

Claim 4.9. $\mathsf{d}_{p_{i-1}h\Omega}(p_i, g) \le r$.

Proof of Claim 4.9. Recall that we can write each λ_i in reduced syllable form to produce an $S(\Gamma)$-geodesic connecting e and g and containing each p_i as a vertex (Lemma 3.5). Thus, Lemma 3.13 says that $\mathsf{d}_{p_{i-1}h\Omega}(p_i, g) \le \mathsf{d}_{p_{i-1}h\Omega}(e, g)$ and $\mathsf{d}_{p_{i-1}h\Omega}(e, g) \le r$ by assumption. ∎

By the equivariance of the gate map (Proposition 2.21 (2)),

$$\mathsf{d}_{h\Omega}(e, \lambda_i) = \mathsf{d}_{p_{i-1}h\Omega}(p_{i-1}, p_i).$$

Claim 4.8 then implies that

$$\mathsf{d}_{p_{i-1}h\Omega}(p_{i-1}, p_i) = \mathsf{d}_{p_{i-1}h\Omega}(e, p_i) \le \mathsf{d}_{p_{i-1}h\Omega}(e, g) + \mathsf{d}_{p_{i-1}h\Omega}(g, p_i).$$

Since $\mathsf{d}_{p_{i-1}h\Omega}(e, g) \le r$ by assumption and $\mathsf{d}_{p_{i-1}h\Omega}(g, p_i) \le r$ by Claim 4.9, we have $\mathsf{d}_{h\Omega}(e, \lambda_i) = \mathsf{d}_{p_{i-1}h\Omega}(p_{i-1}, p_i) \le 2r$ for each $h \in \langle \Lambda_i \rangle$ and $\Omega \subseteq \Lambda_i$. The induction hypothesis now implies that the word length of λ_i in $\langle \Lambda_i \rangle$ is at most $(2^n(2r) + 2)^n$. Thus, we have

$$|g|_{G_\Gamma} \le r(2^{n+1} r + 2)^n \le (2^{n+1} r + 2)^{n+1}$$

because each graphical subgroup is convexly embedded in the word metric d on G_Γ. ∎

The uniqueness axiom allows us to classify boundedness of the hyperbolic spaces $C(g\Lambda)$.

Theorem 4.10. *For any $g \in G_\Gamma$ and any subgraph Λ of Γ containing at least two vertices, the space $C(g\Lambda)$ has infinite diameter if and only if Λ does not split as a join.*

Proof. Recall that if Λ splits as a join, then $\operatorname{diam}(C(g\Lambda)) \leq 2$ by Remark 3.7. Suppose therefore that Λ does not split as a join and let $v_1, \ldots, v_k$ be the vertices of Λ. For each $i \in \{1, \ldots, k\}$, pick $s_i \in S_{v_i}$, where S_{v_i} is the finite generating set for G_{v_i} that we fixed at the beginning of Section 3. Define $\lambda = s_1 \ldots s_k$. For each $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, n\}$, let s_i^j be the jth copy of s_i in the product $(s_1 \ldots s_k)^n = \lambda^n$; that is, $\lambda^n = (s_1^1 \ldots s_k^1)(s_1^2 \ldots s_k^2) \ldots (s_1^n \ldots s_k^n)$.

We claim that for each $n \in \mathbb{N}$, $(s_1^1 \ldots s_k^1) \ldots (s_1^n \ldots s_k^n)$ is a reduced syllable expression for λ^n. Indeed, if $(s_1^1 \ldots s_k^1) \ldots (s_1^n \ldots s_k^n)$ is not reduced, then there exists s_i^j that is combined with some s_i^ℓ $(j \neq \ell)$ after applying some number of commutation relations (Theorem 2.4). However, if s_i^ℓ were to be combined with s_i^j, then s_i would need to commute with each of $s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_k$. This only happens if the vertex v_i is connected to every other vertex of Λ, but this does not happen as Λ does not split as a join. Therefore, $(s_1^1 \ldots s_k^1) \ldots (s_1^n \ldots s_k^n)$ is a reduced syllable expression for λ^n, and we have $|\lambda^n|_{\mathrm{syl}} = kn$ for all $n \in \mathbb{N}$.

To prove $C(g\Lambda)$ has infinite diameter, we use the following claim plus the uniqueness axiom to show that $\mathsf{d}_\Lambda(e, \lambda^n)$ can be made as large as desired by increasing n.

Claim 4.11. For all $\Omega \subsetneq \Lambda$, $h \in \langle \Lambda \rangle$, and $n \geq 2$, $\mathsf{d}_{h\Omega}(e, \lambda^n) \leq 3$.

For now we accept Claim 4.11, deferring its proof until after we have proved that $C(g\Lambda)$ has infinite diameter.

For the purposes of contradiction, assume that there exists $R > 0$ such that $\mathsf{d}_\Lambda(e, \lambda^n) \leq R$ for all $n \in \mathbb{N}$. By Claim 4.11, for every proper subgraph $\Omega \subsetneq \Lambda$ and $h \in \langle \Lambda \rangle$, we have $\mathsf{d}_{h\Omega}(e, \lambda^n) \leq 3$. Applying the uniqueness axiom (Lemma 4.7) to the graph product $\langle \Lambda \rangle = G_\Lambda$, this implies that there exists $D = D(R, |V(\Lambda)|) > 0$ such that $|\lambda^n|_{G_\Lambda} = |\lambda^n|_{G_\Gamma} \leq D$ for all $n \in \mathbb{N}$. However, this is a contradiction as $|\lambda^n|_{G_\Gamma} \geq |\lambda^n|_{\mathrm{syl}} = kn$ for all $n \in \mathbb{N}$. Thus, for each $R > 0$, there exists n_R such that $\mathsf{d}_\Lambda(e, \lambda^{n_R}) > R$. Therefore, $C(\Lambda)$, and hence $C(g\Lambda)$, has infinite diameter. ∎

Proof of Claim 4.11. Let $\Omega \subsetneq \Lambda$ be a proper subgraph and $h \in \langle \Lambda \rangle$. Since $\mathsf{d}_{h\Omega}(e, \lambda^n)$ does not depend on the choice of representative of the coset $h\langle \Omega \rangle$, we can assume that $\operatorname{suffix}_\Omega(h) = e$, and thus $\operatorname{prefix}_\Omega(h^{-1}) = e$.

Recall that $\pi_{h\Omega}(e) = h \cdot \operatorname{prefix}_\Omega(h^{-1})$ and $\pi_{h\Omega}(\lambda^n) = h \cdot \operatorname{prefix}_\Omega(h^{-1}\lambda^n)$ (Remark 3.10). Since $\operatorname{prefix}_\Omega(h^{-1}) = e$, it suffices to prove that $\mathsf{d}_\Omega(e, h^{-1}\lambda^n) \leq 3$. We can also assume that $\operatorname{prefix}_\Omega(h^{-1}\lambda^n) \neq e$.

Proposition 3.11 tells us that all syllables of $\operatorname{prefix}_\Omega(h^{-1}\lambda^n)$ are syllables of λ^n. As $\operatorname{prefix}_\Omega(h^{-1}\lambda^n) \neq e$, there must exist $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, n\}$ such that s_i^j is the first syllable of $(s_1^1 \ldots s_k^1)(s_1^2 \ldots s_k^2) \ldots (s_1^n \ldots s_k^n)$ that is also a syllable of $\operatorname{prefix}_\Omega(h^{-1}\lambda^n)$.

Let $\ell, m \in \{1, \dots, k\}$ be such that $v_\ell \in V(\Lambda \smallsetminus \mathrm{st}(\Omega))$ and $v_m \in V(\Omega)$ is not joined to v_ℓ by an edge. These vertices exist since Λ does not split as a join and thus $\Lambda \not\subseteq \mathrm{st}(\Omega)$. We will show that $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ can be written as a product $p_1 p_2 p_3$ where $\mathrm{supp}(p_2)$ is a single vertex v of Ω and $\mathrm{supp}(p_1), \mathrm{supp}(p_3) \subseteq \Omega \smallsetminus v$. This implies that the $C(\Omega)$-distance between e and $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ is at most 3, which in turn says that $\mathsf{d}_{h\Omega}(e, \lambda^n) \leq 3$.

Suppose that $i < \ell$. Since $v_\ell \notin V(\Omega)$, every syllable of $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ must either be one of $s_i^j, s_{i+1}^j, \dots, s_{\ell-1}^j$ or must commute with s_ℓ^j. As s_m does not commute with s_ℓ, it follows that no s_m^J is a syllable of $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ for $J > j$. Therefore, $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ can contain at most one syllable with support v_m, namely s_m^j. Thus, $\mathrm{prefix}_\Omega(h^{-1}\lambda^n) = p_1 p_2 p_3$ with $\mathrm{supp}(p_1) \subseteq \Omega \smallsetminus v_m$, $\mathrm{supp}(p_2) \subseteq v_m$, and $\mathrm{supp}(p_3) \subseteq \Omega \smallsetminus v_m$. Note that if $\Omega = v_m$, then $\mathrm{prefix}_\Omega(h^{-1}\lambda^n) = p_2 = s_m^j$ and $\mathsf{d}_{h\Omega}(e, \lambda^n) = \mathsf{d}_\Omega(e, s_m^j) = 1$ because $s_m^j \in S_{v_m}$.

The case $i > \ell$ proceeds similarly since every syllable of $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ must either be one of $s_i^j, s_{i+1}^j, \dots, s_k^j, s_1^{j+1}, \dots, s_{\ell-1}^{j+1}$ or must commute with s_ℓ^{j+1}. ■

In Section 5, we use our characterisation of when $C(g\Lambda)$ has infinite diameter to answer two questions of Genevois [16] (Theorems 5.14 and 5.16).

4.4. Bounded geodesic image and large links

As the bounded geodesic image axiom is used to prove large links, we include both in this section.

Lemma 4.12 (Bounded geodesic image). *Let $x, y \in G_\Gamma$ and $[h\Omega] \subsetneq [g\Lambda]$. For any choice of representatives $h\Omega \in [h\Omega]$ and $g\Lambda \in [g\Lambda]$, if $\mathsf{d}_{h\Omega}(x, y) > 0$, then every $C(g\Lambda)$-geodesic γ from $\pi_{g\Lambda}(x)$ to $\pi_{g\Lambda}(y)$ intersects the closed 2-neighbourhood of $\rho^{h\Omega}_{g\Lambda}$.*

Proof. We first need to establish that when $[h\Omega] \sqsubseteq [g\Lambda]$, gating onto $h\langle\Omega\rangle$ is the same as first gating onto $g\langle\Lambda\rangle$ and then gating onto $h\langle\Omega\rangle$. This will allow us to relate $\pi_{g\Lambda}(x)$ and $\pi_{h\Omega}(x)$.

Claim 4.13. If $[h\Omega] \sqsubseteq [g\Lambda]$, then $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x)) = \mathfrak{g}_{h\Omega}(x)$ for all $x \in G_\Gamma$ and for all representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$.

Proof. Let $k \in G_\Gamma$ so that $[k\Omega] = [h\Omega]$ and $[k\Lambda] = [g\Lambda]$. Without loss of generality, we can assume that $x \notin g\langle\Lambda\rangle$.

Suppose that we have $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x)) \neq \mathfrak{g}_{h\Omega}(x)$. Then there is a hyperplane H separating $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x))$ and $\mathfrak{g}_{h\Omega}(x)$. By Proposition 2.21, H also separates $\mathfrak{g}_{g\Lambda}(x)$ and x and cannot cross $g\langle\Lambda\rangle$. However, we know that H crosses $h\langle\Omega\rangle \subseteq h\langle\Lambda\rangle$ and by parallelism (Proposition 3.2) H must also cross $k\langle\Omega\rangle \subseteq k\langle\Lambda\rangle$. But $k\Lambda \| g\Lambda$, so H must also cross $g\langle\Lambda\rangle$. This contradiction means that we must have $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x)) = \mathfrak{g}_{h\Omega}(x)$. ■

Let γ be a $C(g\Lambda)$-geodesic from $\pi_{g\Lambda}(x)$ to $\pi_{g\Lambda}(y)$ and let $p_1, \dots, p_n \in \langle\Lambda\rangle$ so that $\pi_{g\Lambda}(x) = gp_1, gp_2, \dots, gp_n = \pi_{g\Lambda}(y)$ are the vertices of γ. Let α_i be an $S(g\Lambda)$-geodesic from gp_i to gp_{i+1} for each $i \in \{1, \dots, n-1\}$. Let γ' be the path in $S(g\Lambda)$ that is the union of all the α_i.

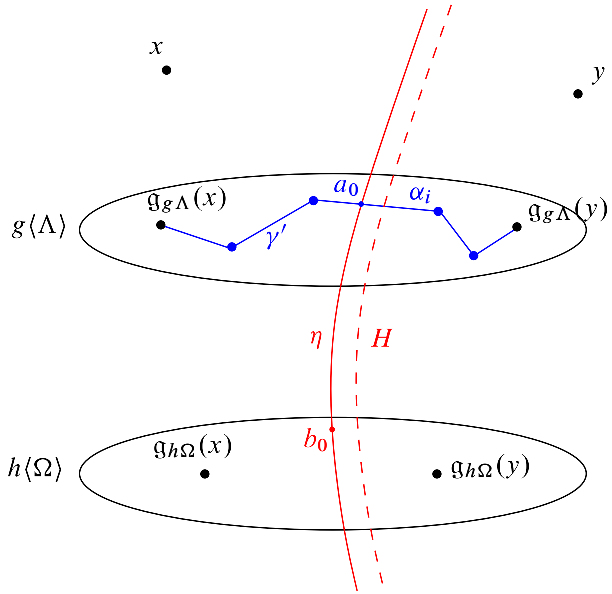

Figure 11. The $S(\Gamma)$-geodesic η connecting $b_0 \in h\langle\Omega\rangle$ and $a_0 \in \alpha_i$ when $d_{h\Omega}(x,y)$ is larger than 0.

Suppose that $\mathsf{d}_{h\Omega}(x,y) > 0$. Then $\mathsf{d}_{\mathrm{syl}}(\mathfrak{g}_{h\Omega}(x), \mathfrak{g}_{h\Omega}(y)) > 0$ and so there is a hyperplane H separating $\mathfrak{g}_{h\Omega}(x) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x))$ and $\mathfrak{g}_{h\Omega}(y) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(y))$ that is labelled by a vertex $w \in V(\Omega)$. The hyperplane H then also separates $\mathfrak{g}_{g\Lambda}(x)$ and $\mathfrak{g}_{g\Lambda}(y)$ by Proposition 2.21. Thus, H must cross one of the segments α_i that make up γ'. Since H crosses both $h\langle\Omega\rangle$ and α_i and H cannot separate $\mathfrak{g}_{g\Lambda}(x)$ from $\mathfrak{g}_{h\Omega}(x) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x))$ nor $\mathfrak{g}_{g\Lambda}(y)$ from $\mathfrak{g}_{h\Omega}(y) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(y))$ (Proposition 2.21 (4)), there exists an $S(\Gamma)$-geodesic, η, from an element $b_0 \in h\langle\Omega\rangle$ to $a_0 \in \alpha_i$ that is labelled by vertices in $\mathrm{lk}(w)$; see Figure 11.

Let $a_1 = \pi_{g\Lambda}(a_0)$ and $b_1 = \pi_{g\Lambda}(b_0)$. Since η was labelled by vertices in $\mathrm{lk}(w)$, Proposition 3.11 tells us we have $\mathrm{supp}(a_1^{-1}b_1) \subseteq \mathrm{lk}(w) \cap \Lambda$, which is a proper subgraph of Λ. Thus, in the subgraph metric, $\mathsf{d}_{g\Lambda}(\alpha_i, \rho^{h\Omega}_{g\Lambda}) \leq 1$ as $a_1 \in \pi_{g\Lambda}(\alpha_i)$ and $b_1 \in \pi_{g\Lambda}(h\langle\Omega\rangle) \subseteq \rho^{h\Omega}_{g\Lambda}$. As α_i is labelled by a proper subgraph of Λ, any subsegment is also labelled by a proper subgraph, hence $\mathsf{d}_{g\Lambda}(ga, gp_{i+1}) \leq 1$ for any vertex ga of α_i. Thus, $\mathsf{d}_{g\Lambda}(ga, \gamma) \leq 1$ and therefore $\mathsf{d}_{g\Lambda}(\gamma, \rho^{h\Omega}_{g\Lambda}) \leq 2$. ∎

We can now use the bounded geodesic image axiom together with the following lemma to prove large links.

Lemma 4.14. *Let* $[g\Lambda], [h\Omega] \in \mathfrak{S}_\Gamma$*. For any representatives* $g\Lambda \in [g\Lambda]$ *and* $h\Omega \in [h\Omega]$*, if* $\mathrm{diam}(\pi_{g\Lambda}(h\langle\Omega\rangle)) > 2$*, then* $[g\Lambda] \sqsubseteq [h\Omega]$.

Proof. If $[g\Lambda] \pitchfork [h\Omega]$ or $[h\Omega] \subsetneq [g\Lambda]$, then $\pi_{g\Lambda}(h\langle\Omega\rangle) \subseteq \rho^{h\Omega}_{g\Lambda}$, which is shown to have diameter at most 2 in Lemmas 3.18 and 3.22. If $[g\Lambda] \perp [h\Omega]$, then $\Lambda \subseteq \mathrm{lk}(\Omega)$. Let $\omega \in \langle\Omega\rangle$.

Then $\mathfrak{g}_{g\Lambda}(h\omega) = g \cdot \mathrm{prefix}_\Lambda(g^{-1}h\omega)$. Assume without loss of generality that

$$\mathrm{suffix}_\Lambda(g) = e \quad \text{and} \quad \mathrm{suffix}_\Omega(h) = e.$$

By Proposition 3.11, all syllables of $\mathrm{prefix}_\Lambda(g^{-1}h\omega)$ are syllables of $h\omega$. Further, since $\Lambda \subseteq \mathrm{lk}(\Omega)$, we have $\mathrm{supp}(\omega) \cap \Lambda = \emptyset$. As $\mathrm{suffix}_\Omega(h) = e$, this implies that

$$\mathrm{prefix}_\Lambda(g^{-1}h\omega) = \mathrm{prefix}_\Lambda(g^{-1}h).$$

Thus, $\pi_{g\Lambda}(h\omega) = g \cdot \mathrm{prefix}_\Lambda(g^{-1}h)$ for all $\omega \in \langle\Omega\rangle$, and so $\mathrm{diam}(\pi_{g\Lambda}(h\langle\Omega\rangle)) = 0$. ■

Lemma 4.15 (Large links). *Let $x, y \in G_\Gamma$ and $n = \mathsf{d}_{k\Pi}(x, y)$, where $k \in G_\Gamma$ and $\Pi \subseteq \Gamma$. There exist $[h_1\Omega_1], \dots, [h_n\Omega_n] \in \mathfrak{S}_\Gamma$ each nested into $[k\Pi]$ so that for any $[g\Lambda] \in \mathfrak{S}_\Gamma$ with $[g\Lambda] \sqsubsetneq [k\Pi]$, if $\mathsf{d}_{g\Lambda}(x, y) > 18$ for some representative of $[g\Lambda]$, then $[g\Lambda] \sqsubseteq [h_i\Omega_i]$ for some $i \in \{1, \dots, n\}$.*

Proof. Let γ be a $C(k\Pi)$-geodesic connecting $\pi_{k\Pi}(x)$ and $\pi_{k\Pi}(y)$, let

$$\pi_{k\Pi}(x) = p_0, p_1, \dots, p_n = \pi_{k\Pi}(y)$$

be the vertices of γ, and let $\lambda_i = p_{i-1}^{-1} p_i$ for each $i \in \{1, \dots, n\}$. For $i \in \{1, \dots, n\}$, define T_i to be $p_{i-1} \cdot \langle \mathrm{supp}(\lambda_i)\rangle$. Note that $p_i \in T_i \cap T_{i+1}$, and $T_i \subseteq k\langle\Pi\rangle$ since $p_{i-1} \in k\langle\Pi\rangle$ and $\mathrm{supp}(\lambda_i) \subsetneq \Pi$. In particular, $[T_i] \sqsubsetneq [k\Pi]$. Note also that $\pi_{k\Pi}(T_i) = T_i$ is contained in the closed 1-neighbourhood of p_i in $C(k\Pi)$, because $\mathrm{supp}(\lambda_i)$ is a proper subgraph of Π.

Next, let $[g\Lambda] \in \mathfrak{S}_\Gamma$ with $[g\Lambda] \sqsubsetneq [k\Pi]$ and suppose that $\mathsf{d}_{g\Lambda}(x, y) > 18$ for some representative $g\Lambda \in [g\Lambda]$. We shall show that $[g\Lambda] \sqsubseteq [T_i]$ for some $i \in \{1, \dots, n\}$. Since we have established the bounded geodesic image axiom (Lemma 4.12), we have $\gamma \cap \mathcal{N}_2(\rho^{g\Lambda}_{k\Pi}) \neq \emptyset$, where $\mathcal{N}_r(A)$ is the closed r-neighbourhood of A in $C(k\Pi)$. Let j be the first number in $\{0, \dots, n\}$ so that $p_j \in \mathcal{N}_4(\rho^{g\Lambda}_{k\Pi})$, and recall that each $\pi_{k\Pi}(T_i) = T_i$ is contained in $\mathcal{N}_1(p_i)$ and $\mathrm{diam}(\rho^{g\Lambda}_{k\Pi}) \leq 2$ (Lemma 3.18). Therefore, if $1 \leq i \leq j$ or $i \geq j + 10$, then

$$\pi_{k\Pi}(T_i) \cap \mathcal{N}_2(\rho^{g\Lambda}_{k\Pi}) = \emptyset$$

and the bounded geodesic image axiom says that $\pi_{g\Lambda}(T_i)$ is a single point.

Since $T_{i-1} \cap T_i \neq \emptyset$ for $i \in \{2, \dots, n\}$ and $x \in T_1$, $y \in T_n$, we have

$$\pi_{g\Lambda}\Big(\bigcup_{i=1}^{j} T_i\Big) = \pi_{g\Lambda}(x) \quad \text{and} \quad \pi_{g\Lambda}\Big(\bigcup_{i=j+10}^{n} T_i\Big) = \pi_{g\Lambda}(y),$$

whenever $j > 0$ and $j + 9 < n$, respectively. This implies that

$$\mathsf{d}_{g\Lambda}(x, y) \leq \sum_{i=j+1}^{\min\{n, j+9\}} \mathrm{diam}\big(\pi_{g\Lambda}(T_i)\big).$$

Since $\mathsf{d}_{g\Lambda}(x, y) > 18$, there must exist $j_0 \in \{j + 1, \dots, \min\{n, j + 9\}\}$ so that $\mathrm{diam}(\pi_{g\Lambda}(T_{j_0})) > 2$. By Lemma 4.14, this implies that $[g\Lambda] \sqsubseteq [T_{j_0}]$. ■

4.5. Partial realisation

We now prove partial realisation, which roughly says that given a collection of pairwise orthogonal $[g_i\Lambda_i] \in \mathfrak{S}_\Gamma$, the hyperbolic spaces $C(g_i\Lambda_i)$ give a coordinate system for G_Γ.

We first prove that we can always represent n mutually orthogonal elements of $\mathfrak{S}_\Gamma$ by the same group element, and similarly for nesting chains. This allows us to simplify arguments involving three or more orthogonal domains by working within a fixed coset.

Proposition 4.16. *Let $[g_1\Lambda_1], \dots, [g_n\Lambda_n] \in \mathfrak{S}_\Gamma$. If either $[g_1\Lambda_1] \sqsubseteq \dots \sqsubseteq [g_n\Lambda_n]$ or $[g_1\Lambda_1], \dots, [g_n\Lambda_n]$ are pairwise orthogonal, then there exists $g \in G_\Gamma$ so that $[g\Lambda_i] = [g_i\Lambda_i]$ for all $i \in \{1, \dots, n\}$.*

Proof. We proceed by induction. The initial case $n = 2$ is true by definition. Suppose that the statement is true for all $n < m$, and consider $n = m$; that is, we have $[g_1\Lambda_1], \dots, [g_m\Lambda_m] \in \mathfrak{S}_\Gamma$ which are either pairwise orthogonal or nested. Then, in particular, $[g_1\Lambda_1], \dots, [g_{m-1}\Lambda_{m-1}]$ are pairwise orthogonal (respectively nested), hence there exists $g \in G_\Gamma$ such that $[g\Lambda_i] = [g_i\Lambda_i]$ for all $i < m$. Since $[g\Lambda_i] = [g_i\Lambda_i]$ if and only if $[\Lambda_i] = [g^{-1}g_i\Lambda_i]$, we can assume that $g = e$ without loss of generality. Then $[\Lambda_i] \perp [g_m\Lambda_m]$ (respectively, $[\Lambda_i] \sqsubseteq [g_m\Lambda_m]$) for each $i < m$, so for each $i < m$ there exists k_i such that $k_i \in \langle \mathrm{st}(\Lambda_i)\rangle$ and $g_m^{-1}k_i \in \langle \mathrm{st}(\Lambda_m)\rangle$. Let h be the shortest prefix of g_m such that $g_m^{-1}h \in \langle \mathrm{st}(\Lambda_m)\rangle$. Since $g_m^{-1}k_i \in \langle \mathrm{st}(\Lambda_m)\rangle$ for each $i \in \{1, \dots, m-1\}$, we know that $\mathrm{supp}(h) \subseteq \mathrm{supp}(k_i) \subseteq \mathrm{st}(\Lambda_i)$ for each $i < m$. Hence $[\Lambda_i] = [h\Lambda_i]$ for each $i < m$ and $[g_m\Lambda_m] = [h\Lambda_m]$. Thus, by induction the statement is true for all n. ∎

Lemma 4.17 (Partial realisation). *Let $\{[g_i\Lambda_i]\}_{i=1}^n$ be a finite collection of pairwise orthogonal elements of $\mathfrak{S}_\Gamma$. For each $i \in \{1, \dots, n\}$, fix a choice of representative $g_i\Lambda_i$ for $[g_i\Lambda_i]$ and let $p_i \in C(g_i\Lambda_i)$. There exists $x \in G_\Gamma$ so that*

- *$\mathsf{d}_{g_i\Lambda_i}(x, p_i) = 0$ for all i;*
- *for each i and each $[h\Omega] \in \mathfrak{S}_\Gamma$, if $[g_i\Lambda_i] \sqsubsetneq [h\Omega]$ or $[h\Omega] \pitchfork [g_i\Lambda_i]$, then for any choice of representative $h\Omega \in [h\Omega]$ we have $\mathsf{d}_{h\Omega}(x, \rho^{g_i\Lambda_i}_{h\Omega}) = 0$.*

Proof. By Proposition 4.16, there exists some $g \in G_\Gamma$ such that $[g_i\Lambda_i] = [g\Lambda_i]$ for all i. Define $p_i' = \mathfrak{g}_{g\Lambda_i}(p_i) = g\lambda_i$, where $\lambda_i \in \langle\Lambda_i\rangle$, and let $x = g\lambda_1\lambda_2 \dots \lambda_n$. Then $\pi_{g\Lambda_i}(x) = g \cdot \mathrm{prefix}_{\Lambda_i}(g^{-1}x) = g\lambda_i = \pi_{g\Lambda_i}(p_i)$ for each i, since orthogonality tells us that the elements λ_i all commute with each other and the subgraphs Λ_i are all disjoint. Therefore, $\mathsf{d}_{g\Lambda_i}(x, p_i) = 0$ for all i, and so by Lemma 3.8, we have $\mathsf{d}_{g_i\Lambda_i}(x, p_i) = 0$ for all i.

Now suppose that $[g\Lambda_i] \sqsubsetneq [h\Omega]$ or $[g\Lambda_i] \pitchfork [h\Omega]$ for some $i \in \{1, \dots, n\}$ and $[h\Omega] \in \mathfrak{S}_\Gamma$. Since $\Lambda_j \subseteq \mathrm{lk}(\Lambda_i) \subseteq \mathrm{st}(\Lambda_i)$ for each $j \neq i$, we have $x = g\lambda_1 \dots \lambda_n \in g\langle \mathrm{st}(\Lambda_i)\rangle$. Thus, $\pi_{h\Omega}(x) \in \pi_{h\Omega}(g\langle \mathrm{st}(\Lambda_i)\rangle) = \rho^{g\Lambda_i}_{h\Omega}$ for any choice of representative $h\Omega$ of $[h\Omega]$. Moreover, we have

$$\rho^{g\Lambda_i}_{h\Omega} = \bigcup_{k\Lambda_i \| g\Lambda_i} \pi_{h\Omega}\big(k\langle\Lambda_i\rangle\big) = \rho^{g_i\Lambda_i}_{h\Omega},$$

since $g_i\Lambda_i \| g\Lambda_i$. This implies that $\mathsf{d}_{h\Omega}(x, \rho^{g_i\Lambda_i}_{h\Omega}) = 0$. ∎

4.6. Consistency

Finally, we prove consistency, which says that given two transverse domains $[g\Lambda]$ and $[h\Omega]$ in $\mathfrak{S}_\Gamma$, each element of G_Γ projects uniformly close to one of the lateral relative projections $\rho^{g\Lambda}_{h\Omega}$ and $\rho^{h\Omega}_{g\Lambda}$.

Our proof shall proceed by contradiction. Assuming that each element of G_Γ projects far from both lateral projections, we can use Lemma 4.14 to show that $[g\Lambda] \sqsubseteq [h\,\mathrm{lk}(w)]$ for each vertex w of Ω, which will imply that $[g\Lambda]\perp[hw]$ for each vertex w of Ω. We then obtain $[g\Lambda]\perp[h\Omega]$ by adapting the proof of Proposition 4.16 to show that we may promote orthogonality with multiple domains to orthogonality with their union. This contradicts $[g\Lambda] \pitchfork [h\Omega]$.

Lemma 4.18. *Let $[g\Lambda_1], \dots, [g\Lambda_{n-1}], [k\Lambda_n] \in \mathfrak{S}_\Gamma$. If $[g\Lambda_i]\perp[k\Lambda_n]$ for all $i < n$, then $[g\bigcup_{i<n}\Lambda_i]\perp[k\Lambda_n]$.*

Proof. Since $[g\Lambda_i]\perp[k\Lambda_n]$ if and only if $[\Lambda_i]\perp[g^{-1}k\Lambda_n]$, we may assume that $g = e$. By orthogonality, for each $i < n$ there exists k_i such that $k_i \in \langle \mathrm{st}(\Lambda_i)\rangle$ and $k^{-1}k_i \in \langle \mathrm{st}(\Lambda_n)\rangle$. Following the proof of Proposition 4.16, let h be the shortest prefix of k such that $k^{-1}h \in \langle \mathrm{st}(\Lambda_n)\rangle$. Then $\mathrm{supp}(h) \subseteq \mathrm{supp}(k_i) \subseteq \mathrm{st}(\Lambda_i)$ for all $i < n$, so $h \in \langle \mathrm{st}(\Lambda_i)\rangle$ for all $i < n$. Therefore, $h \in \langle \bigcap_{i<n} \mathrm{st}(\Lambda_i)\rangle \subseteq \langle \mathrm{st}(\bigcup_{i<n}\Lambda_i)\rangle$, hence $[\bigcup_{i<n}\Lambda_i] = [h\bigcup_{i<n}\Lambda_i]$ and $[k\Lambda_n] = [h\Lambda_n]$. Moreover, by orthogonality, $\Lambda_n \subseteq \mathrm{lk}(\Lambda_i)$ for all $i < n$, hence $\Lambda_n \subseteq \bigcap_{i<n}\mathrm{lk}(\Lambda_i) = \mathrm{lk}(\bigcup_{i<n}\Lambda_i)$. We therefore have $[\bigcup_{i<n}\Lambda_i]\perp[k\Lambda_n]$. ∎

Lemma 4.19 (Consistency). *If $[g\Lambda] \pitchfork [h\Omega]$, then for all $x \in G_\Gamma$ and for any choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$ we have*

$$\min\left\{\mathsf{d}_{h\Omega}\left(\pi_{h\Omega}(x), \rho^{g\Lambda}_{h\Omega}\right), \mathsf{d}_{g\Lambda}\left(\pi_{g\Lambda}(x), \rho^{h\Omega}_{g\Lambda}\right)\right\} \le 2. \tag{$*$}$$

Further, if $[k\Pi] \subsetneq [g\Lambda]$ and either $[g\Lambda] \subsetneq [h\Omega]$ or $[g\Lambda] \pitchfork [h\Omega]$ and $[h\Omega] \not\perp [k\Pi]$, then $\mathsf{d}_{h\Omega}(\rho^{k\Pi}_{h\Omega}, \rho^{g\Lambda}_{h\Omega}) = 0$.

Proof. We prove ($*$) by contradiction. Suppose that

$$\mathsf{d}_{h\Omega}\left(\pi_{h\Omega}(x), \rho^{g\Lambda}_{h\Omega}\right) > 2 \quad \text{and} \quad \mathsf{d}_{g\Lambda}\left(\pi_{g\Lambda}(x), \rho^{h\Omega}_{g\Lambda}\right) > 2.$$

Then we also have

$$\mathsf{d}_{\mathrm{syl}}\left(\mathfrak{g}_{h\Omega}(x), \mathfrak{g}_{h\Omega}\left(g\langle\Lambda\rangle\right)\right) > 2 \quad \text{and} \quad \mathsf{d}_{\mathrm{syl}}\left(\mathfrak{g}_{g\Lambda}(x), \mathfrak{g}_{g\Lambda}\left(h\langle\Omega\rangle\right)\right) > 2.$$

Thus, $\mathfrak{g}_{h\Omega}(x)$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$ are separated by some hyperplane H_w labelled by a vertex w of Ω. By Proposition 2.21 (5), H_w also separates x and $g\langle\Lambda\rangle$. In particular, H_w crosses any $S(\Gamma)$-geodesic segment γ connecting x and $g\langle\Lambda\rangle$. Because of Proposition 2.21 (4), H_w cannot separate $g\langle\Lambda\rangle$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$ as H_w crosses $h\langle\Omega\rangle$. Thus, there exists a combinatorial hyperplane of H_w contained in the same component of $S(\Gamma) \smallsetminus H_w$ as both $g\langle\Lambda\rangle$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$. Let H'_w be this particular combinatorial hyperplane of H_w; see Figure 12.

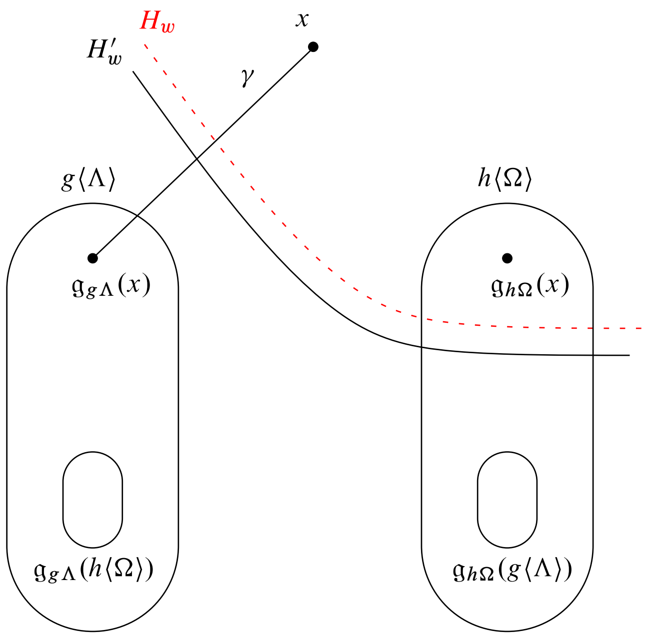

Figure 12. The combinatorial hyperplane H'_w of H_w that is in the same component of $S(\Gamma) \smallsetminus H_w$ as both $g\langle\Lambda\rangle$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$.

We claim that $\operatorname{diam}(\pi_{g\Lambda}(H'_w)) > 2$. By construction, H'_w contains both a vertex of $h\langle\Omega\rangle$ and a vertex of γ. Thus, $\pi_{g\Lambda}(H'_w)$ contains points from both $\pi_{g\Lambda}(h\langle\Omega\rangle)$ and $\pi_{g\Lambda}(\gamma)$. Since $\mathfrak{g}_{g\Lambda}(x)$ is the unique point in $g\langle\Lambda\rangle$ that minimises the $S(\Gamma)$-distance from x to $g\langle\Lambda\rangle$, we have $\mathfrak{g}_{g\Lambda}(\gamma) = \mathfrak{g}_{g\Lambda}(x) \in \pi_{g\Lambda}(H'_w)$. Since

$$\mathsf{d}_{g\Lambda}\big(\pi_{g\Lambda}(x), \pi_{g\Lambda}\big(h\langle\operatorname{st}(\Omega)\rangle\big)\big) = \mathsf{d}_{g\Lambda}\big(\pi_{g\Lambda}(x), \rho^{h\Omega}_{g\Lambda}\big) > 2,$$

and $\pi_{g\Lambda}(H'_w)$ must contain points from both $\pi_{g\Lambda}(x)$ and $\pi_{g\Lambda}(h\langle\Omega\rangle)$, we must have $\operatorname{diam}(\pi_{g\Lambda}(H'_w)) > 2$.

By Remark 2.18, $H'_w \subseteq h\langle\operatorname{lk}(w)\rangle$. Thus, $\operatorname{diam}(\pi_{g\Lambda}(H'_w)) > 2$ implies that

$$\operatorname{diam}\big(\pi_{g\Lambda}\big(h\langle\operatorname{lk}(w)\rangle\big)\big) > 2.$$

Lemma 4.14 then forces

$$[g\Lambda] \sqsubseteq \big[h\operatorname{lk}(w)\big] \sqsubseteq \big[h\operatorname{st}(w)\big].$$

This implies that $\Lambda \subseteq \operatorname{lk}(w)$ and that there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\operatorname{st}(w)] = [h\operatorname{st}(w)]$. Since $\operatorname{st}(\operatorname{st}(w)) = \operatorname{st}(w)$, $[k\operatorname{st}(w)] = [h\operatorname{st}(w)]$ implies that $[kw] = [hw]$. Thus, $[g\Lambda] = [k\Lambda] \perp [kw] = [hw]$. Moreover, since $\mathsf{d}_{h\Omega}(\pi_{h\Omega}(x), \rho^{g\Lambda}_{h\Omega}) > 2$, every vertex of Ω must appear as an edge label for the $S(h\Omega)$-geodesic connecting $\mathfrak{g}_{h\Omega}(x)$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$. Therefore, such a hyperplane H_w exists for every vertex w of Ω, and so

$[g\Lambda] \perp [hw]$ for all $w \in V(\Omega)$. Lemma 4.18 then tells us that $[g\Lambda] \perp [h\Omega]$, contradicting transversality. Hence inequality $(*)$ must hold.

Now suppose that $[k\Pi] \subsetneq [g\Lambda]$ and either $[g\Lambda] \subsetneq [h\Omega]$ or $[g\Lambda] \pitchfork [h\Omega]$ and $[h\Omega] \not\perp [k\Pi]$. Then there exists some element a such that $[k\Pi] = [a\Pi]$ and $[g\Lambda] = [a\Lambda]$. Therefore, $\pi_{h\Omega}(a\langle\Pi\rangle) \subseteq \rho^{k\Pi}_{h\Omega}$ and $\pi_{h\Omega}(a\langle\Lambda\rangle) \subseteq \rho^{g\Lambda}_{h\Omega}$. But $a\langle\Pi\rangle \subseteq a\langle\Lambda\rangle$, so $\mathrm{d}_{h\Omega}(\rho^{k\Pi}_{h\Omega}, \rho^{g\Lambda}_{h\Omega}) = 0$. ■

4.7. Compatibility of the group structure

The results so far show that a graph product G_Γ can be given the structure of a relative HHS. It remains to show that this structure agrees with the group structure of G_Γ.

Lemma 4.20. *The map $\phi : G_\Gamma \times \mathfrak{S}_\Gamma \to \mathfrak{S}_\Gamma$, where $\phi(a, [g\Lambda]) = [ag\Lambda]$, defines a $\sqsubseteq$-, $\perp$-, and $\pitchfork$-preserving action of G_Γ on $\mathfrak{S}_\Gamma$ by bijections such that $\mathfrak{S}_\Gamma$ contains finitely many G_Γ-orbits.*

Proof. Let $\phi_a = \phi(a, \cdot)$. This is well defined, since $[g\Lambda] = [k\Lambda]$ if and only if $[ag\Lambda] = [ak\Lambda]$. Further, since ϕ_a does not alter the subgraph Λ, it preserves the orthogonality, nesting, and transversality relations. Each ϕ_a is also a bijection: if $[ag\Lambda] = [ah\Omega]$, then $\Lambda = \Omega$ and $(ag)^{-1}(ah) = g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$, hence $[g\Lambda] = [h\Omega]$, proving injectivity. Surjectivity holds since we can always write $[g\Lambda] = \phi_a([a^{-1}g\Lambda])$. Finally, there are finitely many G_Γ-orbits; one for each subgraph $\Lambda \subseteq \Gamma$. ■

Lemma 4.21. *For each subgraph $\Lambda \subseteq \Gamma$ and elements $a, g \in G_\Gamma$, there exists an isometry $a_{g\Lambda}: C(g\Lambda) \to C(ag\Lambda)$ satisfying the following for all subgraphs $\Lambda, \Omega \subseteq \Gamma$ and elements $a, b, g, h \in G_\Gamma$.*

- *The isometry $(ab)_{g\Lambda}: C(g\Lambda) \to C(abg\Lambda)$ is equal to the composition*
$$a_{bg\Lambda} \circ b_{g\Lambda}: C(g\Lambda) \to C(abg\Lambda).$$
- *For each $x \in G_\Gamma$, we have $a_{g\Lambda}(\pi_{g\Lambda}(x)) = \pi_{ag\Lambda}(ax)$.*
- *If $[h\Omega] \pitchfork [g\Lambda]$ or $[h\Omega] \subsetneq [g\Lambda]$, then $a_{g\Lambda}(\rho^{h\Omega}_{g\Lambda}) = \rho^{ah\Omega}_{ag\Lambda}$.*

Proof. Let the isometry $a_{g\Lambda}$ be left-multiplication by a; that is, for any $gx \in C(g\Lambda)$, let $a_{g\Lambda}(gx) = agx$. Then

- the equality $(ab)_{g\Lambda} = a_{bg\Lambda} \circ b_{g\Lambda}$ is immediate from our definition;
- we have $a_{g\Lambda}(\pi_{g\Lambda}(x)) = \pi_{ag\Lambda}(ax)$ by Proposition 2.21 (2);
- the final property follows as an immediate consequence of the previous one and the definition of the relative projections. ■

4.8. Graph products are relative HHGs

We now compile the results from Section 4 to obtain the main result of this paper, that any graph product of finitely generated groups is a relative HHG.

Theorem 4.22. *Let G_Γ be a graph product of finitely generated groups. The proto-hierarchy structure $\mathfrak{S}_\Gamma$ from Theorem 3.23 is a relative HHG structure for G_Γ with hierarchy constant* $\max\{18, |V(\Gamma)|\}$ *and uniqueness function*

$$\theta(r) = (2^{|V(\Gamma)|}r + 2)^{|V(\Gamma)|}.$$

Proof. Let $\mathfrak{S}_\Gamma$ be the proto-hierarchy structure for (G_Γ, d) from Theorem 3.23. The work of this section has shown that $\mathfrak{S}_\Gamma$ is a relative HHS structure for (G_Γ, d).

(1) We proved that the spaces associated to the non-$\sqsubseteq$-minimal domains of $\mathfrak{S}_\Gamma$ are $\frac{7}{2}$-hyperbolic in Lemma 4.1.

(2) We proved finite complexity in Lemma 4.5.

(3) We proved the container axiom in Lemma 4.6.

(4) The proof of the uniqueness axiom follows from Lemma 4.7, since if $\mathsf{d}_{C([g\Lambda])}(x, y)$ is uniformly bounded for all $[g\Lambda] \in \mathfrak{S}_\Gamma$, then Lemma 3.8 implies that $\mathsf{d}_{g\Lambda}(x, y)$ has the same uniform bound for all $g \in G_\Gamma$ and $\Lambda \subseteq \Gamma$.

(5) We proved the bounded geodesic image axiom in Lemma 4.12.

(6) We proved the large links axiom in Lemma 4.15.

(7) We proved the consistency axiom in Lemma 4.19.

(8) We proved the partial realisation axiom in Lemma 4.17.

We now verify the remaining axioms required for (G_Γ, d) to be a relative HHG, as laid out in Definition 2.27.

Let $\phi : G_\Gamma \times \mathfrak{S}_\Gamma \to \mathfrak{S}_\Gamma$ be the map $\phi(a, [g\Lambda]) = [ag\Lambda]$. By Lemma 4.20, this is a well-defined G_Γ-action by bijections that preserves the nesting, orthogonality, and transversality relations and has finitely many orbits.

For each $[g\Lambda] \in \mathfrak{S}_\Gamma$, let $\bar{g}\Lambda$ denote the fixed representative of $[g\Lambda]$ such that $C([g\Lambda]) = C(\bar{g}\Lambda)$; see the proto-hierarchy structure in Theorem 3.23. Left multiplication by $a \in G_\Gamma$ gives an isometry $a_{g\Lambda}: C(g\Lambda) \to C(ag\Lambda)$ for each $g \in G_\Gamma$ and each subgraph $\Lambda \subseteq \Gamma$. For each $a \in G_\Gamma$ and $[g\Lambda] \in \mathfrak{S}_\Gamma$, define $\mathbf{a}_{[g\Lambda]}: C(\bar{g}\Lambda) \to C(\overline{ag}\Lambda)$ by $\mathbf{a}_{[g\Lambda]} = \mathfrak{g}_{\overline{ag}\Lambda} \circ a_{\bar{g}\Lambda}$.

Let $a, b \in G_\Gamma$ and $[g\Lambda], [h\Omega] \in \mathfrak{S}_\Gamma$. We now verify the remaining axioms of a relative HHG (Definition 2.27).

- Let $\lambda \in \langle\Lambda\rangle$. To show $(\mathbf{ab})_{[g\Lambda]} = \mathbf{a}_{[bg\Lambda]} \circ \mathbf{b}_{[g\Lambda]}$, we will show that

$$(\mathbf{ab})_{[g\Lambda]}(\bar{g}\lambda) = (\mathbf{a}_{[bg\Lambda]} \circ \mathbf{b}_{[g\Lambda]})(\bar{g}\lambda).$$

Using the last clause of Lemma 3.8, we have

$$(\mathbf{ab})_{[g\Lambda]}(\bar{g}\lambda) = \mathfrak{g}_{\overline{abg}\Lambda}(ab\bar{g}\lambda) = \overline{abg} \cdot p_{ab}\lambda,$$

where $p_{ab} = \mathrm{prefix}_\Lambda((\overline{abg})^{-1} \cdot ab\bar{g})$. Similarly, we have

$$(\mathbf{a}_{[bg\Lambda]} \circ \mathbf{b}_{[g\Lambda]})(\bar{g}\lambda) = \mathbf{a}_{[bg\Lambda]}(\overline{bg} \cdot p_b\lambda) = \overline{abg} \cdot p_a p_b \lambda,$$

where $p_b = \mathrm{prefix}_\Lambda((\overline{bg})^{-1} \cdot b\bar{g})$ and $p_a = \mathrm{prefix}_\Lambda((\overline{abg})^{-1} \cdot a\overline{bg})$. Thus, it suffices to prove that $p_a p_b = p_{ab}$.

Since $\overline{bg}$ and $b\bar{g}$ are both representatives of the parallelism class $[bg\Lambda]$, we have $(\overline{bg})^{-1} \cdot b\bar{g} \in \langle \mathrm{st}(\Lambda)\rangle$. Therefore, $(\overline{bg})^{-1} \cdot b\bar{g} = p_b l_b$, where $l_b \in \langle \mathrm{lk}(\Lambda)\rangle$. Similarly, $(\overline{abg})^{-1} \cdot a\overline{bg} = p_a l_a$, where $l_a \in \langle \mathrm{lk}(\Lambda)\rangle$. Hence the following calculation concludes our argument:

$$\begin{aligned}(\overline{abg})^{-1} \cdot ab\bar{g} &= (\overline{abg})^{-1} \cdot a\overline{bg} \cdot p_b l_b,\\ \mathrm{prefix}_\Lambda\left((\overline{abg})^{-1} \cdot ab\bar{g}\right) &= \mathrm{prefix}_\Lambda\left((\overline{abg})^{-1} \cdot a\overline{bg} \cdot p_b l_b\right),\\ p_{ab} &= \mathrm{prefix}_\Lambda(p_a l_a p_b l_b),\\ p_{ab} &= p_a p_b.\end{aligned}$$

- Let $x \in G_\Gamma$. Since $a\bar{g}\Lambda \| \overline{ag}\Lambda$, we can use Lemma 3.8 and the equivariance of the gate map (Proposition 2.21 (2)) to conclude that

$$\begin{aligned}\mathfrak{g}_{\overline{ag}\Lambda}\left(\mathfrak{g}_{a\bar{g}\Lambda}(ax)\right) &= \mathfrak{g}_{\overline{ag}\Lambda}(ax),\\ \mathfrak{g}_{\overline{ag}\Lambda}\left(a \cdot \mathfrak{g}_{\bar{g}\Lambda}(x)\right) &= \mathfrak{g}_{\overline{ag}\Lambda}(ax),\\ (\mathfrak{g}_{\overline{ag}\Lambda} \circ a_{\bar{g}\Lambda})\left(\pi_{\bar{g}\Lambda}(x)\right) &= \pi_{\overline{ag}\Lambda}(ax),\\ \mathbf{a}_{[g\Lambda]}\left(\pi_{[g\Lambda]}(x)\right) &= \pi_{[ag\Lambda]}(ax).\end{aligned}$$

- Suppose that $[h\Omega] \pitchfork [g\Lambda]$ or $[h\Omega] \subsetneq [g\Lambda]$. Lemmas 3.8, 4.20, and 4.21 imply that $\mathbf{a}_{[g\Lambda]}(\rho^{[h\Omega]}_{[g\Lambda]}) = \rho^{[ah\Omega]}_{[ag\Lambda]}$:

$$\begin{aligned}\mathbf{a}_{[g\Lambda]}(\rho^{[h\Omega]}_{[g\Lambda]}) &= (\mathfrak{g}_{\overline{ag}\Lambda} \circ a_{\bar{g}\Lambda})(\rho^{\bar{h}\Omega}_{\bar{g}\Lambda}) && \text{(definition of } \mathbf{a}_{[g\Lambda]})\\ &= \mathfrak{g}_{\overline{ag}\Lambda}(\rho^{a\bar{h}\Omega}_{a\bar{g}\Lambda}) && \text{(Lemma 4.21)}\\ &= \mathfrak{g}_{\overline{ag}\Lambda}\left(\mathfrak{g}_{a\bar{g}\Lambda}\left(a\bar{h}\langle \mathrm{st}(\Omega)\rangle\right)\right) && \text{(definition of } \rho)\\ &= \mathfrak{g}_{\overline{ag}\Lambda}\left(a\bar{h}\langle \mathrm{st}(\Omega)\rangle\right) && \text{(Lemma 3.8)}\\ &= \mathfrak{g}_{\overline{ag}\Lambda}\left(\overline{ah}\langle \mathrm{st}(\Omega)\rangle\right) && (a\bar{h}\Omega \| \overline{ah}\Omega)\\ &= \rho^{\overline{ah}\Omega}_{\overline{ag}\Lambda}.\end{aligned}$$

■

Behrstock, Hagen, and Sisto show that any relative HHS has a distance formula, which expresses distances in the space as a sum of distances in the projections [3, Theorem 6.10]. As a result, we now have such a distance formula for graph products of finitely generated groups.

Corollary 4.23 (Distance formula for graph products). *Let G_Γ be a graph product of finitely generated groups. There exists $\sigma_0 > 0$ such that for all $\sigma \geq \sigma_0$ there exist $K \geq 1$ and $L \geq 0$ such that for all $g, h \in G_\Gamma$*

$$\frac{1}{K} \sum_{[k\Lambda] \in \mathfrak{S}_\Gamma} \left\{\!\!\left\{ \mathsf{d}_{[k\Lambda]}(g,h) \right\}\!\!\right\}_\sigma - L \leq \mathsf{d}(g,h) \leq K \sum_{[k\Lambda] \in \mathfrak{S}_\Gamma} \left\{\!\!\left\{ \mathsf{d}_{[k\Lambda]}(g,h) \right\}\!\!\right\}_\sigma + L,$$

where we define $\{\!\{N\}\!\}_\sigma = N$ if $N \geq \sigma$ and 0 if $N < \sigma$.

Another key consequence of relative hierarchical hyperbolicity for a group is that the action of the group on the $\sqsubseteq$-maximal space is acylindrical. Thus, we have that the action of G_Γ on $C(\Gamma)$ is acylindrical. Recall that the action of a group G on a metric space X is *acylindrical* if for all $\varepsilon \geq 0$, there exist $R, N \geq 0$ so that if $x, y \in X$ satisfy $\mathsf{d}_X(x, y) \geq R$, then there are at most N elements $g \in G$ such that $\mathsf{d}_X(x, gx) \leq \varepsilon$ and $\mathsf{d}_X(y, gy) \leq \varepsilon$.

Corollary 4.24 (The action on $C(\Gamma)$ is acylindrical). *Let G_Γ be a graph product of finitely generated groups. The action of G_Γ on $C(\Gamma)$ by left multiplication is acylindrical.*

Proof. Behrstock, Hagen, and Sisto proved that if $(G, \mathfrak{S})$ is a (non-relative) HHG and $T \in \mathfrak{S}$ is the $\sqsubseteq$-maximal element, then the action of G on $C(T)$ is acylindrical [5, Theorem 14.3]. However, the argument they employ only uses the hyperbolicity of the space $C(T)$ and not the hyperbolicity of any of the other spaces in the HHG structure. Thus, their argument carries through verbatim if $(G, \mathfrak{S})$ is a relative HHG provided that $\mathfrak{S} \neq \{T\}$. In the case when $\mathfrak{S} = \{T\}$, then $C(T)$ is equivariantly quasi-isometric to a Cayley graph of G with respect to some finite generating set. Thus, G acts on $C(T)$ properly, and hence acylindrically. Applying this to the graph product G_Γ with relative HHG structure $\mathfrak{S}_\Gamma$, we have that G_Γ acts on $C(\Gamma)$ acylindrically. ∎

4.9. The syllable metric is an HHS

Since nearly every argument used in the proof of Theorem 4.22 factors through the syllable metric on the graph product G_Γ, the same arguments show that the syllable metric on G_Γ is itself an HHS. This proves Corollary B stated in the introduction and answers a question of Behrstock, Hagen, and Sisto about the syllable metric on a right-angled Artin group. Note that since we are not working with a word metric on G_Γ in this situation, we do not require the vertex groups to be finitely generated. As the only use of the finite generation of the vertex groups in Theorem 4.22 is to ensure that G_Γ has a word metric, this does not create any additional difficulty.

Corollary 4.25. *Let Γ be a finite simplicial graph, with each vertex v labelled by a group G_v. Then the graph product G_Γ endowed with the syllable metric is an HHS.*

Proof. Define the proto-hierarchy structure for G_Γ as before, except whenever $v \in V(\Gamma)$ and $g \in G_\Gamma$, and define $C(gv)$ to be the graph whose vertices are elements of gG_v and where every pair of vertices is joined by an edge (that is, we endow gG_v with the syllable metric rather than the word metric). The proofs of the HHG axioms then follow as before, with any instance of "word metric" replaced with "syllable metric", and with trivial $\sqsubseteq$-minimal case for the majority of axioms due to such $C(gv)$ having diameter 1. ∎

5. Some applications of hierarchical hyperbolicity

We now give some applications of the relative hierarchical hyperbolicity of graph products. Our main result of this section is Theorem 5.1, which shows that if the vertex groups of a graph product G_Γ are HHGs, then G_Γ is itself a (non-relative) HHG.

We then give a new proof of a theorem of Meier, classifying when a graph product G_Γ with hyperbolic vertex groups is itself hyperbolic. We do this using the relative HHS structure that we just obtained, noting that when the vertex groups are hyperbolic, this is in fact a (non-relative) HHS structure.

Finally, we answer two questions of Genevois regarding the *electrification* $\mathbb{E}(\Gamma)$ of a graph product G_Γ of finite groups [16, Questions 8.3 and 8.4]. The similarity of Genevois' definition of $\mathbb{E}(\Gamma)$ to our own subgraph metric $C(\Gamma)$ allows us to leverage properties of $C(\Gamma)$ to prove statements about $\mathbb{E}(\Gamma)$. In particular, we use Γ to classify when $\mathbb{E}(\Gamma)$ has bounded diameter (Theorem 5.14) and when it is a quasi-line (Theorem 5.16). As Genevois proved that any quasi-isometry between graph products of finite groups induces a quasi-isometry between their electrifications [16, Proposition 1.4], these two theorems provide us with tools for studying quasi-isometric rigidity of graph products of finite groups.

5.1. Graph products of HHGs

Theorem 5.1. *Let G_Γ be a graph product of finitely generated groups. If for each $v \in V(\Gamma)$, the vertex group G_v is an HHG, then G_Γ is an HHG.*

Proof. For each $v \in V(\Gamma)$, let $\mathfrak{R}_{[v]}$ be the HHG structure for G_v and let $\mathfrak{S}_\Gamma$ be the relative HHG structure for G_Γ coming from Theorem 4.22. Fix $E_0 > 0$ to be the maximum of the hierarchy constants for $\mathfrak{S}_\Gamma$ and for each $\mathfrak{R}_{[v]}$. For each $[g\Lambda] \in \mathfrak{S}_\Gamma$, let $\bar{g}\Lambda$ be the fixed representative of $[g\Lambda]$ so that $C([g\Lambda]) = C(\bar{g}\Lambda)$. If $[g\Lambda] = [\Lambda]$, then we choose $\bar{g} = e$.

Let $\mathfrak{S}_\Gamma^{\min} = \{[g\Lambda] \in \mathfrak{S}_\Gamma : \Lambda \text{ is a single vertex of } \Gamma\}$. If Λ is a single vertex v of Γ, then $C([v])$ is the Cayley graph of the vertex group G_v with respect to a finite generating set. Thus, $\mathfrak{R}_{[v]}$ is an HHG structure for $C([v])$. For each $[gv] \in \mathfrak{S}_\Gamma^{\min}$, $\mathfrak{R}_{[v]}$ is also an E_0-HHS structure for $C([gv])$, since $C([gv])$ is isometric to $C([v])$. Let $\mathfrak{R}_{[gv]}$ denote the HHS structure for $C([gv])$ induced by $\mathfrak{R}_{[v]}$. If $U \in \mathfrak{R}_{[v]}$, then we will denote the corresponding element of $\mathfrak{R}_{[gv]}$ by $\bar{g}U$, where $\bar{g}$ is the chosen fixed representative of $[gv]$. Let $\overline{\mathfrak{R}} = \bigcup_{[gv]\in\mathfrak{S}_\Gamma^{\min}} \mathfrak{R}_{[gv]}$, then let $\mathfrak{T}_0 = (\mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}) \cup \overline{\mathfrak{R}}$.

We shall use $\sqsubseteq_\mathfrak{S}$, $\perp_\mathfrak{S}$, and $\pitchfork_\mathfrak{S}$ to denote the nesting, orthogonality, and transversality relations between elements of $\mathfrak{S}_\Gamma$, and $\sqsubseteq_\mathfrak{R}$, $\perp_\mathfrak{R}$, and $\pitchfork_\mathfrak{R}$ to denote the relations between elements of a fixed $\mathfrak{R}_{[gv]}$.

The bulk of our proof of Theorem 5.1 does not use the specifics of the relative HHG structure $\mathfrak{S}_\Gamma$ and instead relies on more general relative HHS properties. Thus, to simplify notation, we will use the capital letters V or V' to denote elements of $\mathfrak{S}_\Gamma^{\min}$ and use $\mathfrak{R}_V$ or $\mathfrak{R}_{V'}$ to denote the corresponding HHS structure on $C(V)$ or $C(V')$. That is, if $V = [gv]$ for a vertex $v \in V(\Gamma)$, then $\mathfrak{R}_V = \mathfrak{R}_{[gv]}$. We will use the capital letters U, W, and Q to denote elements of $\mathfrak{T}_0$. For $U, W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ or $U, W \in \mathfrak{R}_V$ we shall denote the relative projection from U to W in $\mathfrak{S}_\Gamma$ or $\mathfrak{R}_V$ as ρ_W^U. We shall use π_W to denote the projection $G_\Gamma \to 2^{C(W)}$ if $W \in \mathfrak{S}_\Gamma$ and π_W^V to denote the projection $C(V) \to 2^{C(W)}$ if $W \in \mathfrak{R}_V$.

Our proof of Theorem 5.1 proceeds via four claims. First, we prove that the structure $\mathfrak{S}_\Gamma$ can be combined with all of the $\mathfrak{R}_V$ structures in a natural way to produce a proto-hierarchy structure for G_Γ with index set $\mathfrak{T}_0$ (Claim 5.2). This proto-hierarchy structure is not quite an HHS structure, as it satisfies every axiom except the container axiom (Claim 5.3). However, we show that this proto-hierarchy structure has the property that any set of pairwise orthogonal elements of $\mathfrak{T}_0$ has uniformly bounded cardinality (Claim 5.4). This allows us to use the results of the appendix of [1] to upgrade $\mathfrak{T}_0$ to a genuine HHS structure $\mathfrak{T}$. Since the proto-structure will satisfy the equivariance properties of an HHG structure for G_Γ (Claim 5.6), this HHS structure will also be an HHG structure.

Claim 5.2. G_Γ admits an E_1-proto-hierarchy structure with index set $\mathfrak{T}_0$, where $E_1 = E_0^2 + E_0$.

Proof. For $U \in \mathfrak{T}_0$, the associated hyperbolic space $C(U)$ will be the same as the space associated to U in either $\mathfrak{S}_\Gamma$ or $\bar{\mathfrak{R}}$.

Projections. For all $W \in \mathfrak{T}_0$, the projection map will be denoted by $\psi_W : G_\Gamma \to 2^{C(W)}$. If $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$, then $\psi_W = \pi_W$ and if $W \in \mathfrak{R}_V$, then $\psi_W = \pi_W^V \circ \pi_V$. Each ψ_W is $(E_0^2, E_0^2 + E_0)$-coarsely Lipschitz.

Nesting. Let $W, U \in \mathfrak{T}_0$. We define $U \sqsubseteq W$ if one of the following holds:

- $W, U \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \sqsubseteq_{\mathfrak{S}} W$;
- $W, U \in \mathfrak{R}_V$ and $U \sqsubseteq_{\mathfrak{R}} W$;
- $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \in \mathfrak{R}_V$ with $V \sqsubseteq_{\mathfrak{S}} W$.

This definition makes $[\Gamma]$, the $\sqsubseteq_{\mathfrak{S}}$-maximal element of $\mathfrak{S}_\Gamma$, also the $\sqsubseteq$-maximal element of $\mathfrak{T}_0$. For $U, W \in \mathfrak{T}_0$ with $U \sqsubsetneq W$ we denote the relative projection from U to W by β_W^U and define it as follows.

- If $W, U \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \sqsubseteq_{\mathfrak{S}} W$ or $W, U \in \mathfrak{R}_V$ and $U \sqsubseteq_{\mathfrak{R}} W$, then β_W^U is ρ_W^U, the relative projection from U to W in $\mathfrak{S}_\Gamma$ or $\mathfrak{R}_V$ respectively.
- If $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \in \mathfrak{R}_V$ with $V \sqsubseteq_{\mathfrak{S}} W$, then β_W^U is ρ_W^V, the relative projection from V to W in $\mathfrak{S}_\Gamma$.

The diameter of β_W^U is bounded by E_0 in all cases as it always coincides with a relative projection (ρ_W^U or ρ_W^V) from an existing hierarchy structure with constant E_0.

Orthogonality. Let $W, U \in \mathfrak{T}_0$. We define $U \perp W$ if one of the following holds:

- $W, U \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \perp_{\mathfrak{S}} W$;
- $W, U \in \mathfrak{R}_V$ and $U \perp_{\mathfrak{R}} W$;
- $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \in \mathfrak{R}_V$ with $V \perp_{\mathfrak{S}} W$;
- $W \in \mathfrak{R}_{V'}$ and $U \in \mathfrak{R}_V$ where $V \perp_{\mathfrak{S}} V'$.

Transversality. Let $U, W \in \mathfrak{T}_0$. We define $U \pitchfork W$ whenever they are not orthogonal or nested in $\mathfrak{T}_0$. This arises in three different situations, which determine the definition of the relative projections β_U^W and β_W^U.

- Either $U, W \in \mathfrak{S}_\Gamma$ or $U, W \in \mathfrak{R}_V$ and $U \pitchfork_{\mathfrak{S}} W$ or $U \pitchfork_{\mathfrak{R}} W$, respectively. In this case, β_W^U is ρ_W^U, the relative projection from U to W in $\mathfrak{S}_\Gamma$ or $\mathfrak{R}_V$, respectively, and β_U^W is ρ_U^W.
- $W \in \mathfrak{S}_\Gamma$ and $U \in \mathfrak{R}_V$ where $W \pitchfork_{\mathfrak{S}} V$. In this case, β_W^U is ρ_W^V, the relative projection from V to W in $\mathfrak{S}_\Gamma$, and $\beta_U^W = \pi_U^V(\rho_V^W)$.
- $W \in \mathfrak{R}_{V'}$ and $U \in \mathfrak{R}_V$ where $V \pitchfork_{\mathfrak{S}} V'$. In this case, $\beta_W^U = \pi_W^{V'}(\rho_{V'}^V)$ and $\beta_U^W = \pi_U^V(\rho_V^{V'})$.

The projection and transversality axioms of $\mathfrak{R}_V$ and $\mathfrak{S}_\Gamma$ ensure that β_W^U has diameter at most $E_0^2 + E_0$ in all cases. ■

Claim 5.3. $\mathfrak{T}_0$ satisfies all of the axioms of an HHS except for the container axiom.

Proof. Recall that $E_1 > 0$ is the hierarchy constant from the proto-hierarchy structure $\mathfrak{T}_0$. Note that E_1 is larger than E_0, which in turn is larger than the hierarchy constants for $\mathfrak{S}_\Gamma$ and each $\mathfrak{R}_V$.

Hyperbolicity. For all $W \in \mathfrak{T}_0$, the space $C(W)$ is E_1-hyperbolic.

Uniqueness. Let $\kappa \geq 0$ and $\theta\colon [0,\infty) \to [0,\infty)$ be the maximum of the uniqueness functions for $\mathfrak{S}_\Gamma$ and each $\mathfrak{R}_V$. If $x, y \in G_\Gamma$ and $\mathsf{d}(x,y) \geq \theta(\theta(\kappa)+\kappa)$, then there exists $W \in \mathfrak{S}_\Gamma$ such that $\mathsf{d}_W(x,y) \geq \theta(\kappa)+\kappa$ by the uniqueness axiom in $(G_\Gamma, \mathfrak{S}_\Gamma)$. If $W \notin \mathfrak{S}_\Gamma^{\min}$, then W is in $\mathfrak{T}_0$ and the uniqueness axiom is satisfied. If $W \in \mathfrak{S}_\Gamma^{\min}$, then the uniqueness axiom in $(C(W), \mathfrak{R}_W)$ provides $U \in \mathfrak{R}_W$ so that $\mathsf{d}_U(x,y) \geq \kappa$. The uniqueness function for $(G_\Gamma, \mathfrak{T}_0)$ is therefore $\phi(\kappa) = \theta(\theta(\kappa)+\kappa)$.

Finite complexity. The length of a $\sqsubseteq$-chain in $\mathfrak{T}_0$ is at most $2E_1$.

Bounded geodesic image. Let $x, y \in G_\Gamma$ and $U, W \in \mathfrak{T}_0$ with $U \sqsubsetneq W$. If $U, W \in \mathfrak{S}_\Gamma$ or $U, W \in \mathfrak{R}_V$, then the bounded geodesic image axiom from $(G_\Gamma, \mathfrak{S}_\Gamma)$ or $(C(V), \mathfrak{R}_V)$ implies the bounded geodesic image axiom for $(G_\Gamma, \mathfrak{T}_0)$. Suppose, therefore, that $U \in \mathfrak{R}_V$ and $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$. By definition, $V \sqsubseteq_{\mathfrak{S}} W$ and β_W^U coincides with ρ_W^V, the relative projection of V to W in $\mathfrak{S}_\Gamma$. If $\mathsf{d}_U(x,y) > E_1^2 + E_1$, then we have

$$\begin{aligned} E_1^2 + E_1 < \mathsf{d}_U(x,y) &= \mathsf{d}_U\big(\pi_U^V\big(\pi_V(x)\big), \pi_U^V\big(\pi_V(y)\big)\big) \\ &\leq E_1 \mathsf{d}_V\big(\pi_V(x), \pi_V(y)\big) + E_1, \end{aligned}$$

which implies that $E_1 < \mathsf{d}_V(\pi_V(x), \pi_V(y))$. Now the bounded geodesic image axiom in $(G_\Gamma, \mathfrak{S}_\Gamma)$ says that every geodesic in $C(W)$ from $\psi_W(x) = \pi_W(x)$ to $\psi_W(y) = \pi_W(y)$ must pass through the E_1-neighbourhood of $\rho_W^V = \beta_W^U$. Thus, the bounded geodesic image axiom is satisfied for $(G_\Gamma, \mathfrak{T}_0)$.

Large links. Let $W \in \mathfrak{T}_0$ and $x, y \in G_\Gamma$. If $W \in \mathfrak{R}_V$ for some $V \in \mathfrak{S}_\Gamma^{\min}$, then all elements of $\mathfrak{T}_0$ that are nested into W are also elements of $\mathfrak{R}_V$. Thus, the large links axiom in $(C(V), \mathfrak{R}_V)$ immediately implies the large links axiom for $(G_\Gamma, \mathfrak{T}_0)$.

Assume that $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$. The large links axiom for $(G_\Gamma, \mathfrak{S}_\Gamma)$ gives a collection

$$\mathfrak{L} = \{U_1, \dots, U_m\}$$

of elements of $\mathfrak{S}_\Gamma$ nested into W such that m is at most $E_1 \mathsf{d}_W(\pi_W(x), \pi_W(y)) + E_1$, and for all $V \in \mathfrak{S}_W$, either $V \sqsubseteq_{\mathfrak{S}} U_i$ for some i or $\mathsf{d}_V(\pi_V(x), \pi_V(y)) < E_1$. For each $i \in \{1, \dots, m\}$, define $\bar{U}_i$ to be the $\sqsubseteq_{\mathfrak{R}}$-maximal element of $\mathfrak{R}_{U_i}$ if $U_i \in \mathfrak{S}_\Gamma^{\min}$ and define $\bar{U}_i$ to be U_i if $U_i \notin \mathfrak{S}_\Gamma^{\min}$. Let $\bar{\mathfrak{L}} = \{\overline{U_1}, \dots, \overline{U_m}\}$.

If $V \in \mathfrak{S}_\Gamma^{\min}$ is nested into W, but is not nested into an element of $\mathfrak{L}$, then

$$\mathsf{d}_V\big(\pi_V(x), \pi_V(y)\big) < E_1$$

and so

$$\mathsf{d}_Q\big(\psi_Q(x), \psi_Q(y)\big) < E_1^2 + E_1$$

for all $Q \in \mathfrak{R}_V$. Thus, if $\mathsf{d}_Q(\psi_Q(x), \psi_Q(y)) \geq E_1^2 + E_1$ and Q is nested into W, then either $Q \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ or $Q \in \mathfrak{R}_V$ where V is nested into an element of $\mathfrak{L}$ (and so Q is nested into an element of $\bar{\mathfrak{L}}$). If $Q \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$, then Q must be nested into an element of $\mathfrak{L}$ that is not in $\mathfrak{S}_\Gamma^{\min}$ by the large links axiom of $(G_\Gamma, \mathfrak{S}_\Gamma)$, and hence must be nested into an element of $\bar{\mathfrak{L}}$. Thus, $Q \sqsubseteq W$ is nested into an element of $\bar{\mathfrak{L}}$ whenever $\mathsf{d}_Q(\psi_Q(x), \psi_Q(y)) \geq E_1^2 + E_1$.

Consistency. Let $U, W \in \mathfrak{T}_0$ with $U \pitchfork W$ and $x \in G_\Gamma$. Since the relative projections are inherited from $\mathfrak{S}_\Gamma$ and the $\mathfrak{R}_V$, we only need to consider the case where either $W \in \mathfrak{S}_\Gamma$ and $U \in \mathfrak{R}_V$, or $W \in \mathfrak{R}_{V'}$ and $U \in \mathfrak{R}_V$ with $V' \neq V$. Define $Q = W$ if $W \in \mathfrak{S}_\Gamma$ and $Q = V'$ if $W \in \mathfrak{R}_{V'}$. In either case $Q \pitchfork_{\mathfrak{S}} V$.

First assume that $Q = W$ so that $\beta_W^U = \rho_Q^V$ and $\beta_U^W = \pi_U^V(\rho_V^Q)$. If $\mathsf{d}_W(x, \beta_W^U) = \mathsf{d}_Q(x, \rho_Q^V) > E_1$, then the consistency axiom for $(G_\Gamma, \mathfrak{S}_\Gamma)$ says that $\mathsf{d}_V(x, \rho_V^Q) \leq E_1$. The coarse Lipschitzness of the projections then implies that $\mathsf{d}_U(x, \pi_U^V(\rho_V^Q)) = \mathsf{d}_U(x, \beta_U^W) \leq E_1^2 + E_1$.

Now assume that $Q = V'$ so that $\beta_W^U = \pi_W^Q(\rho_Q^V)$ and $\beta_U^W = \pi_U^V(\rho_V^Q)$. If $\mathsf{d}_W(x, \beta_W^U) > E_1^2 + E_1$, then $\mathsf{d}_Q(x, \rho_Q^V) > E_1$. The consistency axiom for $(G_\Gamma, \mathfrak{S}_\Gamma)$ then says that $\mathsf{d}_V(x, \rho_V^Q) \leq E_1$ and we again have

$$\mathsf{d}_U(x, \beta_U^W) = \mathsf{d}_V\big(x, \pi_U^V(\rho_V^Q)\big) \leq E_1^2 + E_1.$$

For the last clause of the consistency axiom, let $Q, U, W \in \mathfrak{T}_0$ with $Q \sqsubsetneq U$. If $U \sqsubsetneq W$, the definition of nesting and relative projection in $\mathfrak{T}_0$ and the consistency axioms in $(G_\Gamma, \mathfrak{S}_\Gamma)$ and the $(C(V), \mathfrak{R}_V)$ ensure that $\mathsf{d}_W(\beta_W^Q, \beta_W^U) \leq E_1^2 + E_1$. Similarly, if $W \in \mathfrak{S}_\Gamma$ with $W \pitchfork U$ and $W \not\perp Q$, then $\mathsf{d}_W(\beta_W^Q, \beta_W^U) \leq E_1^2 + E_1$. Assume that $W \in \mathfrak{R}_V$ for some $V \in \mathfrak{S}_\Gamma^{\min}$, $W \pitchfork U$, and $W \not\perp Q$. If $U, Q \in \mathfrak{R}_{V'}$, then $V' \pitchfork_{\mathfrak{S}} V$ and $\beta_W^U = \beta_W^Q$. If $U, Q \in \mathfrak{S}_\Gamma$, then $U \pitchfork_{\mathfrak{S}} V$ and $Q \pitchfork_{\mathfrak{S}} V$. Thus, the consistency axiom for $(G_\Gamma, \mathfrak{S}_\Gamma)$ provides $\mathsf{d}_V(\rho_V^U, \rho_V^Q) \leq E_1$. Similarly, if $U \in \mathfrak{S}_\Gamma$ and $Q \in \mathfrak{R}_{V'}$, then $U \pitchfork_{\mathfrak{S}} V$, $V' \pitchfork_{\mathfrak{S}} V$, and $\mathsf{d}_V(\rho_V^U, \rho_V^{V'}) \leq E_1$. Hence in both cases $\mathsf{d}_W(\beta_W^U, \beta_W^Q) \leq E_1^2 + E_1$.

Partial realisation. Let $W_1, \dots, W_n$ be pairwise orthogonal elements of $\mathfrak{T}_0$ and $p_i \in C(W_i)$ for each $i \in \{1, \dots, n\}$. Since $(G_\Gamma, \mathfrak{S}_\Gamma)$ satisfies the partial realisation axiom, we can assume at least one W_i is not an element of $\mathfrak{S}_\Gamma$. There exist $V_1, \dots, V_r \in \mathfrak{S}_\Gamma^{\min}$ so that for each $i \in \{1, \dots, n\}$, either $W_i \in \mathfrak{S}_\Gamma$ or there exists a unique $j \in \{1, \dots, r\}$ such that $W_i \in \mathfrak{R}_{V_j}$. For each $j \in \{1, \dots, r\}$, let $\{W_1^j, \dots, W_{k_j}^j\}$ be the elements of $\{W_1, \dots, W_n\}$ that are also elements of $\mathfrak{R}_{V_j}$ and let $\{p_1^j, \dots, p_{k_j}^j\}$ be the subset of $\{p_1, \dots, p_n\}$ satisfying $p_i^j \in C(W_i^j)$ for all $j \in \{1, \dots, r\}$ and $i \in \{1, \dots, k_j\}$. For each $j \in \{1, \dots, r\}$, use partial realisation in $(C(V_j), \mathfrak{R}_{V_j})$ on the points $p_1^j, \dots, p_{k_j}^j$ to produce a point $y_j \in C(V_j)$ so that

- $\mathsf{d}_{W_i^j}(y_j, p_i^j) \leq E_1$ for all $i \in \{1, \dots, k_j\}$;
- for each $i \in \{1, \dots, k_j\}$ and each $U \in \mathfrak{R}_{V_j}$, if $W_i^j \subsetneq U$ or $W_j^i \pitchfork U$, we have

$$\mathsf{d}_U(y_j, \rho_U^{W_i^j}) \leq E_1.$$

Assume, without loss of generality, that $W_m, W_{m+1}, \dots, W_n$ are all of the W_i that are not contained in any of the $\mathfrak{R}_{V_j}$ (it is possible that the set of such W_i is empty). Now applying partial realisation for $(G_\Gamma, \mathfrak{S}_\Gamma)$ to $y_1, \dots, y_r, p_m, \dots, p_n$ produces a point $x \in G_\Gamma$ so that $\psi_{W_i}(x)$ is uniformly close to p_i for each $i \in \{1, \dots, n\}$ and $\psi_U(x)$ is uniformly close to $\beta_U^{W_i}$ whenever $W_i \subsetneq U$ or $U \pitchfork W_i$, for any $U \in \mathfrak{T}_0$. Note that if the set of W_i that are not elements of any of the $\mathfrak{R}_{V_j}$ is empty, then the above applies just to $y_1, \dots, y_r$, but the conclusion still holds. ∎

Claim 5.4. The E_1-proto-hierarchy structure $\mathfrak{T}_0$ has the following property: if $W_1, \dots, W_n \in \mathfrak{T}_0$ are pairwise orthogonal, then $n \leq 4E_1^2 + 2E_1$.

Proof. We first note the following basic lemma from the theory of HHSs.

Lemma 5.5 ([11, Lemma 1.5]). *If $(\mathcal{X}, \mathfrak{S})$ is an E-HHS, then any set of pairwise orthogonal elements of $\mathfrak{S}$ has cardinality at most $2E$.*

Now let $W_1, \dots, W_n \in \mathfrak{T}_0$ be pairwise orthogonal. Without loss of generality, let $W_1, \dots, W_k$ be the elements of $\{W_1, \dots, W_n\}$ that are elements of $\mathfrak{S}_\Gamma$. Since $W_1, \dots, W_k$ is a pairwise orthogonal collection of elements of $\mathfrak{S}_\Gamma$, Lemma 5.5 says that $k \leq 2E_1$.

Let $V_1, \dots, V_m$ be the minimal collection of elements of $\mathfrak{S}_\Gamma^{\min}$ such that if $i \in \{k+1, \dots, n\}$ (i.e., $W_i \notin \mathfrak{S}_\Gamma$), then $W_i \in \mathfrak{R}_{V_j}$ for some $j \in \{1, \dots, m\}$. Minimality implies that for each $j \in \{1, \dots, m\}$, there exists $i \in \{k+1, \dots, n\}$ such that $W_i \in \mathfrak{R}_{V_j}$. Suppose that $W_i \in \mathfrak{R}_{V_j}$ and $W_\ell \in \mathfrak{R}_{V_r}$ with $j \neq r$. Since $W_i \perp W_\ell$ in $\mathfrak{T}_0$, the definition of orthogonality in $\mathfrak{T}_0$ implies that $V_j \perp_{\mathfrak{S}} V_r$. Thus, $V_1, \dots, V_m$ is a pairwise orthogonal collection of elements of $\mathfrak{S}_\Gamma$ and $m \leq 2E_1$ by Lemma 5.5. Similarly, for each $j \in \{1, \dots, m\}$ the set $\{W_i : W_i \in \mathfrak{R}_{V_j}\}$ is a pairwise orthogonal collection of elements of $\mathfrak{R}_{V_j}$ and must have cardinality at most $2E_1$. Putting this together, we have that

$$n \leq k + 2E_1 m \leq 2E_1 + 4E_1^2.$$ ∎

Claim 5.6. The action of G_Γ on $\mathfrak{S}_\Gamma$ induces an action of G_Γ on $\mathfrak{T}_0$ that satisfies axioms (2) and (3) of the definition of an HHG (Definition 2.27).

Proof. The action of G_Γ on $\mathfrak{T}_0$: let $\sigma \in G_\Gamma$ and $W \in \mathfrak{T}_0$. Define $\Phi\colon G \times \mathfrak{T}_0 \to \mathfrak{T}_0$ as follows.

- If $W = [g\Lambda] \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$, then $\Phi(\sigma, [g\Lambda]) = [\sigma g\Lambda]$; i.e., the action is the same as the action of G_Γ on $\mathfrak{S}_\Gamma$.
- If $W = \bar{g}R \in \mathfrak{R}_{[gv]}$ for some $[gv] \in \mathfrak{S}_\Gamma^{\min}$, then $(\overline{\sigma g})^{-1}\sigma\bar{g} \in \mathrm{Stab}_{G_\Gamma}([v])$, where $\overline{\sigma g}$ is the chosen fixed representative of $[\sigma g v] = [\sigma\bar{g}v]$. Since $\mathrm{Stab}_{G_\Gamma}([v]) = \langle \mathrm{st}(v)\rangle$, there exists $l \in \langle \mathrm{lk}(v)\rangle$ and $\hat{\sigma} \in \langle v\rangle$ such that $l\hat{\sigma} = (\overline{\sigma g})^{-1}\sigma\bar{g}$. Because $\mathfrak{R}_{[v]}$ is an HHG structure for $\langle v\rangle = G_v$, there exists $R_\sigma = \hat{\sigma}R \in \mathfrak{R}_{[v]}$ determined by σ and $\bar{g}R$. Define that $\Phi(\sigma, \bar{g}R) = \overline{\sigma g}R_\sigma \in \mathfrak{R}_{[\sigma g v]}$. The following diagram summarises how σ takes elements of $\mathfrak{R}_{[gv]}$ to elements of $\mathfrak{R}_{[\sigma g v]}$:

$$\begin{array}{ccc} \mathfrak{R}_{[gv]} & \xrightarrow{\ \sigma\ } & \mathfrak{R}_{[\sigma g v]} \\ {\scriptstyle \bar{g}^{-1}}\big\downarrow & \nearrow_{\overline{\sigma g}} & \\ \hat{\sigma} \circlearrowright \mathfrak{R}_{[v]} & & \end{array}$$

We now verify that Φ preserves the relations in $\mathfrak{T}_0$. Let $W, U \in \mathfrak{T}_0$. If $W, U \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ or $W, U \in \mathfrak{R}_{[gv]}$ for some $[gv] \in \mathfrak{S}_\Gamma^{\min}$, then Φ preserves the relation between W and U, since the actions of G_Γ on $\mathfrak{S}_\Gamma$ and $G_v = \langle v\rangle$ on $\mathfrak{R}_{[v]}$ preserve the relations in their respective hierarchy structures. If $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{\min}$ and $U \in \mathfrak{R}_{[gv]}$, then $W = [h\Omega]$ and the relation between W and U in $\mathfrak{T}_0$ is the same as the relation between $[h\Omega]$ and $[gv]$ in $\mathfrak{S}_\Gamma$. Thus, Φ preserves the relation between W and U, since the action of G_Γ preserves the relations in $\mathfrak{S}_\Gamma$. Similarly, the same is true in the case where $W \in \mathfrak{R}_{[gv]}$ and $U \in \mathfrak{R}_{[hw]}$ for $[gv] \neq [hw]$ as the relation between W and U in $\mathfrak{T}_0$ is the same as the relation between $[gv]$ and $[hw]$ in $\mathfrak{S}_\Gamma$.

The definition of Φ implies that $\bar{g}R \in \mathfrak{R}_{[gv]}$ is in the G_Γ-orbit of $\bar{h}R' \in \mathfrak{R}_{[hw]}$ if and only if $v = w$ and R is in the G_v-orbit of R'. Thus, the action of G_Γ on $\mathfrak{T}_0$ has finitely many orbits since the actions of G_Γ on $\mathfrak{S}_\Gamma$ and G_v on $\mathfrak{R}_{[v]}$ contain finitely many orbits.

For the remainder of the proof, we shall use σW to denote $\Phi(\sigma, W)$ for all $W \in \mathfrak{T}_0$. This does not conflict with the previous use of the notation as the action of G_Γ on $\mathfrak{T}_0$ agrees with the action of G_Γ on $\mathfrak{S}_\Gamma$ or the action of G_v on $\mathfrak{R}_{[v]}$, when $W \in \mathfrak{S}_\Gamma$ or $\sigma \in \langle v\rangle$ and $W \in \mathfrak{R}_{[v]}$, respectively.

Associated isometries and equivariance with the projection maps: let $\sigma, \tau \in G_\Gamma$ and $W \in \mathfrak{T}_0$. Since the action of G_Γ on $\mathfrak{T}_0$ agrees with the action of G_Γ on $\mathfrak{S}_\Gamma$ for the elements of $\mathfrak{T}_0$ in $\mathfrak{S}_\Gamma$, we can define the isometry

$$\sigma_{[g\Lambda]}\colon C\big([g\Lambda]\big) \to C\big([\sigma g\Lambda]\big)$$

to be the same as the original isometry in $(G_\Gamma, \mathfrak{S}_\Gamma)$; this guarantees that the HHG axioms are satisfied in this case.

If $W \in \mathfrak{R}_{[gv]}$, then $W = \bar{g}R$ for some $R \in \mathfrak{R}_{[v]}$. Now $\sigma W = \overline{\sigma g} R_\sigma$, where R_σ is defined as above. In this case, define the isometry $\sigma_W : C(W) \to C(\sigma W)$ to be the composition

$$C(W) \xrightarrow{(\bar{g}_R)^{-1}} C(R) \xrightarrow{\hat{\sigma}_R} C(R_\sigma) \xrightarrow{\overline{\sigma g}_{R_\sigma}} C(\sigma W),$$

where $\hat{\sigma}_R : C(R) \to C(R_\sigma)$ is the isometry in $\mathfrak{R}_{[v]}$ induced by $\hat{\sigma} \in G_v$, and $\bar{g}_R$ and $\overline{\sigma g}_{R_\sigma}$ are the isometries resulting from identifying $\mathfrak{R}_{[v]}$ with $\mathfrak{R}_{[gv]}$ and $\mathfrak{R}_{[\sigma g v]}$, respectively.

Now, if $\tau \in G_\Gamma$, then $(G_v, \mathfrak{R}_{[v]})$ being an HHG implies that $\hat{\tau}_{R_\sigma} \circ \hat{\sigma}_R = \widehat{\tau\sigma}_R$. Thus, the isometry $(\tau\sigma)_W$ equals the isometry $\tau_{\sigma W} \circ \sigma_W$ for any $W \in \mathfrak{T}_0$. We continue to use the notation set out before Claim 5.2: ψ_* and β_*^* denote the projections and relative projections in $\mathfrak{T}_0$, while π_*^* and ρ_*^* denote the projections and relative projections in $\mathfrak{S}_\Gamma$ and $\mathfrak{R}_{[gv]}$. Since the projection map $\psi_W : G_\Gamma \to 2^{C(W)}$ is equal to $\pi_W^{[gv]} \circ \pi_{[gv]}$, the uniform bound on the distance between $\psi_{\sigma W}(\sigma x)$ and $\sigma_W(\psi_W(x))$ follows from the HHG axioms of $(G_\Gamma, \mathfrak{S}_\Gamma)$ and $(G_v, \mathfrak{R}_{[v]})$. Similarly, since the relative projection β_W^U (where $U \sqsubsetneq W$ or $U \pitchfork W$ in $\mathfrak{T}_0$) is defined using the coarsely equivariant projections and relative projections of $\mathfrak{S}_\Gamma$ and $\mathfrak{R}_{[v]}$, we have that $\sigma_W(\beta_W^U)$ is uniformly close to $\beta_{\sigma W}^{\sigma U}$ whenever $U \sqsubsetneq W$ or $U \pitchfork W$. ∎

We now finish the proof of Theorem 5.1 using the following result.

Theorem 5.7 ([1, Theorem A.1]). *Let G be a finitely generated group and let $\mathfrak{T}_0$ be a proto-hierarchy structure for the Cayley graph of G with respect to some finite generating set. If $\mathfrak{T}_0$ satisfies the following:*

- *all of the axioms of an HHS except the container axiom;*
- *any set of pairwise orthogonal elements of $\mathfrak{T}_0$ has uniformly bounded cardinality;*
- *axioms (2) and (3) of an HHG structure (Definition 2.27);*

then there exists an HHG structure $\mathfrak{T}$ for the group G such that $\mathfrak{T}_0 \subsetneq \mathfrak{T}$ and for all $W \in \mathfrak{T} \smallsetminus \mathfrak{T}_0$, the associated hyperbolic space $C(W)$ is a single point.

Claims 5.3, 5.4, and 5.6 show that the proto-hierarchy structure $\mathfrak{T}_0$ satisfies the requirements of Theorem 5.7. Thus, there exists an HHG structure $\mathfrak{T}$ for G_Γ. ∎

Remark 5.8 (The HHG structure from Theorem 5.7). The proof of Theorem 5.7 produces an explicit HHG structure given the proto-structure $\mathfrak{T}_0$. We will describe that structure briefly now, and direct the reader to the appendix of [1] for full details.

Let $\mathcal{U}$ denote a non-empty set of pairwise orthogonal elements of $\mathfrak{T}_0$ and let $W \in \mathfrak{T}_0$. We say that the pair $(W, \mathcal{U})$ is a *container pair* if the following are satisfied:

- $U \sqsubseteq W$ for all $U \in \mathcal{U}$;
- there exists $Q \sqsubseteq W$ such that $Q \perp U$ for all $U \in \mathcal{U}$.

Let $\mathfrak{D}$ denote the set of all container pairs. We will denote a pair $(W, \mathcal{U}) \in \mathfrak{D}$ by $D_W^{\mathcal{U}}$. The crux of Theorem 5.7 is that the elements of $\mathfrak{D}$ will serve as containers for the elements of $\mathfrak{T}_0$, while the rest of the proto-structure is set up in the minimal way that satisfies all the other axioms.

The HHG structure produced by Theorem 5.7 has index set $\mathfrak{T}_0 \cup \mathfrak{D}$. The hyperbolic spaces, projection maps, relations, and relative projections for elements of $\mathfrak{T}_0$ remain unchanged. The hyperbolic spaces for elements of $\mathfrak{D}$ are single points and the projection maps are the constant maps to these points. The nesting relation involving elements of $\mathfrak{D}$ is defined as follows:

- define $Q \sqsubseteq D_W^{\mathcal{U}}$ if $Q \sqsubseteq W$ in $\mathfrak{T}_0$ and $Q \perp U$ for all $U \in \mathcal{U}$;
- define $D_W^{\mathcal{U}} \sqsubseteq Q$ if $W \sqsubseteq Q$ in $\mathfrak{T}_0$;
- define $D_W^{\mathcal{U}} \sqsubseteq D_T^{\mathcal{R}}$ if $W \sqsubseteq T$ in $\mathfrak{T}_0$ and for all $R \in \mathcal{R}$ either $R \perp W$ or there exists $U \in \mathcal{U}$ with $R \sqsubseteq U$.

Two elements $D_W^{\mathcal{U}}, D_T^{\mathcal{R}} \in \mathfrak{D}$ are orthogonal if $W \perp T$ in $\mathfrak{T}_0$. An element $Q \in \mathfrak{T}_0$ is orthogonal to $D_W^{\mathcal{U}} \in \mathfrak{D}$ if, in $\mathfrak{T}_0$, either $W \perp Q$ or $Q \sqsubseteq U$ for some $U \in \mathcal{U}$. Two elements of $\mathfrak{T}$ are transverse if they are not orthogonal and neither is nested into the other.

Since the associated hyperbolic spaces for elements of $\mathfrak{D}$ are single points, the relative projections onto these elements are just these single points. If $D_W^{\mathcal{U}} \sqsubsetneq Q$ or $Q \pitchfork D_W^{\mathcal{U}}$, then the relative projection $\rho_Q^{D_W^{\mathcal{U}}}$ is defined in one of two ways:

(1) if there exists $U \in \mathcal{U}$ such that $U \sqsubsetneq Q$ or $U \pitchfork Q$, then $\rho_Q^{D_W^{\mathcal{U}}}$ is the union of all ρ_Q^U for $U \in \mathcal{U}$ with $U \sqsubsetneq Q$ or $U \pitchfork Q$;

(2) if there does not exist $U \in \mathcal{U}$ such that $U \sqsubsetneq Q$ or $U \pitchfork Q$, then the definition of the relations given above forces $Q \pitchfork D_W^{\mathcal{U}}$ and $W \pitchfork Q$. In this case, $\rho_Q^{D_W^{\mathcal{U}}} = \rho_Q^W$.

5.2. Meier's condition for hyperbolicity

We now recover a theorem of Meier classifying hyperbolicity of graph products. We do this by applying Behrstock, Hagen, and Sisto's bounded orthogonality condition for HHSs.

Theorem 5.9 ([6, Corollary 2.16]). *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. The following are equivalent.*

- *$\mathcal{X}$ is hyperbolic.*
- *(Bounded orthogonality.) There exists a constant $D \geq 0$ such that*

$$\min\big(\operatorname{diam}\big(C(U)\big), \operatorname{diam}\big(C(V)\big)\big) \leq D$$

for all $U, V \in \mathfrak{S}$ satisfying $U \perp V$.

Theorem 5.10 (Meier's criterion for hyperbolicity of graph products; [24]). *Let Γ be a finite simplicial graph with hyperbolic groups associated to its vertices. Let Γ_F be the subgraph spanned by the vertices associated with finite groups. Then G_Γ is hyperbolic if and only if the following conditions hold:*

(i) *there are no edges connecting two vertices of $\Gamma \smallsetminus \Gamma_F$;*

(ii) *if v is a vertex of $\Gamma \smallsetminus \Gamma_F$, then $\operatorname{lk}(v)$ is a complete graph;*

(iii) *Γ_F does not contain any induced squares.*

Proof. We show hyperbolicity via the bounded orthogonality condition, noting that since each of the vertex groups is hyperbolic, the graph product G_Γ is an HHS. We call the vertices of Γ_F the *finite vertices* of Γ and the vertices of $\Gamma \smallsetminus \Gamma_F$ the *infinite vertices* of Γ.

($\Rightarrow$) Suppose that we have bounded orthogonality. Then the following hold.

(i) Suppose that two infinite vertices v, w are connected by an edge. Then $[v]\perp[w]$ and $C(v)$, $C(w)$ have infinite diameter as they are the infinite groups G_v, G_w with the word metric. This contradicts bounded orthogonality.

(ii) Suppose that $\mathrm{lk}(v)$ is incomplete for some vertex v of $\Gamma \smallsetminus \Gamma_F$. Then there exist some vertices x, y in $\mathrm{lk}(v)$ with no edge between them. Moreover, $[v]\perp[x \cup y]$, $C(v)$ has infinite diameter as v is an infinite vertex, and $C(x \cup y)$ has infinite diameter since $\mathsf{d}_{x\cup y}(e,(g_x g_y)^n) = 2n$ for elements $g_x \in G_x \smallsetminus \{e\}$, $g_y \in G_y \smallsetminus \{e\}$. This again contradicts bounded orthogonality.

(iii) Suppose that Γ_F contains a square with vertices v, x, w, y, where v, w and x, y are non-adjacent. Then $[v \cup w]\perp[x \cup y]$ and both $C(v \cup w)$ and $C(x \cup y)$ have infinite diameter as in case (ii). Once again, this contradicts bounded orthogonality.

($\Leftarrow$) Conversely, suppose that conditions (i), (ii), and (iii) are satisfied, and set $D = \max\{2, |G_v| : v \in V(\Gamma_F)\}$. Moreover, suppose that $[g\Lambda], [h\Omega] \in \mathfrak{S}$ satisfy $[g\Lambda]\perp[h\Omega]$.

Suppose that $\mathrm{diam}(C(g\Lambda)) > D$. Then Theorem 4.10 tells us that either Λ consists of a single infinite vertex or Λ contains at least two vertices and does not split as a join.

If Λ consists of a single infinite vertex, then conditions (i) and (ii) tell us that $\mathrm{lk}(\Lambda) \supseteq \Omega$ is a complete graph consisting of finite vertices, hence either Ω is a single finite vertex or Ω splits as a join. In both cases, $\mathrm{diam}(C(h\Omega)) \leq D$.

If Λ contains at least two vertices and does not split as a join, then, in particular, it contains two non-adjacent vertices v and w. As $\Omega \subseteq \mathrm{lk}(\Lambda)$, every vertex of Ω is connected to both v and w. Since v and w are non-adjacent, condition (ii) implies that $\Omega \subseteq \Gamma_F$. If either v or w is an infinite vertex, condition (ii) implies that Ω is a complete graph, and if both v and w are finite vertices, condition (iii) implies that Ω is a complete graph. That is, Ω either consists of a single finite vertex or splits as a join. In both cases, $\mathrm{diam}(C(h\Omega)) \leq D$. Thus, the bounded orthogonality condition holds. ∎

5.3. Genevois' minsquare electrification

We now use our characterisation of when $C(g\Lambda)$ has infinite diameter (Theorem 4.10) to answer two questions of Genevois [16, Questions 8.3 and 8.4] regarding the *electrification* of G_Γ, defined as follows.

Definition 5.11. Let Γ be a simplicial graph. An induced subgraph $\Lambda \subseteq \Gamma$ is called *square-complete* if every induced square in Γ sharing two non-adjacent vertices with Λ is a subgraph of Λ. A subgraph is *minsquare* if it is a minimal square-complete subgraph containing at least one induced square.

The *electrification* $\mathbb{E}(\Gamma)$ of a graph product G_Γ is the graph whose vertices are elements of G_Γ and where two vertices g and h are joined by an edge if $g^{-1}h$ is an element of a vertex group or $g^{-1}h \in \langle\Lambda\rangle$ for some minsquare subgraph Λ of Γ. We use $\mathrm{d}_{\mathbb{E}}(g,h)$ to denote the distance in $\mathbb{E}(\Gamma)$ between $g,h \in G_\Gamma$.

Genevois' interest in the electrification arises from the fact that it forms a quasi-isometry invariant whenever the vertex groups of a graph product are all finite, as is the case for right-angled Coxeter groups.

Theorem 5.12 ([16, Proposition 1.4]). *Let G_Γ and G_Λ be graph products of finite groups. Any quasi-isometry $G_\Gamma \to G_\Lambda$ induces a quasi-isometry between $\mathbb{E}(\Gamma)$ and $\mathbb{E}(\Lambda)$.*

For graph products of finite groups, we classify when $\mathbb{E}(\Gamma)$ has bounded diameter and when $\mathbb{E}(\Gamma)$ is a quasi-line. These classifications answer Questions 8.3 and 8.4 of [16] in the affirmative. The core idea behind both proofs is the same: when Γ is not minsquare, the electrification $\mathbb{E}(\Gamma)$ sits between the syllable metric $S(\Gamma)$ and the subgraph metric $C(\Gamma)$; that is, we obtain $\mathbb{E}(\Gamma)$ from $S(\Gamma)$ by adding edges and then obtain $C(\Gamma)$ from $\mathbb{E}(\Gamma)$ by adding more edges. This means that large distances in $C(\Gamma)$, which we can detect with Theorem 4.10, will persist in $\mathbb{E}(\Gamma)$. We start with a lemma that we use in both classifications to reduce to the case where Γ does not split as a join.

Lemma 5.13. *If Γ splits as a join and contains a proper minsquare subgraph, then Γ splits as a join $\Gamma = \Gamma_1 \bowtie \Gamma_2$, where Γ_1 contains every minsquare subgraph of Γ and Γ_2 is a complete graph. In this case, $\mathbb{E}(\Gamma)$ is the 1-skeleton of $\mathbb{E}(\Gamma_1) \times \mathbb{E}(\Gamma_2)$.*

Proof. Suppose that Γ contains a proper minsquare subgraph Λ and splits as a join $\Gamma = \Omega_1 \bowtie \Omega_2$. We first show that Γ splits as a (possibly different) join $\Gamma_1 \bowtie \Gamma_2$, where Γ_1 contains the minsquare subgraph Λ. If Λ is a subgraph of either Ω_1 or Ω_2 we are done. Otherwise, Λ contains vertices of both Ω_1 and Ω_2. By minimality of Λ, there must exist a square of Λ containing vertices of both Ω_1 and Ω_2. Moreover, since Ω_1 and Ω_2 form a join, this square must arise in the form of two pairs of disjoint vertices $v_i, w_i \in V(\Omega_i)$, $i = 1,2$. Then any vertex v of $\Omega_1 \smallsetminus \Lambda$ must be connected to every vertex w of $\Lambda \cap \Omega_1$, else v, w, v_2, w_2 form an induced square, contradicting square-completeness of Λ. Similarly, any vertex of $\Omega_2 \smallsetminus \Lambda$ must be connected to every vertex of $\Lambda \cap \Omega_2$. This then gives a decomposition of Γ as a join of the minsquare subgraph Λ and the graph $\Gamma \smallsetminus \Lambda$.

We have shown that Γ splits as a join $\Gamma_1 \bowtie \Gamma_2$ with $\Lambda \subseteq \Gamma_1$. We now show that Γ_2 must be a complete graph. Since Λ is minsquare, there exists an induced square S in $\Lambda \subseteq \Gamma_1$. Let v_1, w_1 be two disjoint vertices of S, and suppose that there exists a pair of disjoint vertices v_2, w_2 in Γ_2. Since Γ is a join of Γ_1 and Γ_2 and $\Lambda \subseteq \Gamma_1$, the vertices v_1, w_1, v_2, w_2 define an induced square that shares two opposite vertices with Λ, but is not contained in Λ. This would contradict square-completeness of Λ. Therefore, Γ_2 must be complete.

Finally, we show that every other minsquare subgraph of Γ must also be contained in Γ_1. Let $\Omega \subseteq \Gamma$ be minsquare. If four vertices v_1, v_2, v_3, v_4 of Ω form an induced square of Γ, then each v_i must be contained in Γ_1, since any v_i that Γ_2 contains must be connected to all v_j in Γ_1, but Γ_2 cannot contain a pair of disjoint vertices since it is complete. Thus, the minimality of Ω implies that Ω must be contained in Γ_1 (otherwise $\Omega \cap \Gamma_1$ would be a proper square-complete subgraph of Ω).

Since Γ splits as a join $\Gamma_1 \bowtie \Gamma_2$, it follows that $S(\Gamma)$ is the 1-skeleton of $S(\Gamma_1) \times S(\Gamma_2)$ and since the only minsquare subgraphs of Γ are the minsquare subgraphs of Γ_1, $\mathbb{E}(\Gamma)$ is the 1-skeleton of $\mathbb{E}(\Gamma_1) \times \mathbb{E}(\Gamma_2)$ by construction. ∎

We now show that $\mathbb{E}(\Gamma)$ is bounded only in the obvious cases.

Theorem 5.14. *The electrification $\mathbb{E}(\Gamma)$ is bounded if and only if Γ is either minsquare, complete or splits as a join of a minsquare subgraph and a complete graph.*

Proof. We first show that if Γ is minsquare, complete, or splits as the join of a minsquare subgraph and a complete graph, then the electrification is bounded. If Γ is minsquare, then $\mathbb{E}(\Gamma)$ has diameter 1 by definition. Let x, y be vertices of $\mathbb{E}(\Gamma)$, so that $x^{-1}y \in G_\Gamma$. If Γ is a complete graph on n vertices, then all vertex groups of Γ commute, so we can write $x^{-1}y = s_1 \dots s_n$, where $\mathrm{supp}(s_i) = v_i \in V(\Gamma)$ and $v_i \neq v_j$ for all $i \neq j$. Thus, $\mathsf{d}_{\mathbb{E}}(x, y) \leq n$, and hence $\mathbb{E}(\Gamma)$ is bounded. If Γ splits as a join of a minsquare subgraph Γ_1 and a complete graph Γ_2 on n vertices, then $G_\Gamma \cong \langle \Gamma_1 \rangle \times \langle \Gamma_2 \rangle$ and so we can write $x^{-1}y = g_1 g_2$, where $g_i \in \langle \Gamma_i \rangle$. Therefore, $\mathsf{d}_{\mathbb{E}}(x, y) \leq n + 1$, hence $\mathbb{E}(\Gamma)$ is bounded.

We now assume that $\mathbb{E}(\Gamma)$ is bounded and prove that this implies that Γ either is complete, minsquare or splits as a join of a minsquare subgraph and a complete graph. The proof will proceed by induction on the number of vertices of Γ. The base case is immediate as Γ is complete and $\mathbb{E}(\Gamma)$ has diameter 1 when Γ is a single vertex. Assume that the conclusion holds whenever the defining graph has at most $n-1$ vertices. Let G_Γ be a graph product of groups where Γ contains $n \geq 2$ vertices.

Claim 5.15. If $\mathbb{E}(\Gamma)$ is bounded and Γ is neither complete nor minsquare, then Γ must split as a join and must contain a proper minsquare subgraph.

Proof. Suppose that Γ does not split as a join. By Theorem 4.10, $C(\Gamma)$ is therefore unbounded. Since Γ is not minsquare, $\mathbb{E}(\Gamma)$ can be obtained from $C(\Gamma)$ by removing some edges. In particular, if $C(\Gamma)$ has infinite diameter, then so does $\mathbb{E}(\Gamma)$. This implies that if Γ is not minsquare and does not split as a join, then $\mathbb{E}(\Gamma)$ is unbounded, contradicting our assumption.

Now suppose that Γ does not contain any proper minsquare subgraphs. Then $\mathbb{E}(\Gamma)$ is simply $S(\Gamma)$. Since Γ is not complete, there exist two disjoint vertices $v, w \in V(\Gamma)$. Therefore, $\mathsf{d}_{\mathbb{E}}(e, (g_v g_w)^m) = \mathsf{d}_{\mathrm{syl}}(e, (g_v g_w)^m) = 2m$ for any $g_v \in G_v \setminus \{e\}$ and $g_w \in G_w \setminus \{e\}$, hence $\mathbb{E}(\Gamma)$ is unbounded, a contradiction. ∎

Assume that Γ is neither complete nor minsquare, so that Γ must contain a strict minsquare subgraph Λ and splits as a join by Claim 5.15. By Lemma 5.13, Γ must split as a join of Γ_1 and Γ_2, where Γ_2 is complete and $\mathbb{E}(\Gamma)$ is the 1-skeleton of $\mathbb{E}(\Gamma_1) \times \mathbb{E}(\Gamma_2)$. Thus, $\mathbb{E}(\Gamma)$ having bounded diameter implies that $\mathbb{E}(\Gamma_1)$ must also have bounded diameter. Since Γ_1 contains at most $n-1$ vertices, the induction hypothesis then implies that Γ_1 is either minsquare, complete, or splits as a join of a minsquare subgraph and a complete graph. Since $\Lambda \subseteq \Gamma_1$ contains a square, Γ_1 cannot be complete. Thus, Γ_1 is either minsquare or a join of Λ with a complete graph Ω. Hence Γ splits either as a join of the minsquare subgraph Γ_1 and the complete graph Γ_2 or as a join of the minsquare subgraph Λ and the complete graph $\Omega \bowtie \Gamma_2$. ∎

Finally, we show that $\mathbb{E}(\Gamma)$ being a quasi-line coincides with G_Γ being virtually cyclic. The key step of the proof is to produce two elements of G_Γ that act as independent loxodromic elements on $C(\Gamma)$. This creates more than two directions to escape to infinity in $C(\Gamma)$, which then gives more than two directions to escape to infinity in $\mathbb{E}(\Gamma)$.

Theorem 5.16. *Let G_Γ be a graph product of finite groups. The electrification $\mathbb{E}(\Gamma)$ is a quasi-line if and only if G_Γ is virtually cyclic.*

Proof. A graph product of finite groups G_Γ is virtually cyclic if and only if either Γ is a pair of disjoint vertices each with vertex group $\mathbb{Z}_2$ or Γ splits as a join $\Gamma_1 \bowtie \Gamma_2$, where Γ_1 is a pair of disjoint vertices each with vertex group $\mathbb{Z}_2$ and Γ_2 is a complete graph (this follows from [2, Lemma 3.1]). Thus, if G_Γ is virtually cyclic, then $\mathbb{E}(\Gamma) = S(\Gamma)$ is a quasi-line by construction.

Let us now assume that G_Γ is not virtually cyclic. If Γ is either minsquare, complete or the join of a minsquare graph and a complete graph, then $\mathbb{E}(\Gamma)$ has bounded diameter by Theorem 5.14 and is therefore not a quasi-line. Let us therefore assume that Γ is not minsquare, not complete, and does not split as a join of a minsquare graph and a complete graph.

First assume that Γ does not split as a join at all. Since the action of G_Γ on $C(\Gamma)$ by left multiplication is acylindrical (Corollary 4.24), a result of Osin [25, Theorem 1.1] says that G_Γ must satisfy exactly one of the following: G_Γ has bounded orbits in $C(\Gamma)$, G_Γ is virtually cyclic or G_Γ contains two elements that act loxodromically and independently on $C(\Gamma)$. Since Γ does not split as a join, the proof of Theorem 4.10 implies that G_Γ does not have bounded orbits in $C(\Gamma)$. Further, G_Γ is not virtually cyclic by assumption. Thus, there exist $g, h \in G_\Gamma$ such that $n \mapsto \pi_\Gamma(g^n)$ and $n \mapsto \pi_\Gamma(h^n)$ are bi-infinite quasi-geodesics in $C(\Gamma)$ whose images, $\pi_\Gamma(\langle g\rangle)$ and $\pi_\Gamma(\langle h\rangle)$, have infinite Hausdorff distance from each other. Now, since Γ is not minsquare, $C(\Gamma)$ is obtained from $\mathbb{E}(\Gamma)$ by adding edges and therefore

$$\mathsf{d}_\Gamma(x, y) \leq \mathsf{d}_\mathbb{E}(x, y) \quad \text{for all } x, y \in G_\Gamma.$$

Hence the subsets $\langle g\rangle$ and $\langle h\rangle$ in $\mathbb{E}(\Gamma)$ are also the images of bi-infinite quasi-geodesics that have infinite Hausdorff distance from each other. This implies that $\mathbb{E}(\Gamma)$ is not a quasi-line, as any two bi-infinite quasi-geodesics in a quasi-line have finite Hausdorff distance.

Now assume that Γ splits as a join. If Γ contains no minsquare subgraph, then $\mathbb{E}(\Gamma) = S(\Gamma)$. Since the vertex groups are all finite, $S(\Gamma)$ is quasi-isometric to the word metric on G_Γ and hence $S(\Gamma) = \mathbb{E}(\Gamma)$ is not a quasi-line, because we assumed that G_Γ is not virtually cyclic. Thus, we can assume that Γ contains a minsquare subgraph Λ. By applying Lemma 5.13 iteratively, we have that Γ splits as a join $\Gamma = \Gamma_1 \bowtie \Gamma_2$ such that

- Γ_1 either does not split as a join or is minsquare;
- Γ_2 is a complete graph;
- $\mathbb{E}(\Gamma)$ is the 1-skeleton of $\mathbb{E}(\Gamma_1) \times \mathbb{E}(\Gamma_2)$.

Recall our assumption that Γ does not split as a join of a minsquare graph and a complete graph, hence Γ_1 cannot be minsquare and thus must not split as a join by the first item above. Further, $\langle \Gamma_1 \rangle$ is not virtually cyclic since it is a finite index subgroup of G_Γ, which is not virtually cyclic. Thus, we can apply the previous case to conclude that $\mathbb{E}(\Gamma_1)$ is not a quasi-line and hence $\mathbb{E}(\Gamma)$ is not a quasi-line. ■

Acknowledgements. The authors would like to thank Mark Hagen for his useful comments and suggestions for applications, Anthony Genevois for his comments on an earlier draft of the paper, and their advisor Jason Behrstock for many valuable conversations and his support throughout. The authors would also like to thank an anonymous referee for comments that improved the paper.

References

[1] C. Abbott, J. Behrstock, and M. G. Durham, Largest acylindrical actions and stability in hierarchically hyperbolic groups. With an appendix by Daniel Berlyne and Jacob Russell. *Trans. Amer. Math. Soc. Ser. B* **8** (2021), 66–104 Zbl 07319368 MR 4215647

[2] H. Baik, B. Petri, and J. Raimbault, Subgroup growth of virtually cyclic right-angled Coxeter groups and their free products. *Combinatorica* **39** (2019), no. 4, 779–811 Zbl 1449.20036 MR 4015351

[3] J. Behrstock, M. Hagen, and A. Sisto, Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.* **299** (2019), no. 2, 257–338 Zbl 07062864 MR 3956144

[4] J. Behrstock, M. F. Hagen, and A. Sisto, Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)* **114** (2017), no. 5, 890–926 Zbl 1431.20028 MR 3653249

[5] J. Behrstock, M. F. Hagen, and A. Sisto, Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.* **21** (2017), no. 3, 1731–1804 Zbl 1439.20043 MR 3650081

[6] J. Behrstock, M. F. Hagen, and A. Sisto, Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.* **170** (2021), no. 5, 909–996 Zbl 07369844 MR 4255047

[7] J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, Geometry and rigidity of mapping class groups. *Geom. Topol.* **16** (2012), no. 2, 781–888 Zbl 1281.20045 MR 2928983

[8] F. Berlai and B. Robbio, A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.* **14** (2020), no. 4, 1127–1203 Zbl 07362555 MR 4186470

[9] J. F. Brock, The Weil–Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.* **16** (2003), no. 3, 495–535 Zbl 1059.30036 MR 1969203

[10] M. G. Durham, The augmented marking complex of a surface. *J. Lond. Math. Soc. (2)* **94** (2016), no. 3, 933–969 Zbl 1360.30038 MR 3614935
[11] M. G. Durham, M. F. Hagen, and A. Sisto, Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.* **21** (2017), no. 6, 3659–3758 Zbl 1439.20044 MR 3693574
[12] M. G. Durham, M. F. Hagen, and A. Sisto, Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.* **24** (2020), no. 2, 1051–1073 Zbl 1443.20071 MR 4153656
[13] A. Eskin, H. Masur, and K. Rafi, Large-scale rank of Teichmüller space. *Duke Math. J.* **166** (2017), no. 8, 1517–1572 Zbl 1373.32012 MR 3659941
[14] A. Genevois, *Cubical-like geometry of quasi-median graphs and applications to geometric group theory*. Ph.D. thesis, Université d'Aix-Marseille, 2017
[15] A. Genevois, Automorphisms of graph products of groups and acylindrical hyperbolicity. 2020, arXiv:1807.00622
[16] A. Genevois, Quasi-isometrically rigid subgroups in right-angled Coxeter groups. *Algebr. Geom. Topol.* **22** (2022), no. 2, 657–708 MR 4464462
[17] A. Genevois and A. Martin, Automorphisms of graph products of groups from a geometric perspective. *Proc. Lond. Math. Soc. (3)* **119** (2019), no. 6, 1745–1779 Zbl 1461.20008 MR 4295519
[18] E. Green, *Graph products of groups*. Ph.D. thesis, University of Warwick, Coventry, 1990
[19] M. Gromov, Hyperbolic groups. In *Essays in group theory*, pp. 75–263, Math. Sci. Res. Inst. Publ. 8, Springer, New York, 1987 Zbl 0634.20015 MR 919829
[20] S.-H. Kim and T. Koberda, The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.* **24** (2014), no. 2, 121–169 Zbl 1342.20042 MR 3192368
[21] J. F. Manning, Geometry of pseudocharacters. *Geom. Topol.* **9** (2005), 1147–1185 Zbl 1083.20038 MR 2174263
[22] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* **138** (1999), no. 1, 103–149 Zbl 0941.32012 MR 1714338
[23] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.* **10** (2000), no. 4, 902–974 Zbl 0972.32011 MR 1791145
[24] J. Meier, When is the graph product of hyperbolic groups hyperbolic? *Geom. Dedicata* **61** (1996), no. 1, 29–41 Zbl 0874.20026 MR 1389635
[25] D. Osin, Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.* **368** (2016), no. 2, 851–888 Zbl 1380.20048 MR 3430352
[26] K. Rafi, A combinatorial model for the Teichmüller metric. *Geom. Funct. Anal.* **17** (2007), no. 3, 936–959 Zbl 1129.30031 MR 2346280

Received 22 June 2020.

Daniel Berlyne
The Graduate Center, City University of New York, New York, NY 10016, USA;
dberlyne@gradcenter.cuny.edu

Jacob Russell
Department of Mathematics, Rice University, Houston, TX 77005, USA;
jacob.russell@rice.edu