6-2021

Hierarchical Hyperbolicity of Graph Products and Graph Braid Groups

Daniel James Solomon Berlyne
The Graduate Center, City University of New York

Hierarchical hyperbolicity of graph products and graph braid groups

by

Daniel James Solomon Berlyne

A dissertation submitted to the Graduate Faculty in Mathematics in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York.

2021

Hierarchical hyperbolicity of graph products and graph braid groups

by

Daniel James Solomon Berlyne

This manuscript has been read and accepted for the Graduate Faculty in Mathematics in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date	**Jason Behrstock** Chair of Examining Committee
Date	**Ara Basmajian** Executive Officer

Jason Behrstock

Ilya Kapovich

Joseph Maher

Supervisory Committee

The City University of New York

Abstract

Hierarchical hyperbolicity of graph products and graph braid groups

by

Daniel James Solomon Berlyne

Advisor: Jason Behrstock

This thesis comprises three original contributions by the author concerning hierarchical hyperbolicity, a coarse geometric tool developed by Behrstock, Hagen, and Sisto to provide a common framework for studying aspects of non-positive curvature in a wide variety of groups and spaces.

We show that any graph product of finitely generated groups is hierarchically hyperbolic relative to its vertex groups. We apply this to answer two questions of Genevois about the electrification of a graph product of finite groups. We also answer two questions of Behrstock, Hagen, and Sisto: we show that the syllable metric on a 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. To achieve this, we develop a technique that allows an almost hierarchically hyperbolic structure to be promoted to a hierarchically hyperbolic structure. This technique has found independent use in work of Abbott, Behrstock, and Durham, where it is used to significantly streamline their proofs.

We then turn to graph braid groups, using their structure as fundamental groups of special cube complexes to endow them with a natural hierarchically hyperbolic structure. By expressing this structure in terms of the graph, we obtain characterisations of when these groups are hyperbolic or acylindrically hyperbolic. We also conjecture and partially prove a similar characterisation of relative hyperbolicity.

Acknowledgments

I would like to thank my advisor, Jason Behrstock, who has provided guidance throughout my PhD, patiently answering all my questions while also teaching me valuable lessons on how to thrive as an independent mathematician. I remember well the day you approached me after attending my talk on Teichmüller theory and suggested I speak with you. Despite my background being primarily in complex analysis, you made the transition to geometric group theory easy and exciting. Your support and your advocacy on my behalf have been invaluable. I also extend my thanks to Ilya Kapovich and Joseph Maher for serving on my thesis committee alongside Jason.

I would also like to thank my past and current collaborators. You have made mathematics an even more enjoyable and meaningful pursuit than it already was. The papers I wrote with Jacob Russell have been formative to my identity as a mathematician, while my current venture with Carolyn Abbott, Alex Rasmussen, and Thomas Ng has already provided many inspiring and eye-opening conversations of its own. It is fair to say that these joint works have vastly changed my perspective on how mathematics ought to be undertaken.

I would like to thank Jason Behrstock, Mark Hagen, Carolyn Abbott, Ilya Kapovich, and Olga Kharlampovich for writing letters of recommendation for me. You were instrumental in my success in finding a job. I would especially like to thank Mark for inviting me to speak at Bristol and spending so much time talking about mathematics with me and introducing me to Bristol's mathematics community. Your hospitality made me feel very welcome, and I am sure that my postdoc in Bristol is in no small part due to your efforts.

I am tremendously grateful for all the friends I made along my journey as a graduate student. I thank Jacob Russell, Ivan Levcovitz, Hai Yu, GB Pignatti, Laura Lopez-Cruz, AJ Stas, Chris Natoli, Alice Kwon, and Kieran O'Reilly for teaching and learning with me, and for brightening my experience at the Graduate Center. I have many fond memories of our experiences in the shoebox. I am also very fortunate to have met so many wonderful people through conferences, such as Harry Petyt, Bruno Robbio, Davide Spriano, Yvon Verberne, and Kasia Jankiewicz.

I would like to thank Jason Behrstock, Abhijit Champanerkar, Ilya Kofman, and Joseph Maher for countless enjoyable Korean lunches, both down in Midtown and up near Columbia. The seminars you run and the community you have fostered around them have been a cornerstone to my graduate experience. You have taught me a great deal about hierarchical hyperbolicity, 3–manifolds, knot theory, and random walks, all of which have influenced the research I conduct today.

I am grateful to my Master's advisor, Saul Schleimer, for introducing me to geometric group theory, even if I didn't know it at the time. You planted the seed which turned my interest in Teichmüller space and Beltrami differentials into an appreciation of mapping class groups and eventually into a love of geometric group theory as a whole.

I am thankful to my teachers at the Manchester Grammar School, who opened my mind to the sheer breadth of mathematics and its interconnectedness. I am especially thankful to Neil Sheldon, who seemed to be able to teach everything under the sun, from probability to differential equations. I will always remember the summer of 2010 when you taught us an entire numerical analysis course in one week in the confines of a cramped IT lab.

But I am most thankful to my kind and loving family, who provided the environment which made all of this possible. My mother and father, Jayne and Brian, deserve a great deal of credit for my success, having nurtured my curiosity from a young age and having done all they could to allow me to spread my wings. My brother, Josh, has instigated many an interesting philosophical conversation, and has humoured my many ramblings about alternative systems of logic and non-Euclidean geometries. Nathan Somers, while not family by blood, has been a close friend, ally, and confidante throughout most of my life. You have seen me through many highs and lows, and have provoked in me a love of linguistics and its mathematical connections.

And Miranda, my dearest love, closest companion, and the newest addition to my family, has been by my side at all times throughout this most challenging leg of my journey. You have taught me confidence and resolve, and have given me the strength to carry on in the face of Sisyphean adversity. You have pushed me when I have faltered, held me when I have buckled, and picked me up when I have fallen. Without you I am not sure I would be in this position today. I thank you sincerely.

The results in Chapter 3 appear as an appendix to [ABD21], and the results in Chapter 4 appear in [BR20].

Contents

List of Figures

To Jayne and Brian. For encouraging me to follow my passions no matter where they might lead, and supporting me unconditionally every step of the way.

We shall not cease from exploration

And the end of all our exploring

Will be to arrive where we started

And know the place for the first time.

– T. S. Eliot.

Chapter 1

Introduction

While groups have long been used as a tool for studying geometry—for example, the study of Kleinian groups dates back to Klein and Poincaré in the 1880s—the use of geometry to study groups is a much more recent endeavour. This area, known as *geometric group theory*, was studied by a number of authors in the 20th century such as Dehn, Milnor–Švarc, and Bass–Serre, but did not become recognisable as a field in its own right until the publication of two seminal papers by Gromov [Gro87, Gro93]. In these, he detailed a programme for the study of group theory from the point of view of a collection of natural discrete metrics associated to any finitely generated group. These so-called *word metrics* measure distance between two elements g and h of a group G by counting the number of elements required to express the difference $g^{-1}h$ as a product of generators of G and their inverses.

One must beware that a group can have many word metrics, depending on the choice of generating set. However, so long as these generating sets are finite, the distances with respect to any two word metrics are equivalent up to a multiplicative and additive constant; we say the metric spaces are *quasi-isometric*. By restricting ourselves to geometric properties that are invariant under quasi-isometry, we therefore obtain invariants of the group itself, independent of the choice of generating set.

The discreteness of this geometry bars one from utilising the traditional techniques of

Riemannian geometry. Nonetheless, there are ways of developing analogues of classical Riemannian concepts which make sense in this discrete setting. This is primarily achieved by studying the *Cayley graphs* of a finitely generated group G, defined by representing elements of G by vertices and connecting two vertices by an edge whenever the corresponding elements differ by a generator. By assigning each edge a length of 1, we see that these Cayley graphs are quasi-isometric to the group G with the corresponding word metrics. Moreover, the Cayley graphs can be viewed as geodesic metric spaces, granting access to a greater range of geometric tools.

For example, one of Gromov's most influential contributions to geometric group theory was through the development of the concept of *hyperbolicity* of a group. By mimicking the 'thin' appearance of triangles in hyperbolic Riemannian manifolds, one is able to develop a version of hyperbolicity that can be applied to any geodesic metric space, and thus in particular to Cayley graphs. Specifically, a geodesic metric space is defined to be *δ–hyperbolic* if every geodesic triangle satisfies the condition that each side is contained in the union of the δ–neighbourhoods of the other two sides. A group is then defined to be hyperbolic if its Cayley graphs are hyperbolic.

Surprisingly, this simple condition is sufficient to capture many aspects of negative curvature seen in the Riemannian setting. For example, one can develop a discrete analogue of volume by counting the number of group elements contained in a ball of a given radius. In a hyperbolic group, this volume is seen to grow exponentially with the radius of the ball. Finitely presented hyperbolic groups can also be shown to satisfy a linear isoperimetric inequality, with the appropriate analogues of perimeter and area.

Hyperbolic groups are seen in a wide variety of places. For example, finite groups, free groups, small cancellation groups, and fundamental groups of closed hyperbolic manifolds are hyperbolic. In fact, hyperbolicity is in some sense ubiquitous among groups; by studying random groups, Gromov and Ol'shanskii show that almost every finitely presented group is hyperbolic [Gro93, Ol'92]. Despite this, there are many interesting and important classes of groups that are not hyperbolic in the strictly homogeneous sense required by δ–hyperbolicity, but do exhibit some hyperbolic behaviour. There have therefore been many attempts to generalise the definition of hyperbolicity in order to capture this.

One of the first such generalisations was that of *relative hyperbolicity*, suggested by Gromov himself and improved upon by Farb, Bowditch, and others [Gro87, Far98, Bow12, DS05, Osi06]. Another example is *acylindrical hyperbolicity*, developed by Osin by building upon ideas of Sela and Bowditch [Osi16, Sel97, Bow08]. The primary focus of this thesis, however, shall be *hierarchical hyperbolicity*, developed by Behrstock, Hagen, and Sisto [BHS17b, BHS19] as a way of describing hyperbolic behaviour in quasi-geodesic metric spaces via machinery akin to that introduced for mapping class groups by Masur and Minsky [MM99, MM00].

The work of Behrstock, Hagen, and Sisto originally focused on developing such machinery for right-angled Artin groups, but hierarchical hyperbolicity also provides a common framework in which to study a wide variety of other groups and spaces. Prominent examples include virtually cocompact special groups and most CAT(0) cube complexes [BHS17b], fundamental groups of closed 3–manifolds with no Nil or Sol components in their prime decomposition [BHS19], Teichmüller space with either the Teichmüller metric or the Weil–Petersson metric ([Raf07, Dur16, EMR17] and [Bro03, Beh06, BKMM12] respectively), and graph products of hyperbolic groups [BR18].

Hierarchical hyperbolicity has deep geometric consequences for a space, including a Masur and Minsky style distance formula [BHS19], a quadratic isoperimetric inequality [BHS19], rank rigidity and Tits alternative theorems [DHS17, DHS19], control over top-dimensional quasi-flats [BHS17c], and bounds on the asymptotic dimension [BHS17a]. Moreover, the rich structure afforded by hierarchical hyperbolicity also allows one to detect the presence of other forms of hyperbolicity, such as δ–hyperbolicity [BHS17c], relative hyperbolicity [Rus20], and acylindrical hyperbolicity [BHS17b].

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 geometry of the entire space to be recovered from the projections to the hyperbolic spaces $C(W)$. Furthermore, this structure encodes non-positive curvature occurring in the space; we obtain information about hyperbolic aspects of $\mathcal{X}$ through the projections to the hyperbolic spaces $C(W)$, while flats (copies of $\mathbb{Z}^n$) in $\mathcal{X}$ are encoded by the orthogonality relation.

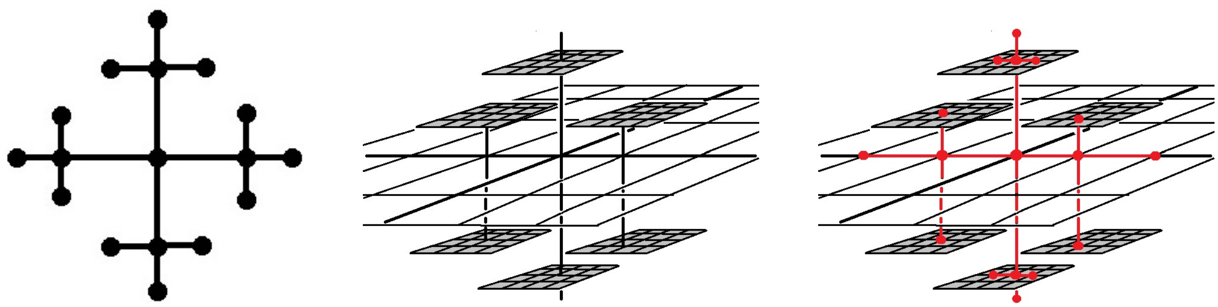

Figure 1.1: The right-angled Artin group $(\mathbb{Z} \times \mathbb{Z}) * \mathbb{Z}$ can be endowed with a hierarchically hyperbolic structure. Its Cayley graph (centre) contains both Euclidean planes and copies of the Cayley graph of the hyperbolic group $\mathbb{Z} * \mathbb{Z}$ (left).

1.1 Contributions of the author

The original contributions of this thesis focus on three topics, based on three pieces of work the author has produced during his graduate studies. Two of these are joint projects with Jacob Russell, which have resulted in a joint paper [BR20] and an appendix to a paper of Carolyn Abbott, Jason Behrstock, and Matthew Durham [ABD21]. The three pieces of work are summarised below.

1.1.1 Hierarchical hyperbolicity of graph products

In Chapter 4 we construct an explicit hierarchically hyperbolic structure for any graph product, generalising the standard hierarchically hyperbolic structures on right-angled Artin groups. 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

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

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

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

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 in [BHS17b] using the collection of induced subgraphs of the defining graph Γ, in the following way. 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 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 hierarchichally hyperbolic structure to arbitrary graph products. It is important to note, however, that arbitrary graph products *cannot* 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 [BHS19]. 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 ([BR20, Theorem 4.22])**.** *Let Γ be a finite simplicial graph, with each vertex v labelled by a non-trivial 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 [BHS19] and is explored further in [BHS17a]. 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 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 4.2.23 and 4.2.24.

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 minimum 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) [KK14]. 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 [BHS19]. We give a positive answer to this question, not just for right-angled Artin groups but for all graph products.

Corollary B ([BR20, Corollary 4.25]). *Let Γ be a finite simplicial graph, with each vertex v labelled by a non-trivial group G_v. Then 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 [Gen17, GM18] 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 [BHS19, Question D].

Theorem C ([BR20, Theorem 5.1]). *Let Γ be a finite simplicial graph, with each vertex v labelled by a non-trivial 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 groups that implies Theorem C when the vertex groups satisfy some natural, but non-trivial, additional hypotheses [BR18]. 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 to be the graph whose vertices correspond to elements of G_Γ, and where $g, h \in G_\Gamma$ are joined by an edge in $\mathbb{E}(\Gamma)$ whenever $g^{-1}h \in G_\Lambda \leqslant G_\Gamma$ 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 [Gen19b]. 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 [Gen19b].

Theorem D ([BR20, Theorems 5.14, 5.16]). *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 [Mei96].

1.1.2 Almost hierarchical hyperbolicity implies hierarchical hyperbolicity

The concept of an *almost hierarchically hyperbolic space* was introduced by Abbott, Behrstock, and Durham as a generalisation of hierarchical hyperbolicity [ABD21]. This *a priori* broader class of spaces is obtained by relaxing the container axiom, which requires that for each element W in a hierarchically hyperbolic structure $\mathfrak{S}$ there is a corresponding element of $\mathfrak{S}$ which contains everything orthogonal to W.

The proof of Theorem C requires the hierarchically hyperbolic structures of the vertex groups to be stitched into the fabric of the relatively hierarchically hyperbolic structure given by Theorem A, by first removing the nesting-minimal domains in the relative structure and then replacing them with the structures of the vertex groups. This causes one notable problem, however; this new structure does not satisfy the container axiom, providing us only with an almost hierarchically hyperbolic structure.

In Chapter 3 we fix this problem by showing that any almost hierarchically hyperbolic structure can be upgraded to a genuine hierarchically hyperbolic structure by adding a collection of dummy domains to serve as containers.

Theorem E ([ABD21, Theorem A.1])**.** *Any almost hierarchically hyperbolic space can be endowed with the structure of a hierarchically hyperbolic space.*

This allows us to complete the proof of Theorem C, while also streamlining proofs in [ABD21] by removing the need to navigate the subtleties of almost hierarchical hyperbolicity.

1.1.3 Non-positive curvature in graph braid groups

Given a topological space X, one can construct the *configuration space* $C_n(X)$ of n particles on X by taking the direct product of n copies of X and removing the diagonal. Informally, this space tracks the movement of the particles through X; removing the diagonal ensures the particles do not collide. One then obtains the *unordered configuration space* $UC_n(X)$ by taking the quotient by the action of the symmetric group on the coordinates of X^n. Finally, the *braid group* $B_n(X)$ is defined to be the fundamental group of $UC_n(X)$.

Braid groups have been a popular object of study since they were first introduced by Artin in 1926 [Art26]. Originally, these were studied geometrically as knots; see Figure 1.2. One can obtain the configuration space interpretation from this geometric model by taking horizontal cross-sections, each of which gives an arrangement of particles on a disc. Each cross-section can be thought of as a snapshot in time, tracking the locations of the particles as they weave between each other. This configuration space approach was first introduced by Fox [Fox62].

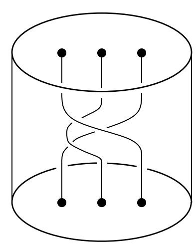

Figure 1.2: A 3–braid.

Classically, the space X is taken to be a disc, as in the above example. However, one may also study braid groups of other spaces. Taking X to be a manifold, Birman showed that braid groups are trivial in dimensions 3 and higher [Bir69, Theorem 1], therefore much work on braid groups is concentrated on the case where X is a surface. However, by weakening the manifold assumption, one also obtains interesting braid groups in dimension 1, namely those of graphs.

These so-called *graph braid groups* were first developed by Abrams [Abr00], who showed that they can be expressed as fundamental groups of non-positively curved cube complexes. Results of Genevois show that these cube complexes are in fact *special* [Gen19a], in the sense of Haglund and Wise [HW08]. By applying Behrstock–Hagen–Sisto's result that special cube complexes are hierarchically hyperbolic, it follows that $B_n(\Gamma)$ is a hierarchically hyperbolic group, but this structure does not admit a nice description in terms of Γ.

In Chapter 5, we construct an explicit hierarchically hyperbolic structure on a graph braid group $B_n(\Gamma)$. By expressing this structure in terms of the graph Γ, we are able to characterise when a graph braid group exhibits other aspects of non-positive curvature in terms of properties of Γ. For example, we are able to apply Behrstock–Hagen–Sisto's bounded orthogonality criterion [BHS17c, Corollary 2.16] to classify hyperbolicity of graph braid groups. A version of this theorem was first proved by Genevois [Gen19a, Theorem 4.1].

Theorem F. *A graph braid group $B_n(\Gamma)$ is hyperbolic if and only one of the following holds.*

(1) $n = 1$.

(2) $n = 2$ *and* Γ *does not contain two disjoint cycle subgraphs.*

(3) $n = 3$ *and* Γ *does not contain two disjoint cycle subgraphs, nor does it contain a disjoint star subgraph and cycle subgraph.*

(4) $n \geqslant 4$ *and* Γ *does not contain two disjoint subgraphs, each of which is a star or a cycle.*

We are able to recover another theorem of Genevois [Gen19a] by applying Behrstock–Hagen–Sisto's criteria for acylindrical hyperbolicity in hierarchically hyperbolic groups.

Theorem G. *Let* Γ *be a finite connected graph. The graph braid group* $B_n(\Gamma)$ *is either cyclic or acylindrically hyperbolic.*

Finally, we conjecture and partially prove a similar classification result for relative hyperbolicity of graph braid groups. This would answer a question of Genevois, generalising his characterisation of toral relative hyperbolicity [Gen19a]. We achieve this by applying

Russell's isolated orthogonality criterion [Rus20], which allows one to determine if a hierarchically hyperbolic space is relatively hyperbolic. Moreover, by adapting techniques developed by Levcovitz in his classification of right-angled Coxeter groups [Lev20], we show that one can simultaneously characterise when a graph braid group is *strongly thick*. Introduced by Behrstock–Druţu–Mosher as an obstruction to relative hyperbolicity [BDM09] and further developed by Behrstock–Druţu [BD14], thickness measures the complexity of coarse intersection patterns of non-negatively curved regions of a space. Thickness comes in various orders, with each order being a quasi-isometry invariant. It is conjectured that all hierarchically hyperbolic groups are thick or relatively hyperbolic; we seek to confirm this in the case of graph braid groups.

Following [Lev20], we introduce a sequence of hypergraphs which encode collections of mutually orthogonal domains arising in the hierarchically hyperbolic structure of a graph braid group $B_n(\Gamma)$. By analysing connectedness properties of these hypergraphs via the so-called *hypergraph index* and applying Russell's isolated orthogonality criterion, we claim that it is possible to determine whether the graph braid group is relatively hyperbolic. By construction, our hypergraphs show that any graph braid group which does not satisfy the isolated orthogonality criterion is in fact strongly thick, and moreover we obtain an upper bound on the order of thickness.

Conjecture H. Let Γ be a finite connected graph and let $n \geqslant 1$, $k \geqslant 0$ be integers.

(1) If $B_n(\Gamma)$ has hypergraph index k, then $B_n(\Gamma)$ is strongly thick of order at most k. In particular, $B_n(\Gamma)$ is not relatively hyperbolic.

(2) If $B_n(\Gamma)$ has hypergraph index ∞, then $B_n(\Gamma)$ is relatively hyperbolic.

Behrstock and Druţu show that one can obtain a lower bound on the order of strong thickness of a space by studying its *divergence* [BD14]. Levcovitz uses this to give a precise characterisation of orders of strong thickness of right-angled Coxeter groups, employing disc diagram techniques to measure divergence [Lev20]. By adapting these techniques to the setting of graph braid groups, we conjecture that the results of Conjecture H can be strengthened even further.

Conjecture I. $B_n(\Gamma)$ has hypergraph index k if and only if it is strongly thick of order k.

1.2 Outline of the thesis

In Chapter 2 we lay the foundations for the main results in Chapters 3, 4, and 5. We begin with a quick overview of the basic concepts from coarse geometry (Section 2.1), including a glossary of some important graph theoretic terms (Section 2.1.1), followed by a brief summary of relative hyperbolicity and thickness (Section 2.2) as well as acylindrical hyperbolicity (Section 2.3). We then give a more in-depth analysis of the geometry of cube complexes (Section 2.4), including a new proof of hyperbolicity of the contact graph (Theorem 2.4.16). Cube complexes are an important prerequisite for our approach to graph braid groups (Section 2.5) and a motivation for the study of quasi-median graphs (Section 2.6). Quasi-median graphs in turn form the basis for the geometry of graph products (Section 2.6.1), which are a generalisation of right-angled Artin groups (Section 2.4.1) and right-angled Coxeter groups (Section 2.4.2). We conclude the chapter by formally introducing the concept of hierarchical hyperbolicity (Section 2.7), which underpins the majority of the results of this thesis. We spend some time developing this theory, reviewing methods of detecting other forms of hyperbolicity in hierarchically hyperbolic spaces (Section 2.7.1) and introducing two important variants—relative hierarchical hyperbolicity (Section 2.7.2) and almost hierarchical hyperbolicity (Section 2.7.3)—as well as delving into the details of hierarchically hyperbolic structures on certain CAT(0) cube complexes (Section 2.7.4), which will be key to understanding hierarchically hyperbolic structures on both graph products and graph braid groups.

Chapter 3 is devoted primarily to the proof of Theorem E, showing that any almost hierarchically hyperbolic structure can be promoted to a hierarchically hyperbolic structure. We conclude the chapter with a description of how this result is applied in the paper of Abbott, Behrstock, and Durham.

In Chapter 4 we set about proving our theorems on graph products. In Section 4.1, we set up our proof of the relative hierarchical hyperbolicity of graph products by defining the necessary spaces, projections, and relations. In Section 4.2, we show the spaces, projections, and relations defined in Section 4.1 satisfy the axioms of a relative HHG (or non-relative HHS in the case of the syllable metric). This completes the proofs of Theorem A and Corollary B. Section 4.3 is devoted to applications of this hierarchically hyperbolic structure. We start by proving graph products of HHGs are HHGs (Theorem C) in Section 4.3.1, which requires the technical results shown in Chapter 3. In Section 4.3.2, we record our proof of Meier's hyperbolicity criteria and in Section 4.3.3, we classify when Genevois' electrification has infinite diameter and when it is a quasi-line, proving Theorem D.

Finally, Chapter 5 contains our results on graph braid groups. We spend Section 5.1 constructing an explicit hierarchically hyperbolic structure for graph braid groups, first by studying the cubical structure and how it relates to properties of the graph (Section 5.1.1) and then translating this into a hierarchically hyperbolic structure (Section 5.1.2). We then apply this hierarchically hyperbolic structure in Section 5.2 to detect other forms of hyperbolicity. In particular, we classify when a graph braid group is hyperbolic or acylindrically hyperbolic (Section 5.2.1), proving Theorems F and G. We then introduce the notion of the hypergraph index in Section 5.2.2 and use this to partially prove Conjecture H, which characterises when a graph braid group is relatively hyperbolic or thick.

Chapter 2

Background

Notation 2.0.1. We set the following conventions for the notation we shall use.

- Given a metric space X, $B_x(R)$ denotes the ball of radius R centred at the point x.
- Given a metric space X and a subspace $A \subseteq X$, $N_R(A)$ denotes the closed R–neighbourhood of the subspace A.
- The identity element of a group is denoted e.
- $\mathbb{N}$ denotes the positive integers.
- The set of vertices of a graph Γ is denoted $V(\Gamma)$ and the set of edges is denoted $E(\Gamma)$.

2.1 Coarse geometry

In this section we shall review some basic tools of geometric group theory which will be used throughout this thesis. The first and foremost tool is the *Cayley graph* of a group, which is the primary method of studying a group's geometry. This graph allows one to put a natural discrete metric on any finitely generated group. Due to the discrete nature of this form of geometry, it is often referred to as *coarse geometry*.

Definition 2.1.1 (Cayley graph, word metric)**.** Let G be a group and let S be a generating set for G. The *Cayley graph* $\mathrm{Cay}(G, S)$ of G with respect to S is the graph whose vertices are the elements of G, and where two vertices g and h are connected by an edge if and only if $g^{-1}h = s$ for some $s \in S \cup S^{-1}$. In this case, we label the edge by s. The *word metric* d_S on G with respect to S is the graph metric on $\mathrm{Cay}(G, S)$; that is, $\mathsf{d}_S(g, h)$ is equal to the shortest distance between the vertices g and h in $\mathrm{Cay}(G, S)$, where each edge has length 1.

The geometry of the Cayley graph $\mathrm{Cay}(G, S)$ is intimately related to algebraic properties of the group G. For example, the distance between two elements $g, h \in G$ in the word metric d_S is equal to the minimum number of elements required to express $g^{-1}h$ as a product of elements of $S \cup S^{-1}$. One must beware that a group can have many Cayley graphs, depending on the choice of generating set, meaning this geometry is not fully determined by the choice of group G. However, so long as these generating sets are finite, the distances in any two Cayley graphs are equivalent up to a multiplicative and additive constant; we say the Cayley graphs are *quasi-isometric*. By restricting ourselves to geometric properties of Cayley graphs that are invariant under quasi-isometry, we therefore obtain genuine invariants of the group itself.

Definition 2.1.2 (Coarsely Lipschitz, quasi-isometry)**.** Let (X, d_X) and (Y, d_Y) be metric spaces and let $L \geqslant 1$, $C \geqslant 0$. A function $f : X \to Y$ is said to be (L, C)–*coarsely Lipschitz* if for all $x, x' \in X$, we have

$$\mathsf{d}_Y(f(x), f(x')) \leqslant L\mathsf{d}_X(x, x') + C.$$

A function $f : X \to Y$ is an (L, C)–*quasi-isometric embedding* if for all $x, x' \in X$, we have

$$\frac{1}{L}\mathsf{d}_X(x, x') - C \leqslant \mathsf{d}_Y(f(x), f(x')) \leqslant L\mathsf{d}_X(x, x') + C.$$

If in addition $Y \subseteq N_C(f(X))$, we say f is an (L, C)–*quasi-isometry*. Two metric spaces

(X, d_X) and (Y, d_Y) are said to be *quasi-isometric* if there exists a quasi-isometry $f : X \to Y$ between them.

In light of this, the appropriate coarse geometric analogue of a geodesic is a *quasi-geodesic*, defined as a quasi-isometric embedding of a closed interval $I \subseteq \mathbb{R}$.

Definition 2.1.3 (Quasi-geodesic)**.** Let X be a metric space. An (L, C)–*quasi-geodesic* in X is an (L, C)–quasi-isometric embedding $\gamma : I \to X$ for some closed interval $I \subseteq \mathbb{R}$. In particular, if $L = 1$ and $C = 0$ then γ is a geodesic. We say X is an (L, C)–*quasi-geodesic space* if for any two points $x, x' \in X$ there exists an (L, C)–quasi-geodesic $\gamma : [0, l] \to X$ such that $\gamma(0) = x$ and $\gamma(l) = x'$.

This leads to a natural coarse version of convexity.

Definition 2.1.4 (Quasi-convex)**.** A subset A of a metric space X is (L, C)–*quasi-convex* if any two points in A can be connected in $N_C(A)$ by an (L, L)–quasi-geodesic.

Example 2.1.5. The following three examples of quasi-isometries will be critical to our understanding of the coarse geometry of groups.

(1) The group $\mathbb{Z}^n$ with the word metric induced by its standard generating set is quasi-isometric to Euclidean n–space. For this reason, we often refer to quasi-isometrically embedded copies of $\mathbb{Z}^n$ in a group G as *flats* in the group.

(2) A finitely generated group G is quasi-isometric to any finite index subgroup $H \leqslant G$. For this reason, we shall often speak of a group *virtually* having a property. This simply means there is a finite index subgroup which has this property.

(3) **(Milnor–Švarc Lemma.)** Let G be a group acting properly discontinuously and co-compactly on a proper geodesic metric space X. Then G is finitely generated and quasi-isometric to X. We say G acts *geometrically* on X.

Note that the fact that any two Cayley graphs of a finitely generated group G are quasi-isometric follows as a consequence of the Milnor–Švarc lemma, by considering the action of G on these two metric spaces.

With the notion of quasi-isometry in mind, we can now begin producing geometric group invariants by finding properties which are invariant under quasi-isometry. One such invariant is a coarse version of hyperbolicity, originally introduced by Gromov in [Gro87] and now a cornerstone of geometric group theory.

Definition 2.1.6 (δ–hyperbolic)**.** Let X be a (quasi-)geodesic space and let $\delta > 0$. We say X is *δ–hyperbolic* if every (quasi-)geodesic triangle in X satisfies the property that any side is contained in the δ–neighbourhood of the union of the other two sides.

Hyperbolic spaces have a number of nice properties. For example, hyperbolic spaces have a natural notion of *boundary*.

Definition 2.1.7 (Gromov boundary)**.** Let X be a proper geodesic hyperbolic space and fix a point $O \in X$. Define two geodesic rays $\gamma_1 : [0, \infty) \to X$ and $\gamma_2 : [0, \infty) \to X$ with $\gamma_1(0) = \gamma_2(0) = O$ to be equivalent if there exists some $K \geqslant 0$ such that $\mathsf{d}_X(\gamma_1(t), \gamma_2(t)) \leqslant K$ for all $t \geqslant 0$. Denote the equivalence class of a geodesic ray γ by $[\gamma]$. The *Gromov boundary* ∂X of X is defined to be the set of equivalence classes $\partial X = \{[\gamma] \mid \gamma \text{ is a geodesic ray in } X\}$.

Given a group G acting on a hyperbolic space X, one can study the behaviour of orbits of elements $g \in G$.

Definition 2.1.8 (Loxodromic)**.** An element g is said to be *loxodromic* if the map $\mathbb{Z} \to X$ defined by $n \mapsto g^n x$ is a quasi-isometry for some (equivalently, any) $x \in X$. In particular, the orbits of g in X form quasi-geodesics, with precisely two limit points in ∂X. We say two loxodromic elements $g, h \in G$ are *independent* if they do not share any limit points.

2.1.1 Graph terminology

Graphs prove to be highly useful tools in geometric group theory, even outside of the specific example of the Cayley graph, by giving us powerful ways of keeping track of combinatorial data. As such, it shall be useful to review some concepts from graph theory.

Definition 2.1.9 (Simplicial)**.** Let Γ be a graph. We say Γ is *simplicial* if no two edges connect the same pair of vertices and no edge both begins and ends at the same vertex. That is, every edge of Γ may be expressed uniquely as an unordered pair $\{v, w\}$ of distinct vertices of Γ corresponding to its endpoints.

Definition 2.1.10 (Induced subgraph)**.** Let Γ be a graph and let $\Lambda \subseteq \Gamma$ be a subgraph of Γ. We say Λ is an *induced* subgraph of Γ if the edges of Λ are precisely the edges of Γ whose endpoints are vertices of Λ.

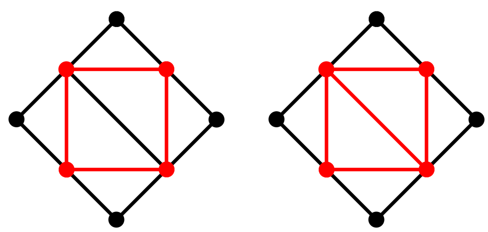

Figure 2.1: The red square on the left is not an induced subgraph, as it does not contain the diagonal edge. The subgraph on the right is an induced subgraph.

Definition 2.1.11 (Cycle, tree)**.** Let Γ be a graph. A *cycle* in Γ is a sequence of distinct vertices $v_1, \ldots, v_n$ such that v_i and v_{i+1} are connected by an edge for each $1 \leqslant i < n$, and v_n is connected to v_1 by an edge. A *cycle* also refers to the subgraph of Γ given by the union of these vertices and edges. We say Γ is a *tree* if it is connected and contains no cycles.

Definition 2.1.12 (Star graph)**.** A connected graph Γ is a *star graph* if it consists of one vertex of valence $m \geqslant 3$ and m vertices of valence 1.

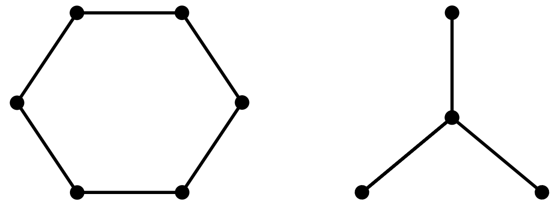

Figure 2.2: A cycle and a star graph.

Definition 2.1.13 (Complete, clique)**.** A simplicial graph Γ is *complete* if every pair of vertices of Γ is connected by an edge. An induced subgraph Λ of a simplicial graph Γ is called a *clique* if Λ is a complete subgraph of Γ.

Definition 2.1.14 (Star, link, and join)**.** Let Γ be a simplicial graph and Λ an induced subgraph of Γ. The *link* of Λ, denoted $\mathrm{lk}(\Lambda)$, is the subgraph of Γ induced by the set of vertices of $\Gamma \setminus \Lambda$ that are connected to every vertex of Λ. The *star* of Λ, denoted $\mathrm{st}(\Lambda)$, is $\Lambda \cup \mathrm{lk}(\Lambda)$. We say Λ is a *join* if it can be written as $\Lambda = \Lambda_1 \sqcup \Lambda_2$ where Λ_1 and Λ_2 are non-empty induced 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 $\Lambda_1 \star \Lambda_2$. In particular, $\mathrm{st}(\Lambda)$ is the join $\Lambda \star \mathrm{lk}(\Lambda)$.

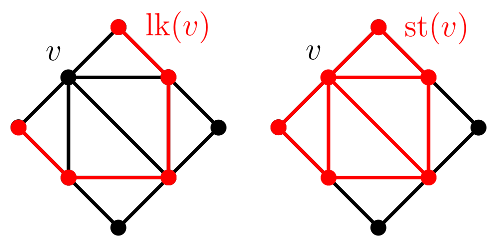

Figure 2.3: The link and star of a vertex.

Generalising the idea that edges in a simplicial graph are represented by unordered pairs of vertices, one can develop a notion of a *hypergraph*, which has ‘hyperedges’ consisting of unordered collections of any number of vertices.

Definition 2.1.15 (Hypergraph)**.** A *hypergraph* Γ is a set of vertices $V(\Gamma)$ and a set of hyperedges $\mathcal{E}(\Gamma)$, where each hyperedge $E \in \mathcal{E}(\Gamma)$ is a subset of $V(\Gamma)$.

2.2 Relative hyperbolicity and thickness

2.2.1 Relative hyperbolicity

There are many interesting and important classes of groups that are not hyperbolic in the homogeneous sense of δ–hyperbolicity, but do still exhibit some hyperbolic behaviour. One of the first attempts to capture this was the notion of *relative hyperbolicity*, first introduced by Gromov in his seminal paper and since improved upon by Bowditch, Farb, and others [Gro87, Far98, Bow12, DS05, Osi06]. Roughly, a group is relatively hyperbolic if it is δ–hyperbolic outside of some isolated collection of peripheral subgroups.

Although we shall mostly be using relative hyperbolicity as a black box, a definition is included here for the sake of completeness. Note that there are a number of equivalent definitions; see [Hru10, Sis12] for surveys of these and proofs of their equivalence. The version of the definition given below is due to Farb.

Definition 2.2.1 (Relatively hyperbolic; [Far98])**.** Let G be a finitely generated group with Cayley graph Γ and let $\mathcal{P}$ be a finite collection of finitely generated subgroups of G. The *coned-off Cayley graph* $\hat{\Gamma} = \hat{\Gamma}(\mathcal{P})$ of G with respect to $\mathcal{P}$ is obtained from Γ by adding a vertex v_{gP} for each coset gP of each subgroup $P \in \mathcal{P}$, and adding edges connecting v_{gP} to each vertex of gP in Γ. The group G is *hyperbolic relative to* $\mathcal{P}$ if there exists some $\delta \geqslant 0$ such that $\hat{\Gamma}$ is δ–hyperbolic and the pair $(G, \mathcal{P})$ satisfies the *bounded coset penetration* property, defined below.

Definition 2.2.2 (Bounded coset penetration; [Far98])**.** Let G be a finitely generated group with Cayley graph Γ and let $\mathcal{P}$ be a finite collection of finitely generated subgroups of G. A path w in Γ is a word in the generators of G. By searching w for maximal subwords z contained in cosets gP for $P \in \mathcal{P}$, one can obtain a path $\hat{w}$ of $\hat{\Gamma}$ by replacing subpaths given by such subwords z with two edges connected to the corresponding cone point v_{gP}. If $\hat{w}$ passes through the cone point v_{gP}, we say w *penetrates* the coset gP. If $\hat{w}$ is a geodesic of $\hat{\Gamma}$, we say w is a *relative geodesic* of Γ. If $\hat{w}$ is an L–quasi-geodesic of $\hat{\Gamma}$, we say w is a *relative L–quasi-geodesic* of Γ. The path w is a *path without backtracking* if w never returns to a coset that it penetrates.

The pair $(G, \mathcal{P})$ satisfies the *bounded coset penetration* property if for all $L \geqslant 1$ there exists a constant $c = c(L)$ so that for every pair u, v of relative L–quasi-geodesics without backtracking and with the same endpoints, the following conditions hold.

(1) If u penetrates a coset gP but v does not penetrate gP, then u travels a Γ–distance of at most c in gP.

(2) If u and v both penetrate a coset gP, then the vertices of Γ where u and v first enter gP are at Γ–distance at most c from each other. Similarly, the vertices of Γ where u and v last exit gP are at distance at most c from each other.

A classical example of a relatively hyperbolic group is the fundamental group of a complete, finite-volume, cusped hyperbolic manifold. This is hyperbolic relative to its cusp subgroups [Far98, Theorem 4.11]. Other examples include free products of finitely-generated groups and non-uniform lattices in rank 1 symmetric spaces.

Importantly, relative hyperbolicity is a quasi-isometry invariant property [Dru09, Theorem 1.2]. In fact, relatively hyperbolic groups are quasi-isometrically rigid in the following stronger sense.

Theorem 2.2.3 (Rigidity of relatively hyperbolic groups; [BDM09, Theorem 4.8])**.** *Let G be a finitely generated group and suppose G is hyperbolic relative to a finite collection $\mathcal{P}$ of finitely generated subgroups, none of which are relatively hyperbolic. If G' is a finitely generated group quasi-isometric to G, then G' is hyperbolic relative to some finite collection $\mathcal{P}'$ of finitely generated subgroups, each of which is quasi-isometric to some subgroup in $\mathcal{P}$.*

Remark 2.2.4. Although we concern ourselves primarily with relative hyperbolicity of groups, a notion of relative hyperbolicity also exists for metric spaces in general; see e.g.

[DS05].

2.2.2 Thickness and divergence

Thickness was introduced by Behrstock–Druţu–Mosher as an obstruction to relative hyperbolicity [BDM09], and further developed by Behrstock–Druţu [BD14]. Roughly, thickness measures the complexity of coarse intersection patterns of non-negatively curved regions of a space. We shall use the strong version of thickness, defined in [BD14].

Definition 2.2.5 (Strongly thick)**.** A metric space X is *strongly thick of order* 0 if none of its asymptotic cones have cut points and every point in X is uniformly close to a bi-infinite uniform quasi-geodesic. In particular, X is strongly thick of order 0 if X is quasi-isometric to a product of two infinite-diameter metric spaces, contains a bi-infinite quasigeodesic, and admits a cocompact group action.

X is *strongly thick of order at most* n if there exist $C \geqslant 0$, an index set I, and a collection $\{P_\alpha\}_{\alpha \in I}$ of quasi-convex subsets of X satisfying the following three conditions.

(1) **(Thick pieces.)** Each P_α is strongly thick of order at most $n-1$.

(2) **(Coarse covering.)** $X \subseteq N_C(\bigcup_{\alpha \in I} P_\alpha)$.

(3) **(Thick chaining.)** For each pair $P_\alpha, P_{\alpha'}$ there is a sequence $P_\alpha = P_0, P_1, \ldots, P_k = P_{\alpha'}$ such that $N_C(P_i) \cap N_C(P_{i+1})$ has infinite diameter for each $0 \leqslant i \leqslant k-1$.

Each order of thickness is a quasi-isometry invariant [BDM09, Remark 7.2]. In this way, thickness and relative hyperbolicity may be used in tandem to aid in distinguishing quasi-isometry classes of groups and spaces.

Remark 2.2.6 (The relatively hyperbolic–thick dichotomy)**.** It is important to note that there is not always a strict dichotomy between relative hyperbolicity and thickness (see [BDM09, Section 7.1]). However, in many cases, there is. For example, the dichotomy has been proven for mapping class groups, 3–manifold groups, and Artin groups [BDM09], Teichmüller space [BM08], Coxeter groups [BHS17d], and free-by-cyclic groups [Hag19].

When studying groups, it is often more convenient to use the following algebraic version of thickness. Note that algebraic thickness of order n implies metric thickness of order n, by [BDM09, Proposition 7.6].

Definition 2.2.7 (Strongly algebraically thick)**.** A finitely generated group G is *strongly algebraically thick of order* 0 if none of its asymptotic cones have cut points and every point in G is uniformly close to a bi-infinite uniform quasi-geodesic. In particular, G is strongly algebraically thick of order 0 if it is quasi-isometric to a product of two infinite-diameter metric spaces.

G is *strongly algebraically thick of order at most* n if there exists a finite index set I and a collection $\{H_\alpha\}_{\alpha\in I}$ of quasi-convex subgroups of G satisfying the following three conditions.

(1) **(Thick pieces.)** Each H_α is strongly algebraically thick of order at most $n-1$.

(2) **(Coarse covering.)** $\bigcup_{\alpha\in I} H_\alpha$ generates a finite-index subgroup of G.

(3) **(Thick chaining.)** For each pair $H_\alpha, H_{\alpha'}$ there is a sequence $H_\alpha = H_0, H_1, \ldots, H_k = H_{\alpha'}$ such that $H_i \cap H_{i+1}$ has infinite diameter for each $0 \leqslant i \leqslant k-1$.

In practice, it can be difficult to show that a group or space has a specific order of thickness; by virtue of the definition, we often initially only obtain an upper bound on the order. To facilitate finding the exact order of thickness, it is often convenient to compute the *divergence* of the group or space.

Definition 2.2.8 (Divergence)**.** Let (X, d_X) be a geodesic metric space, and let $0 < \delta < 1$ and $\gamma \geqslant 0$. Given $a, b, c \in X$ with $\mathsf{d}_X(c, \{a, b\}) = r > 0$, define $\mathrm{div}_\gamma(a, b, c; \delta)$ to be the infimum of the lengths of paths in X connecting a to b and avoiding the ball $B_c(\delta r - \gamma)$. If no such path exists, define $\mathrm{div}_\gamma(a, b, c; \delta) = \infty$. The *divergence function* $\mathrm{Div}_\gamma(n; \delta)$ of X is defined to be the supremum of all $\mathrm{div}_\gamma(a, b, c; \delta)$ with $\mathsf{d}_X(a, b) \leqslant n$.

Behrstock and Druţu show that the rate of divergence gives a lower bound on the order of strong thickness.

Theorem 2.2.9 (Divergence bounds thickness; [BD14, Corollary 4.17])**.** *Let X be a geodesic metric space. If the divergence $\mathrm{Div}_\gamma(x;\delta)$ is at least polynomial of order $n+1$ for every $0<\delta<\frac{1}{54}$ and every $\gamma \geqslant 0$, then X is strongly thick of order at least n.*

2.3 Acylindrical hyperbolicity

Another generalisation of δ–hyperbolicity is the notion of *acylindrical hyperbolicity*. This version of hyperbolicity was developed by Osin [Osi16], who built upon ideas of Sela and Bowditch to bring it to its current form [Sel97, Bow08].

Definition 2.3.1 (Acylindrical)**.** The action of a group G on a metric space X is *acylindrical* if for all $E \geqslant 0$, there exist $R, N \geqslant 0$ so that if $x, y \in X$ satisfy $\mathsf{d}_X(x,y) \geqslant R$, then there are at most N elements $g \in G$ such that $\mathsf{d}_X(x,gx) \leqslant E$ and $\mathsf{d}_X(y,gy) \leqslant E$. We say G is *acylindrically hyperbolic* if it admits a non-elementary acylindrical action on a hyperbolic space X by isometries. (An action is non-elementary if it has two independent loxodromic elements.)

One notable consequence of acylindricity is the following classification due to Osin.

Theorem 2.3.2 ([Osi16, Theorem 1.1])**.** *Let G be a group acting acylindrically on a hyperbolic space. Then G satisfies exactly one of the following.*

(1) G has bounded orbits.

(2) G is virtually cyclic and contains a loxodromic element.

(3) G contains infinitely many independent loxodromic elements.

2.4 Cube complexes

Definition 2.4.1 (Cube complex)**.** Let $n \geqslant 0$. An *n–cube* is a Euclidean cube $[-\frac{1}{2}, \frac{1}{2}]^n$. A *face* of a cube is a subcomplex obtained by restricting one or more of the coordinates to $\pm\frac{1}{2}$. A *cube complex* is a CW complex where each cell is a cube and the attaching maps are given by isometries along faces.

We will often refer to the 0–cubes of a cube complex X as *vertices*, the 1–cubes as *edges*, and the 2–cubes as *squares*.

Definition 2.4.2 (Non-positively curved, CAT(0), cubical group)**.** Let X be a cube complex. The *link* $\mathrm{link}(v)$ of a vertex v of X is a simplicial complex defined as follows.

- The vertices of $\mathrm{link}(v)$ are the edges of X that are incident at v.
- n vertices of $\mathrm{link}(v)$ span an n–simplex if the corresponding edges of X are faces of a common cube.

The complex $\mathrm{link}(v)$ is said to be *flag* if n vertices $v_1, \ldots, v_n$ of $\mathrm{link}(v)$ span an n–simplex if and only if v_i and v_j are connected by an edge for all $i \neq j$. A cube complex X is *non-positively curved* if the link of each vertex of X is flag and contains no bigons (that is, no loops consisting of two edges). A cube complex X is *CAT(0)* if it is non-positively curved and simply connected. A group is said to be *cubical* if it acts geometrically on a CAT(0) cube complex.

Henceforth, all the cube complexes we use shall be non-positively curved, unless otherwise specified. The link condition tells us that a non-positively curved cube complex X is determined by its 1–skeleton $X^{(1)}$. In general, we shall therefore work in this 1–skeleton, where we use the graph metric, denoted d_X. The resulting metric space is a *median graph*, as shown in [Che00, Theorem 6.1].

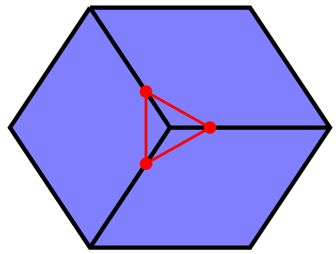

Figure 2.4: This 2–dimensional cube complex is not non-positively curved because the link of the central vertex (shown in red) is not flag.

Definition 2.4.3 (Median graph)**.** Let Γ be a graph with graph metric d. The resulting metric graph is a *median graph* if for any three distinct vertices u, v, w of Γ, there exists a unique vertex $m = m(u, v, w)$ such that

$$\mathsf{d}(u, v) = \mathsf{d}(u, m) + \mathsf{d}(m, v),$$
$$\mathsf{d}(v, w) = \mathsf{d}(v, m) + \mathsf{d}(m, w),$$
$$\mathsf{d}(w, u) = \mathsf{d}(w, m) + \mathsf{d}(m, u).$$

The vertex m is the *median* of u, v, and w.

Many important properties of cube complexes may be proved from this combinatorial point of view of median graphs; see [Che00, Section 6.2] for an overview of such arguments, as well as their relation to the more geometric point of view of *hyperplanes*, which is the perspective we shall take.

Definition 2.4.4 (Mid-cube, hyperplane, combinatorial hyperplane, carrier)**.** Let X be a cube complex. A *mid-cube* of a cube $C \cong [-\frac{1}{2}, \frac{1}{2}]^n$ of X is obtained by restricting one of the coordinates of C to 0. Each mid-cube has two isometric associated faces of C, obtained by restricting this coordinate to $\pm\frac{1}{2}$ instead of 0. A *hyperplane* H of X is a maximal connected union of mid-cubes. A *combinatorial hyperplane* associated to H is a maximal connected union of faces associated to mid-cubes of H. The *closed (resp. open) carrier* of H is the union of all closed (resp. open) cubes of X which contain mid-cubes of H.

Convention 2.4.5. We will almost always be using the closed version of the carrier of a hyperplane H, therefore we will refer to this version as simply the 'carrier' of H. We denote this closed version of the carrier by $N(H)$.

A result of Chepoi tells us that carriers and combinatorial hyperplanes form convex subcomplexes of X [Che00, Proposition 6.6]. The combinatorial hyperplanes associated to a given hyperplane are also 'parallel', in the following sense.

Definition 2.4.6 (Parallelism and separation)**.** Let H be a hyperplane of a cube complex X. We say that H is *dual* to an edge E of X if H contains a mid-cube which intersects E. We say that H *crosses* a subcomplex Y of X if there exists some edge E of Y that is dual to H. We say that H crosses another hyperplane H' if it crosses a combinatorial hyperplane associated to H'. We say H *separates* two subcomplexes Y, Y' of X if Y and Y' are contained in two different connected components of $X \smallsetminus H$. Two subcomplexes Y, Y' of X are *parallel* if each hyperplane H of X crosses Y if and only if it crosses Y'.

Parallelism of combinatorial hyperplanes can in fact be generalised to a much stronger result which characterises parallelism of any convex subcomplexes.

Lemma 2.4.7 ([BHS17b, Lemma 2.4])**.** *Let K and K' be convex subcomplexes of a CAT(0) cube complex X. The following are equivalent.*

(1) K and K' are parallel.

(2) There is a cubical isometric embedding $i : K \times [0, l] \to X$ such that $i(K \times \{0\}) = K$ and $i(K \times \{l\}) = K'$, and for each $x \in K$, $i(\{x\} \times [0, l])$ is a geodesic segment in X whose dual hyperplanes are precisely those separating K from K'.

Remark 2.4.8. Parallelism defines an equivalence relation on the edges of a cube complex X. In particular, two edges are in the same equivalence class (or *parallelism class*) if and only if they are dual to the same hyperplane. Therefore, one may instead consider hyperplanes of X to be parallelism classes of edges of X. One may also define an orientation $\overrightarrow{H}$ on a hyperplane H by taking an equivalence class of oriented edges. In this case, we say that $\overrightarrow{H}$ is dual to the oriented edges in this class.

Definition 2.4.9 (Osculation)**.** We say H *directly self-osculates* if there exist two oriented edges $\overrightarrow{E_1}, \overrightarrow{E_2}$ dual to $\overrightarrow{H}$ that have the same initial or terminal vertex but do not span a square. We say two hyperplanes H_1, H_2 *inter-osculate* if they intersect and there exist dual edges E_1, E_2 of H_1, H_2, respectively, such that E_1 and E_2 share a common endpoint but do not span a square. See Figure 2.5.

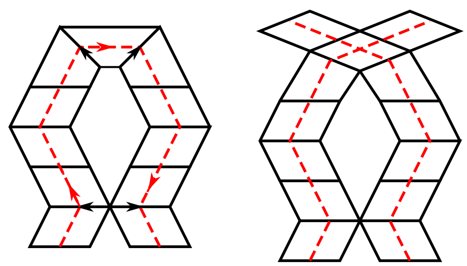

Figure 2.5: A directly self-osculating hyperplane and a pair of inter-osculating hyperplanes.

Definition 2.4.10 (Special)**.** A non-positively curved cube complex X is said to be *special* if its hyperplanes satisfy the following properties.

(1) Hyperplanes are two-sided; that is, the open carrier of a hyperplane H is homeomorphic to $H \times (-\frac{1}{2}, \frac{1}{2})$.

(2) Hyperplanes of X do not self-intersect.

(3) Hyperplanes of X do not directly self-osculate.

(4) Hyperplanes of X do not inter-osculate.

The following fact is an easy but noteworthy consequence of the above definitions.

Fact 2.4.11. *The universal cover of a non-positively curved cube complex is a CAT(0) cube complex. In particular, the universal cover of a special cube complex is a CAT(0) cube complex.*

Conversely, a result of Haglund and Wise tells us that a CAT(0) cube complex is itself special.

Proposition 2.4.12 ([HW08, Example 3.3])**.** *Any CAT(0) cube complex is special.*

The hyperplanes of a CAT(0) cube complex X enjoy many useful properties beyond the ones in the definition of 'special', as outlined below. We shall see in Section 2.7.4 that these hyperplane properties translate into powerful geometric properties, as well as algebraic properties of $\pi_1(Y)$ in the case where X is the universal cover of a special cube complex Y.

Proposition 2.4.13 (Hyperplane properties; [Sag95, Che00, Hag14])**.** *Let X be a CAT(0) cube complex.*

(1) Any hyperplane H of X is itself a CAT(0) cube complex. Moreover, its hyperplanes are of the form $H \cap J$, where $J \neq H$ is a hyperplane of X.

(2) Each hyperplane separates X into two connected components.

(3) If H is a hyperplane of X, then its carrier $N(H)$ and any combinatorial hyperplane associated to H are convex in X.

(4) If H is a hyperplane of X, then each connected component of $X \setminus H$ is convex in X.

(5) A continuous path γ in $X^{(1)}$ is a geodesic if and only if γ intersects each hyperplane at most once. In particular, if x and y are two vertices of X, then $\mathsf{d}_X(x, y)$ is equal to the number of hyperplanes of X separating x and y.

(6) If two hyperplanes H_1 and H_2 cross, then there exists a unique pair of edges E_1, E_2 dual to H_1 and H_2, respectively, such that E_1 and E_2 span a square of X.

(7) Let K be a connected convex subcomplex of X. Any pair of hyperplanes of X that cross each other and also cross K must necessarily cross inside K (that is, the square given by Property 6 is contained in K).

Proof. Properties (1) and (2) are proved in [Sag95, Theorems 4.11, 4.10]. Property (3) is shown in the proof of [Che00, Proposition 6.6]. Property (4) follows immediately from Properties (2) and (5). Property (7) is given in [Hag14, Lemma 2.13].

Property (5) follows from convexity of combinatorial hyperplanes. Indeed, let γ be a geodesic in X and suppose a hyperplane H crosses γ more than once. By orienting γ, we can order the edges of γ that are dual to H. Let E and E' be two consecutive dual edges with respect to this order. Let x, y and x', y' be the vertices of E and E', respectively, such that x, x' are contained in a combinatorial hyperplane H_+ and y, y' are contained in a combinatorial hyperplane H_-. By convexity, the subsegment of γ connecting x and x' must be contained in H_+ and the subsegment of γ connecting y and y' must be contained in H_-. However, this implies that $H_+ \cap H_- \neq \varnothing$, contradicting two-sidedness of H.

Property (6) follows from Property (5). Indeed, suppose two hyperplanes H_1 and H_2 of X cross. Then there exists some edge E_1 dual to H_1 that is contained in a combinatorial hyperplane associated to H_2. By definition, every edge of a combinatorial hyperplane spans a square with an edge dual to the associated hyperplane. Thus, there exists some edge E_2 dual to H_2 that spans a square with E_1. Moreover, by Property (5), H_1 cannot cross H_2 more than once, as it would have to cross a combinatorial hyperplane and thus a geodesic more than once. Therefore, the edges E_1 and E_2 must be unique. $\square$

As discussed in [BHS17b, Section 2.1], CAT(0) cube complexes admit projections to convex subcomplexes, which satisfy particularly nice geometric properties, as summarised below. We call such projections *gate maps*.

Proposition 2.4.14 (Gate map; [BHS17b, Section 2.1]). *Let X be a CAT(0) cube complex. For each convex subcomplex $K \subseteq X$, there exists a map $\mathfrak{g}_K : X \to K$ satisfying the following* *properties.*

(1) *For all* $x, y \in X$, $\mathsf{d}_X(\mathfrak{g}_K(x), \mathfrak{g}_K(y)) \leqslant \mathsf{d}_X(x, y)$.

(2) *For all* $x \in X$, $\mathfrak{g}_K(x)$ *is the unique vertex of* K *such that* $\mathsf{d}_X(x, \mathfrak{g}_K(x)) = \mathsf{d}_X(x, K)$.

(3) *Any hyperplane of* X *that separates* x *from* $\mathfrak{g}_K(x)$ *separates* x *from* K.

(4) *If* $x, y \in X$ *and* H *is a hyperplane of* X *separating* $\mathfrak{g}_K(x)$ *and* $\mathfrak{g}_K(y)$, *then* H *separates* x *and* y, *so that* x *and* $\mathfrak{g}_K(x)$ *(resp.* y *and* $\mathfrak{g}_K(y)$*) are contained in the same connected component of* $X \smallsetminus H$.

Another important tool in the study of CAT(0) cube complexes is the *contact graph*, a structural invariant originally introduced by Hagen [Hag12] which encodes the hyperplane structure of a cube complex while ignoring the CAT(0) geometry. A number of coarse properties of cube complexes can be inferred from simple properties of the contact graph; in fact, we will see in Section 2.7.4 that the contact graph plays an important role in developing hierarchically hyperbolic structures for CAT(0) cube complexes.

Definition 2.4.15 (Contact graph)**.** Let X be a CAT(0) cube complex. The contact graph of X, denoted $C_0(X)$, is the simplicial graph whose vertex set is the set $\mathcal{H}$ of hyperplanes of X, and where two vertices $H_1, H_2 \in \mathcal{H}$ are connected by an edge if the (closed) carriers $N(H_1), N(H_2)$ intersect. We say that H_1 and H_2 *contact.*

One key property of the contact graph is that it is δ–hyperbolic. In fact, Hagen showed it is a quasi-tree [Hag14, Theorem 4.1]. Below, we give a simple alternate proof of hyperbolicity of $C_0(X)$, relying only on the hyperplane properties of CAT(0) cube complexes detailed above. In particular, we obtain an explicit hyperbolicity constant of $\delta = \frac{5}{2}$.

Theorem 2.4.16 (The contact graph is hyperbolic)**.** *Let* X *be a CAT(0) cube complex. The contact graph* $C_0(X)$ *is* $\frac{5}{2}$*–hyperbolic.*

Proof. Let $x, y, z \in C_0(X)$ be three distinct points, and let $\gamma_1, \gamma_2, \gamma_3$ be three $C_0(X)$–geodesics connecting the pairs $\{y, z\}$, $\{z, x\}$, and $\{x, y\}$, respectively. Without loss of generality, we may assume x, y, z are vertices of $C_0(X)$. Indeed, if x lies in the interior of an edge e, then the geodesic triangle $\Delta = \gamma_1 \cup \gamma_2 \cup \gamma_3$ must contain at least one endpoint of e. If Δ contains both endpoints of e, then x may be replaced with one of the endpoints without affecting Δ, since $C_0(X)$ is a graph. If Δ contains only one endpoint v of e, then the path $\gamma_1 \cup \gamma_2 \cup \gamma_3$ backtracks at the point x. Thus, we may replace x with v to obtain a new geodesic triangle Δ' that is δ–hyperbolic if and only if Δ is δ–hyperbolic.

The points x, y, z therefore correspond to hyperplanes H_x, H_y, H_z of X. Moreover, the geodesics $\gamma_1, \gamma_2, \gamma_3$ correspond to minimal-length sequences $H_y = H_1^1, \ldots, H_1^{n_1} = H_z$, $H_z = H_2^1, \ldots, H_2^{n_2} = H_x$, $H_x = H_3^1, \ldots, H_3^{n_3} = H_y$ of hyperplanes of X such that each pair of consecutive hyperplanes has intersecting carriers. We wish to show that the geodesic triangle Δ is $\frac{5}{2}$–slim, that is, each of the geodesics γ_i is contained in the $\frac{5}{2}$–neighbourhood of the other two geodesics. Let p be a point on γ_3. We wish to show that p is within a distance of $\frac{5}{2}$ of some point on $\gamma_1 \cup \gamma_2$. It is sufficient to assume p is a vertex of γ_3 and show that p is within a distance of 2 of some vertex of $\gamma_1 \cup \gamma_2$. In particular, it suffices to show there exists some hyperplane H that crosses both H_p and H_i^j for some $i \in \{1, 2\}$ and $j \in \{1, \ldots, n_i\}$, where H_p is the hyperplane of X corresponding to the vertex p of $C_0(X)$.

Claim 2.4.17. There exists a hyperplane H of X that crosses both H_p and H_i^j for some $i \in \{1, 2\}$ and $j \in \{1, \ldots, n_i\}$.

Proof of claim. Consider the hyperplanes $H_1^1, \ldots, H_1^{n_1} = H_2^1, \ldots, H_2^{n_2} = H_3^1, \ldots, H_3^{n_3} = H_1^1$ corresponding to the vertices of the geodesic triangle Δ. We say two hyperplanes H_i^j, H_k^l are adjacent if the corresponding vertices are adjacent in Δ; that is, if $i = k$ and $|j - l| = 1$, or if $i = 1$, $j = n_1$, $k = 2$, $l = 1$ or $i = 2$, $j = n_2$, $k = 3$, $l = 1$ or $i = 3$, $j = n_3$, $k = 1$, $l = 1$. In particular, if two hyperplanes are adjacent then their carriers intersect. For each pair of adjacent hyperplanes H_i^j, H_k^l, pick a vertex $q \in N(H_i^j) \cap N(H_k^l)$. We obtain a sequence

$q_1^1, \ldots, q_1^{n_1} = q_2^1, \ldots, q_2^{n_2} = q_3^1, \ldots q_3^{n_3} = q_1^1$ of vertices of X. We say q_i^j and q_k^l are adjacent if the corresponding hyperplanes are adjacent. We can then connect pairs of adjacent q_i^j with geodesics in $N(H_i^j)$, by convexity of carriers (Proposition 2.4.13(3)). Let β_i^j be the geodesic from q_i^j to q_i^{j+1}, and denote the geodesic contained in $N(H_p)$ by β_p. The union of the geodesics β_i^j forms a loop β in X; see Figure 2.6.

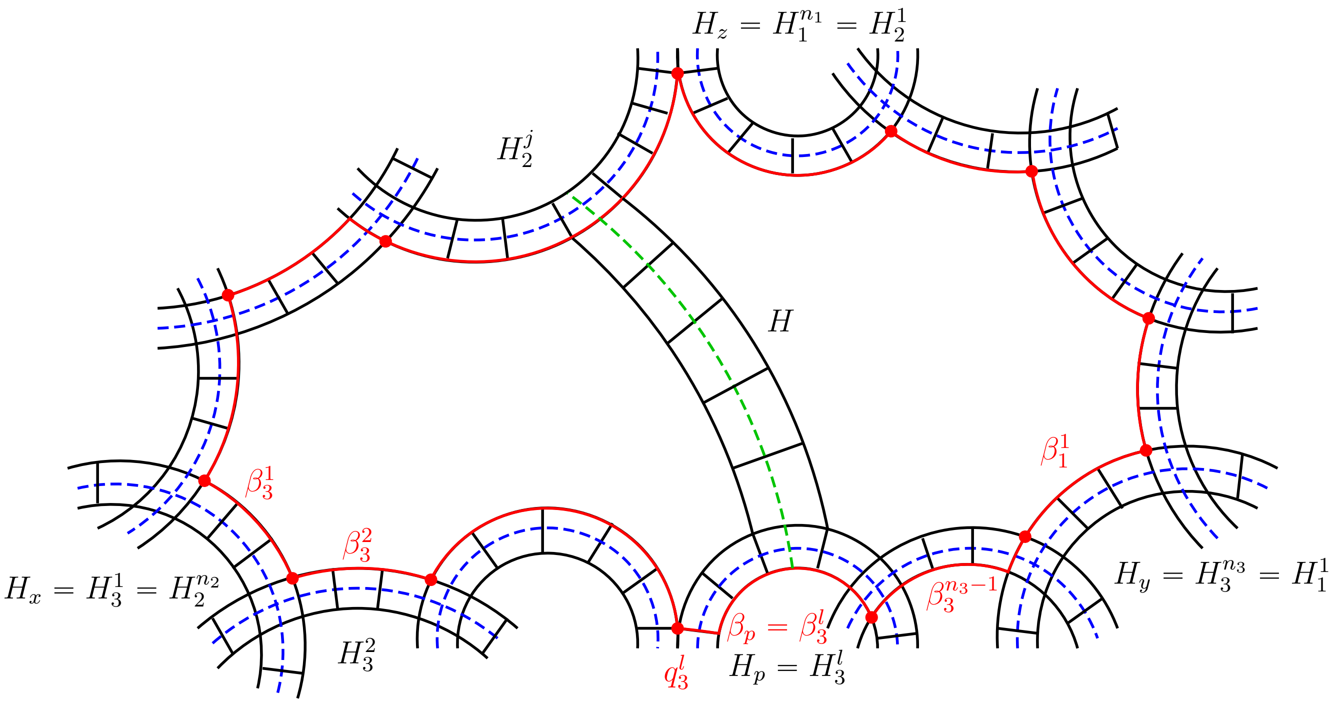

Figure 2.6: The geodesic triangle Δ in $C_0(X)$ gives a loop in X constructed from the corresponding hyperplanes. We wish to show there is a hyperplane H that crosses both H_p and some H_i^j for $i = 1$ or 2.

Suppose $H_p = H_3^l$, so that $\beta_p = \beta_3^l$. Note that since β_3^l is contained in $N(H_3^l)$, all edges of β_3^l are either dual to H_3^l or contained in a combinatorial hyperplane associated to H_3^l. Let $\mathcal{E} = \{E_1, \ldots, E_n\}$ be the edges of β_3^l that are not dual to H_3^j for $j = l-1$, l, $l+1$ and such that for each k, the maximal-dimensional cube of $N(H_3^l)$ containing E_k does not contain any edges of $N(H_3^{l-1})$ or $N(H_3^{l+1})$. Furthermore, suppose the edges E_k are ordered according to the orientation of β_3^l. Let H_k be the hyperplane dual to E_k. Since all of the edges E_k are contained in a combinatorial hyperplane associated to H_3^l, it follows that H_k crosses H_3^l for all k. It remains to show that there exists some k such that H_k also crosses some H_i^j for $i \in \{1, 2\}$. It suffices to show that H_k crosses some β_i^j for $i \in \{1, 2\}$.

First consider H_1. Since H_1 separates X into two connected components (Proposition 2.4.13(2)) and H_1 is dual to an edge of the loop β, it follows that H_1 must also be dual to another edge of β. Furthermore, since each hyperplane crosses each geodesic at most once (Proposition 2.4.13(5)), H_1 must cross some $\beta_i^j \neq \beta_3^l$. If $i = 1$ or 2, then we are done. Suppose therefore that $i = 3$ and $j \neq l$. We then break our analysis down into the following cases.

Case 1: $j \leqslant l - 3$ or $j \geqslant l + 2$. Note that because E_1 is the first edge of $\mathcal{E}$, it follows that $N(H_1)$ must contain a vertex of $N(H_3^{l-1}) \cap N(H_3^l)$; that is, H_1 contacts H_3^{l-1}. Therefore, if $j \geqslant l + 2$, then we obtain a sequence of contacting hyperplanes H_3^{l-1}, H_1, H_3^j which is shorter than the sequence $H_3^{l-1}, H_3^l, \ldots, H_3^j$. If $j \leqslant l - 3$, then we obtain a sequence of contacting hyperplanes H_3^l, H_1, H_3^j which is shorter than the sequence $H_3^l, H_3^{l-1}, \ldots, H_3^j$. In both cases, this contradicts our assumption that γ_3 is a geodesic in $C_0(X)$. Therefore, $l - 2 \leqslant j \leqslant l + 1$. See Figure 2.7.

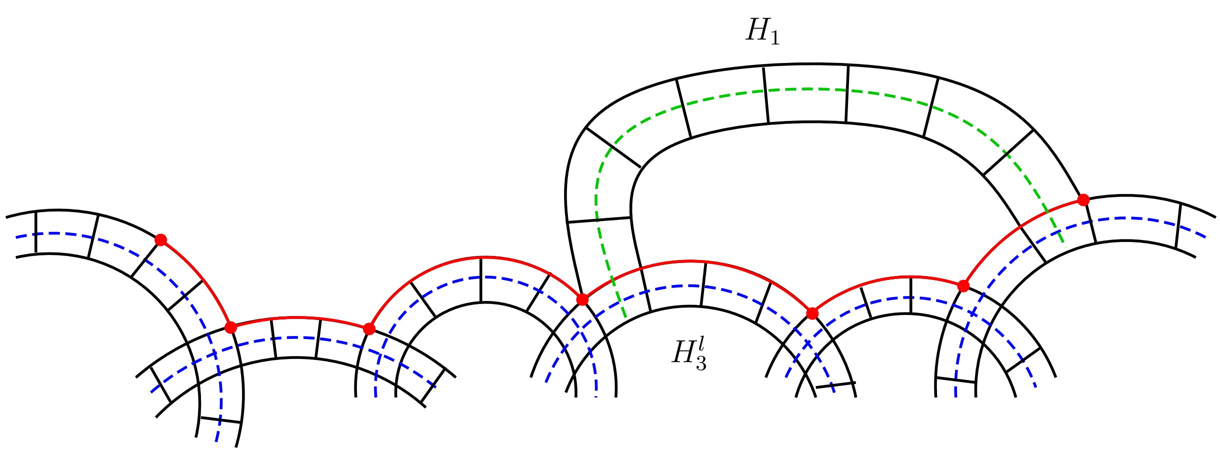

Figure 2.7: If H_1 crosses H_3^j for $j \geqslant l + 2$ then this creates a shortcut in the contact graph, as H_1 contacts both H_3^{l-1} and H_3^{l+2}.

Case 2: $j = l - 1$**.** Recall that $N(H_1)$ contains a vertex v of $N(H_3^{l-1}) \cap N(H_3^l)$. Let E be the edge dual to H_1 containing v, and let E' be the edge dual to H_3^{l-1} containing v. Note that E is contained in $N(H_3^l)$, and E is contained in the same cube as E_1. If E and E' span a square, then E is also contained in $N(H_3^{l-1})$, contradicting our assumption that the maximal-dimensional cube of $N(H_3^l)$ containing E_1 does not contain any edges of $N(H_3^{l-1})$ or $N(H_3^{l+1})$. On the other hand, if E and E' do not span a square then H_1 and H_3^{l-1} inter-osculate, contradicting our assumption that X is CAT(0) (and therefore special by Proposition 2.4.12); see Figure 2.8. Therefore, $j \neq l - 1$.

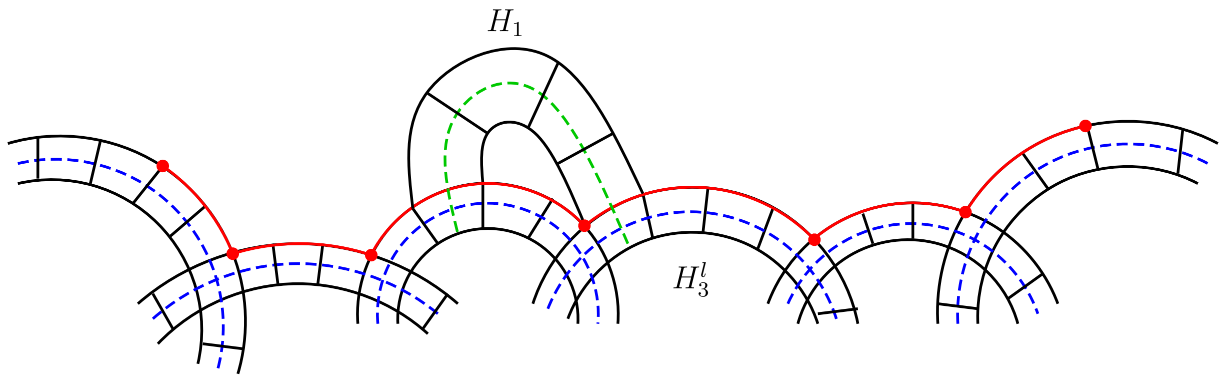

Figure 2.8: If H_1 crosses H_3^{l-1} then H_1 and H_3^{l-1} inter-osculate.

Case 3: $j = l + 1$**.** Suppose $\mathcal{E} = \{E_1\}$. Because E_1 is the last edge of $\mathcal{E}$, it follows that $N(H_1)$ must contain a vertex v of $N(H_3^l) \cap N(H_3^{l+1})$. Let E be the edge dual to H_1 containing v, and let E' be the edge dual to H_3^{l+1} containing v. Note that E is contained in $N(H_3^l)$, and E is contained in the same cube as E_1. If E and E' span a square, then E is also contained in $N(H_3^{l+1})$, contradicting our assumption that the maximal-dimensional cube of $N(H_3^l)$ containing E_1 does not contain any edges of $N(H_3^{l-1})$ or $N(H_3^{l+1})$. On the other hand, if E and E' do not span a square then H_1 and H_3^{l+1} inter-osculate, contradicting our assumption that X is special. Thus, we may assume that $\mathcal{E}$ contains more than one edge. In this case, we have a loop in $C_0(X)$ given by the contacting hyperplanes H_3^l, H_3^{l+1}, H_1. We can therefore construct a loop α in X in a similar manner to β; pick vertices w_1 in $N(H_1) \cap N(H_3^l)$, w_2 in

$N(H_3^l) \cap N(H_3^{l+1})$, and w_3 in $N(H_3^{l+1}) \cap N(H_1)$ and connect them via geodesics $\alpha_1, \alpha_2, \alpha_3$ in $N(H_3^l), N(H_3^{l+1}), N(H_1)$. Moreover, without loss of generality we may assume that w_1 lies on β_3^l, $w_2 = q_3^{l+1}$, w_3 lies on β_3^{l+1}, and $\alpha_1 \subseteq \beta_3^l$, $\alpha_2 \subseteq \beta_3^{l+1}$; see Figure 2.9. Consider the hyperplane H_2, which is dual to an edge of α_1. Note that H_2 must cross another edge of α by Proposition 2.4.13(2). Further, H_2 cannot cross another edge of α_1 by Proposition 2.4.13(5) and H_2 cannot cross α_3 since this would imply H_2 and H_1 inter-osculate. Thus, H_2 crosses an edge of α_2. We may repeat this analysis on each of the hyperplanes H_k in turn, noting that H_k either crosses α_1 a second time (contradicting Proposition 2.4.13(5)), or crosses H_{k-1} (causing an inter-osculation), or crosses α_2. We conclude that H_k must cross α_2 for all k. However, if H_n crosses α_2 then this causes H_n and H_3^{l+1} to inter-osculate, contradicting specialness of X. Therefore, we must have $j \neq l + 1$.

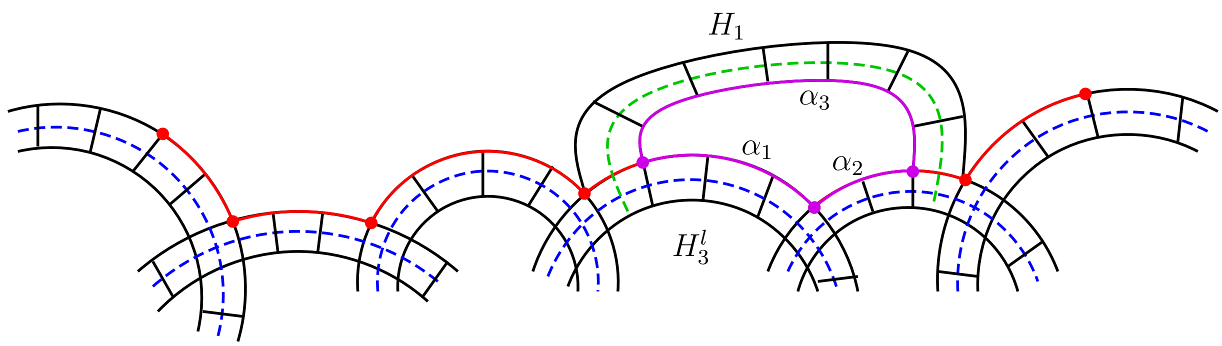

Figure 2.9: If H_1 crosses H_3^{l+1} then we can analyse the behaviour of hyperplanes crossing the new loop α.

Case 4: H_1 does not exist. We must also consider the case that the collection $\mathcal{E}$ is empty; that is, all edges of β_3^l are either dual to some H_3^j with $j = l - 1$, l, or $l + 1$, or are contained in a maximal-dimensional cube of $N(H_3^l)$ that also contains an edge of $N(H_3^{l-1})$ or $N(H_3^{l+1})$. In this case, either H_3^{l-1} and H_3^{l+1} contact (contradicting γ_3 being a geodesic) or H_3^l separates the endpoints of β_3^l and we can repeat the above case analysis for H_3^l; see Figure 2.10. This time, the case of $j = l - 2$ is ruled out too since H_3^l contacts H_3^{l+1}. Therefore, H_3^l must cross H_i^j for $i \in \{1, 2\}$, concluding our argument.

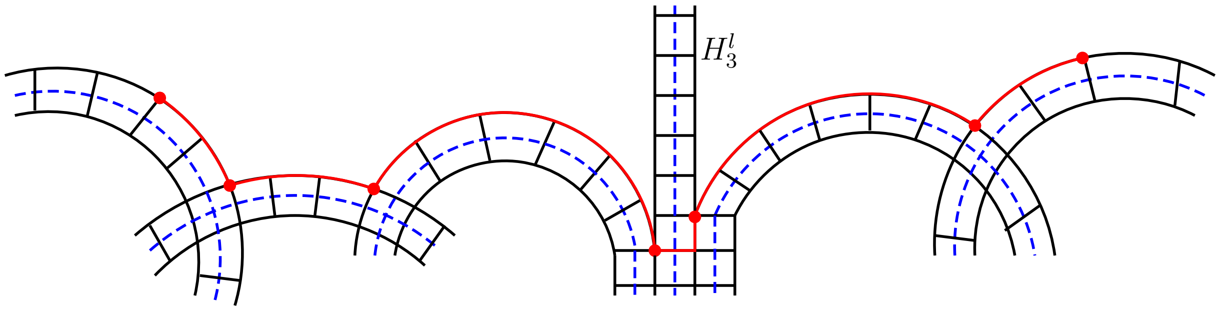

Figure 2.10: If H_1 does not exist, then H_3^l must cross H_i^j for $i = 1$ or 2.

Notice that if H_1 exists, the only case that does not result in a contradiction is therefore when $j = l - 2$. That is, H_1 crosses H_3^{l-2}. Moreover, each H_k contacts H_{k-1}. Thus, by inducting on k and repeating this case analysis with H_{k-1} in place of H_3^{l-1}, we see that each H_k either crosses some H_i^j with $i \in \{1, 2\}$, concluding the proof, or crosses H_3^{l-2}. Furthermore, E_n is the last edge of $\mathcal{E}$, so $N(H_n)$ must share a vertex with $N(H_3^{l+1})$; that is, H_n contacts H_3^{l+1}. Thus, if H_n crosses H_3^{l-2} then we obtain a sequence of contacting hyperplanes $H_3^{l+1}, H_n, H_3^{l-2}$ which is shorter than the sequence $H_3^{l+1}, H_3^l, H_3^{l-1}, H_3^{l-2}$. This contradicts our assumption that γ_3 is a geodesic in $C_0(X)$. Hence, H_n cannot cross H_3^{l-2} and so there must exist some k such that H_k crosses H_i^j for some $i \in \{1, 2\}$. See Figure 2.11. □

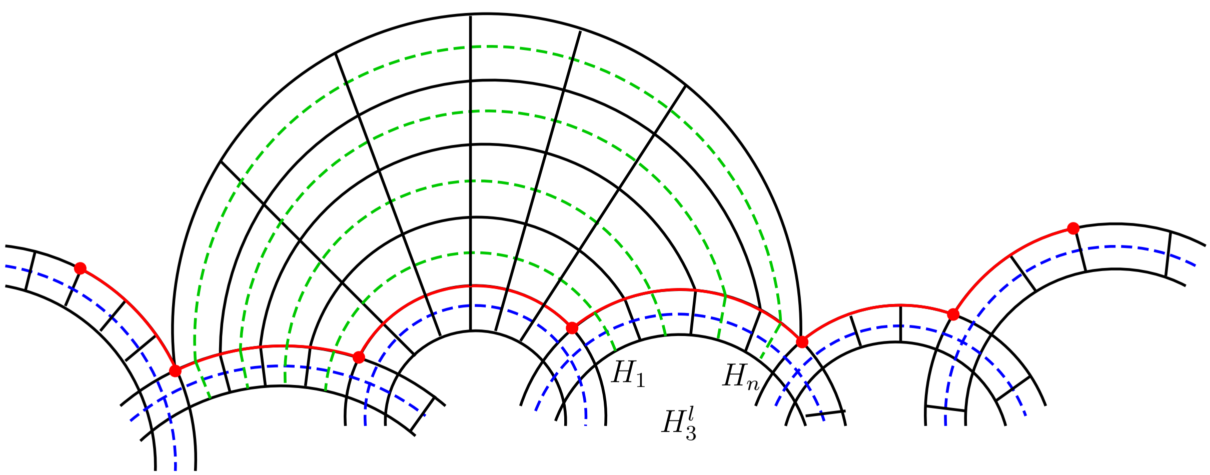

Figure 2.11: If H_n crosses H_3^{l-2} then this creates a shortcut in the contact graph, as H_n contacts both H_3^{l+1} and H_3^{l-2}.

This concludes the proof of Theorem 2.4.16. □

2.4.1 Right-angled Artin groups

Right-angled Artin groups form one of the foundational examples of cubical groups.

Definition 2.4.18 (Right-angled Artin group)**.** Let Γ be a finite simplicial graph. The *right-angled Artin group* A_Γ is defined as

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

We call Γ the *defining graph* of A_Γ.

A right-angled Artin group may be expressed as the fundamental group of a cube complex called the *Salvetti complex*, constructed as follows.

Definition 2.4.19 (Salvetti complex)**.** The *Salvetti complex* S_Γ of the right-angled Artin group A_Γ is the cube complex defined as follows.

- S_Γ has a single vertex.
- S_Γ has an edge E_v for each vertex v of Γ.
- Edges $E_{v_1}, \ldots, E_{v_n}$ span an n–cube if $v_1, \ldots, v_n$ span an n–clique of Γ.

Note that since S_Γ has just a single vertex, all edges of S_Γ must form loops and each n–cube is an n–torus. It therefore follows immediately that A_Γ is the fundamental group of S_Γ. Moreover, S_Γ is a special cube complex (see [HW08, Example 3.3]). In fact, one can show that any special cube complex X embeds in the Salvetti complex of some right-angled Artin group A_Γ via a local isometry [HW08, Theorem 1.1], and therefore $\pi_1(X)$ embeds in A_Γ itself.

Theorem 2.4.20 ([HW08, Theorem 1.1])**.** *Let X be a special cube complex. Then $\pi_1(X)$ embeds in a right-angled Artin group.*

Note that if Γ is disconnected, then A_Γ can be expressed as a free product of the right-angled Artin groups defined on each of the connected components of Γ. In particular, A_Γ is hyperbolic relative to these free factors. On the other hand, if Γ splits as a join, then A_Γ splits as a direct product, where the factors are the right-angled Artin groups defined on the components of the join. In this case, A_Γ is thick of order 0.

Behrstock–Druţu–Mosher go further by showing that if Γ is connected and has at least 2 vertices, then A_Γ is thick of order at most 1 [BDM09, Corollary 10.8]. Moreover, Behrstock and Charney show that if Γ also does not split as a join, then the divergence of A_Γ is at least quadratic [BC12, Corollary 4.8]. Combining these two results, we have a complete characterisation of relative hyperbolicity and thickness in right-angled Artin groups.

Theorem 2.4.21 (Characterisation of relative hyperbolicity and thickness in right-angled Artin groups; [BDM09, Corollary 10.8],[BC12, Corollary 4.8])**.** *Let Γ be a finite simplicial graph with at least 2 vertices.*

- *If Γ is disconnected, then A_Γ is freely decomposable, and in particular is hyperbolic relative to its free factors.*
- *If Γ splits as a join, then A_Γ is strongly thick of order 0.*
- *Otherwise, A_Γ is strongly thick of order 1.*

2.4.2 Right-angled Coxeter groups

By taking the definition of a right-angled Artin group and adding extra relations that require all generators to have order 2, one obtains a *right-angled Coxeter group.*

Definition 2.4.22 (Right-angled Coxeter group)**.** Let Γ be a finite simplicial graph. The *right-angled Coxeter group* W_Γ is defined as

$$W_\Gamma = \langle V(\Gamma) \mid [v,w] = e \ \forall \ \{v,w\} \in E(\Gamma), \ v^2 = e \ \forall \ v \in V(\Gamma) \rangle.$$

Right-angled Coxeter groups form another salient example of cubical groups. This time, their cubical structure is given by the *Davis complex*.

Definition 2.4.23 (Davis complex)**.** The *Davis complex* Σ_Γ of the right-angled Coxeter group W_Γ is the cube complex defined as follows.

- The 1–skeleton of Σ_Γ is the Cayley graph of W_Γ. Thus, each edge of Σ_Γ is labelled by a vertex of Γ.

- Pairwise adjacent edges $E_1, \ldots, E_n$ span an n–cube if they are labelled by distinct vertices $v_1, \ldots, v_n$ which span an n–clique of Γ.

The Davis complex Σ_Γ is a CAT(0) cube complex upon which W_Γ acts geometrically [Dav08]. It may in some sense be considered to be analogous to the universal cover of the Salvetti complex in the right-angled Artin group case; indeed, Haglund and Wise show that right-angled Coxeter groups have a finite-index subgroup which is the fundamental group of a special cube complex [HW10]. Furthermore, by studying these two CAT(0) cube complexes, Davis and Januszkiewicz show that every right-angled Artin group can be realised as a finite-index subgroup of a right-angled Coxeter group.

Theorem 2.4.24 ([DJ00])**.** *For every right-angled Artin group there exists a right-angled Coxeter group which contains it as a subgroup of finite index.*

One immediate consequence of this theorem is that the spectrum of possible rates of divergence for right-angled Coxeter groups must at the very least contain those of right-angled Artin groups. Behrstock and Charney's results therefore tell us that there exist right-angled Coxeter groups with linear, quadratic, and infinite divergence [BC12, Corollary 4.8].

Levcovitz takes this further, obtaining a complete classification of divergence in right-angled Coxeter groups, as well as a classification of relative hyperbolicity and thickness akin to Theorem 2.4.21 [Lev20, Theorem A, Corollary B]. This is achieved by studying a graph invariant called the *hypergraph index*.

Definition 2.4.25 (Hypergraph index)**.** Let Γ be a finite simplicial graph, and let Λ be a subgraph of Γ which splits as a join $\Lambda = \Lambda_1 \star \Lambda_2$. We say Λ is a *wide* subgraph if Λ_1 and Λ_2 each contain two non-adjacent vertices. We say Λ is a *strip* subgraph if Λ_1 consists of two non-adjacent vertices and Λ_2 is a clique. Let Φ be the collection of all maximal wide subgraphs of Γ and let Ψ be the collection of all maximal strip subgraphs.

Inductively define a collection of hypergraphs $\Delta_i = \Delta_i(\Gamma)$ as follows. Let Δ_0 be the hypergraph with vertex set $V(\Gamma)$ and hyperedge set $\{V(\Lambda) \mid \Lambda \in \Phi \cup \Psi\}$. For each $i \geqslant 0$, define an equivalence relation $\equiv_i$ on the hyperedges of Δ_i by setting $E \equiv_i E'$ for $E, E' \in \mathcal{E}(\Delta_i)$ if there exists a sequence of hyperedges $E = E_1, E_2, \ldots, E_n = E'$ in $\mathcal{E}(\Delta_i)$ such that for each $1 \leqslant j < n$, $E_j \cap E_{j+1}$ contains a pair of vertices which are non-adjacent in Γ. Now define Δ_{i+1} to be the hypergraph with vertex set $V(\Gamma)$ and where $E \subseteq V(\Gamma)$ is a hyperedge of Δ_{i+1} if and only if $E = E_1 \cup \cdots \cup E_m$ for some maximal collection $\{E_1, \ldots, E_m\}$ of $\equiv_i$–equivalent hyperedges of Δ_i.

We define the *hypergraph index* of Γ to be the smallest integer $k \geqslant 0$ such that there exists a hyperedge $E \in \mathcal{E}(\Delta_k)$ with $E = V(\Gamma)$. If no such k exists, then we say the hypergraph index of Γ is ∞.

Theorem 2.4.26 (Characterisation of relative hyperbolicity and thickness in right-angled Coxeter groups; [Lev20, Theorem A, Corollary B])**.** *Let Γ be a finite simplicial graph.*

- *If Γ has infinite hypergraph index, then W_Γ is relatively hyperbolic. Moreover, W_Γ has exponential divergence if it is one-ended, and infinite divergence otherwise.*

- *If Γ has hypergraph index $k \geqslant 0$, then W_Γ is strongly thick of order k and has polynomial divergence of degree $k + 1$.*

2.5 Graph braid groups

In this section we introduce graph braid groups and show they can be expressed as fundamental groups of special cube complexes.

Consider a finite collection of particles lying on a finite metric graph Γ. The *configuration space* of these particles on Γ is the collection of all possible ways the particles can be arranged on the graph with no two particles occupying the same location. As we move through the configuration space, the particles move along Γ, without colliding. If we do not distinguish between each of the different particles, we call this an *unordered* configuration space. A *graph braid group* is the fundamental group of an unordered configuration space. More precisely:

Definition 2.5.1 (Graph braid group)**.** Let Γ be a finite graph, and let n be a positive integer. The *topological configuration space* $C_n^{\text{top}}(\Gamma)$ is defined as

$$C_n^{\text{top}}(\Gamma) = \Gamma^n \smallsetminus D^{\text{top}},$$

where $D^{\text{top}} = \{(x_1, \ldots, x_n) \in \Gamma^n \mid x_i = x_j \text{ for some } i \neq j\}$. The *unordered topological configuration space* $UC_n^{\text{top}}(\Gamma)$ is then defined as

$$UC_n^{\text{top}}(\Gamma) = C_n^{\text{top}}(\Gamma)/S_n,$$

where the symmetric group S_n acts on $C_n^{\text{top}}(\Gamma)$ by permuting its coordinates. We define the *graph braid group* $B_n(\Gamma, S)$ as

$$B_n(\Gamma, S) = \pi_1(UC_n^{\text{top}}(\Gamma), S),$$

where $S \in UC_n^{\text{top}}(\Gamma)$ is a fixed base point.

The base point S in our definition represents an initial configuration of the particles on the graph Γ. As the particles are unordered, they may always be moved along Γ into any other desired initial configuration, so long as the correct number of particles are present in each connected component of Γ. In particular, if Γ is connected, then the graph braid group $B_n(\Gamma, S)$ is independent of the choice of base point, and may therefore be denoted simply $B_n(\Gamma)$. The following result of Genevois shows that we can always express $B_n(\Gamma, S)$ as a product of braid groups of connected graphs; thus, we will often be able to work under the assumption that Γ is connected without loss of generality.

Lemma 2.5.2 ([Gen19a, Lemma 3.5]). *Let $n > 1$, let Γ be a finite graph, and suppose $\Gamma = \Gamma_1 \sqcup \Gamma_2$. Then*

$$UC_n^{top}(\Gamma) \simeq \bigsqcup_{k=0}^{n} UC_k^{top}(\Gamma_1) \times UC_{n-k}^{top}(\Gamma_2).$$

Moreover, if $S \in UC_n^{top}(\Gamma)$ has k particles in Γ_1 and $n-k$ particles in Γ_2, then

$$B_n(\Gamma, S) \cong B_k(\Gamma_1, S \cap \Gamma_1) \times B_{n-k}(\Gamma_2, S \cap \Gamma_2),$$

where $S \cap \Gamma_1$ (resp. $S \cap \Gamma_2$) denotes the configuration of the k (resp. $n-k$) particles of S lying in Γ_1 (resp. Γ_2).

Note that in some sense, the space $UC_n^{\mathrm{top}}(\Gamma)$ is almost a cube complex. Indeed, Γ^n is a cube complex, but removing the diagonal breaks the structure of some of its cubes. By expanding the diagonal slightly, we are able to fix this by ensuring that we are always removing whole cubes.

Definition 2.5.3 (Combinatorial configuration space). Let Γ be a finite graph, and let n be a positive integer. For each $x \in \Gamma$, the *carrier* $c(x)$ of x is the lowest dimensional simplex of Γ containing x. The *combinatorial configuration space* $C_n(\Gamma)$ is defined as

$$C_n(\Gamma) = \Gamma^n \smallsetminus D,$$

where $D = \{(x_1, \ldots, x_n) \in \Gamma^n \mid c(x_i) \cap c(x_j) \neq \varnothing \text{ for some } i \neq j\}$. The *unordered combina-*

torial configuration space $UC_n(\Gamma)$ is then

$$UC_n(\Gamma) = C_n(\Gamma)/S_n.$$

The *reduced graph braid group* $RB_n(\Gamma, S)$ is defined as

$$RB_n(\Gamma, S) = \pi_1(UC_n(\Gamma), S),$$

where $S \in UC_n(\Gamma)$ is a fixed base point.

Removing this new version of the diagonal tells us that two particles cannot occupy the same edge of Γ. This effectively discretises the motion of the particles to jumps between vertices, as each particle must fully traverse an edge before another particle may enter.

Observe that $C_n(\Gamma)$ is the union of all products $c(x_1)\times\cdots\times c(x_n)$ satisfying $c(x_i)\cap c(x_j) = \varnothing$ for all $i \neq j$. Since the carrier is always a vertex or a closed edge, this defines an n–dimensional cube complex, and moreover $C_n(\Gamma)$ is compact with finitely many hyperplanes, as Γ is a finite graph. It follows that $UC_n(\Gamma)$ is also a compact cube complex with finitely many hyperplanes. Indeed, we have the following useful description of the cube complex structure, due to Genevois [Gen19a].

- The vertices of $UC_n(\Gamma)$ are the subsets S of $V(\Gamma)$ with size $|S| = n$.
- Two vertices S_1 and S_2 of $UC_n(\Gamma)$ are connected by an edge if their symmetric difference $S_1 \triangle S_2$ is a pair of adjacent vertices of Γ. We therefore label each edge E of $UC_n(\Gamma)$ with a closed edge e of Γ.
- A collection of m edges of $UC_n(\Gamma)$ with a common endpoint span an m–cube if their labels are pairwise disjoint.

Abrams showed that if Γ has more than n vertices, then $UC_n(\Gamma)$ is connected if and only if Γ is connected.

Theorem 2.5.4 ([Abr00, Theorem 2.6]). *Let Γ be a finite graph. If Γ has more than n vertices, then $UC_n(\Gamma)$ is connected if and only if Γ is connected.*

Furthermore, Abrams showed that if we subdivide edges of Γ sufficiently (i.e. add 2–valent vertices to the middle of edges) to give a new graph Γ', then $UC_n^{\text{top}}(\Gamma')$ deformation retracts onto $UC_n(\Gamma')$ [Abr00, Theorem 2.1]. As $UC_n^{\text{top}}(\Gamma')$ does not distinguish vertices from other points on the graph, we have $UC_n^{\text{top}}(\Gamma') = UC_n^{\text{top}}(\Gamma)$, implying that $B_n(\Gamma, S) \cong RB_n(\Gamma', S)$. This allows us to consider $B_n(\Gamma, S)$ as the fundamental group of the cube complex $UC_n(\Gamma')$.

Prue and Scrimshaw later improved upon the constants in Abrams' result to give the following theorem [PS14, Theorem 3.2].

Theorem 2.5.5 ([Abr00],[PS14]). *Let $n \in \mathbb{N}$ and let Γ be a finite graph with at least n vertices. The unordered topological configuration space $UC_n^{top}(\Gamma)$ deformation retracts onto the unordered combinatorial configuration space $UC_n(\Gamma)$ if the following conditions hold.*

(1) Every path between distinct vertices of Γ of valence $\geqslant 3$ has length at least $n-1$.

(2) Every homotopically essential loop in Γ has length at least $n+1$.

Remark 2.5.6. Note that Prue and Scrimshaw's version of the theorem only deals with the case where Γ is connected. However, the disconnected case follows easily by deformation retracting each connected component of $UC_n(\Gamma)$, noting that Lemma 2.5.2 tells us that each component can be expressed as a product of cube complexes $UC_k(\Lambda)$, where $k \leqslant n$ and Λ is a connected component of Γ.

Another foundational result of Abrams states that the cube complex $UC_n(\Gamma)$ is non-positively curved [Abr00, Theorem 3.10]. Furthermore, Genevois proved that $UC_n(\Gamma)$ admits a *special colouring* [Gen19a, Proposition 3.7]. We shall omit the details of his theory of special colourings, and direct the reader to [Gen21] for further details. The key result is that a cube complex X admits a special colouring if and only if there exists a special cube complex Y

such that $Y^{(2)} = X^{(2)}$ [Gen21, Lemma 3.2]. Furthermore, Genevois constructs Y by taking $X^{(2)}$ and inductively attaching m–cubes C whenever a copy of $C^{(m-1)}$ is present in the complex, for $m \geqslant 3$; this ensures non-positive curvature of Y. Since in our case $X = UC_n(\Gamma)$ is already non-positively curved, this means $Y = X$. Thus, $B_n(\Gamma)$ is the fundamental group of the special cube complex $UC_n(\Gamma')$.

In summary, we have the following result.

Corollary 2.5.7 (Graph braid groups are special; [Abr00, Gen19a, Gen21])**.** *Let $n > 1$ and let Γ be a finite, connected graph. Then $B_n(\Gamma) \cong RB_n(\Gamma')$, where Γ' is obtained from Γ by subdividing edges. In particular, $B_n(\Gamma)$ is the fundamental group of the connected compact special cube complex $UC_n(\Gamma')$.*

2.6 Quasi-median graphs

The notion of a *quasi-median graph* was originally introduced by Mulder in [Mul80] as a generalisation of median graphs (see Definition 2.4.3) and further developed by Bandelt–Mulder–Wilkeit [BMW94]. However, these were not studied in the context of geometric group theory until recently, when Genevois explored their coarse geometry, drawing on analogies with the cubical geometry of median graphs [Gen17]. In this section we shall study this geometry, paying particular attention to its application to graph products.

There are numerous equivalent definitions of a quasi-median graph (see [BMW94, Theorem 1]). We shall focus on the one given below.

Definition 2.6.1 (Quasi-median graph)**.** A connected simplicial metric graph X is *weakly modular* if it satisfies the following two conditions.

(1) **(Triangle condition.)** Let u, v, w be vertices of X such that v is adjacent to w and both v and w are at distance $k \geqslant 2$ from u. Then there exists a vertex x of X that is adjacent to v and w and at distance $k - 1$ from u. See Figure 2.13.

(2) **(Quadrangle condition.)** Let u, v, w, z be vertices of X such that both v and w are adjacent to z and at distance $k-1 \geqslant 2$ from u, and z is at distance k from u. Then there exists a vertex x of X that is adjacent to v and w and at distance $k-2$ from u. See Figure 2.13.

We say X is *quasi-median* if it is weakly modular and does not contain K_4^- or $K_{3,2}$ as induced subgraphs; see Figure 2.12.

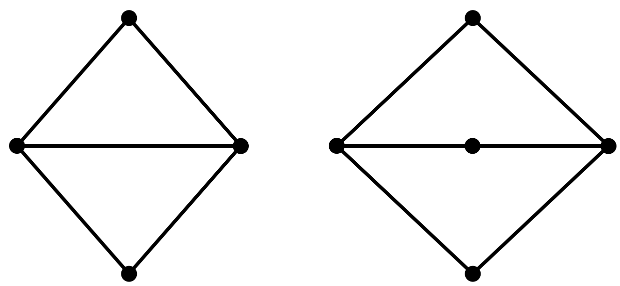

Figure 2.12: The graphs K_4^- (left) and $K_{3,2}$ (right).

Note that as a consequence of this definition, every cycle in a quasi-median graph is contained in a union of 3–cycles and 4–cycles. Indeed, given an n–cycle with $n \geqslant 5$, one can apply the triangle condition (if n is odd) or the quadrangle condition (if n is even) in order to cover this n–cycle with cycles of strictly shorter length. By applying these conditions inductively, one achieves the desired result; see Figure 2.13. Moreover, a result of Genevois says that a quasi-median graph is a median graph precisely when it does not contain any cycles of length 3.

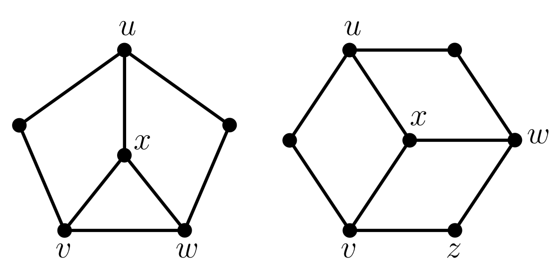

Figure 2.13: The triangle and quadrangle conditions break up cycles into triangles and squares.

Proposition 2.6.2 ([Gen17, Corollary 2.92])**.** *A graph is median if and only if it is quasi-median and contains no* 3*–cycles.*

In this way, the appropriate generalisation of a cube complex in the context of quasi-median graphs is that of a *prism* complex. Just as a cube may be viewed as a product of 1–simplices, a prism is defined to be a product of simplices of any dimension. In the following section, we shall explore the geometry of such complexes for the specific example of graph products. For a more general treatment of this quasi-median geometry, the interested reader is directed to [Gen17].

2.6.1 Graph products

Definition 2.6.3 (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 non-trivial group G_v. Then the *graph product* G_Γ is the group

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

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

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 2.6.4 (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 $\mathsf{d}_{syl}(\cdot,\cdot)$ and say $\mathsf{d}_{syl}(g,h)$ is the *syllable distance* between g and h. When convenient, we will use $|g|_{syl}$ to denote $\mathsf{d}_{syl}(e,g)$ 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 2.6.5 (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 $s_1 \dots s_n$ is a *syllable expression* for g. If $s_1 \dots s_n$ is a syllable expression for g and $n = \mathsf{d}_{syl}(e,g)$, then we say $s_1 \dots s_n$ 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 2.6.6 (Reduction algorithm for graph products; [Gre90, 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+i} 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.*

When each of the vertex groups of the graph product is finitely generated, Theorem 2.6.6 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 2.6.7 (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 2.6.6, 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 $m_i \leqslant m$ and σ_i is a permutation of $\{1, \dots, m\}$. Further, if $i \neq k$, then $\{\sigma_i(1), \dots, \sigma_i(m_i)\} \cap \{\sigma_k(1), \dots, \sigma_k(m_k)\} = \varnothing$. Thus, $\sum_{i=1}^{n} |s_i| \leqslant \sum_{i=1}^{n} m_i \leqslant m$. However, $m = |g| \leqslant \sum_{i=1}^{n} |s_i|$ by definition, so $|g| = \sum_{i=1}^{n} |s_i|$. □

Another critical consequence of Theorem 2.6.6 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 2.6.8 (Syllables and support of an element)**.** Let G_Γ be a graph product of groups 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 $\mathrm{supp}(g)$ to denote the induced subgraph of Γ spanned by the vertices $\{v_1, \dots, v_n\}$, where $s_i \in G_{v_i}$. We call $\mathrm{supp}(g)$ the *support* of g.

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

Definition 2.6.9 (Graphical subgroups)**.** Let G_Γ be a graph product with vertex groups $\{G_v : v \in V(\Gamma)\}$ and let $\Lambda \subseteq \Gamma$ be an induced 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.

Convention 2.6.10. Whenever we consider a subgraph $\Lambda \subseteq \Gamma$, we will assume that Λ is an induced subgraph of Γ.

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

Definition 2.6.11. 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 2.6.4 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 2.6.12. 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 $\mathsf{d}_{syl}(k, h)$.

Remark 2.6.13. The graph theoretic properties of subgraphs Λ of Γ have important algebraic significance in the context of the graph product G_Γ itself. 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: $\langle \mathrm{st}(\Lambda) \rangle \cong \langle \Lambda \rangle \times \langle \mathrm{lk}(\Lambda) \rangle$. 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 a quasi-median graph [Gen17, Proposition 8.2]. Moreover, he showed that the only non-cubical behaviour arises from the vertex groups.

Proposition 2.6.14 ([Gen17, Lemmas 8.5, 8.8])**.** *Two adjacent edges of $S(\Gamma)$ are the 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 the edges of an induced square if and only if they are labelled by elements of adjacent vertex groups. In this case, the 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 2.6.15 (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 $S(\Gamma)$.

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

Definition 2.6.16 (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 $2 \leqslant k \leqslant n$, add edges connecting the barycentre of each k–simplex in T to the barycentres of each of its $(k-1)$–sub-simplices; see Figure 2.14. The complex

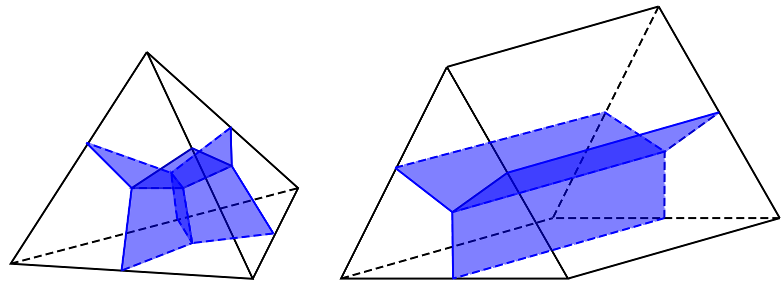

Figure 2.14: The mid-prism of a 3–simplex and a mid-prism of the product of a 2–simplex and a 1–simplex.

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 $K_i = T_1 \times \dots T_{i-1} \times K(T_i) \times T_{i+1} \times \cdots \times T_m$.

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, 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 2.6.17 (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 the *hyperplane* dual to the equivalence class $[E]$ is the union of all mid-prisms that intersect edges of $[E]$; see Figure 2.15. 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 H *crosses* γ or $g\langle\Lambda\rangle$. We say 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 $S(\Gamma) \smallsetminus H$.

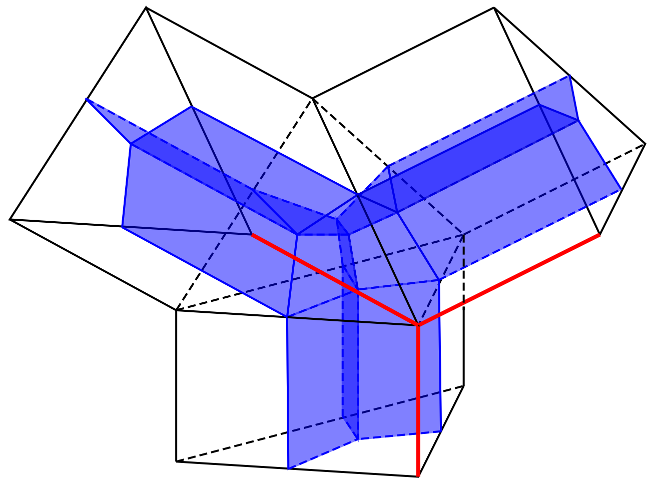

Figure 2.15: A hyperplane (blue) inside its carrier, and an associated combinatorial hyperplane (red).

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 2.6.18 (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 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 2.15.

Remark 2.6.19 (Labelling hyperplanes)**.** Proposition 2.6.14 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 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 a fact that will be exploited repeatedly in our proofs in Chapter 4.

Genevois established that the hyperplanes of $S(\Gamma)$ maintain many of the fundamental properties from the cubical setting (cf. Proposition 2.4.13).

Proposition 2.6.20 (Properties of hyperplanes; [Gen17, 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 2.6.21. 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 Chapter 4. Again, these gate maps share many properties with their cubical counterparts (cf. Proposition 2.4.14).

Proposition 2.6.22 (Gate onto graphical subgroups; [GM18, 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 $k, h \in G_\Gamma$, $\mathsf{d}_{syl}(\mathfrak{g}_{g\Lambda}(h), \mathfrak{g}_{g\Lambda}(k)) \leqslant \mathsf{d}_{syl}(h, k)$.

(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 $\mathsf{d}_{syl}(x, \mathfrak{g}_{g\Lambda}(x)) = \mathsf{d}_{syl}(x, g\langle\Lambda\rangle)$.

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

Lemma 2.6.24 (Algebraic description of the gate map)**.** *For all $\Lambda \subseteq \Gamma$ and $g \in G_\Gamma$, there exists $p \in$ prefix$(g) \cap \langle\Lambda\rangle$ so that $\mathfrak{g}_\Lambda(g) = p$. Further, p is the element of* prefix$(g) \cap \langle\Lambda\rangle$ *with the largest syllable length.*

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

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

Since $p \in \langle\Lambda\rangle$, this implies $x = p$, as x is the unique element of $\langle\Lambda\rangle$ which minimises the syllable distance of g to $\langle\Lambda\rangle$ (Proposition 2.6.22(3)). □

Definition 2.6.25. Denote the element p of $\text{prefix}(g) \cap \langle\Lambda\rangle$ with largest syllable length by $\text{prefix}_\Lambda(g)$, and define $\text{suffix}_\Lambda(g) = (\text{prefix}_\Lambda(g^{-1}))^{-1}$.

2.7 Hierarchical hyperbolicity

Hierarchical hyperbolicity is a quasi-isometry invariant property which combines elements of both coarse Euclidean and hyperbolic geometry. This version of non-positive curvature was devised by Behrstock, Hagen, and Sisto by axiomatising Masur and Minsky's treatment of mapping class groups using subsurface projections and curve graphs [MM99, MM00, BHS17b]. The geometric information of a hierarchically hyperbolic space X (commonly abbreviated 'HHS') is encoded in a collection of projections onto hyperbolic spaces associated to X. These projections are arranged via a partial order called *nesting*, and flats (quasi-isometrically embedded copies of $\mathbb{Z}^n$) are encoded via a combinatorial relation between the projections called *orthogonality*. Due to the extra structure endowed by the projections and relations, one must be careful to distinguish a hierarchically hyperbolic *space* from a hierarchically hyperbolic *group*. A hierarchically hyperbolic group (commonly abbreviated 'HHG') is not merely a group whose Cayley graph is an HHS; the hierarchy structure must also be equivariant with respect to the group action.

Examples of HHGs include mapping class groups [MM99, MM00, Beh06, BKMM12] and fundamental groups of special cube complexes [BHS17b]; see Section 2.7.4 for more details on the HHG structure of the latter. Examples of HHSs include Teichmüller space with the Teichmüller or Weil–Petersson metric ([Raf07, Dur16, EMR17] and [Bro03, Beh06, BKMM12] respectively) and fundamental groups of closed 3–manifolds with no Nil or Sol components [BHS19, Theorem 10.1]. Note that these 3–manifold groups are conjectured not to be HHGs in general [BHS19, Remark 10.2].

We break the definition of an HHG given by Behrstock, Hagen, and Sisto in [BHS19] 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 hierarchically hyperbolic space. We then define a *hierarchically hyperbolic group* to be a group whose Cayley graph is an HHS in such a way that the HHS structure agrees with the group structure.

Definition 2.7.1 (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 $(C(W), \mathsf{d}_W)$ 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 $\operatorname{diam}(\pi_W(x)) \leqslant E$. Moreover, each π_W is (E, E)–coarsely Lipschitz and $C(W) \subseteq N_E(\pi_W(\mathcal{X}))$ 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 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 $\rho_W^V \subseteq C(W)$ with $\operatorname{diam}(\rho_W^V) \leqslant E$.

(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 $\mathfrak{S}_W^{\perp}$ 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 V, W are *transverse*, denoted $V \pitchfork W$. Moreover, for all $V, W \in \mathfrak{S}$ with $V \pitchfork W$ there are non-empty sets $\rho_W^V \subseteq C(W)$ and $\rho_V^W \subseteq C(V)$ 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_W^V the *relative projection* from V to W. The number E is called the *hierarchy constant* for $\mathfrak{S}$.

Definition 2.7.2 (Hierarchically hyperbolic space)**.** An E–proto-hierarchy structure $\mathfrak{S}$ on a quasi-geodesic space $\mathcal{X}$ is an *E–hierarchically hyperbolic space structure* (E–HHS structure) on $\mathcal{X}$ if it satisfies the following additional axioms.

(1) **(Hyperbolicity.)** For each $W \in \mathfrak{S}$, $C(W)$ 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 $Q \in \mathfrak{S}_W \smallsetminus \{W\}$ such that $V \sqsubseteq Q$ whenever $V \in \mathfrak{S}_W \cap \mathfrak{S}_U^{\perp}$. We call Q the *container of U in W*.

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

$\mathsf{d}_V(\pi_V(x), \pi_V(y)) \geqslant E$, 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 $\mathfrak{L} = \{V_1, \ldots, V_m\} \subseteq \mathfrak{S}_W \smallsetminus \{W\}$ such that m is at most $E\mathsf{d}_W(\pi_W(x), \pi_W(y)) + E$, and for all $U \in \mathfrak{S}_W \smallsetminus \{W\}$, either $U \in \mathfrak{S}_{V_i}$ for some i, or $\mathsf{d}_U(\pi_U(x), \pi_U(y)) \leqslant E$.

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

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

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 $\mathsf{d}_W(\rho_W^U, \rho_W^V) \leqslant E$.

(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:

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

We call a quasi-geodesic space $\mathcal{X}$ an *E–hierarchically hyperbolic space* (E–HHS) if there exists an E–hierarchically hyperbolic space structure on $\mathcal{X}$. We use the pair $(\mathcal{X}, \mathfrak{S})$ to denote a hierarchically hyperbolic space equipped with the specific HHS structure $\mathfrak{S}$.

Definition 2.7.3 (Hierarchically hyperbolic group)**.** 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 G is an *E–hierarchically hyperbolic group* (E–HHG) if:

(1) The space X admits an 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 $V \subsetneq W$, then $g_W(\rho^V_W)$ and ρ^{gV}_{gW} are at most E–far apart in $C(gW)$.

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

Remark 2.7.4. Since the property of being an HHS is invariant under quasi-isometry [BHS19, Propositon 1.10], one may also obtain an HHG structure for a group G by finding an HHS $(X, \mathfrak{S})$ such that G acts on X geometrically and satisfies Properties (2) and (3). In particular, the Milnor–Švarc Lemma tells us that this HHG structure is given by composing the projections in $(X, \mathfrak{S})$ with an orbit map $G \to X$, $g \mapsto g \cdot x$ for some fixed $x \in X$.

We may deduce a number of additional useful properties as a direct consequence of these axioms, many of which shall prove essential in Chapter 3. For example, Durham–Hagen–Sisto show that the partial realisation axiom implies that the relative projections ρ^W_Q and ρ^V_Q of orthogonal domains $W, V \in \mathfrak{S}$ coarsely coincide. Note, ρ^W_Q and ρ^V_Q are both defined when $W \pitchfork Q$ or $W \subsetneq Q$ and $V \pitchfork Q$ or $V \subsetneq Q$.

Lemma 2.7.5 ([DHS17, Lemma 1.5])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an E–HHS. If $W, V \in \mathfrak{S}$ with $W \perp V$, and $Q \in \mathfrak{S}$ with ρ^W_Q and ρ^V_Q both defined, then $\mathsf{d}_Q(\rho^W_Q, \rho^V_Q) \leqslant 2E$.*

One may also deduce a strengthened version of the partial realisation axiom, called the *realisation theorem*, which characterises which tuples in the product $\prod_{V \in \mathfrak{S}} C(V)$ are coarsely the image of a point in $\mathcal{X}$. Essentially, it says if a tuple $(b_V) \in \prod_{V \in \mathfrak{S}} C(V)$ satisfies the consistency and bounded geodesic image axioms of an almost HHS, then there exists a point $x \in \mathcal{X}$ such that $\pi_V(x)$ is uniformly close to b_V for each $V \in \mathfrak{S}$. While it is straightforward to state what it means for a tuple to satisfy the consistency axiom—either $\mathsf{d}_W(b_W, \rho_W^V)$ or $\mathsf{d}_V(b_V, \rho_V^W)$ is less than E whenever $V \pitchfork W$—it is more opaque as to how a tuple can satisfy the bounded geodesic image axiom. For this we need the following map from $C(W)$ to $C(V)$ when $V \subsetneq W$.

Definition 2.7.6 (Downward relative projection)**.** Let $\mathfrak{S}$ be an E–HHS structure for $\mathcal{X}$. For each $W \in \mathfrak{S}$ and $p \in C(W)$, pick a point $x_{p,W} \in \mathcal{X}$ so that $\pi_W(x_{p,W})$ is within E of p. If $V, W \in \mathfrak{S}$ with $V \subsetneq W$, then define the map $\rho_V^W : C(W) \to 2^{C(V)}$ by $\rho_V^W(p) = \pi_V(x_{p,W})$. We call the map ρ_V^W a *downward relative projection* from W to V.

With the downward relative projection, we can formulate the necessary conditions for a tuple $(b_V)_{V \in \mathfrak{S}}$ to be realised by point in $\mathcal{X}$. In the following, one should think of the first condition as saying that the tuple satisfies the consistency axiom and the second as saying the tuple satisfies the bounded geodesic image axiom.

Definition 2.7.7 (Consistent tuple)**.** Let $(\mathcal{X}, \mathfrak{S})$ be an HHS and let $b_V \in C(V)$ for each $V \in \mathfrak{S}$. For each $\kappa \geqslant 0$, the tuple $(b_V)_{V \in \mathfrak{S}}$ is κ*–consistent* if:

(1) whenever $V \pitchfork W$, $\min\{\mathsf{d}_W(b_W, \rho_W^V), \mathsf{d}_V(b_V, \rho_V^W)\} \leqslant \kappa$;

(2) whenever $V \subsetneq W$, $\min\{\mathsf{d}_W(b_W, \rho_W^V), \operatorname{diam}(b_V \cup \rho_V^W(b_W))\} \leqslant \kappa$.

Given $x \in \mathcal{X}$, the tuple $(\pi_V(x))_{V \in \mathfrak{S}}$ is always consistent; properties (1) and (2) follow from the consistency and bounded geodesic image axioms for $(\mathcal{X}, \mathfrak{S})$, respectively. Conversely, the realisation theorem says that all consistent tuples are coarsely the image of point in $\mathcal{X}$.

Lemma 2.7.8 (Projections of points are consistent, [BHS19, Proposition 1.11])**.** *Let $\mathfrak{S}$ be an E–HHS structure for $\mathcal{X}$. If $x \in \mathcal{X}$, then $(\pi_V(x))_{V \in \mathfrak{S}}$ is a $3E$–consistent tuple.*

Theorem 2.7.9 (The realisation of consistent tuples, [BHS19, Theorem 3.1]). *Let $(\mathcal{X}, \mathfrak{S})$ be an HHS. There exists a function $\tau\colon [0,\infty) \to [0,\infty)$ so that if $(b_V)_{V\in\mathfrak{S}}$ is a κ–consistent tuple, then there exists $x \in \mathcal{X}$ so that $\mathsf{d}_V(x, b_V) \leqslant \tau(\kappa)$ for all $V \in \mathfrak{S}$.*

Furthermore, the relative projections of an HHS also satisfy the inequalities in the definition of a consistent tuple.

Lemma 2.7.10 (ρ–consistency, [BHS19, Proposition 1.8]). *Let $\mathfrak{S}$ be an E–HHS structure for $\mathcal{X}$ and $V, W, Q \in \mathfrak{S}$. Suppose $W \pitchfork Q$ or $W \subsetneq Q$ and $W \pitchfork V$ or $W \subsetneq V$. Then we have the following.*

(1) If $Q \pitchfork V$, then $\min\{\mathsf{d}_Q(\rho^W_Q, \rho^V_Q), \mathsf{d}_V(\rho^Q_V, \rho^W_V)\} \leqslant 2E$.

(2) If $Q \sqsubseteq V$, then $\min\{\mathsf{d}_V(\rho^Q_V, \rho^W_V), \operatorname{diam}(\rho^W_Q \cup \rho^V_Q(\rho^W_V))\} \leqslant 2E$.

One remarkable and far less straightforward consequence of the HHS axioms is the existence of a Masur–Minsky style distance formula, which allows for distances in an HHS to be expressed as a sum of distances in the associated hyperbolic spaces.

Theorem 2.7.11 (Distance formula; [BHS19, Theorem 4.5]). *Let $(X, \mathfrak{S})$ be an HHS. There exists $\sigma_0 > 0$ such that for all $\sigma \geqslant \sigma_0$ there exist $K \geqslant 1$ and $L \geqslant 0$ such that for all $x, y \in X$,*

$$\frac{1}{K}\sum_{U\in\mathfrak{S}} \{\!\{\mathsf{d}_U(x,y)\}\!\}_\sigma - L \leqslant \mathsf{d}_X(x,y) \leqslant K\sum_{U\in\mathfrak{S}} \{\!\{\mathsf{d}_U(x,y)\}\!\}_\sigma + L$$

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

2.7.1 Detecting other forms of hyperbolicity in HHSs

One powerful tool of hierarchical hyperbolicity is the ability to extract information regarding other generalisations of hyperbolicity from the HHS structure. In this section, we shall outline methods of detecting δ–hyperbolicity, acylindrical hyperbolicity, and relative hyperbolicity in HHSs.

Morally, detecting forms of hyperbolicity (or lack thereof) amounts to analysing patterns of orthogonality occurring within the HHS structure. For example, a theorem of Behrstock–Hagen–Sisto tells us that an HHS is hyperbolic precisely when there is no non-trivial orthogonality.

Theorem 2.7.12 (Characterisation of hyperbolicity; [BHS17c, 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* $D \geqslant 0$ *such that*

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

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

In the case of acylindrical hyperbolicity, a result of Behrstock–Hagen–Sisto tells us that every HHG G acts acylindrically on its $\sqsubseteq$–maximal hyperbolic space; however, G is only acylindrically hyperbolic when this hyperbolic space is unbounded.

Theorem 2.7.13 (Criteria for acylindrical hyperbolicity; [BHS17b, Corollary 14.4])**.** *Let* $(G, \mathfrak{S})$ *be an HHG and let* S *be the unique* $\sqsubseteq$*–maximal element of* $\mathfrak{S}$*. Then* G *acts acylindrically on* $C(S)$*. In particular, if* $C(S)$ *is unbounded and* G *is not virtually cyclic, then* G *is acylindrically hyperbolic.*

A result of Russell tells us that if collections of intersecting non-negatively curved regions in an HHG can be isolated from each other (the *isolated orthogonality* criterion), then the HHG is relatively hyperbolic and these isolated collections form the peripheral subgroups [Rus20, Theorem 1.1].

Definition 2.7.14 (Isolated orthogonality)**.** Let $(X, \mathfrak{S})$ be an HHS and let S be the $\sqsubseteq$–maximal element of $\mathfrak{S}$. We say $(X, \mathfrak{S})$ has *isolated orthogonality* if there exists a collection of domains $\mathcal{I} \subseteq \mathfrak{S} \smallsetminus \{S\}$ satisfying the following conditions.

(1) For all $V, W \in \mathfrak{S}$ with $V \perp W$, there exists $U \in \mathcal{I}$ such that $V, W \sqsubseteq U$.

(2) $\mathfrak{S}_{U_1} \cap \mathfrak{S}_{U_2} = \varnothing$ for all pairs of distinct $U_1, U_2 \in \mathcal{I}$.

Theorem 2.7.15 (Criterion for relative hyperbolicity; [Rus20, Theorem 1.1])**.** *A hierarchically hyperbolic space* $(G, \mathfrak{S})$ *is relatively hyperbolic if* $\mathfrak{S}$ *has isolated orthogonality.*

A characterisation of thickness is currently unknown for HHGs in their full generality. In fact, it is not even known if HHGs satisfy a strict dichotomy between relative hyperbolicity and thickness, although this is conjectured to be true, and has been shown in many cases. For example, Behrstock–Druţu–Mosher show that mapping class groups are either hyperbolic or thick of order 1 [BDM09, Theorem 8.1], and Brock–Masur build upon results of Behrstock–Druţu–Mosher to show that Teichmüller space with the Weil–Petersson metric is either hyperbolic, relatively hyperbolic, or thick of order 1 [BDM09, Theorem 12.5][BM08, Theorems 1, 6]. Characterisations of thickness and relative hyperbolicity also exist for right-angled Artin and Coxeter groups; see Theorems 2.4.21 and 2.4.26.

2.7.2 Relative HHSs

In [BHS19], Behrstock, Hagen, and Sisto introduce *relative HHSs*, a broader class of spaces obtained by relaxing the hyperbolicity axiom of an HHS. In Chapter 4, we show that all graph products satisfy this version of hierarchical hyperbolicity.

Definition 2.7.16 (Relative HHS/HHG)**.** Let $\mathfrak{S}$ be an E–proto-hierarchy structure for a quasi-geodesic space $\mathcal{X}$. We say $\mathfrak{S}$ is a *relatively E–hierarchically hyperbolic space structure* for $\mathcal{X}$ if $\mathfrak{S}$ satisfies all of the HHS axioms in Definition 2.7.1 except the hyperbolicity axiom, and instead satisfies the following weaker version.

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

If $\mathfrak{S}$ is a relative HHS structure for $\mathcal{X}$, we say the pair $(\mathcal{X}, \mathfrak{S})$ is a *relatively E–hierarchically hyperbolic space*. Furthermore, if $\mathcal{X}$ is the Cayley graph of a finitely generated group G with respect to some finite generating set, and G satisfies conditions (2)–(3) of Definition 2.7.3, then $\mathfrak{S}$ is called a *relatively E–hierarchically hyperbolic group structure* on G, and we say the pair $(G, \mathfrak{S})$ is a *relatively E–hierarchically hyperbolic group*.

Remark 2.7.17. Despite this weakening of the hyperbolicity axiom, relative hierarchical hyperbolicity still retains many of the important properties of hierarchical hyperbolicity; for example, relative HHSs satisfy the distance formula given in Theorem 2.7.11, and relative HHGs are acylindrically hyperbolic under the conditions of Theorem 2.7.13.

2.7.3 Almost HHSs

Just as relative HHSs are obtained by weakening the hyperbolicity axiom, one can obtain an *almost HHS* by weakening the container axiom. The first and foremost consequence of the container axiom is that every HHS structure has 'finite rank', i.e., a uniform bound on the size of any pairwise orthogonal collection of domains.

Lemma 2.7.18 ([BHS19, Lemma 2.1])**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an E–hierarchically hyperbolic space. If $W_1, \dots, W_n$ is a pairwise orthogonal collection of elements of $\mathfrak{S}$, then $n \leqslant E$.*

Observing that many consequences of being a hierarchically hyperbolic space structure still apply when the container axiom is replaced with the conclusion of Lemma 2.7.18, Abbott, Behrstock, and Durham coined the term *almost HHS structure* to describe such spaces.

Definition 2.7.19 (Almost HHS)**.** Let $\mathfrak{S}$ be an E–proto-hierarchy structure for a quasi-geodesic space $\mathcal{X}$. We say $\mathfrak{S}$ is an *almost E–hierarchically hyperbolic space structure* for $\mathcal{X}$ if $\mathfrak{S}$ satisfies all of the HHS axioms in Definition 2.7.1 except the container axiom, and instead satisfies the following restriction on the orthogonality relation.

(3′) **(Finite rank)** If $W_1, \dots, W_n$ is a pairwise orthogonal collection of elements of $\mathfrak{S}$, then $n \leqslant E$.

If $\mathfrak{S}$ is an almost HHS structure for $\mathcal{X}$, we say the pair $(\mathcal{X}, \mathfrak{S})$ is an *almost hierarchically hyperbolic space.*

Remark 2.7.20. Lemma 2.7.5, Lemma 2.7.8, Theorem 2.7.9, and Lemma 2.7.10 also hold in the almost HHS setting, since the only use of the container axiom in their proofs is Lemma 2.7.18.

Using this weakened axiom, Abbott, Behrstock, and Durham were able to show several impressive results, including a complete characterisation of contracting quasigeodesics in HHGs [ABD21]. In Chapter 3, we show that every almost HHS structure can in fact be upgraded to an HHS structure, simplifying Abbott–Behrstock–Durham's proofs significantly. This also has implications for Chapter 4, where it enables us to show that any graph product of HHGs can be endowed with an HHG structure.

2.7.4 HHS structures on CAT(0) cube complexes

We shall conclude this chapter by providing explicit examples of HHS structures for CAT(0) cube complexes, which in turn will be key in constructing HHG structures on graph braid groups in Chapter 5. As shown in [BHS17b], HHS structures can be put on CAT(0) cube complexes by studying *factor systems.*

Definition 2.7.21 (Factor system)**.** A *factor system* $\mathfrak{F}$ on a CAT(0) cube complex X is a collection of subcomplexes of X with the following properties.

(1) $X \in \mathfrak{F}$.

(2) Each $F \in \mathfrak{F}$ is non-empty and convex.

(3) There exists $\Delta \geqslant 1$ such that each $x \in X^{(0)}$ is contained in at most Δ elements of $\mathfrak{F}$.

(4) $\mathfrak{F}$ contains all non-trivial convex subcomplexes parallel to combinatorial hyperplanes of X.

(5) There exists $\xi \geqslant 0$ such that for all $F, F' \in \mathfrak{F}$, either $\mathfrak{g}_F(F') \in \mathfrak{F}$ or $\operatorname{diam}(\mathfrak{g}_F(F')) < \xi$.

Behrstock, Hagen, and Sisto show that a special cube complex C with finitely many hyperplanes has a canonical factor system on its universal cover X. This factor system is obtained by considering the collection of all subgraphs of the *crossing graph* of C.

Definition 2.7.22 (Crossing graph)**.** Let X be a cube complex. The *crossing graph* Ξ of X is defined as follows.

- The vertices of Ξ are the hyperplanes of X.
- Two vertices of Ξ are connected by an edge if the corresponding hyperplanes cross in X.

Let $\mathcal{H}$ be the collection of hyperplanes of C, let Ξ be the crossing graph of C, and let $\mathcal{R}$ be the collection of all subgraphs of Ξ. Note that there is a one-to-one correspondence between the vertices of Ξ and the elements of $\mathcal{H}$, therefore each $\Omega \in \mathcal{R}$ corresponds to a subset $\mathcal{A} \subseteq \mathcal{H}$, by taking $\mathcal{A} = \Omega^{(0)}$.

Given two edges E, E' of C, write $E \sim_\Omega E'$ if there is a path γ in $C^{(1)}$ from E to E' such that every edge of γ (including E and E') is dual to some hyperplane in $\Omega^{(0)}$. Let $[E]_\Omega$ denote the equivalence class of E with respect to $\sim_\Omega$. Define C_Ω to be the collection of induced subcomplexes of C whose 1–skeleton is $[E]_\Omega$ for some edge E.

Theorem 2.7.23 ([BHS17b, Corollary 8.9])**.** *Let C be a special cube complex with finitely many hyperplanes, and let X be its universal cover. Let Ξ be the crossing graph of C and let $\mathcal{R}$ be the collection of all subgraphs of Ξ. The collection of all lifts of subcomplexes in $\bigcup_{\Omega \in \mathcal{R}} C_\Omega$ forms a factor system for X.*

Denote this factor system by $\mathfrak{F}$. A proto-hierarchy structure $\mathfrak{S}$ can then be constructed for X as follows. See [BHS17b, Remark 13.2] for full details on how $\mathfrak{S}$ forms an HHS structure.

Define an equivalence relation $||$ on $\mathfrak{F}$ as follows. Given two subcomplexes $F, F' \in \mathfrak{F}$, write $F||F'$ if each hyperplane of X crosses F if and only if it crosses F'. We say F and F' are *parallel*, and call the equivalence class $[F]$ of F its *parallelism class*. We then define the index set $\mathfrak{S}$ for the proto-hierarchy structure by choosing exactly one element of $\mathfrak{F}$ in each parallelism class.

Given $F \in \mathfrak{S}$, define $C_0(F)$ to be the graph whose vertices are the hyperplanes of F, and where two vertices are connected by an edge if the carriers of the corresponding hyperplanes intersect. We call $C_0(F)$ the *contact graph* of F. We then take $C(F)$ to be the *factored contact graph*, defined as follows.

Let $F \in \mathfrak{F}$ and define $\mathfrak{F}_F = \{F' \cap F \mid F' \in \mathfrak{F}\}$. For each $Y \in \mathfrak{F}_F$, we can then consider $C_0(Y)$ as a subgraph of $C_0(F)$. For each parallelism class $[Y]$, if $Y \in \mathfrak{F}_F \smallsetminus \{F\}$ and Y is either parallel to a combinatorial hyperplane or has diameter at least ξ, then cone off $C_0(Y)$ in $C_0(F)$. That is, for each such $[Y]$, add a vertex v_Y to $C_0(F)$ and add edges connecting v_Y to every vertex of $C_0(Y) \subseteq C_0(F)$. The resulting graph is the *factored contact graph* $C(F)$.

The proto-hierarchy structure is then given as follows.

(1) **(Projections.)** Given $F \in \mathfrak{S}$, define the projection map $\pi_F : X \to 2^{C(F)}$ as $\pi_F = i_F \circ \rho_F \circ \mathfrak{g}_F$, where $\mathfrak{g}_F : X \to F$ is the gate map onto the subcomplex $F \subseteq X$, $\rho_F : F \to 2^{C_0(F)}$ is defined by taking $\rho_F(x)$ to be the maximal collection of hyperplanes of F whose carriers all contain x, and $i_F : 2^{C_0(F)} \to 2^{C(F)}$ is the inclusion map.

(2) **(Nesting.)** Given $F, F' \in \mathfrak{S}$, we say F is *nested* in F' (written $F \sqsubseteq F'$) if there exists a subcomplex $K \subseteq F'$ such that $F||K$. If $F \sqsubsetneq F'$, then the set $\mathcal{H}$ of hyperplanes crossing F is the same as the set of hyperplanes crossing K, and therefore $\mathcal{H}$ is a subset of the

set of hyperplanes crossing F'. It follows that $C_0(F) \subseteq C_0(F')$. The upwards relative projection is then defined to be $\rho^F_{F'} = \pi_{F'}(F) = i_{F'}(C_0(F))$.

(3) **(Orthogonality.)** Given $F, F' \in \mathfrak{S}$, we say F is *orthogonal* to F' (written $F \perp F'$) if there exist $K || F$ and $K' || F'$ such that there is a cubical isometric embedding $F \times F' \to X$ where K is the image of $F \times \{\mathfrak{g}_{F'}(K)\}$ and K' is the image of $\{\mathfrak{g}_F(K')\} \times F'$.

(4) **(Transversality.)** If $F, F' \in \mathfrak{S}$ are not nested or orthogonal, then we say they are *transverse* (written $F \pitchfork F'$). The lateral relative projection from F to F' is defined as $\rho^F_{F'} = \pi_{F'}(F)$, and $\rho^{F'}_F$ is defined in the same way.

Theorem 2.7.24 (Special groups are HHGs; [BHS17b, Proposition B, Remark 13.2])**.** *Let X be a special cube complex with finitely many hyperplanes. Then its universal cover is a hierarchically hyperbolic space, and $\pi_1(X)$ is a hierarchically hyperbolic group.*

Corollary 2.7.25 (Graph braid groups are HHGs)**.** *Let Γ be a finite graph. Then $B_n(\Gamma, S)$ is a hierarchically hyperbolic group for all $n \geqslant 1$ and for all $S \in UC_n(\Gamma)$.*

Proof. The case where Γ is connected follows from such graph braid groups being fundamental groups of special cube complexes with finitely many hyperplanes (Corollary 2.5.7), which are HHGs by Theorem 2.7.24. If Γ has more than one connected component, then we can express $B_n(\Gamma, S)$ as a product of graph braid groups of connected components of Γ by Lemma 2.5.2. Thus, $B_n(\Gamma, S)$ is a product of hierarchically hyperbolic groups, whence it follows that $B_n(\Gamma, S)$ is itself a hierarchically hyperbolic group, by a combination theorem of Behrstock, Hagen, and Sisto [BHS17b, Corollary 8.26]. □

Chapter 3

Almost HHSs are HHSs

The main result of this section establishes that all almost HHS structures can be promoted to HHS structures by adding dummy domains to serve as the containers.

Theorem 3.0.1 (Almost HHSs are HHSs)**.** *Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS. There exists an HHS structure $\mathfrak{T}$ for $\mathcal{X}$ so that $\mathfrak{S} \subseteq \mathfrak{T}$, and if $W \in \mathfrak{T} \smallsetminus \mathfrak{S}$ then the associated hyperbolic space for W is a single point.*

If $(\mathcal{X}, \mathfrak{S})$ is an almost HHS, then the only HHS axiom that is not satisfied is the container axiom. The most obvious way to address this is to add an extra element to $\mathfrak{S}$ every time we need a container. That is, if $V, W \in \mathfrak{S}$ with $V \sqsubseteq W$ and there exists some $Q \sqsubseteq W$ with $Q \perp V$, then we add a domain D_W^V to serve as the container for V in W, i.e., every Q nested into W and orthogonal to V will be nested into D_W^V. However, this approach is perilous as once a domain Q is nested into D_W^V, we may now need a container for Q in D_W^V! To avoid this, we add domains $D_W^{\mathcal{V}}$ where $\mathcal{V}$ is a pairwise orthogonal set of domains nested into W; that is, $D_W^{\mathcal{V}}$ contains all domains Q that are nested into W and orthogonal to all $V \in \mathcal{V}$. This allows for all the needed containers to be added at once, avoiding an iterative process.

Remark 3.0.2. To prove Theorem 3.0.1, we shall require Lemma 2.7.5, Lemma 2.7.8, Theorem 2.7.9, and Lemma 2.7.10. Each of these results was originally proved in the setting of hierarchically hyperbolic spaces, but as noted in Remark 2.7.20 they continue to hold in the almost HHS setting.

Proof of Theorem 3.0.1. Let $(\mathcal{X}, \mathfrak{S})$ be an almost HHS. Let $\mathcal{V}$ denote a non-empty set of pairwise orthogonal elements of $\mathfrak{S}$ and let $W \in \mathfrak{S}$. We say the pair $(W, \mathcal{V})$ is a *container pair* if the following are satisfied:

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

Let $\mathfrak{D}$ denote the set of all container pairs. We will denote a pair $(W, \mathcal{V}) \in \mathfrak{D}$ by $D_W^{\mathcal{V}}$.

Let $\mathfrak{T} = \mathfrak{S} \cup \mathfrak{D}$. If $D_W^{\mathcal{V}} \in \mathfrak{D}$, then the associated hyperbolic space, $C(D_W^{\mathcal{V}})$, will be a single point.

Claim 3.0.3. $\mathcal{X}$ admits a proto-hierarchy structure with index set $\mathfrak{T}$.

Proof. Since $(\mathcal{X}, \mathfrak{S})$ is an almost HHS, we can continue to use the spaces, projections, and relations of $\mathfrak{S}$. Thus, it suffices to verify the axioms for elements of $\mathfrak{D}$ and relations involving elements of $\mathfrak{D}$.

Projections: For $D_W^{\mathcal{V}} \in \mathfrak{D}$, the projection map is just the constant map to the single point in $C(D_W^{\mathcal{V}})$.

Nesting: Let $Q \in \mathfrak{S}$ and $D_W^{\mathcal{V}}, D_T^{\mathcal{R}} \in \mathfrak{D}$.

- Define $Q \sqsubseteq D_W^{\mathcal{V}}$ if $Q \sqsubseteq W$ in $\mathfrak{S}$ and $Q \perp V$ for all $V \in \mathcal{V}$.
- Define $D_W^{\mathcal{V}} \sqsubseteq Q$ if $W \sqsubseteq Q$ in $\mathfrak{S}$.
- Define $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$ if $W \sqsubseteq T$ in $\mathfrak{S}$ and for all $R \in \mathcal{R}$ either $R \perp W$ or there exists $V \in \mathcal{V}$ with $R \sqsubseteq V$.

These definitions ensure $\sqsubseteq$ is still a partial order and maintain the $\sqsubseteq$–maximal element of $\mathfrak{S}$ as the $\sqsubseteq$–maximal element of $\mathfrak{T}$.

Since the hyperbolic spaces associated to elements of $\mathfrak{D}$ are points, define $\rho^{Q}_{D^{\mathcal{V}}_W} = C(D^{\mathcal{V}}_W)$ for every $Q \in \mathfrak{T}$ and $D^{\mathcal{V}}_W \in \mathfrak{D}$ with $Q \sqsubsetneq D^{\mathcal{V}}_W$. If $D^{\mathcal{V}}_W \in \mathfrak{D}$ and $Q \in \mathfrak{S}$ with $D^{\mathcal{V}}_W \sqsubsetneq Q$, then $V \sqsubsetneq Q$ in $\mathfrak{S}$ for each $V \in \mathcal{V}$. Thus we define $\rho^{D^{\mathcal{V}}_W}_Q = \bigcup_{V \in \mathcal{V}} \rho^V_Q$. Lemma 2.7.5 ensures that $\rho^{D^{\mathcal{V}}_W}_Q$ has diameter at most $4E$.

Orthogonality: Two elements $D^{\mathcal{V}}_W, D^{\mathcal{R}}_T \in \mathfrak{D}$ are orthogonal if $W \perp T$ in $\mathfrak{S}$. Let $Q \in \mathfrak{S}$ and $D^{\mathcal{V}}_W \in \mathfrak{D}$. Define $Q \perp D^{\mathcal{V}}_W$ if, in $\mathfrak{S}$, either $Q \perp W$ or $Q \sqsubseteq V$ for some $V \in \mathcal{V}$.

Transversality: An element of $\mathfrak{T}$ is transverse to an element of $\mathfrak{D}$ whenever it is not nested or orthogonal. Since the hyperbolic spaces associated to elements of $\mathfrak{D}$ are points, we only need to define the relative projections from an element of $\mathfrak{D}$ to an element of $\mathfrak{S}$. Let $D^{\mathcal{V}}_W \in \mathfrak{D}$ and $Q \in \mathfrak{S}$ and suppose $D^{\mathcal{V}}_W \pitchfork Q$. This implies $W \not\perp Q$ and $W \not\sqsubseteq Q$. We define $\rho^{D^{\mathcal{V}}_W}_Q$ based on the $\mathfrak{S}$–relation between Q and the elements of $\mathcal{V}$.

- If $Q \perp V$ for all $V \in \mathcal{V}$, then $Q \not\sqsubseteq W$ as $Q \sqsubseteq W$ would imply $Q \sqsubseteq D^{\mathcal{V}}_W$. Thus we must have $Q \pitchfork W$, so we define $\rho^{D^{\mathcal{V}}_W}_Q = \rho^W_Q$.

- If $V \pitchfork Q$ or $V \sqsubsetneq Q$ for some $V \in \mathcal{V}$, then ρ^V_Q exists and we define $\rho^{D^{\mathcal{V}}_W}_Q$ to be the union of all the ρ^V_Q for $V \in \mathcal{V}$ with $V \pitchfork Q$ or $V \sqsubsetneq Q$. Lemma 2.7.5 ensures $\rho^{D^{\mathcal{V}}_W}_Q$ has diameter at most $4E$ in this case.

- If $Q \sqsubseteq V$ for some V, then $Q \perp D^{\mathcal{V}}_W$ which contradicts $Q \pitchfork D^{\mathcal{V}}_W$, so this case does not occur. □

We now prove that $(\mathcal{X}, \mathfrak{T})$ is a hierarchically hyperbolic space. This will complete the proof of Theorem 3.0.1. By abuse of notation, let E be the largest of the constants for the proto-structure of $\mathfrak{T}$.

Hyperbolicity: For all elements of $\mathfrak{D}$ the associated spaces are points and thus hyperbolic. For elements of $\mathfrak{S}$, the associated spaces are hyperbolic since $\mathfrak{S}$ is an almost HHS structure.

Finite complexity: First consider a nesting chain of the form $D^{\mathcal{V}_1}_W \sqsubsetneq D^{\mathcal{V}_2}_W \sqsubsetneq \ldots \sqsubsetneq D^{\mathcal{V}_n}_W$.

Claim 3.0.4. The length of $D_W^{\mathcal{V}_1} \subsetneq D_W^{\mathcal{V}_2} \subsetneq \ldots \subsetneq D_W^{\mathcal{V}_n}$ is bounded by $E^2 + E$.

Proof. For each $V \in \bigcup_{i=1}^n \mathcal{V}_i$, we have $V \sqsubseteq W$ and hence $V \not\perp W$. As $D_W^{\mathcal{V}_{i-1}} \subsetneq D_W^{\mathcal{V}_i}$ for each $i \in \{2, \ldots, n\}$, every element of $\mathcal{V}_i$ must therefore be nested into an element of $\mathcal{V}_{i-1}$. Denote the elements of $\mathcal{V}_i$ by $V_1^i, \ldots V_{k_i}^i$. Since each $\mathcal{V}_i$ is a pairwise orthogonal subset of $\mathfrak{S}$, we have $k_i \leqslant E$ for each $i \in \{1, \ldots, n\}$ by the finite rank axiom of an almost HHS (Definition 2.7.19). We define a $\mathcal{V}$*–nesting chain* to be a maximal chain of the form $V_{j_m}^m \sqsubseteq V_{j_{m-1}}^{m-1} \sqsubseteq \ldots \sqsubseteq V_{j_1}^1$ for some $m \in \{1, \ldots, n\}$ and $j_i \in \{1, \ldots, k_i\}$, with $i \in \{1, \ldots, m\}$. Since the elements of $\mathcal{V}_i$ are pairwise orthogonal for each $i \in \{1, \ldots, n\}$, if $V_{j_m}^m$ is the $\sqsubseteq$–minimal element of a $\mathcal{V}$–nesting chain, then $V_{j_m}^m$ is nested into exactly one element of $\mathcal{V}_i$ for each $i \leqslant m$. This implies that each $\mathcal{V}$-nesting chain is determined by its $\sqsubseteq$–minimal element. Further, the set of $\sqsubseteq$–minimal elements of $\mathcal{V}$–nesting chains is pairwise orthogonal. By the finite rank axiom of an almost HHS, this implies there exist at most E $\mathcal{V}$–nesting chains.

In order for $D_W^{\mathcal{V}_i} \neq D_W^{\mathcal{V}_{i+1}}$, either $k_{i+1} < k_i$ or there exists $j_i \in \{1, \ldots, k_i\}$, $j_{i+1} \in \{1, \ldots, k_{i+1}\}$ such that $V_{j_{i+1}}^{i+1} \subsetneq V_{j_i}^i$. Thus, every step up the chain $D_W^{\mathcal{V}_1} \subsetneq D_W^{\mathcal{V}_2} \subsetneq \ldots \subsetneq D_W^{\mathcal{V}_n}$ results in either a strict decrease in k_i (the cardinality of $\mathcal{V}_i$) to k_{i+1} (the cardinality of $\mathcal{V}_{i+1}$) or a strict decrease within one of the $\mathcal{V}$–nesting chains. Note that k_i may increase when we encounter a strict decrease in one of the $\mathcal{V}$–nesting chains, since multiple elements of $\mathcal{V}_{i+1}$ may be nested into the same element of $\mathcal{V}_i$. However, this may only happen at most $E - k_1$ times, as there are at most E $\mathcal{V}$–nesting chains. Hence, the length of $D_W^{\mathcal{V}_1} \subsetneq D_W^{\mathcal{V}_2} \subsetneq \ldots \subsetneq D_W^{\mathcal{V}_n}$ is bounded by E plus the total number of times a strict decrease can occur across all of the $\mathcal{V}$–nesting chains.

Each $\mathcal{V}$–nesting chain $V_{j_m}^m \sqsubseteq V_{j_{m-1}}^{m-1} \sqsubseteq \ldots \sqsubseteq V_{j_1}^1$ contains at most E distinct elements of $\mathfrak{S}$ by the finite complexity of $\mathfrak{S}$. Finite rank implies there are at most E different $\mathcal{V}$–nesting chains, thus the number of steps of the chain $D_W^{\mathcal{V}_1} \subsetneq D_W^{\mathcal{V}_2} \subsetneq \ldots \subsetneq D_W^{\mathcal{V}_n}$ where there is a strict decrease within one of the $\mathcal{V}$–nesting chains is at most E^2. This bounds the length of $D_W^{\mathcal{V}_1} \subsetneq D_W^{\mathcal{V}_2} \subsetneq \ldots \subsetneq D_W^{\mathcal{V}_n}$ by $E^2 + E$. □

We now consider a nesting chain of the form $D_{W_1}^{\mathcal{V}_1} \sqsubsetneq D_{W_2}^{\mathcal{V}_2} \sqsubsetneq \ldots \sqsubsetneq D_{W_n}^{\mathcal{V}_n}$. In this case, $W_1 \sqsubseteq W_2 \sqsubseteq \ldots \sqsubseteq W_n$, but not all of these nestings must be proper. Let $1 = i_1 < i_2 < \cdots < i_k$ be the minimal subset of $\{1, \ldots, n\}$ such that if $i_j \leqslant i < i_{j+1}$, then $W_{i_j} = W_i$. Thus $W_{i_1} \sqsubsetneq W_{i_2} \sqsubsetneq \ldots \sqsubsetneq W_{i_k}$ and $k \leqslant E$. Claim 3.0.4 established that $|i_j - i_{j+1}| \leqslant E^2 + E$, so $n \leqslant k(E^2 + E) \leqslant E^3 + E^2$, that is, any $\sqsubsetneq$–chain of elements of $\mathfrak{D}$ has length at most $E^3 + E^2$.

Finally, since any $\sqsubsetneq$–chain of elements of $\mathfrak{T}$ can be partitioned into a $\sqsubsetneq$–chain of elements of $\mathfrak{D}$ and a $\sqsubsetneq$–chain of elements of $\mathfrak{S}$, any $\sqsubsetneq$–chain in $\mathfrak{T}$ has length at most $E^3 + E^2 + E$.

Containers: Let $W, V \in \mathfrak{S}$ with $V \sqsubsetneq W$ and $\{Q \in \mathfrak{T}_W : Q \perp V\} \neq \varnothing$, i.e., $(W, \{V\})$ is a container pair. In this case, the container of V in W for $\mathfrak{T}$ is $D_W^{\{V\}}$.

We now show containers exist for situations involving elements of $\mathfrak{D}$. We split this into three subcases.

Case 1: $\mathbf{D_W^{\mathcal{V}} \in \mathfrak{D}}$ and $\mathbf{Q \in \mathfrak{S}}$ with $\mathbf{D_W^{\mathcal{V}} \sqsubseteq Q}$. Since $(W, \mathcal{V})$ is a container pair, there exists $P \in \mathfrak{S}$ with $P \sqsubseteq W$ and $P \perp V$ for all $V \in \mathcal{V}$. Suppose that $D_W^{\mathcal{V}}$ requires a container in Q, that is, there is an element of $\mathfrak{T}$ that is orthogonal to $D_W^{\mathcal{V}}$ and nested in Q. We verify that $(Q, \{P\})$ is a container pair and $D_Q^{\{P\}}$ is a container of $D_W^{\mathcal{V}}$ in Q.

If $T \in \mathfrak{S}$ with $T \perp D_W^{\mathcal{V}}$ and $T \sqsubseteq Q$, then $T \perp W$ or $T \sqsubseteq V$ for some $V \in \mathcal{V}$. In either case, we have $T \perp P$, so $(Q, \{P\})$ is a container pair and $T \sqsubseteq D_Q^{\{P\}}$. If $D_T^{\mathcal{R}} \in \mathfrak{D}$ with $D_T^{\mathcal{R}} \sqsubseteq Q$ and $D_T^{\mathcal{R}} \perp D_W^{\mathcal{V}}$, then $T \sqsubseteq Q$ and $T \perp W$. Since $P \sqsubseteq W$, this implies $T \perp P$ and so $(Q, \{P\})$ is again a container pair, and $D_T^{\mathcal{R}} \sqsubseteq D_Q^{\{P\}}$.

Case 2: $\mathbf{D_W^{\mathcal{V}}, D_T^{\mathcal{R}} \in \mathfrak{D}}$ where $\mathbf{D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}}$. Since $(W, \mathcal{V})$ is a container pair, there exists $P \in \mathfrak{S}$ so that $P \sqsubseteq W$ and $P \perp V$ for all $V \in \mathcal{V}$. Since $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$, it follows that for all $R \in \mathcal{R}$, either $R \perp W$ or there exists $V \in \mathcal{V}$ so that $R \sqsubseteq V$. In both cases, $R \perp P$. Thus $\mathcal{P} = \mathcal{R} \cup \{P\}$ is a pairwise orthogonal collection of elements of $\mathfrak{S}$. Suppose that $D_W^{\mathcal{V}}$ requires a container in $D_T^{\mathcal{R}}$, that is, there is an element of $\mathfrak{T}$ that is orthogonal to $D_W^{\mathcal{V}}$ and nested in $D_T^{\mathcal{R}}$. We verify that $(T, \mathcal{P})$ is a container pair and $D_T^{\mathcal{P}} \sqsubsetneq D_T^{\mathcal{R}}$ is a container for $D_W^{\mathcal{V}}$ in $D_T^{\mathcal{R}}$.

If $Q \in \mathfrak{S}$ satisfies $Q \sqsubseteq D_T^{\mathcal{R}}$ and $D_W^{\mathcal{V}} \perp Q$, then $Q \sqsubseteq T$ and we have either $Q \perp W$ or

$Q \sqsubseteq V$ for some $V \in \mathcal{V}$. In both cases, $Q \perp P$. Further, we must have $Q \perp R$ for each $R \in \mathcal{R}$ as $Q \sqsubseteq D_T^{\mathcal{R}}$. Thus $(T, \mathcal{P})$ is a container pair and $Q \sqsubseteq D_T^{\mathcal{P}}$. On the other hand, if $D_Q^{\mathcal{Z}} \in \mathfrak{D}$ satisfies $D_Q^{\mathcal{Z}} \perp D_W^{\mathcal{V}}$ and $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{R}}$, then $Q \perp W$, $Q \sqsubseteq T$, and for each $R \in \mathcal{R}$ either $Q \perp R$ or there exists $Z \in \mathcal{Z}$ with $R \sqsubseteq Z$. Since $(Q, \mathcal{Z})$ is a container pair, there exists $U \in \mathfrak{S}$ such that $U \sqsubseteq Q$ and $U \perp Z$ for all $Z \in \mathcal{Z}$. Since $Q \perp W$, we also have $U \perp P$ as $U \sqsubseteq Q$ and $P \sqsubseteq W$. For each $R \in \mathcal{R}$, either $R \perp Q$ or there exists $Z \in \mathcal{Z}$ with $R \sqsubseteq Z$. In both cases, $R \perp U$. Thus, U is orthogonal to all elements of $\mathcal{P} = \mathcal{R} \cup \{P\}$ and moreover $U \sqsubseteq Q \sqsubseteq T$, so $(T, \mathcal{P})$ is a container pair. Furthermore, $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{P}} = D_T^{\mathcal{R} \cup \{P\}}$ since $D_Q^{\mathcal{Z}} \sqsubseteq D_T^{\mathcal{R}}$ and $P \perp Q$. We have therefore shown that $D_T^{\mathcal{P}}$ is a container for $D_W^{\mathcal{V}}$ in $D_T^{\mathcal{R}}$.

Case 3: $\mathbf{D_T^{\mathcal{R}} \in \mathfrak{D}}$ and $\mathbf{Q \in \mathfrak{S}}$ with $\mathbf{Q \sqsubseteq D_T^{\mathcal{R}}}$. This implies $\mathcal{Q} = \mathcal{R} \cup \{Q\}$ is a pairwise orthogonal set of elements of $\mathfrak{S}$. Further, suppose that Q requires a container in $D_T^{\mathcal{R}}$, that is, there is an element of $\mathfrak{T}$ that is orthogonal to Q and nested in $D_T^{\mathcal{R}}$. We verify that $(T, \mathcal{Q})$ is a container pair and $D_T^{\mathcal{Q}}$ is a container for Q in $D_T^{\mathcal{R}}$.

Suppose there exists $V \in \mathfrak{S}$ with $V \sqsubseteq D_T^{\mathcal{R}}$ and $V \perp Q$. Then $V \sqsubseteq T$ and V is orthogonal to all the elements of $\mathcal{R} \cup \{Q\}$. Thus $(T, \mathcal{Q})$ is a container pair, so $D_T^{\mathcal{Q}}$ exists and $V \sqsubseteq D_T^{\mathcal{Q}}$. Now suppose there exists $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$ such that $Q \perp D_W^{\mathcal{V}}$. Since $(W, \mathcal{V})$ is a container pair, there exists $U \in \mathfrak{S}$ with $U \sqsubseteq W$ and U orthogonal to each element of $\mathcal{V}$. As $Q \perp D_W^{\mathcal{V}}$, we have $Q \perp W$ or $Q \sqsubseteq V$ for some $V \in \mathcal{V}$. In both cases, $Q \perp U$. Therefore U is orthogonal to every element of $\mathcal{Q}$, and moreover $U \sqsubseteq W \sqsubseteq T$ since $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{R}}$. Thus $(T, \mathcal{Q})$ is a container pair and $U \sqsubseteq D_T^{\mathcal{Q}}$. Now, for each $R \in \mathcal{R}$, either $R \perp W$ or $R \sqsubseteq V$ for some $V \in \mathcal{V}$. Since $\mathcal{Q} = \mathcal{R} \cup \{Q\}$ and $Q \perp W$, this implies $D_W^{\mathcal{V}} \sqsubseteq D_T^{\mathcal{Q}}$. Thus, $(T, \mathcal{Q})$ is a container pair and $D_T^{\mathcal{Q}}$ is a container for Q in $D_T^{\mathcal{R}}$.

Uniqueness, bounded geodesic image, large links: Since the only elements of $\mathfrak{T}$ whose associated spaces are not points are in $\mathfrak{S}$, these axioms for $(\mathcal{X}, \mathfrak{T})$ follow from the fact they hold in $(\mathcal{X}, \mathfrak{S})$.

Consistency: Since the only elements of $\mathfrak{T}$ whose associated spaces are not points are in $\mathfrak{S}$, the first inequality in the consistency axiom for $(\mathcal{X}, \mathfrak{S})$ implies the same inequality for $(\mathcal{X}, \mathfrak{T})$. To verify the final clause of the consistency axiom, we need to check that if $Q, R, T \in \mathfrak{T}$ such that $Q \sqsubsetneq R$ with ρ_T^R and ρ_T^Q both defined, then $\mathsf{d}_T(\rho_T^Q, \rho_T^R)$ is uniformly bounded in terms of E. We can assume $T \in \mathfrak{S}$ as $C(T)$ has diameter zero otherwise. We can further assume at least one of Q and R is an element of $\mathfrak{D}$, as we have the consistency axiom for elements of $\mathfrak{S}$.

Case 1: $\mathbf{Q \sqsubsetneq R \sqsubsetneq T}$.

- Assume $Q \in \mathfrak{S}$ and $R = D_W^{\mathcal{V}} \in \mathfrak{D}$. Fix $V \in \mathcal{V}$. Since $D_W^{\mathcal{V}} = R \sqsubseteq T$ and $\rho_T^{D_W^{\mathcal{V}}} = \bigcup_{U \in \mathcal{V}} \rho_T^U$, we have $\rho_T^V \subseteq \rho_T^{D_W^{\mathcal{V}}} = \rho_T^R$. Since $V \perp Q$, Lemma 2.7.5 says $\mathsf{d}_T(\rho_T^R, \rho_T^Q) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^Q) \leqslant 2E$.

- Assume $Q = D_W^{\mathcal{V}} \in \mathfrak{D}$ and $R \in \mathfrak{S}$. Fix $V \in \mathcal{V}$. In this case, $\rho_T^V \subseteq \rho_T^Q$ since $D_W^{\mathcal{V}} = Q \sqsubsetneq T$. Since $D_W^{\mathcal{V}} = Q \sqsubsetneq R$, we have $V \sqsubsetneq W \sqsubseteq R$. Thus, the consistency axiom for $\mathfrak{S}$ says $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^R) \leqslant E$.

- Assume $Q = D_W^{\mathcal{V}} \in \mathfrak{D}$ and $R = D_{W'}^{\mathcal{V}'}$. Thus $W \sqsubseteq W' \sqsubsetneq T$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho_T^W, \rho_T^{W'}) \leqslant E$. Fix $V \in \mathcal{V}$ and $V' \in \mathcal{V}'$. Consistency in $\mathfrak{S}$ also implies $\mathsf{d}_T(\rho_T^V, \rho_T^W) \leqslant E$ and $\mathsf{d}_T(\rho_T^{V'}, \rho_T^{W'}) \leqslant E$. Since $\rho_T^V \subseteq \rho_T^Q$ and $\rho_T^{V'} \subseteq \rho_T^R$, we have $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^{V'}) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^W) + \operatorname{diam}(\rho_T^W) + \mathsf{d}_T(\rho_T^W, \rho_T^{W'}) + \operatorname{diam}(\rho_T^{W'}) + \mathsf{d}_T(\rho_T^{W'}, \rho_T^{V'}) \leqslant 5E$.

Case 2: $\mathbf{Q \sqsubsetneq R}$, $\mathbf{R \pitchfork T}$, and $\mathbf{Q \not\perp T}$. In this case we have either $Q \pitchfork T$ or $Q \sqsubsetneq T$.

- Assume $Q \in \mathfrak{S}$ and $R = D_W^{\mathcal{V}} \in \mathfrak{D}$. Since $D_W^{\mathcal{V}} = R \pitchfork T$, we cannot have $T \sqsubseteq V$ for any $V \in \mathcal{V}$ (this would imply $D_W^{\mathcal{V}} \perp T$). If $V \perp T$ for all $V \in \mathcal{V}$, then $W \pitchfork T$ (as shown in the proof of transversality in Claim 3.0.3) and $\rho_T^R = \rho_T^{D_W^{\mathcal{V}}} = \rho_T^W$. Since $Q \sqsubseteq R = D_W^{\mathcal{V}}$, we have $Q \sqsubseteq W$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho_T^Q, \rho_T^R) = \mathsf{d}_T(\rho_T^Q, \rho_T^W) \leqslant E$. If instead there exists $V \in \mathcal{V}$ so that $T \pitchfork V$ or $V \sqsubsetneq T$, then $\rho_T^V \subseteq \rho_T^{D_W^{\mathcal{V}}} = \rho_T^R$. Since $Q \sqsubseteq R = D_W^{\mathcal{V}}$, we have $Q \perp V$ and Lemma 2.7.5 gives $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^Q, \rho_T^V) \leqslant 2E$.

- Assume $Q = D_W^{\mathcal{V}} \in \mathfrak{D}$ and $R \in \mathfrak{S}$. As before, $T \not\sqsubseteq V$ for all $V \in \mathcal{V}$. First assume there exists $V \in \mathcal{V}$ so that $V \pitchfork T$ or $V \subsetneq T$. This occurs when either $D_W^{\mathcal{V}} = Q \subsetneq T$ or $Q \pitchfork T$ and not every element of $\mathcal{V}$ is orthogonal to T. In both cases, $\rho_T^V \subseteq \rho_T^{D_W^{\mathcal{V}}} = \rho_T^Q$ and consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^R) \leqslant 2E$ because $V \sqsubseteq W \subsetneq R$. Now assume $T \perp V$ for all $V \in \mathcal{V}$. This can only occur when $D_W^{\mathcal{V}} = Q \pitchfork T$. In this case, $W \pitchfork T$ and $\rho_T^Q = \rho_T^{D_W^{\mathcal{V}}} = \rho_T^W$. Since $W \subsetneq R$, consistency in $\mathfrak{S}$ implies $\mathsf{d}_T(\rho_T^R, \rho_T^Q) \leqslant \mathsf{d}_T(\rho_T^R, \rho_T^W) \leqslant E$.

- Assume $Q = D_W^{\mathcal{V}} \in \mathfrak{D}$ and $R = D_{W'}^{\mathcal{V}'} \in \mathfrak{D}$. As before, $T \not\sqsubseteq V$ for all $V \in \mathcal{V} \cup \mathcal{V}'$. If $\rho_T^R = \rho_T^{W'}$, then we have the first case of transversality laid out in the proof of Claim 3.0.3, that is, $W' \pitchfork T$ and $V' \perp T$ for all $V' \in \mathcal{V}$. Thus, if $\rho_T^R = \rho_T^{W'}$, then the result reduces to the previous bullet, replacing R with W'. We can therefore assume $\rho_T^R \neq \rho_T^{W'}$, meaning we have the second case of transversality where there exists $V' \in \mathcal{V}'$ so that V' is either transverse to or properly nested into T.

  Suppose $\rho_T^Q \neq \rho_T^W$ too. This implies there also exists $V \in \mathcal{V}$ so that V is either transverse to or properly nested into T. Furthermore, $\rho_T^V \subseteq \rho_T^Q$ and $\rho_T^{V'} \subseteq \rho_T^R$. Now, $D_W^{\mathcal{V}} \sqsubseteq D_{W'}^{\mathcal{V}'}$ implies $V' \perp W$ or V' is nested into an element of $\mathcal{V}$. If $V' \perp W$, then $V \perp V'$ and Lemma 2.7.5 implies $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^{V'}) \leqslant 2E$. If V' is nested into an element of $\mathcal{V}$, then either $V' \sqsubseteq V$ or $V' \perp V$ since $\mathcal{V}$ is a pairwise orthogonal subset of $\mathfrak{S}$. By applying consistency in $\mathfrak{S}$ when $V' \sqsubseteq V$ or Lemma 2.7.5 when $V' \perp V$, we have $\mathsf{d}_T(\rho_T^Q, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^V, \rho_T^{V'}) \leqslant 2E$.

  Now suppose $\rho_T^Q = \rho_T^W$. Then $D_W^{\mathcal{V}} \sqsubseteq D_{W'}^{\mathcal{V}'}$ implies $V' \perp W$ or V' is nested into W. Applying Lemma 2.7.5 if $V' \perp W$, or consistency in $\mathfrak{S}$ if $V' \sqsubseteq W$, we again obtain $\mathsf{d}_T(\rho_T^Q, \rho_T^R) = \mathsf{d}_T(\rho_T^W, \rho_T^R) \leqslant \mathsf{d}_T(\rho_T^W, \rho_T^{V'}) \leqslant 2E$.

Partial realisation: Let $T_1, \ldots, T_n$ be pairwise orthogonal elements of $\mathfrak{T}$, and let $p_i \in C(T_i)$ for each $i \in \{1, \ldots, n\}$. Without loss of generality, assume $T_1, \ldots, T_k \in \mathfrak{S}$ and

$T_{k+1}, \dots, T_n \in \mathfrak{D}$ where $k \in \{0, \dots, n\}$. If $k = 0$ (resp. $k = n$), then each $T_i \in \mathfrak{D}$ (resp. each $T_i \in \mathfrak{S}$).

For $i \in \{k+1, \dots, n\}$, let $T_i = D^{\mathcal{V}_i}_{W_i}$ and let q_i be any point in $\rho^{D^{\mathcal{V}_i}_{W_i}}_{W_i} \subseteq C(W_i)$. Since $T_1, \dots, T_n$ are pairwise orthogonal, it follows that $W_{k+1}, \dots, W_n$ are pairwise orthogonal too, and for each $j \in \{1, \dots, k\}$, T_j is either nested into an element of some $\mathcal{V}_i$ or orthogonal to all $W_{k+1}, \dots, W_n$. Without loss of generality, assume that $T_1, \dots, T_l$ are nested into elements of $\mathcal{V}_{m+1} \cup \cdots \cup \mathcal{V}_n$ and $T_{l+1}, \dots, T_k, W_{k+1}, \dots, W_n$ are pairwise orthogonal, where $0 \leqslant n - m \leqslant l \leqslant k$. If $l = 0$, then $n = m$ and each T_j is orthogonal to every W_i. Otherwise, for each $j \in \{1, \dots, l\}$, T_j is nested in some W_i for $i \in \{m+1, \dots, n\}$. In both cases, $T_1, \dots, T_k, W_{k+1}, \dots, W_m$ are pairwise orthogonal elements of $\mathfrak{S}$. We can therefore use the partial realisation axiom in $\mathfrak{S}$ on the points $p_1, \dots, p_k, q_{k+1}, \dots, q_m$ to produce a point $x \in \mathcal{X}$ with the following properties:

(1) $\mathsf{d}_{T_i}(x, p_i) \leqslant E$ for $i \in \{1, \dots, k\}$;

(2) $\mathsf{d}_{W_i}(x, q_i) \leqslant E$ for $i \in \{k+1, \dots, m\}$;

(3) for all $i \in \{1, \dots, k\}$ if $Q \pitchfork T_i$ or $T_i \subsetneq Q$, then $\mathsf{d}_Q(x, \rho^{T_i}_Q) \leqslant E$;

(4) for all $i \in \{k, \dots, m\}$ if $Q \pitchfork W_i$ or $W_i \subsetneq Q$, then $\mathsf{d}_Q(x, \rho^{W_i}_Q) \leqslant E$.

Now, for $Q \in \mathfrak{S}$, define $b_Q \in C(Q)$ as follows. Let $\mathcal{V} = \bigcup_{i=k+1}^{n} \mathcal{V}_i$ and $\mathcal{V}_Q = \{V \in \mathcal{V} : V \pitchfork Q \text{ or } V \subsetneq Q\}$. If $\mathcal{V}_Q \neq \varnothing$, then define b_Q to be any point in $\bigcup_{V \in \mathcal{V}_Q} \rho^V_Q$. Since $\mathcal{V}$ is a collection of pairwise orthogonal elements of $\mathfrak{S}$, the diameter of $\bigcup_{V \in \mathcal{V}_Q} \rho^V_Q$ is at most $2E$ by Lemma 2.7.5. If either $Q \sqsubseteq V$ for some $V \in \mathcal{V}$ or $Q \perp V$ for all $V \in \mathcal{V}$ then define $b_Q = \pi_Q(x)$. Since $\mathcal{V}$ is a collection of pairwise orthogonal elements of $\mathfrak{S}$, these two cases encompass all elements of $\mathfrak{S}$.

Claim 3.0.5. The tuple $(b_Q)_{Q \in \mathfrak{S}}$ is $3E$–consistent.

Proof. Let $R, Z \in \mathfrak{S}$. Recall that if $b_Z = \pi_Z(x)$ and $b_R = \pi_R(x)$, then the $3E$–consistency inequalities for b_R and b_Z are satisfied by Lemma 2.7.8. Thus we can assume that there exists $V \in \mathcal{V}$ so that either $V \sqsubsetneq Z$ or $V \pitchfork Z$. Fix $V \in \mathcal{V}$ so that $b_Z \in \rho_Z^V$. We need to verify the consistency inequalities when $R \pitchfork Z$, $R \sqsubsetneq Z$, and $Z \sqsubsetneq R$.

Consistency when R⋔Z: Assume $R \pitchfork Z$. If $R \perp V$, $V \sqsubseteq R$, or $R \sqsubseteq V$ then either Lemma 2.7.5 or consistency in $\mathfrak{S}$ implies $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) \leqslant 2E$. Since $b_Z \in \rho_Z^V$, we have $\mathsf{d}_Z(b_Z, \rho_Z^R) \leqslant 3E$. Thus, we can assume $R \pitchfork V$ so that $\mathcal{V}_R$ is non-empty, that is $b_R \in \bigcup_{U \in \mathcal{V}_R} \rho_R^U$ and so b_R is within $2E$ of ρ_R^V. Now, if $\mathsf{d}_Z(b_Z, \rho_Z^R) > 3E$, then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) > 2E$. Thus ρ–consistency (Lemma 2.7.10) implies $\mathsf{d}_R(\rho_R^V, \rho_R^Z) \leqslant E$. It follows that $\mathsf{d}_R(b_R, \rho_R^Z) \leqslant 3E$ by the triangle inequality.

Consistency when R ⊊ Z: Assume $R \sqsubsetneq Z$. As before, if $R \perp V$, $V \sqsubseteq R$, or $R \sqsubseteq V$ then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) \leqslant 2E$ and we have $\mathsf{d}_Z(b_Z, \rho_Z^R) \leqslant 3E$. Thus, we can assume $R \pitchfork V$ so that b_R is within $2E$ of ρ_R^V. Now, if $\mathsf{d}_Z(b_Z, \rho_Z^R) > 3E$, then $\mathsf{d}_Z(\rho_Z^V, \rho_Z^R) > 2E$, and ρ–consistency implies $\operatorname{diam}(\rho_R^V \cup \rho_R^Z(\rho_Z^V)) \leqslant E$. However, this implies $\operatorname{diam}(b_R \cup \rho_R^Z(b_Z)) \leqslant 3E$ since $b_Z \in \rho_Z^V$ and $\mathsf{d}_R(b_R, \rho_R^V) \leqslant 2E$.

Consistency when Z ⊊ R: Assume $Z \sqsubsetneq R$. If R is orthogonal to all elements of $\mathcal{V}$, then $R \perp V$ implies $V \perp Z$ which contradicts the assumption that $V \sqsubsetneq Z$ or $V \pitchfork Z$. On the other hand, if there exists $V' \in \mathcal{V}$ so that $R \sqsubseteq V'$, then either $V \perp R$ or $R \sqsubseteq V = V'$. But this implies either $V \perp Z$ or $Z \sqsubsetneq V$, both of which give a contradiction if $V \pitchfork Z$ or $V \sqsubsetneq Z$. There must therefore be an element of $\mathcal{V}$ that is either properly nested in or transverse to R, and we can repeat the same argument as in the previous case, switching the roles of R and Z. □

Let $y \in \mathcal{X}$ be the point produced by applying the realisation theorem (Theorem 2.7.9) in $\mathfrak{S}$ to the tuple (b_Q). We claim y is a partial realisation point for $p_1, \ldots, p_n$ in $\mathfrak{T}$. Since $C(D_{W_i}^{\mathcal{V}_i})$ is a single point, y satisfies the first requirement of the partial realisation axiom in $\mathfrak{T}$ for $p_{k+1}, \ldots, p_n$. For $i \leqslant k$, T_i is either nested into an element of $\mathcal{V}_{m+1} \cup \cdots \cup \mathcal{V}_n$ or orthogonal to all $W_{k+1}, \ldots, W_n$. This implies T_i is either nested into an element of $\mathcal{V}$ or orthogonal to all elements of $\mathcal{V}$. In both cases $b_{T_i} = \pi_{T_i}(x)$, and we have that $\pi_{T_i}(y)$ is uniformly close to $\pi_{T_i}(x)$, which is in turn E–close to p_i by item (1).

Now, let $Q \in \mathfrak{S}$ with $Q \pitchfork T_i$ or $T_i \subsetneq Q$ for some $i \in \{1, \ldots, n\}$. We verify $\mathsf{d}_Q(y, \rho_Q^{T_i})$ is uniformly bounded when $i \leqslant k$ and $i > k$ separately.

Assume $i \leqslant k$, so that $T_i \in \mathfrak{S}$. If $i \leqslant k$ and $b_Q = \pi_Q(x)$, then $\mathsf{d}_Q(y, \rho_Q^{T_i})$ is bounded by item (3). If $i \leqslant k$ and $b_Q \neq \pi_Q(x)$, then $b_Q \in \rho_Q^V$ for some $V \in \mathcal{V}$ and T_i is either orthogonal to or nested into V. If $T_i \perp V$ then $\mathsf{d}_Q(b_Q, \rho_Q^{T_i}) \leqslant 3E$ by Lemma 2.7.5. If $T_i \subsetneq V$ then $\mathsf{d}_Q(b_Q, \rho_Q^{T_i}) \leqslant 2E$ by consistency. The result then follows from the triangle inequality since $\pi_Q(y)$ is uniformly close to b_Q.

Now assume $i > k$, so that $T_i = D_{W_i}^{\mathcal{V}_i} \in \mathfrak{D}$. If $D_{W_i}^{\mathcal{V}_i} \subsetneq Q$, then $\rho_Q^V \subseteq \rho_Q^{D_{W_i}^{\mathcal{V}_i}}$ for all $V \in \mathcal{V}_i$. Since b_Q is within $2E$ of any ρ_Q^V for $V \in \mathcal{V}_i$, this bounds $\mathsf{d}_Q(y, \rho_Q^{D_{W_i}^{\mathcal{V}_i}})$ uniformly. On the other hand, if $D_{W_i}^{\mathcal{V}_i} \pitchfork Q$, then either $Q \perp V$ for all $V \in \mathcal{V}_i$ or there exists $V \in \mathcal{V}_i$ so that $V \pitchfork Q$ or $V \subsetneq Q$. In the latter case, $\rho_Q^V \subseteq \rho_Q^{D_{W_i}^{\mathcal{V}_i}}$ and we are finished since b_Q is within $2E$ of ρ_Q^V, giving a uniform bound on the distance from $\pi_Q(y)$ to $\rho_Q^{D_{W_i}^{\mathcal{V}_i}}$. In the former case, we must have $W_i \pitchfork Q$ and $\rho_Q^{D_{W_i}^{\mathcal{V}_i}}$ is equal to $\rho_Q^{W_i}$. If $b_Q = \pi_Q(x)$ than we are done by item (4). Otherwise, there exists $V' \in \mathcal{V} \smallsetminus \mathcal{V}_i$ so that $V' \pitchfork Q$ or $V' \subsetneq Q$ and $b_Q \in \rho_Q^{V'}$. Since $V' \perp W_i$, it follows that $\rho_Q^{V'}$ is within $2E$ of $\rho_Q^{W_i}$. Thus b_Q, and hence $\pi_Q(y)$, is uniformly close to $\rho_Q^{W_i} = \rho_Q^{D_{W_i}^{\mathcal{V}_i}}$. This completes the proof of Theorem 3.0.1. □

Remark 3.0.6. (Almost HHGs are HHGs) If G is a group and $\mathfrak{S}$ is an almost HHS structure for the Cayley graph of G, then we say $\mathfrak{S}$ is an *almost HHG structure* for G if it satisfies items (2) and (3) of the definition of an HHG. The above proof shows that if $(G, \mathfrak{S})$ is an almost HHG, then the structure $\mathfrak{T}$ from Theorem 3.0.1 is an HHG structure for G.

We conclude this section with a noteworthy application of Theorem 3.0.1. In [ABD21], Abbott, Behrstock, and Durham sought to show every hierarchically hyperbolic group admits an HHG structure with the following property.

Definition 3.0.7 (Unbounded products)**.** We say an (almost) hierarchically hyperbolic group $(G, \mathfrak{S})$ has *unbounded products* if there exists $B \geqslant 0$ so that one of the following holds for each $V \in \mathfrak{S}$.

(1) V is the unique $\sqsubseteq$–maximal element of $\mathfrak{S}$.

(2) $\operatorname{diam}(C(W)) \leqslant B$ whenever $W \sqsubseteq V$.

(3) There exists $W \in \mathfrak{S}$ so that $V \perp W$ and $\operatorname{diam}(C(W)) = \infty$.

In [ABD21], it was originally shown every HHG admits an *almost* HHG structure with unbounded products, and moreover you can verify that this structure satisfies the container axiom if the original HHG satisfies an additional hypothesis called *clean containers.* With Theorem 3.0.1, we were able to tie off this loose end in the theory of hierarchically hyperbolic groups and show all HHGs admit a structure with unbounded products.

Corollary 3.0.8. *If $(G, \mathfrak{S})$ is a hierarchically hyperbolic group, then there exists an HHG structure $\mathfrak{T}$ for G with unbounded products.*

Proof. In the proofs of [ABD21, Theorem 3.7, Corollary 3.8], it is shown that every HHG admits an almost HHG structure with unbounded products. Thus, G admits an almost HHG structure $\mathfrak{T}_0$ with unbounded products. Further, from the proof of [ABD21, Theorem 3.7], $\mathfrak{T}_0$ has the property that for every non–$\sqsubseteq$–maximal domain $V \in \mathfrak{T}_0$, there exist $W, Q \in \mathfrak{T}_0$ so that $W \sqsubseteq V$, $Q \perp V$ and $\operatorname{diam}(C(W)) = \operatorname{diam}(C(Q)) = \infty$. Let $\mathfrak{T}$ be the HHG structure obtained from $\mathfrak{T}_0$ using Theorem 3.0.1. We need only verify unbounded products for elements of $\mathfrak{T} \setminus \mathfrak{T}_0$. Using the notation of Theorem 3.0.1, let $D_W^{\mathcal{V}} \in \mathfrak{T} \setminus \mathfrak{T}_0$. By construction of $\mathfrak{T}_0$, there exists $R \sqsubseteq D_W^{\mathcal{V}}$ with $\operatorname{diam}(C(R)) = \infty$, so item (2) does not hold, and moreover $D_W^{\mathcal{V}}$ is not the $\sqsubseteq$–maximal element of $\mathfrak{T}$. However, $V \perp D_W^{\mathcal{V}}$ for all $V \in \mathcal{V}$, and by construction of $\mathfrak{T}_0$, there exists $T \in \mathfrak{T}_0$ so that $T \sqsubseteq V$ and $\operatorname{diam}(C(T)) = \infty$. Since $T \perp D_W^{\mathcal{V}}$, item (3) holds. □

Chapter 4

Hierarchical hyperbolicity of graph products

In this chapter we show that all graph products of finitely generated groups can be endowed with a relative HHG structure (Theorem 4.2.22), generalising results of Behrstock–Hagen–Sisto for right-angled Artin groups [BHS17b]. We build the proto-hierarchy structure for a graph product in Section 4.1 and spend Section 4.2 verifying this structure satisfies the axioms of a relative HHS and respects the group structure. We also show that any graph product has a (non-relative) HHS structure with respect to the syllable metric (Theorem 4.2.25), answering a question of Behrstock–Hagen–Sisto. Furthermore, in the particular case where all the vertex groups are themselves HHGs, the graph product can be endowed with a (non-relative) HHG structure with respect to the word metric (Theorem 4.3.1). This answers a second question of Behrstock–Hagen–Sisto.

In Section 4.3, we give some applications of our theorems. We give a new proof of a theorem of Meier, classifying when a graph product of hyperbolic groups is itself hyperbolic (Theorem 4.3.6). We also answer two questions of Genevois regarding a quasi-isometry invariant called the *electrification* of a graph product of finite groups (Theorems 4.3.10 and 4.3.12).

4.1 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 4.1.1 and then define the relations and relative projections in Section 4.1.2.

4.1.1 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 [BHS17b].

Definition 4.1.1 (Parallelism and the index set for a graph product)**.** Let G_Γ be a graph product. For an induced subgraph $\Lambda \subseteq \Gamma$, we shall use $g\Lambda$ to denote the coset $g\langle\Lambda\rangle$ for ease of notation. We say $g\Lambda$ and $h\Lambda$ are *parallel* if $g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$ and write $g\Lambda \parallel h\Lambda$. Let $[g\Lambda]$ denote the equivalence class of $g\Lambda$ under the parallelism relation $\parallel$. 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 $S(\Gamma)$. Recall that these hyperplanes, introduced in Definition 2.6.17, are generalisations of those in cube complexes.

Proposition 4.1.2 (Parallel cosets have the same hyperplanes)**.** *Let $\Lambda \subseteq \Gamma$ and $g, h \in G_\Gamma$. If $g\langle\Lambda\rangle \parallel h\langle\Lambda\rangle$, then every hyperplane of $S(\Gamma)$ crossing $g\langle\Lambda\rangle$ must also cross $h\langle\Lambda\rangle$.*

Proof. Since $g\langle\Lambda\rangle \parallel h\langle\Lambda\rangle$, $g^{-1}h \in \langle \mathrm{st}(\Lambda)\rangle$ and there exists $\lambda \in \langle\Lambda\rangle$ and $l \in \langle \mathrm{lk}(\Lambda)\rangle$ such that $g^{-1}h = \lambda l$. 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 4.1.3 (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 $\mathsf{d}_\Gamma(g,h)$ 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 $\mathrm{supp}(\lambda_i)$ a proper subgraph of Γ for each $i \in \{1, \dots, n\}$.

If $g = \lambda_1 \dots \lambda_n$ where $\mathrm{supp}(\lambda_i)$ 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 4.1.4. When Γ has at least 2 vertices, $S(\Gamma)$ is obtained from $\mathrm{Cay}(G_\Gamma, S)$ 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 $\mathsf{d}_\Gamma \leqslant \mathsf{d}_{syl} \leqslant \mathsf{d}$, where d is the word metric on G_Γ induced by S.

In a reduced subgraph expression $g = \lambda_1 \ldots \lambda_n$ we may assume $\mathrm{suffix}_{\Lambda_{i+1}}(\lambda_1 \ldots \lambda_i) = e$ for each $i \in \{1, \ldots, n-1\}$ by removing any non-trivial suffix from the end of $\lambda_1 \ldots \lambda_i$ 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 4.1.5. *If Γ contains at least 2 vertices, then for each $g \in G_\Gamma$, there exist $\lambda_1, \ldots, \lambda_n \in G_\Gamma$ with $\mathrm{supp}(\lambda_i) = \Lambda_i \subsetneq \Gamma$ such that the following hold.*

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

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

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

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

Proof. We begin by noting how the final conclusion of the lemma follows from the main conclusion. Let $\lambda_1 \ldots \lambda_n$ be a reduced subgraph expression for $x^{-1}y$ that satisfies (3). For each $i \in \{1, \ldots, n\}$, let $s_1^i \ldots s_{m_i}^i$ be a reduced syllable expression for λ_i. Since $|x^{-1}y|_{syl} = |\lambda_1 \ldots \lambda_n|_{syl} = \sum_{j=1}^{n} |\lambda_j|_{syl}$, it follows that $(s_1^1 \ldots s_{m_1}^1) \ldots (s_1^n \ldots s_{m_n}^n)$ 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 $(s_1^1 \ldots s_{m_1}^1) \ldots (s_1^n \ldots s_{m_n}^n)$, and this implies the element $\lambda_1 \ldots \lambda_i$ appears as a vertex of η for each $i \in \{1, \ldots, n\}$. Translating by x gives $\gamma = x\eta$ as the desired geodesic.

We now prove 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 $\text{supp}(g)$ is a proper subgraph of Γ and the conclusion is trivially true.

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

Let $\sigma = \text{suffix}_{\Omega_n}(\omega_1 \dots \omega_{n-1})$. 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, $(s_1^1 \dots s_{m_1}^1) \dots (s_1^n \dots s_{m_{n-1}}^{n-1})$ is a reduced syllable expression for $\omega_1 \dots \omega_{n-1}$ as $|\omega_1 \dots \omega_{n-1}|_{syl} = \sum_{j=1}^{n-1} |\omega_j|_{syl}$. 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 $\text{suffix}_{\Omega_n}(\omega_1' \dots \omega_{n-1}') = e$.

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

Now let $\lambda_n = \omega_n'$ and $\Lambda_i = \text{supp}(\lambda_i)$ 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 $\Lambda_i = \text{supp}(\lambda_i)$ 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 $\text{suffix}_{\Lambda_{i+1}}(\lambda_1 \dots \lambda_i) = e$.

(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 $|\lambda_1 \dots \lambda_{n-1}|_{syl} = \sum_{j=1}^{n-1} |\lambda_j|_{syl}$, we know $(t_1^1 \dots t_{k_1}^1) \dots (t_1^{n-1} \dots t_{k_{n-1}}^{n-1})$ is a reduced syllable expression for $\lambda_1 \dots \lambda_{n-1}$. Thus, 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 2.6.6 implies 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 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 $|\lambda_1 \dots \lambda_n|_{syl} = |\lambda_1|_{syl} + \dots + |\lambda_n|_{syl}$ 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 4.1.6. 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 $\mathsf{d}_{g\Lambda}(\cdot,\cdot)$.

Remark 4.1.7. If $\Lambda \subseteq \Gamma$ is a join $\Lambda = \Lambda_1 \star \Lambda_2$, 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 $C(\Lambda)$, 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 2.6.22 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 4.1.8 (Gates to parallelism classes are well defined)**.** *If $g\Lambda \parallel h\Lambda$, 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 $g\Lambda \parallel h\Lambda$, 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 2.6.22, 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 4.1.2. As this is a contradiction, we must have that $\mathfrak{g}_{h\Lambda}(x) = \mathfrak{g}_{h\Lambda}(\mathfrak{g}_{g\Lambda}(x))$.

Note, if $g\lambda \in g\langle\Lambda\rangle$, then equivariance (Proposition 2.6.22(2)) plus the prefix description of the gate map (Lemma 2.6.24) imply

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

Since $h^{-1}g \in \langle \mathrm{st}(\Lambda)\rangle$, we can write $h^{-1}g = pl$, where $p \in \langle\Lambda\rangle$ and $l \in \langle \mathrm{lk}(\Lambda)\rangle$. Therefore $\mathfrak{g}_{h\Lambda}(g\lambda) = h \cdot \mathrm{prefix}_{\Lambda}(pl\lambda) = hp\lambda$, 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 4.1.9. 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 4.1.10. Combining the prefix description of the gate map (Lemma 2.6.24) with equivariance (Proposition 2.6.22.(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 $\pi_{g\Lambda}(x) = g \cdot \mathrm{prefix}_{\Lambda}(g^{-1}x)$ 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 $\mathrm{suffix}_{\Lambda}(h) = \lambda$. 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 4.1.11. *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* $\text{suffix}_\Lambda(g) = e$ *and $h \in G_\Gamma$, then every syllable of* $\text{prefix}_\Lambda(g^{-1}h) = \mathfrak{g}_\Lambda(g^{-1}h)$ *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 $s_1, \dots, s_n$ 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 $s_i \in G_{v_i}$.

By Proposition 2.6.20(4) and Proposition 2.6.22(5), since each H_i separates $\mathfrak{g}_{g\Lambda}(x)$ and $\mathfrak{g}_{g\Lambda}(y)$, each H_i must also cross γ. For $i \in \{1, \dots, n\}$, let E_i be the edge of γ dual to H_i. Note, 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 2.6.20(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 4.1. Without

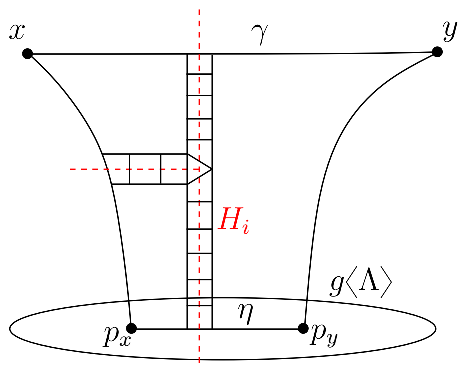

Figure 4.1: If the hyperplane H_i encounters a triangle of $S(\Gamma)$ between η and γ, then a branch of H_i must cross an $S(\Gamma)$–geodesic from x to p_x (shown) or from y to p_y.

loss of generality, assume H_i crosses an $S(\Gamma)$–geodesic connecting x and $p_x = \mathfrak{g}_{g\Lambda}(x)$. This means 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

2.6.20(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 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 $\mathfrak{g}_\Lambda(g^{-1}) = \mathrm{prefix}_\Lambda(g^{-1}) = e$. Thus, we can apply the above with $x = g^{-1}$ and $y = g^{-1}h$ to conclude 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 $h, k \in G_\Gamma$, 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 4.1.12 (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 $\mathsf{d}_{g\Lambda}(x,y) \leqslant \mathsf{d}(x,y)$ for all $x, y \in G_\Gamma$. First assume 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 $|p_x^{-1}p_y|$ 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 4.1.11, s must be a syllable of $x^{-1}y$, that is, s appears in a reduced syllable expression for $x^{-1}y$. Recall, 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 2.6.7). Thus $|x^{-1}y| \geqslant |s| = |p_x^{-1}p_y|$.

Now assume Λ contains at least 2 vertices. By Proposition 2.6.22(1), we have

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

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

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

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 4.1.13 (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*

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

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 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 2.6.22(5). However, this is impossible as H cannot cross γ twice (Proposition 2.6.20(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 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 4.1.5, so that 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 2.6.22(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 2.6.20(4)), it follows that any hyperplane that crosses α (resp. β) cannot cross β (resp. α). By Remark 2.6.21, 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 $\mathsf{d}_{g\Lambda}(x, q) \leqslant \mathsf{d}_{g\Lambda}(x, y)$. The proof for $\mathsf{d}_{g\Lambda}(q, y) \leqslant \mathsf{d}_{g\Lambda}(x, y)$ is nearly identical with β replacing α. Let $E_1, \ldots, E_k$ be the edges of α and let H_j be the hyperplane that crosses E_j for $j \in \{1, \ldots, k\}$. 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 4.2.

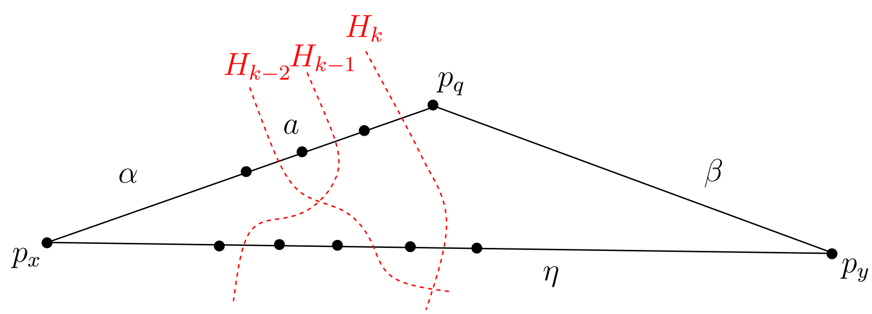

Figure 4.2: The hyperplanes H_{k-2} and H_{k-1} cross between α and η because the vertex a is separated from every vertex of η by either H_{k-2} or H_{k-1}. Even though H_{k-2} and H_k cross, they do not cross *between* α and η.

Claim 4.1.14. There exists an $S(\Gamma)$–geodesic α' that connects p_x and p_q such that no two of $H_1, \ldots, H_k$ cross between α' and η.

Proof. Let $\alpha_1 = \alpha$ and let K_i be the number of times two of $H_1, \ldots, H_k$ cross between α_i and η. Note, $K_1 \leqslant \frac{k(k-1)}{2}$. If $K_1 = 0$ we are done. Otherwise, there exists $j \in \{1, \ldots, k\}$ 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 2.6.20(5) tells us the edges E_{j-1} and E_j are labelled by elements of adjacent vertex groups. By Proposition 2.6.14, 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 4.3.

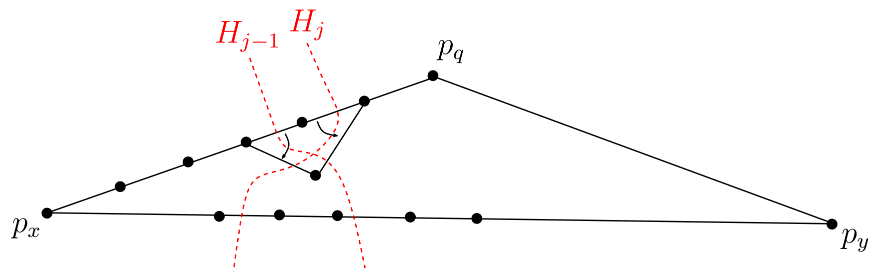

Figure 4.3: The edges E_{j-1} and E_j can be replaced with the other two edges of the square S to obtain a new $S(\Gamma)$–geodesic with $K_2 = K_1 - 1$.

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 $H_1, \ldots, H_k$ cross between α_2 and η is one less than the number of times two of $H_1, \ldots, H_k$ crossed between α_1 and η. Reindex $H_1, \ldots, H_k$ such that H_j crosses the jth edge of α_2.

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

Let α' be as in Claim 4.1.14 and reindex $H_1, \ldots, H_k$ so that H_j crosses the jth edge of α' for each $j \in \{1, \ldots, k\}$. Since H_j crosses η for each $j \in \{1, \ldots, k\}$, the labels for the edges of α' are a subset of the labels of η. Further, since no two of $H_1, \ldots, H_k$ 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 $p_x\lambda_1 \ldots \lambda_i$ for each $i \in \{1, \ldots, m\}$, this implies that we can write $p_x^{-1}p_q = \lambda'_1 \ldots \lambda'_m$, where $\mathrm{supp}(\lambda'_i) \subseteq \mathrm{supp}(\lambda_i)$ for each $i \in \{1, \ldots, m\}$. 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 $\mathsf{d}_{g\Lambda}(x, q) \leqslant \mathsf{d}_{g\Lambda}(x, y)$. □

4.1.2 The relations

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

Definition 4.1.15 (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 4.1.1. We say $[g\Lambda] \sqsubseteq [h\Omega]$ if $\Lambda \subseteq \Omega$ and there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\Omega] = [h\Omega]$.

Lemma 4.1.16. *The relation* $\sqsubseteq$ *is a partial order.*

Proof. The only property that requires checking is transitivity, that is, if $[g_1\Lambda_1] \sqsubseteq [g_2\Lambda_2]$ and $[g_2\Lambda_2] \sqsubseteq [g_3\Lambda_3]$, 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 $[g_1\Lambda_1] = [a\Lambda_1]$, $[a\Lambda_2] = [g_2\Lambda_2] = [b\Lambda_3]$, $[g_3\Lambda_3] = [b\Lambda_3]$, that is, $g_1^{-1}a \in \langle \mathrm{st}(\Lambda_1)\rangle$, $g_2^{-1}a, g_2^{-1}b \in \langle \mathrm{st}(\Lambda_2)\rangle$, $g_3^{-1}b \in \langle \mathrm{st}(\Lambda_3)\rangle$. Thus $g_1^{-1}a = l_1\lambda_1$, $g_2^{-1}a = l_2\lambda_2$, $g_2^{-1}b = l_2'\lambda_2'$, $g_3^{-1}b = l_3\lambda_3$ 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 $\mathrm{lk}(\Lambda_2) \subseteq \mathrm{lk}(\Lambda_1)$,

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

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 4.1.17 (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

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

The equality between $\bigcup_{k\Lambda \| h\Lambda} \pi_{h\Omega}\big(k\langle\Lambda\rangle\big)$ and $\pi_{h\Omega}\left(g\langle \mathrm{st}(\Lambda)\rangle\right)$ is a consequence of the definition that $k\Lambda \parallel g\Lambda$ 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 $k\Lambda \parallel g\Lambda$. 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 $g\Lambda \parallel gl\Lambda$.

Lemma 4.1.18 (Upwards relative projections have bounded diameter)**.** *If $[g\Lambda] \subsetneq [h\Omega]$, then for any choice of representatives $g\Lambda \in [g\Lambda]$ and $h\Omega \in [h\Omega]$, we have* $\mathrm{diam}\left(\rho^{g\Lambda}_{h\Omega}\right) \leqslant 2$.

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 4.1.7, and hence $\mathrm{diam}\left(\rho^{g\Lambda}_{h\Omega}\right) \leqslant 2$. 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 $\mathrm{st}(\Lambda) \cap \Omega = \Omega$. 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 equivariance (Proposition 2.6.22(2)) and the prefix description of the gate map (Lemma 2.6.24), 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 $\mathfrak{g}_{a\Omega}(a\lambda l) = a\lambda l_0$, where $l_0 = \mathrm{prefix}_{\Omega}(l) \in \langle \mathrm{lk}(\Lambda) \cap \Omega\rangle$ and so $\mathrm{supp}(\lambda l_0) \subseteq \Lambda \cup (\mathrm{lk}(\Lambda) \cap \Omega) = \mathrm{st}(\Lambda) \cap \Omega \subsetneq \Omega$. Moreover, by Lemma 4.1.8, $\mathfrak{g}_{h\Omega}(a\lambda l) = \mathfrak{g}_{h\Omega}(\mathfrak{g}_{a\Omega}(a\lambda l)) = \mathfrak{g}_{h\Omega}(a\lambda l_0)$.

Since $a\Omega \parallel h\Omega$, 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 $p = \mathrm{prefix}_{\Lambda}(h^{-1}a)$ (Lemma 4.1.8). Since $\mathrm{supp}(\lambda l_0) \subsetneq \Omega$, this implies $\mathfrak{g}_{h\Omega}(a\lambda l_0) = hpa^{-1} \cdot a\lambda l_0 = hp\lambda l_0$. Therefore, given two arbitrary elements

$a\lambda l, a\lambda' l' \in a\langle \mathrm{st}(\Lambda)\rangle$, we have $(\mathfrak{g}_{h\Omega}(a\lambda l))^{-1}\mathfrak{g}_{h\Omega}(a\lambda' l') = l_0^{-1}\lambda^{-1}\lambda' l_0'$, where $\mathrm{supp}(l_0^{-1}\lambda^{-1}\lambda' l_0') \subseteq \mathrm{st}(\Lambda) \cap \Omega \subsetneq \Omega$. This implies the $C(h\Omega)$–diameter of $\pi_{h\Omega}(g\langle \mathrm{st}(\Lambda)\rangle) = \rho_{h\Omega}^{g\Lambda}$ is at most 1 in this case. □

Next we deal with the orthogonality relation.

Definition 4.1.19 (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 4.1.1. We say $[g\Lambda] \perp [h\Omega]$ 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 4.1.20 (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 $[g\Lambda] \sqsubseteq [h\Omega]$ and $[h\Omega]\perp[k\Pi]$. Then $\Lambda \subseteq \Omega \subseteq \mathrm{lk}(\Pi)$, and there exist $a, b \in G_\Gamma$ such that $[a\Lambda] = [g\Lambda]$, $[a\Omega] = [h\Omega] = [b\Omega]$ and $[b\Pi] = [k\Pi]$. 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 4.1.21 (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 $[g\Lambda] \pitchfork [h\Omega]$. 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

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

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

Lemma 4.1.22. *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* $\mathrm{diam}\big(\pi_{g\Lambda}(h\langle \mathrm{st}(\Omega)\rangle)\big) \leqslant 2$ *and* $\mathrm{diam}\big(\pi_{h\Omega}(g\langle \mathrm{st}(\Lambda)\rangle)\big) \leqslant 2$.

Proof. We provide the proof for $\mathrm{diam}\big(\pi_{g\Lambda}(h\langle \mathrm{st}(\Omega)\rangle)\big) \leqslant 2$. 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 $\Lambda_1 \star \Lambda_2$, then $\mathsf{d}_{g\Lambda}(p_x, p_y) \leqslant \mathrm{diam}(C(g\Lambda)) \leqslant 2$ by Remark 4.1.7.

Now suppose 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 $\mathrm{supp}(p_x^{-1}p_y) \neq \Lambda$.

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

First assume $\Lambda \nsubseteq \mathrm{lk}(\Omega)$ and $\Lambda \nsubseteq \Omega$. Since Λ does not split as a join, if $\Lambda = \mathrm{st}(\Omega) \cap \Lambda$, then Λ would need to be a subgraph of either Ω or $\mathrm{lk}(\Omega)$. As this is impossible in this case, we must have that $\mathrm{st}(\Omega) \cap \Lambda \neq \Lambda$. By Proposition 4.1.11, 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 $\mathrm{supp}(p_x^{-1}p_y) \subseteq \mathrm{st}(\Omega) \cap \Lambda \neq \Lambda$ as desired.

Now assume $\Lambda \subseteq \mathrm{lk}(\Omega)$ 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 $\mathrm{supp}(p_x^{-1}p_y) = \Lambda$.

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 4.5.

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, a hyperplane in $S(\Gamma)$ crosses a geodesic segment if and only if it separates the end points of the segment (Proposition 2.6.20(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 2.6.22(4). In particular, no H_i crosses η. Thus, by Remark 2.6.21, 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 4.4. 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 2.6.20(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, $t_i \in \langle \mathrm{lk}(\Lambda)\rangle$.

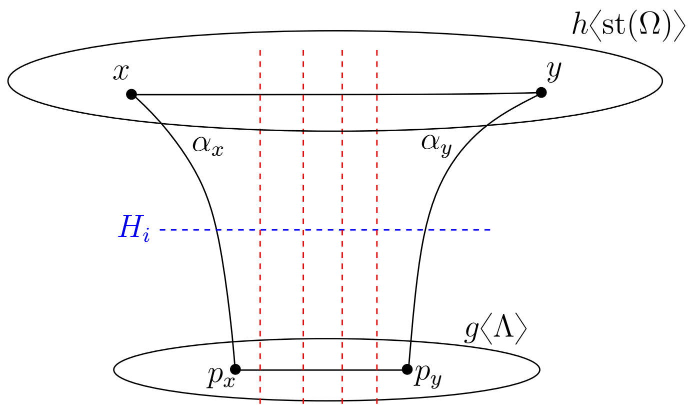

Figure 4.4: Any hyperplane that crosses α_x and α_y must cross all of the hyperplanes separating p_x and p_y.

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 $i < j$ and that H_i crosses γ, but H_j crosses α_y. As shown in Figure 4.5, this forces H_i to cross H_j, which implies that t_i and t_j commute by Proposition 2.6.20(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 $l_x \in \langle \mathrm{lk}(\Lambda) \rangle$.

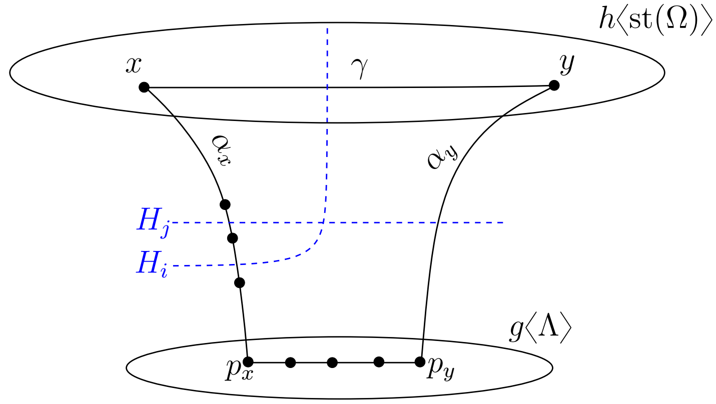

Figure 4.5: The hyperplane H_i crosses α_x and γ while H_j crosses α_x and α_y. Since H_i appears before H_j along α_x, H_i must cross H_j.

Now, since $x \in h\langle \mathrm{st}(\Omega) \rangle$, we have $h^{-1}x \in \langle \mathrm{st}(\Omega) \rangle$, which implies $[h\Omega] = [x\Omega]$. 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 $[g\Lambda] = [p_x\Lambda]$. 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

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

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 we must have $\mathrm{supp}(p_x^{-1}p_y) \neq \Lambda$ as desired. $\square$

4.1.3 The proto-hierarchy structure

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

Theorem 4.1.23. *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 4.1.1.*
- *The space $C([g\Lambda])$ associated to $[g\Lambda]$ is the space $C(g\Lambda)$ from Definition 4.1.3, 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 4.1.9 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 4.1.17, 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 4.1.21, 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 4.1.12. Nesting is shown to be a partial order in Lemma 4.1.16. The upward relative projection has diameter at most 2 by Lemma 4.1.18. Lemma 4.1.20 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 4.1.22. □

4.2 Graph products are relative HHGs

In this section, we complete our proof that graph products of finitely generated groups are relative HHGs (Theorem 4.2.22) by proving the eight remaining HHS axioms and showing that the group structure is compatible with our hierarchy structure. In Section 4.2.1, we prove hyperbolicity of $C(g\Lambda)$ whenever Λ contains at least two vertices. Section 4.2.2 is devoted to proving the finite complexity and containers axioms. Section 4.2.3 deals with the uniqueness axiom, and in Section 4.2.4, we the prove the bounded geodesic image and large links axioms. In Section 4.2.5, we verify partial realisation, and Section 4.2.6 deals with the consistency axiom. Finally, in Section 4.2.7, compatibility of the relative HHS structure with the group structure is checked.

We also obtain some auxiliary results along the way: in Section 4.2.1, we show that not only are the spaces $C(g\Lambda)$ hyperbolic whenever Λ contains at least 2 vertices, but they are also quasi-trees; and in Section 4.2.3, 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 a hierarchically hyperbolic space. This is true even when the vertex groups are not finitely generated. However, until then we will continue to assume 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.

4.2.1 Hyperbolicity

Lemma 4.2.1 (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 4.2.2. The hyperbolicity of $C(g\Lambda)$ can also be deduced from [Gen18, 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 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 $p_1^i, \ldots, p_{m_i}^i$ 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_j^i and p_{j+1}^i with an $S(g\Lambda)$–geodesic α_j^i. Since α_j^i is labelled by vertices of $\mathrm{supp}((p_j^i)^{-1}p_{j+1}^i)$, which is a proper subgraph of Λ, the $C(g\Lambda)$–distance between any vertex of α_j^i and p_j^i or p_{j+1}^i is at most 1. It therefore suffices to show that given any vertex p_j^1 of γ_1, either α_{j-1}^1 or α_j^1 is $C(g\Lambda)$–distance 1 from some α_t^i with $i = 2$ or 3. See Figure 4.6.

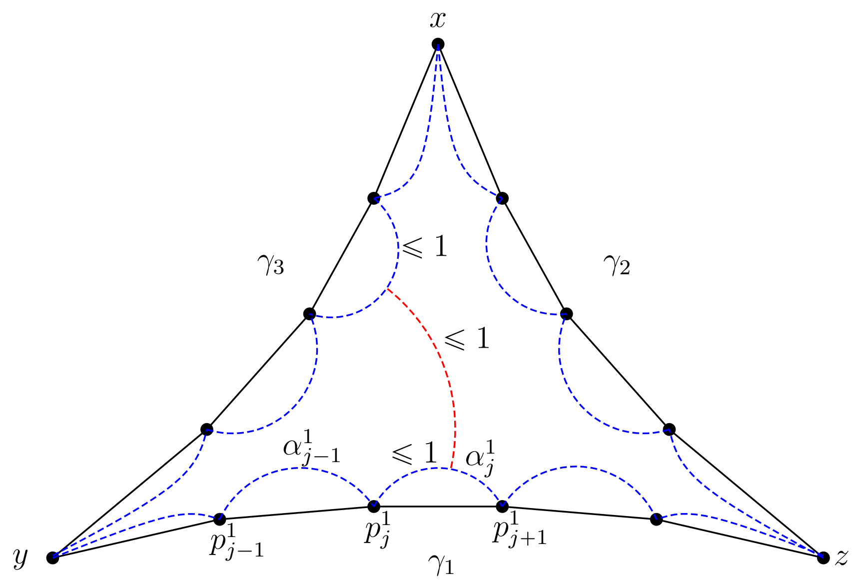

Figure 4.6: For each edge of the $C(g\Lambda)$–geodesic triangle, we construct an $S(g\Lambda)$–geodesic segment α_j^i between its endpoints (shown in blue). To show the triangle is $\frac{7}{2}$–slim, it then suffices to show that for each j, $\alpha_{j-1}^1 \cup \alpha_j^1$ is $C(g\Lambda)$–distance 1 from some α_t^i with $i \neq 1$.

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 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 4.1.7 and hence is clearly $\frac{7}{2}$–hyperbolic. Take a vertex p_j^1 of γ_1. If p_j^1 is one of the first or last

4 vertices of γ_1, then it is at distance at most 3 from γ_2 or γ_3. 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, 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 $\mathrm{lk}(w) \cap \Lambda \neq \varnothing$, 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.

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 2.6.21). Further, Proposition 2.6.20(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 4.7. 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 $|k - r| \leqslant 2$. As we know that H_0 crosses $\alpha^1_{j-1} \cup \alpha^1_j$, this implies 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 2.6.20(4)), we may therefore suppose without loss of generality that $r = j$ and $j < k \leqslant j+2$ (the cases where $j-2 \leqslant k < j$ or $r = j-1$ proceed similarly).

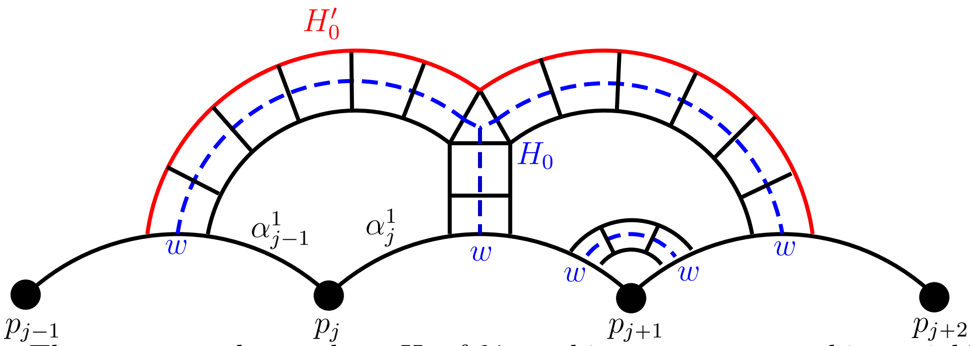

Figure 4.7: The outermost hyperplane H_0 of $\mathcal{H}_w$ and its outermost combinatorial hyperplane H_0'.

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 $\mathsf{d}_{g\Lambda}(p^1_j, p^1_{k+1}) = 1$ 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 $\mathsf{d}_{g\Lambda}(p^1_l, p^1_{k+1}) \leqslant 2$, 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.

If H_1 crosses α^1_s with $|s-j| \geqslant 3$, 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 $|s-k| \geqslant 3$, then similarly we obtain a contradiction. Assume therefore that $|s-j| \leqslant 2$ and

$|s-k| \leqslant 2$. 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 4.8.

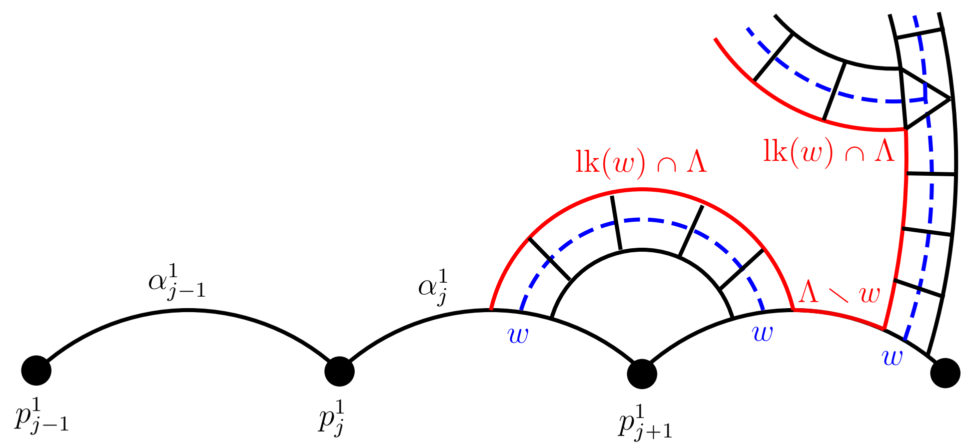

Figure 4.8: By following a sequence of combinatorial hyperplanes, we obtain a path labelled by $\Lambda \smallsetminus w$ (shown in red) that must eventually leave γ'_1 and cross $\gamma'_2 \cup \gamma'_3$.

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 $|t-s| \geqslant 3$, and the second case gives $|t-j| \geqslant 3$, 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 4.2.3 (Bottleneck criterion [Man05, Theorem 4.6])**.** *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 4.2.4. *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 4.2.4 proceeds similarly to the proof of Lemma 4.2.1, 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 $[g\Lambda]$ 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 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 4.2.1 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. $\square$

4.2.2 Finite complexity and containers

Lemma 4.2.5 (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 $|V(\Gamma)|$. $\square$

Lemma 4.2.6 (Containers)**.** *Let $[h\Omega] \subsetneq [g\Lambda]$ 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 $[k\Pi] \sqsubseteq [a(\mathrm{lk}(\Omega) \cap \Lambda)] \subsetneq [a\Lambda]$ 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 $\mathrm{supp}(k^{-1}b) \subseteq \mathrm{st}(\Pi)$, $\mathrm{supp}(b^{-1}a) \subseteq \mathrm{st}(\Omega)$, $\mathrm{supp}(k^{-1}c) \subseteq \mathrm{st}(\Pi)$, $\mathrm{supp}(c^{-1}a) \subseteq \mathrm{st}(\Lambda)$. Writing $p = \mathrm{prefix}_{\mathrm{st}(\Pi)}(k^{-1}a)$, we have $p^{-1}k^{-1}a = s$, where $\mathrm{prefix}_{\mathrm{st}(\Pi)}(s) = e$. 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 $p^{-1}k^{-1}a \in \langle \mathrm{st}(\Omega) \rangle$. 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, 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 $\mathrm{st}(\Omega) \cap \mathrm{st}(\Lambda) = \Omega \cup \mathrm{lk}(\Lambda) \cup (\mathrm{lk}(\Omega) \cap \Lambda)$. 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. □

4.2.3 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 4.1.

Lemma 4.2.7 (Uniqueness). *Let G_Γ be a graph product of finitely generated groups. There exists a function $\theta\colon [0,\infty) \to [0,\infty)$, depending only on the number of vertices of Γ, so that for all $g \in G_\Gamma$, if $\mathsf{d}_{h\Lambda}(e,g) \leqslant r$ for all $h \in G_\Gamma$ and subgraphs $\Lambda \subseteq \Gamma$, then $|g|_{G_\Gamma} \leqslant \theta(r)$.*

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

Since $\mathsf{d}_\Gamma(e,g) \leqslant r$, there exist proper subgraphs $\Lambda_i \subsetneq \Gamma$ and elements λ_i with $\mathrm{supp}(\lambda_i) = \Lambda_i$ so that $g = \lambda_1 \dots \lambda_m$ and $\mathsf{d}_\Gamma(e,g) = m \leqslant r$. We shall see that $\mathsf{d}_{h\Omega}(e,g) \leqslant r$ implies $\mathsf{d}_{h\Omega}(e,\lambda_i)$ 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 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 $\Gamma = \Lambda_1 \star \Lambda_2$, 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 $\mathfrak{g}_{h\Omega}(g) = h \cdot \mathrm{prefix}_\Omega(h^{-1}g) = h \cdot \mathrm{prefix}_\Omega(h^{-1}\lambda_i) = \mathfrak{g}_{h\Omega}(\lambda_i)$. Therefore

$\mathsf{d}_{h\Omega}(e, \lambda_i) = \mathsf{d}_{h\Omega}(e, g) \leqslant r$ and by induction there exists $D = D(n, r)$ so that $|\lambda_i|_{G_\Gamma} \leqslant D$ for $i = 1, 2$. Thus, $|g|_{G_\Gamma} \leqslant 2D$, which depends only on r and the number of vertices of Γ.

Suppose Γ does not split as a join, and define $p_0 = e$ and $p_i = \lambda_1 \cdots \lambda_i$ 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 4.1.5, we can assume that $\text{suffix}_{\Lambda_i}(p_{i-1}) = e$ 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 $\mathsf{d}_{h\Omega}(e, \lambda_i)$ 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 $\text{suffix}_\Omega(h) = e$. To achieve the bound on $\mathsf{d}_{h\Omega}(e, \lambda_i)$, we use the following two claims plus the assumption that $\mathsf{d}_{h\Omega}(e, g) \leqslant r$.

Claim 4.2.8. $\pi_{p_{i-1}h\Omega}(p_{i-1}) = \pi_{p_{i-1}h\Omega}(e)$.

Proof. By equivariance of the gate map and the prefix description of the gate map (Lemma 2.6.24),

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

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

Claim 4.2.9. $\mathsf{d}_{p_{i-1}h\Omega}(p_i, g) \leqslant r$.

Proof of Claim 4.2.9. Recall, 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 4.1.5). Thus, Lemma 4.1.13 says $\mathsf{d}_{p_{i-1}h\Omega}(p_i, g) \leqslant \mathsf{d}_{p_{i-1}h\Omega}(e, g)$, and $\mathsf{d}_{p_{i-1}h\Omega}(e, g) \leqslant r$ by assumption. $\square$

By the equivariance of the gate map (Proposition 2.6.22(2)), $\mathsf{d}_{h\Omega}(e,\lambda_i) = \mathsf{d}_{p_{i-1}h\Omega}(p_{i-1},p_i)$. Claim 4.2.8 then implies

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

Since $\mathsf{d}_{p_{i-1}h\Omega}(e,g) \leqslant r$ by assumption and $\mathsf{d}_{p_{i-1}h\Omega}(g,p_i) \leqslant r$ by Claim 4.2.9, we have $\mathsf{d}_{h\Omega}(e,\lambda_i) = \mathsf{d}_{p_{i-1}h\Omega}(p_{i-1},p_i) \leqslant 2r$ for each $h \in \langle \Lambda_i \rangle$ and $\Omega \subseteq \Lambda_i$. The induction hypothesis now implies there exists $D = D(n,r)$ such that the word length of λ_i in $\langle \Lambda_i \rangle$ is at most D. Since each graphical subgroup is convexly embedded in the word metric d on G_Γ, this implies $|g|_{G_\Gamma} \leqslant rD$, which depends only on r and the number of vertices of Γ. □

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

Theorem 4.2.10. *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, if Λ splits as a join, then $\operatorname{diam}(C(g\Lambda)) \leqslant 2$ by Remark 4.1.7. Suppose therefore that Λ does not split as a join and let $v_1, \dots, v_k$ be the vertices of Λ. For each $i \in \{1, \dots, k\}$, 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 4.1. Define $\lambda = s_1 \dots s_k$. For each $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, n\}$, let s_i^j be the jth copy of s_i in the product $(s_1 \dots s_k)^n = \lambda^n$, that is, $\lambda^n = (s_1^1 \dots s_k^1)(s_1^2 \dots s_k^2) \dots (s_1^n \dots s_k^n)$.

We claim that for each $n \in \mathbb{N}$, $(s_1^1 \dots s_k^1)(s_1^2 \dots s_k^2) \dots (s_1^n \dots s_k^n)$ is a reduced syllable expression for λ^n. Indeed, if $(s_1^1 \dots s_k^1)(s_1^2 \dots s_k^2) \dots (s_1^n \dots s_k^n)$ 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. However, if s_i^ℓ were to be combined with s_i^j, then s_i would need to commute with each of $s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_k$. 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 $(s_1^1 \dots s_k^1)(s_1^2 \dots s_k^2) \dots (s_1^n \dots s_k^n)$ is a reduced syllable expression for λ^n, and we have $|\lambda^n|_{syl} = kn$ for all $n \in \mathbb{N}$.

To prove $C(\Lambda)$ 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 4.2.11. For all $\Omega \subsetneq \Lambda$, $h \in \langle\Lambda\rangle$ and $n \geqslant 2$, $\mathsf{d}_{h\Omega}(e, \lambda^n) \leqslant 3$.

For now we accept Claim 4.2.11, deferring its proof until after we have proved $C(\Lambda)$ has infinite diameter.

For the purposes of contradiction, assume there exists $R > 0$ such that $\mathsf{d}_\Lambda(e, \lambda^n) \leqslant R$ for all $n \in \mathbb{N}$. By Claim 4.2.11, for every proper subgraph $\Omega \subsetneq \Lambda$ and $h \in \langle\Lambda\rangle$, we have $\mathsf{d}_{h\Omega}(e, \lambda^n) \leqslant 3$. Applying the uniqueness axiom (Lemma 4.2.7) to the graph product $\langle\Lambda\rangle = G_\Lambda$, this implies there exists $D = D(R, |V(\Lambda)|) > 0$ such that $|\lambda^n|_{G_\Lambda} = |\lambda^n|_{G_\Gamma} \leqslant D$ for all $n \in \mathbb{N}$. However, this is a contradiction as $|\lambda^n|_{G_\Gamma} \geqslant |\lambda^n|_{syl} = kn$ 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 4.2.11. 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 $\mathrm{suffix}_\Omega(h) = e$, and thus $\mathrm{prefix}_\Omega(h^{-1}) = e$.

Recall, $\pi_{h\Omega}(e) = h \cdot \mathrm{prefix}_\Omega(h^{-1})$ and $\pi_{h\Omega}(\lambda^n) = h \cdot \mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ (Remark 4.1.10). Since $\mathrm{prefix}_\Omega(h^{-1}) = e$, it suffices to prove that $\mathsf{d}_\Omega(e, h^{-1}\lambda^n) \leqslant 3$. We can also assume that $\mathrm{prefix}_\Omega(h^{-1}\lambda^n) \neq e$.

By Proposition 4.1.11, all syllables of $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ are syllables of λ^n. As $\mathrm{prefix}_\Omega(h^{-1}\lambda^n) \neq e$, there must exist $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, n\}$ such that s_i^j is the first syllable of $(s_1^1 \dots s_k^1)(s_1^2 \dots s_k^2) \dots (s_1^n \dots s_k^n)$ that is also a syllable of $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$.

Let $\ell, m \in \{1, \dots, k\}$ be such that $v_\ell \in \Lambda \smallsetminus \mathrm{st}(\Omega)$ and $v_m \in \Omega$ is not joined to v_ℓ by an edge. These vertices exist since Λ does not split as a join and thus $\Lambda \neq \mathrm{st}(\Omega)$. We will show that $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ can be written as a product $p_1p_2p_3$ 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 the $C(\Omega)$–distance between e and $\mathrm{prefix}_\Omega(h^{-1}\lambda^n)$ is at most 3, which in turn says $\mathsf{d}_{h\Omega}(e, \lambda^n) \leqslant 3$.

Suppose $i < \ell$. Since $v_\ell \notin \Omega$, every syllable of $\text{prefix}_\Omega(h^{-1}\lambda^n)$ must either be one of $s_i^j, s_{i+1}^j, \ldots, s_{\ell-1}^j$ 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 $\text{prefix}_\Omega(h^{-1}\lambda^n)$ for $J > j$. Therefore $\text{prefix}_\Omega(h^{-1}\lambda^n)$ can contain at most one syllable with support v_m, namely s_m^j. Thus $\text{prefix}_\Omega(h^{-1}\lambda^n) = p_1 p_2 p_3$ with $\text{supp}(p_1) \subseteq \Omega \smallsetminus v_m$, $\text{supp}(p_2) \subseteq v_m$, and $\text{supp}(p_3) \subseteq \Omega \smallsetminus v_m$. Note, if $\Omega = v_m$, then $\text{prefix}_\Omega(h^{-1}\lambda^n) = p_2 = s_m^j$ 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 because every syllable of $\text{prefix}_\Omega(h^{-1}\lambda^n)$ must either be one of $s_i^j, s_{i+1}^j, \ldots, s_k^j, s_1^{j+1}, \ldots, s_{\ell-1}^{j+1}$ or must commute with s_ℓ^{j+1}. □

In Section 4.3, we use our characterisation of when $C(g\Lambda)$ has infinite diameter to answer two questions of Genevois [Gen19b] (Theorems 4.3.10 and 4.3.12).

4.2.4 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 4.2.12 (Bounded geodesic image)**.** *Let $x, y \in G_\Gamma$ and $[h\Omega] \sqsubsetneq [g\Lambda]$. 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 $\rho_{g\Lambda}^{h\Omega}$.*

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 4.2.13. 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 $x \notin g\langle\Lambda\rangle$.

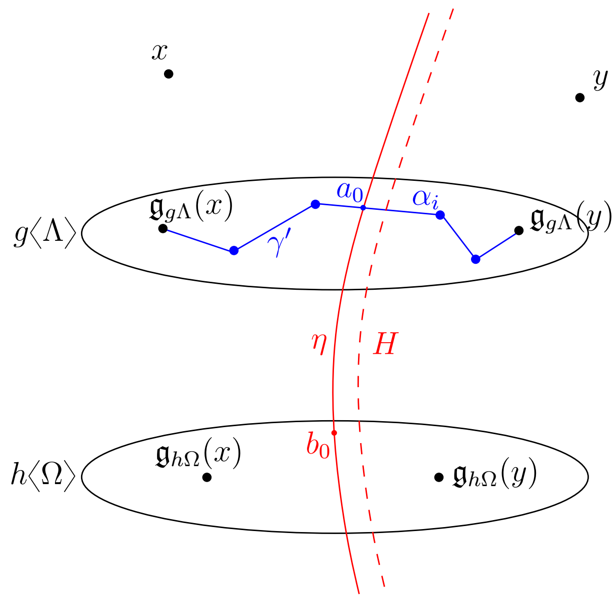

Figure 4.9: The $S(\Gamma)$–geodesic η connecting $b_0 \in h\langle\Omega\rangle$ and $a_0 \in \alpha_i$ when $d_{h\Omega}(x, y) > 0$.

Suppose $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x)) \neq \mathfrak{g}_{h\Omega}(x)$. Then there is a hyperplane H separating $\mathfrak{g}_{h\Omega}(\mathfrak{g}_{g\Lambda}(x))$ and $\mathfrak{g}_{h\Omega}(x)$. By Proposition 2.6.22, 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 4.1.2) H must also cross $k\langle\Omega\rangle \subseteq k\langle\Lambda\rangle$. But $k\Lambda \parallel g\Lambda$, so H must also cross $g\langle\Lambda\rangle$. This contradiction means 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 $p_1, \ldots, p_n \in \langle\Lambda\rangle$ so that $\pi_{g\Lambda}(x) = gp_1, gp_2, \ldots, gp_n = \pi_{g\Lambda}(y)$ are the vertices of γ. Let α_i be an $S(g\Lambda)$–geodesic from gp_i to gp_{i+1} for each $i \in \{1, \ldots, n-1\}$. Let γ' be the path in $S(g\Lambda)$ that is the union of all the α_i.

Suppose $\mathsf{d}_{h\Omega}(x, y) > 0$. Then $\mathsf{d}_{syl}(\mathfrak{g}_{h\Omega}(x), \mathfrak{g}_{h\Omega}(y)) > 0$ 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 2.6.22. 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 $\mathfrak{g}_{h\Omega}(x)$ nor $\mathfrak{g}_{g\Lambda}(y)$ from $\mathfrak{g}_{h\Omega}(y)$, 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 4.9.

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 4.1.11 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, $\mathsf{d}_{g\Lambda}(\alpha_i, \rho_{g\Lambda}^{h\Omega}) \leqslant 1$ as $a_1 \in \pi_{g\Lambda}(\alpha_i)$ and $b_1 \in \pi_{g\Lambda}(h\langle\Omega\rangle) \subseteq \rho_{g\Lambda}^{h\Omega}$. As α_i is labelled by a proper subgraph of Λ, any subsegment is also labelled by a proper subgraph, hence $\mathsf{d}_{g\Lambda}(ga, gp_{i+1}) \leqslant 1$ for any vertex ga of α_i. Thus, $\mathsf{d}_{g\Lambda}(ga, \gamma) \leqslant 1$ and therefore $\mathsf{d}_{g\Lambda}(\gamma, \rho_{g\Lambda}^{h\Omega}) \leqslant 2$. □

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

Lemma 4.2.14. *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 $\pi_{g\Lambda}\big(h\langle\Omega\rangle\big) \subseteq \rho_{g\Lambda}^{h\Omega}$, which is shown to have diameter at most 2 in Lemmas 4.1.18 and 4.1.22. 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 $\mathrm{suffix}_\Lambda(g) = e$ and $\mathrm{suffix}_\Omega(h) = e$. By Proposition 4.1.11, all syllables of $\mathrm{prefix}_\Lambda(g^{-1}h\omega)$ are syllables of $h\omega$. Further, since $\Lambda \subseteq \mathrm{lk}(\Omega)$, we have $\mathrm{supp}(\omega) \cap \Lambda = \varnothing$. As $\mathrm{suffix}_\Omega(h) = e$, this implies $\mathrm{prefix}_\Lambda(g^{-1}h\omega) = \mathrm{prefix}_\Lambda(g^{-1}h)$. Thus $\pi_{g\Lambda}(h\omega) = g \cdot \mathrm{prefix}_\Lambda(g^{-1}h)$ for all $\omega \in \langle\Omega\rangle$, and so $\pi_{g\Lambda}(h\langle\Omega\rangle)$ has diameter 0. □

Lemma 4.2.15 (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 $[h_1\Omega_1], \ldots, [h_n\Omega_n] \in \mathfrak{S}_\Gamma$ each nested into $[k\Pi]$ so that for any $[g\Lambda] \in \mathfrak{S}_\Gamma$ with $[g\Lambda] \subsetneq [k\Pi]$, if $\mathsf{d}_{g\Lambda}(x,y) > 18$ for some representative of $[g\Lambda]$, then $[g\Lambda] \sqsubseteq [h_i\Omega_i]$ for some $i \in \{1, \ldots, n\}$.*

Proof. Let γ be a $C(k\Pi)$–geodesic connecting $\pi_{k\Pi}(x)$ and $\pi_{k\Pi}(y)$, let $\pi_{k\Pi}(x) = p_0, p_1, \cdots, p_n = \pi_{k\Pi}(y)$ be the vertices of γ, and let $\lambda_i = p_{i-1}^{-1}p_i$ for each $i \in \{1, \ldots, n\}$. For $i \in \{1, \ldots, n\}$, define T_i to be $p_{i-1} \cdot \langle\mathrm{supp}(\lambda_i)\rangle$. Note that $p_i \in T_{i-1} \cap T_i$, and $T_i \subseteq k\langle\Pi\rangle$ since $p_{i-1} \in k\langle\Pi\rangle$

and $\operatorname{supp}(\lambda_i) \subsetneq \Pi$. In particular, $[T_i] \subsetneq [k\Pi]$. 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 $\operatorname{supp}(\lambda_i)$ is a proper subgraph of Π.

Next, let $[g\Lambda] \in \mathfrak{S}_\Gamma$ with $[g\Lambda] \subsetneq [k\Pi]$ and suppose $\mathsf{d}_{g\Lambda}(x, y) > 18$ for some representative $g\Lambda \in [g\Lambda]$. We shall show $[g\Lambda] \sqsubseteq [T_i]$ for some $i \in \{1, \ldots, n\}$. Since we have established the bounded geodesic image axiom (Lemma 4.2.12), we have $\gamma \cap N_2(\rho^{g\Lambda}_{k\Pi}) \neq \varnothing$. Let j be the first number in $\{0, \ldots, n\}$ so that $p_j \in N_4(\rho^{g\Lambda}_{k\Pi})$, and recall that each $\pi_{k\Pi}(T_i) = T_i$ is contained in $N_1(p_i)$ and $\operatorname{diam}(\rho^{g\Lambda}_{k\Pi}) \leqslant 2$ (Lemma 4.1.18). Therefore, if $1 \leqslant i \leqslant j$ or $i \geqslant j + 10$, then $\pi_{k\Pi}(T_i) \cap N_2(\rho^{g\Lambda}_{k\Pi}) = \varnothing$ and the bounded geodesic image axiom says $\pi_{g\Lambda}(T_i)$ is a single point.

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

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

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

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

Since $\mathsf{d}_{g\Lambda}(x, y) > 18$, there exists $j_0 \in \{j+1, \ldots, \min\{n, j+9\}\}$ so that $\operatorname{diam}\big(\pi_{g\Lambda}(T_{j_0})\big) > 2$. By Lemma 4.2.14, this implies $[g\Lambda] \sqsubseteq [T_{j_0}]$. $\square$

4.2.5 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 4.2.16. *Let $[g_1\Lambda_1], \ldots, [g_n\Lambda_n] \in \mathfrak{S}_\Gamma$. If either $[g_1\Lambda_1] \sqsubseteq \ldots \sqsubseteq [g_n\Lambda_n]$ or*

$[g_1\Lambda_1], \ldots, [g_n\Lambda_n]$ *are pairwise orthogonal, then there exists* $g \in G_\Gamma$ *so that* $[g\Lambda_i] = [g_i\Lambda_i]$ *for all* $i \in \{1, \ldots, n\}$.

Proof. We proceed by induction. The initial case $n = 2$ is true by definition. Suppose the statement is true for all $n < m$, and consider $n = m$, that is, we have $[g_1\Lambda_1], \ldots, [g_m\Lambda_m] \in \mathfrak{S}_\Gamma$ which are either pairwise orthogonal or nested. Then in particular $[g_1\Lambda_1], \ldots, [g_{m-1}\Lambda_{m-1}]$ 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 $g = e$ 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 $i \in \{1, \ldots, m-1\}$, we know $\mathrm{supp}(h) \subseteq \mathrm{supp}(k_i) \subseteq \mathrm{st}(\Lambda_i)$ 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 4.2.17 (Partial realisation)**.** *Let* $\{[g_i\Lambda_i]\}_{i=1}^n$ *be a finite collection of pairwise orthogonal elements of* $\mathfrak{S}_\Gamma$*. For each* $i \in \{1, \ldots, n\}$*, 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 4.2.16 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 $x = g\lambda_1\lambda_2 \ldots \lambda_n$. 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 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 4.1.8, we have $\mathsf{d}_{g_i\Lambda_i}(x, p_i) = 0$ for all i.

Now, suppose $[g\Lambda_i] \sqsubsetneq [h\Omega]$ or $[g\Lambda_i] \pitchfork [h\Omega]$ for some $i \in \{1, \ldots, n\}$ 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 $x = g\lambda_1 \ldots \lambda_n \in g\langle \mathrm{st}(\Lambda_i) \rangle$. 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 $g_i\Lambda_i \parallel g\Lambda_i$. This implies $\mathsf{d}_{h\Omega}(x, \rho^{g_i\Lambda_i}_{h\Omega}) = 0$. ☐

4.2.6 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 4.2.14 to show that $[g\Lambda] \sqsubseteq [h\,\mathrm{lk}(w)]$ for each vertex w of Ω, which will imply $[g\Lambda]\perp[hw]$ for each vertex w of Ω. We then obtain $[g\Lambda]\perp[h\Omega]$ by adapting the proof of Proposition 4.2.16 to show that we may promote orthogonality with multiple domains to orthogonality with their union. This contradicts $[g\Lambda]\pitchfork[h\Omega]$.

Lemma 4.2.18. *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 4.2.16, 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 4.2.19 (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*

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

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 $\mathsf{d}_{h\Omega}(\pi_{h\Omega}(x), \rho^{g\Lambda}_{h\Omega}) > 2$ and $\mathsf{d}_{g\Lambda}(\pi_{g\Lambda}(x), \rho^{h\Omega}_{g\Lambda}) > 2$. Then we also have

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

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 2.6.22(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 2.6.22(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 4.10.

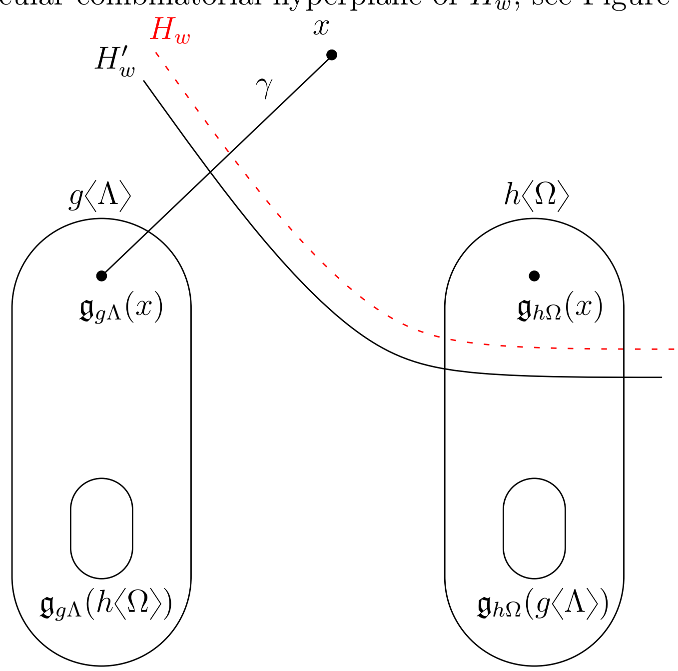

Figure 4.10: The combinatorial hyperplane H'_w of H_w that is in the same component of $S(\Gamma) \smallsetminus H_w$ as both $g\langle\Lambda\rangle$ and $\mathfrak{g}_{h\Omega}(g\langle\Lambda\rangle)$.

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 $\mathsf{d}_{g\Lambda}(\pi_{g\Lambda}(x), \pi_{g\Lambda}(h\langle\mathrm{st}(\Omega)\rangle)) = \mathsf{d}_{g\Lambda}(\pi_{g\Lambda}(x), \rho^{h\Omega}_{g\Lambda}) > 2$, 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 $\mathrm{diam}(\pi_{g\Lambda}(H'_w)) > 2$.

By Remark 2.6.19, we have $H'_w \subseteq h\langle\mathrm{lk}(w)\rangle$. Thus, $\mathrm{diam}(\pi_{g\Lambda}(H'_w)) > 2$ implies that $\mathrm{diam}(\pi_{g\Lambda}(h\langle\mathrm{lk}(w)\rangle)) > 2$. Lemma 4.2.14 then forces $[g\Lambda] \sqsubseteq [h\,\mathrm{lk}(w)] \sqsubseteq [h\,\mathrm{st}(w)]$. This implies $\Lambda \subseteq \mathrm{lk}(w)$ and that there exists $k \in G_\Gamma$ such that $[k\Lambda] = [g\Lambda]$ and $[k\,\mathrm{st}(w)] = [h\,\mathrm{st}(w)]$. Since $\mathrm{st}(\mathrm{st}(w)) = \mathrm{st}(w)$, $[k\,\mathrm{st}(w)] = [h\,\mathrm{st}(w)]$ implies $[kw] = [hw]$. 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 4.2.18 then tells us $[g\Lambda] \perp [h\Omega]$, contradicting transversality. Hence inequality $(*)$ must hold.

Now suppose $[k\Pi] \subsetneq [g\Lambda]$ 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 $[k\Pi] = [a\Pi], [g\Lambda] = [a\Lambda]$. 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 $\mathsf{d}_{h\Omega}(\rho^{k\Pi}_{h\Omega}, \rho^{g\Lambda}_{h\Omega}) = 0$. □

4.2.7 Compatibility of the group structure

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

Lemma 4.2.20. *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 4.2.21. *For each subgraph $\Lambda \subseteq \Gamma$ and elements $a, g \in G_\Gamma$, there exists an isometry $a_{g\Lambda} \colon C(g\Lambda) \to C(ag\Lambda)$ satisfying the following for all subgraphs $\Lambda, \Omega \subseteq \Gamma$ and elements $a, b, g, h \in G_\Gamma$.*

- *The map $(ab)_{g\Lambda} \colon C(g\Lambda) \to C(abg\Lambda)$ is equal to the map $a_{bg\Lambda} \circ b_{g\Lambda} \colon C(g\Lambda) \to C(abg\Lambda)$.*
- *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 2.6.22(2).
- The final property follows as an immediate consequence of the previous one and the definition of the relative projections. □

4.2.8 Graph products are relative HHGs

We now compile the results from Section 4.2 to obtain the main result of this chapter, that any graph product of finitely generated groups is a relative HHG.

Theorem 4.2.22. *Let G_Γ be a graph product of finitely generated groups. The proto-hierarchy structure $\mathfrak{S}_\Gamma$ from Theorem 4.1.23 is a relatively hierarchically hyperbolic group structure for G_Γ with hierarchy constant* $\max\{18, |V(\Gamma)|\}$.

Proof. Let $\mathfrak{S}_\Gamma$ be the proto-hierarchy structure for (G_Γ, d) from Theorem 4.1.23. 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 4.2.1.

(2) We proved finite complexity in Lemma 4.2.5.

(3) We proved the container axiom in Lemma 4.2.6.

(4) The proof of the uniqueness axiom follows from Lemma 4.2.7, since if $\mathsf{d}_{C([g\Lambda])}(x, y)$ is uniformly bounded for all $[g\Lambda] \in \mathfrak{S}_\Gamma$, then Lemma 4.1.8 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 4.2.12.

(6) We proved the large links axiom in Lemma 4.2.15.

(7) We proved the consistency axiom in Lemma 4.2.19.

(8) We proved the partial realisation axiom in Lemma 4.2.17.

We now verify the remaining axioms required for (G_Γ, d) to be a relative HHG, as laid out in Definition 2.7.3.

Let $\phi : G_\Gamma \times \mathfrak{S}_\Gamma \to \mathfrak{S}_\Gamma$ be the map $\phi(a, [g\Lambda]) = [ag\Lambda]$. By Lemma 4.2.20, this is a well-defined G_Γ–action by bijections that preserves the nesting, orthogonality, and transversality relations and has finitely many orbits. We will use $a \cdot [g\Lambda]$ to denote $\phi(a, [g\Lambda]) = [ag\Lambda]$.

For each $[g\Lambda] \in \mathfrak{S}_\Gamma$, let $\overline{g}\Lambda$ denote the fixed representative of $[g\Lambda]$ such that $C([g\Lambda]) = C(\overline{g}\Lambda)$; see the proto-hierarchy structure in Theorem 4.1.23. Left multiplication by $a \in G_\Gamma$ gives an isometry $a_{g\Lambda} \colon C(g\Lambda) \to C(ag\Lambda)$ 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 $\mathbf{a}_{[g\Lambda]} \colon C(\overline{g}\Lambda) \to C(\overline{ag}\Lambda)$ by $\mathbf{a}_{[g\Lambda]} = \mathfrak{g}_{\overline{ag}\Lambda} \circ a_{\overline{g}\Lambda}$.

Let $a, b \in G_\Gamma$ and $[g\Lambda], [h\Omega] \in \mathfrak{S}_\Gamma$. We now verify the remaining axioms of a relatively hierarchically hyperbolic group (Definition 2.7.3).

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

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

Using the last clause of Lemma 4.1.8, we have

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

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

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

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

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

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

- Let $x \in G_\Gamma$. Since $a\overline{g}\Lambda \parallel \overline{ag}\Lambda$, we can use Lemma 4.1.8 and equivariance of the gate map (Proposition 2.6.22(2)) to conclude:

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

- Suppose $[h\Omega] \pitchfork [g\Lambda]$ or $[h\Omega] \subsetneq [g\Lambda]$. Lemmas 4.1.8, 4.2.20, and 4.2.21 imply $\mathfrak{a}_{[g\Lambda]}\left(\rho^{[h\Omega]}_{[g\Lambda]}\right) = \rho^{[ah\Omega]}_{[ag\Lambda]}$:

$$\begin{aligned}
\mathfrak{a}_{[g\Lambda]}\left(\rho^{[h\Omega]}_{[g\Lambda]}\right) &= \left(\mathfrak{g}_{\overline{ag}\Lambda} \circ a_{\overline{g}\Lambda}\right)\left(\rho^{\overline{h}\Omega}_{\overline{g}\Lambda}\right) && \text{(Definition of } \mathbf{a}_{[g\Lambda]}) \\
&= \mathfrak{g}_{\overline{ag}\Lambda}\left(\rho^{a\overline{h}\Omega}_{a\overline{g}\Lambda}\right) && \text{(Lemma 4.2.21)} \\
&= \mathfrak{g}_{\overline{ag}\Lambda}\left(\mathfrak{g}_{a\overline{g}\Lambda}(a\overline{h}\langle \mathrm{st}(\Omega)\rangle)\right) && \text{(Definition of } \rho) \\
&= \mathfrak{g}_{\overline{ag}\Lambda}(a\overline{h}\langle \mathrm{st}(\Omega)\rangle) && \text{(Lemma 4.1.8)} \\
&= \mathfrak{g}_{\overline{ag}\Lambda}(\overline{ah}\langle \mathrm{st}(\Omega)\rangle) && (a\overline{h}\Omega \parallel \overline{ah}\Omega) \\
&= \rho^{\overline{ah}\Omega}_{\overline{ag}\Lambda}. && \square
\end{aligned}$$

Behrstock, Hagen, and Sisto show that any relatively hierarchically hyperbolic space has a distance formula, which expresses distances in the space as a sum of distances in the projections [BHS19, Theorem 6.10]. As a result, we now have such a distance formula for graph products of finitely generated groups.

Corollary 4.2.23 (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 $\sigma \geqslant \sigma_0$ there exist $K \geqslant 1$ and*

$L \geqslant 0$ *such that for all* $g, h \in G_\Gamma$

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

where we define $\{\!\{N\}\!\}_\sigma = N$ *if* $N \geqslant \sigma$ *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.

Corollary 4.2.24 (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) hierarchically hyperbolic group and $T \in \mathfrak{S}$ is the $\sqsubseteq$–maximal element, then the action of G on $C(T)$ is acylindrical [BHS17b, 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 $\mathfrak{S} \neq \{T\}$. 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. □

4.2.9 The syllable metric is an HHS

Since nearly every argument used in the proof of Theorem 4.2.22 factors through the syllable metric on the graph product G_Γ, the same arguments show that the syllable metric on G_Γ is itself a hierarchically hyperbolic space. 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 4.2.22 is to ensure that G_Γ has a word metric, this does not create any additional difficulty.

Corollary 4.2.25. *Let* Γ *be a finite simplicial graph, with each vertex* v *labelled by a non-trivial group* G_v*. Then the graph product* G_Γ *endowed with the syllable metric is a hierarchically hyperbolic space.*

Proof. Define the proto-hierarchy structure for G_Γ as before, except whenever $v \in V(\Gamma)$ and $g \in G_\Gamma$, 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. $\square$

4.3 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 4.3.1, which uses our results from Chapter 3 to show 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 [Gen19b, Questions 8.3, 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 4.3.10) and when it is a quasi-line (Theorem 4.3.12). As Genevois proved that any quasi-isometry between graph products of finite groups induces a quasi-isometry between their electrifications [Gen19b, Proposition 1.4], these two theorems provide us with tools for studying quasi-isometric rigidity of graph products of finite groups.

4.3.1 Graph products of HHGs

Theorem 4.3.1. *Let G_Γ be a graph product of finitely generated groups. If for each $v \in V(\Gamma)$, the vertex group G_v is a hierarchically hyperbolic group, then G_Γ is a hierarchically hyperbolic group.*

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 4.2.22. 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 $\overline{g}\Lambda$ be the fixed representative of $[g\Lambda]$ so that $C([g\Lambda]) = C(\overline{g}\Lambda)$. If $[g\Lambda] = [\Lambda]$, then we choose $\overline{g} = e$.

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 $\overline{g}U$ where $\overline{g}$ 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 4.3.1 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 4.3.1 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 4.3.2). This proto-hierarchy structure is not quite a hierarchically hyperbolic space structure, as it satisfies every axiom except the container axiom (Claim 4.3.3). 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 4.3.4). This allows us to use the results of Chapter 3 to upgrade $\mathfrak{T}_0$ to a genuine HHS structure $\mathfrak{T}$. Since the proto-structure will satisfy the equivariance properties of a hierarchically hyperbolic group structure for G_Γ (Claim 4.3.5), this HHS structure will also be a hierarchically hyperbolic group structure.

Claim 4.3.2. *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 $\overline{\mathfrak{R}}$.

Projections: For all $W \in \mathfrak{T}_0$, the projection map will be denoted $\psi_W \colon G_\Gamma \to 2^{C(W)}$. 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 $U \sqsubseteq_{\mathfrak{R}} W$;

- $W \in \mathfrak{S}_\Gamma \setminus \mathfrak{S}_\Gamma^{min}$ 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 $U \subsetneq W$ we denote the relative projection from U to W by β_W^U and define it as follows.

- If $W, U \in \mathfrak{S}_\Gamma \setminus \mathfrak{S}_\Gamma^{min}$ 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 $W \in \mathfrak{S}_\Gamma \setminus \mathfrak{S}_\Gamma^{min}$ 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 is 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:

- $W, U \in \mathfrak{S}_\Gamma \setminus \mathfrak{S}_\Gamma^{min}$ and $U \perp_{\mathfrak{S}} W$;

- $W, U \in \mathfrak{R}_V$ and $U \perp_{\mathfrak{R}} W$;

- $W \in \mathfrak{S}_\Gamma \setminus \mathfrak{S}_\Gamma^{min}$ 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 4.3.3. $\mathfrak{T}_0$ satisfies all of the axioms of a hierarchically hyperbolic space except for the container axiom.

Proof. Recall, $E_1 > 0$ is the hierarchy constant from the proto-hierarchy structure $\mathfrak{T}_0$. Note 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 $\kappa \geqslant 0$ 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 $\mathsf{d}(x,y) \geqslant \theta(\theta(\kappa)+\kappa)$, then there exists $W \in \mathfrak{S}_\Gamma$ such that $\mathsf{d}_W(x,y) \geqslant \theta(\kappa)+\kappa$ 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 $\mathsf{d}_U(x,y) \geqslant \kappa$. 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

$$E_1^2 + E_1 < \mathsf{d}_U(x,y) = \mathsf{d}_U(\pi_U^V(\pi_V(x)), \pi_U^V(\pi_V(y))) \leqslant E_1 \mathsf{d}_V(\pi_V(x), \pi_V(y)) + E_1,$$

which implies $E_1 < \mathsf{d}_V(\pi_V(x), \pi_V(y))$. Now, the bounded geodesic image axiom in $(G_\Gamma, \mathfrak{S}_\Gamma)$ says 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 $W \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{min}$. The large links axiom for $(G_\Gamma, \mathfrak{S}_\Gamma)$ gives a collection $\mathfrak{L} = \{U_1, \ldots, U_m\}$ 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 $i \in \{1, \ldots, m\}$, define $\overline{U_i}$ to be the $\sqsubseteq_{\mathfrak{R}}$–maximal element of $\mathfrak{R}_{U_i}$ if $U_i \in \mathfrak{S}_\Gamma^{min}$ and define $\overline{U}_i$ to be U_i if $U_i \notin \mathfrak{S}_\Gamma^{min}$. Let $\overline{\mathfrak{L}} = \{\overline{U_1}, \ldots, \overline{U_m}\}$.

If $V \in \mathfrak{S}_\Gamma^{min}$ is nested into W, but is not nested into an element of $\mathfrak{L}$, then $\mathsf{d}_V(\pi_V(x), \pi_V(y)) < E_1$ and so

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

for all $Q \in \mathfrak{R}_V$. Thus, if $\mathsf{d}_Q(\psi_Q(x), \psi_Q(y)) \geqslant E_1^2 + E_1$ and Q is nested into W, then either $Q \in \mathfrak{S} \smallsetminus \mathfrak{S}_\Gamma^{min}$ 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 $\overline{\mathfrak{L}}$). 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 $\overline{\mathfrak{L}}$. Thus, $Q \sqsubseteq W$ is nested into an element of $\overline{\mathfrak{L}}$ whenever $\mathsf{d}_Q(\psi_Q(x), \psi_Q(y)) \geqslant E_1^2 + E_1$.

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 $Q = W$ 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 $\mathsf{d}_V(x, \rho_V^Q) \leqslant E_1$. The coarse Lipschitzness of the projections then implies $\mathsf{d}_U(x, \pi_U^V(\rho_V^Q)) = \mathsf{d}_U(x, \beta_U^W) \leqslant E_1^2 + E_1$.

Now assume $Q = V'$ 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 $\mathsf{d}_V(x, \rho_V^Q) \leqslant E_1$ and we again have $\mathsf{d}_U(x, \beta_U^W) = \mathsf{d}_V(x, \pi_U^V(\rho_V^Q)) \leqslant E_1^2 + E_1$.

For the last clause of the consistency axiom, let $Q, U, W \in \mathfrak{T}_0$ with $Q \subsetneq U$. If $U \subsetneq W$, 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 $\mathsf{d}_W(\beta_W^Q, \beta_W^U) \leqslant E_1^2 + E_1$. Similarly, if $W \in \mathfrak{S}_\Gamma$ with $W \pitchfork U$ and $W \not\perp Q$, then $\mathsf{d}_W(\beta_W^Q, \beta_W^U) \leqslant E_1^2 + E_1$. Assume $W \in \mathfrak{R}_V$ 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 $\mathsf{d}_V(\rho_V^U, \rho_V^Q) \leqslant E_1$. 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 $\mathsf{d}_V(\rho_V^U, \rho_V^{V'}) \leqslant E_1$. Hence in both cases $\mathsf{d}_W(\beta_W^U, \beta_W^Q) \leqslant E_1^2 + E_1$.

Partial realisation: Let $W_1, \ldots, W_n$ be pairwise orthogonal elements of $\mathfrak{T}_0$ and $p_i \in C(W_i)$ for each $i \in \{1, \ldots, n\}$. 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 $V_1, \ldots, V_r \in \mathfrak{S}_\Gamma^{min}$ so that for each $i \in \{1, \ldots, n\}$, either $W_i \in \mathfrak{S}_\Gamma$ or there exists a unique $j \in \{1, \ldots, r\}$ such that $W_i \in \mathfrak{R}_{V_j}$. For each $j \in \{1, \ldots, r\}$, let $\{W_1^j, \ldots, W_{k_j}^j\}$ be the elements of $\{W_1, \ldots, W_n\}$ that are also elements of $\mathfrak{R}_{V_j}$ and let $\{p_1^j, \ldots, p_{k_j}^j\}$ be the subset of $\{p_1, \ldots, p_n\}$ satisfying $p_i^j \in C(W_i^j)$ for all $j \in \{1, \ldots, r\}$ and $i \in \{1, \ldots, k_j\}$. Using partial realisation for each of the $(C(V_j), \mathfrak{R}_{V_j})$ on the points $p_1^j, \ldots, p_{k_j}^j$ produces a set of points $y_1, \ldots y_r$ so that for each $j \in \{1, \ldots, r\}$:

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

Assume, without loss of generality, that $W_m, W_{m+1}, \ldots, W_n$ are all of the W_i that are not contained in any of the $\mathfrak{R}_{V_j}$ (it is possible the set of such W_i is empty). Now, applying partial realisation for $(G_\Gamma, \mathfrak{S}_\Gamma)$ to $y_1, \ldots, y_r, p_m, \ldots, p_n$ produces a point $x \in G_\Gamma$ so that $\psi_{W_i}(x)$ is uniformly close to p_i for each $i \in \{1, \ldots, n\}$ 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, 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 $y_1, \ldots, y_r$, but the conclusion still holds. □

Claim 4.3.4. The E_1–proto-hierarchy structure $\mathfrak{T}_0$ has the following property: if $W_1, \ldots, W_n \in \mathfrak{T}_0$ are pairwise orthogonal, then $n \leqslant E_1^2 + E_1$.

Proof. Let $W_1, \ldots, W_n \in \mathfrak{T}_0$ be pairwise orthogonal. Without loss of generality, let $W_1, \ldots W_k$ be the elements of $\{W_1, \ldots, W_n\}$ that are elements of $\mathfrak{S}_\Gamma$. Since $W_1, \ldots, W_k$ is a pairwise orthogonal collection of elements of $\mathfrak{S}_\Gamma$, Lemma 2.7.18 says $k \leqslant E_1$.

Let $V_1, \ldots, V_m$ be the minimal collection of elements of $\mathfrak{S}_\Gamma^{min}$ such that if $i \in \{k+1, \ldots, n\}$ (i.e. $W_i \notin \mathfrak{S}_\Gamma$), then $W_i \in \mathfrak{R}_{V_j}$ for some $j \in \{1, \ldots, m\}$. Minimality implies that for each $j \in \{1, \ldots, m\}$, there exists $i \in \{k+1, \ldots, n\}$ such that $W_i \in \mathfrak{R}_{V_j}$. Suppose $W_i \in \mathfrak{R}_{V_j}$ and $W_\ell \in \mathfrak{R}_{V_r}$ with $j \neq r$. Since $W_i \perp W_\ell$ in $\mathfrak{T}_0$, then definition of orthogonality in $\mathfrak{T}_0$ implies that $V_j \perp_{\mathfrak{S}} V_r$. Thus, $V_1, \ldots, V_m$ is a pairwise orthogonal collection of elements of $\mathfrak{S}_\Gamma$ and $m \leqslant E_1$ by Lemma 2.7.18. Similarly, for each $j \in \{1, \ldots, m\}$ 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 E_1. Putting this together, we have that $n \leqslant k + E_1 m \leqslant E_1 + E_1^2$. □

Claim 4.3.5. 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 a hierarchically hyperbolic group (Definition 2.7.3).

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 $W = \overline{g}R \in \mathfrak{R}_{[gv]}$ for some $[gv] \in \mathfrak{S}_\Gamma^{min}$, then $(\overline{\sigma g})^{-1}\sigma\overline{g} \in \mathrm{Stab}_{G_\Gamma}([v])$, where $\overline{\sigma g}$ is the chosen fixed representative of $[\sigma gv] = [\sigma\overline{g}v]$. 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 $l\hat{\sigma} = (\overline{\sigma g})^{-1}\sigma\overline{g}$. 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 $\overline{g}R$. Define $\Phi(\sigma, \overline{g}R) = \overline{\sigma g}R_\sigma \in \mathfrak{R}_{[\sigma gv]}$. The following commutative diagram summarises how σ takes elements of $\mathfrak{R}_{[gv]}$ to elements of $\mathfrak{R}_{[\sigma gv]}$.

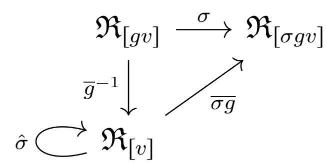

We now verify that Φ preserves the relations in $\mathfrak{T}_0$. Let $W, U \in \mathfrak{T}_0$. If $W, U \in \mathfrak{S}_\Gamma \smallsetminus \mathfrak{S}_\Gamma^{min}$ or $W, U \in \mathfrak{R}_{[gv]}$ for some $[gv] \in \mathfrak{S}_\Gamma^{min}$, then Φ preserves the relation between W and U, since 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 $\overline{g}R \in \mathfrak{R}_{[gv]}$ is in the G_Γ–orbit of $\overline{h}R' \in \mathfrak{R}_{[hw]}$ 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 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 $\sigma_{[g\Lambda]}\colon C([g\Lambda]) \to C([\sigma g\Lambda])$ to be the same as the original isometry in $(G_\Gamma, \mathfrak{S}_\Gamma)$; this guarantees the HHG axioms are satisfied in this case.

If $W \in \mathfrak{R}_{[gv]}$, then $W = \overline{g}R$ 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 $\sigma_W\colon C(W) \to C(\sigma W)$ to be the composition

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

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

Now, if $\tau \in G_\Gamma$, then $(G_v, \mathfrak{R}_{[v]})$ being an HHG implies $\hat{\tau}_{R_\sigma} \circ \hat{\sigma}_R = \widehat{\tau\sigma}_R$. 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 4.3.2: 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 $\psi_W\colon G_\Gamma \to 2^{C(W)}$ 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 $U \subsetneq W$ 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 $U \subsetneq W$ or $U \pitchfork W$. □

We now conclude the proof of Theorem 4.3.1 by noting that Claims 4.3.3, 4.3.4, and 4.3.5 show that the proto-hierarchy structure $\mathfrak{T}_0$ defines an almost HHG structure on G_Γ. Thus, there exists an HHG structure $\mathfrak{T}$ for G_Γ by Theorem 3.0.1 and Remark 3.0.6. □

4.3.2 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 hierarchically hyperbolic spaces (Theorem 2.7.12).

Theorem 4.3.6 (Meier's criterion for hyperbolicity of graph products; [Mei96])**.** *Let Γ be a finite simplicial graph with hyperbolic groups associated to its vertices. Let Γ_F be the induced 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 $\mathrm{lk}(v)$ is a complete graph.

(iii) Γ_F does not contain any induced squares.

Proof. We show hyperbolicity via the bounded orthogonality condition of Theorem 2.7.12, 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 we have bounded orthogonality. Then:

(i) Suppose two infinite vertices v, w are connected by an edge. Then $[v]\bot[w]$ 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 $\mathrm{lk}(v)$ 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, $[v]\bot[x \cup y]$, $C(v)$ 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 $g_x \in G_x \smallsetminus \{1\}, g_y \in G_y \smallsetminus \{1\}$. This again contradicts bounded orthogonality.

(iii) Suppose Γ_F contains a square with vertices v, x, w, y, where v, w and x, y are non-adjacent. Then $[v \cup w]\bot[x \cup y]$ 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 conditions (i)–(iii) hold and let $D = \max\{2, |G_v| : v \in \Gamma_F\}$. Moreover, suppose $[g\Lambda], [h\Omega] \in \mathfrak{S}$ satisfy $[g\Lambda] \perp [h\Omega]$.

Suppose $\operatorname{diam}(C(g\Lambda)) > D$. Then Theorem 4.2.10 tells us that either Λ consists of a single infinite vertex or Λ contains at least 2 vertices and does not split as a join.

If Λ consists of a single infinite vertex, then conditions (i) and (ii) tell us that $\operatorname{lk}(\Lambda) \supseteq \Omega$ 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, $\operatorname{diam}(C(h\Omega)) \leqslant D$.

If Λ contains at least 2 vertices and does not split as a join, then in particular it contains two non-adjacent vertices v and w. As $\Omega \subseteq \operatorname{lk}(\Lambda)$, every vertex of Ω is connected to both v and w. Condition (iii) implies that every pair of vertices of Ω must be connected by an edge, and condition (i) then implies that $\Omega \subseteq \Gamma_F$. That is, Ω either consists of a single finite vertex or splits as a join. In both cases, $\operatorname{diam}(C(h\Omega)) \leqslant D$. Thus the bounded orthogonality condition holds. □

4.3.3 Genevois' minsquare electrification.

We now use our characterisation of when $C(g\Lambda)$ has infinite diameter (Theorem 4.2.10) to answer two questions of Genevois [Gen19b, Questions 8.3, 8.4] regarding the *electrification* of G_Γ, defined as follows.

Definition 4.3.7. 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 $\mathsf{d}_{\mathbb{E}}(g,h)$ 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 4.3.8 ([Gen19b, 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 [Gen19b] 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 large distances in $C(\Gamma)$, which we can detect with Theorem 4.2.10, 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 4.3.9. *If Γ splits as a join and contains a proper minsquare subgraph, then Γ splits as a join $\Gamma = \Gamma_1 \star \Gamma_2$ 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 Γ contains a proper minsquare subgraph Λ and splits as a join $\Gamma = \Omega_1 \star \Omega_2$. We first show Γ splits as a (possibly different) join $\Gamma_1 \star \Gamma_2$, 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 $\Gamma_1 \star \Gamma_2$ 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 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 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 $\Gamma_1 \star \Gamma_2$, 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 4.3.10. *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 $\operatorname{supp}(s_i) = v_i \in V(\Gamma)$ and $v_i \neq v_j$ for all $i \neq j$. Thus $\mathsf{d}_{\mathbb{E}}(x, y) \leqslant n$, 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 $\mathsf{d}_{\mathbb{E}}(x, y) \leqslant n + 1$, hence $\mathbb{E}(\Gamma)$ is bounded.

We now assume $\mathbb{E}(\Gamma)$ is bounded and prove this implies Γ is either 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 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 $n \geqslant 2$ vertices.

Claim 4.3.11. 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 Γ does not split as a join. By Theorem 4.2.10, $C(\Gamma)$ is therefore unbounded. Since Γ is not minsquare, $\mathbb{E}(\Gamma)$ is $C(\Gamma)$ with some edges removed, so if $C(\Gamma)$ has infinite diameter then so does $\mathbb{E}(\Gamma)$. That is, if Γ is not minsquare and does not split as a join then $\mathbb{E}(\Gamma)$ is unbounded, contradicting our assumption.

Now suppose Γ does not contain any proper minsquare subgraphs. Then $\mathbb{E}(\Gamma)$ is simply G_Γ with the syllable metric. Since Γ is not complete, there exist two disjoint vertices $v, w \in V(\Gamma)$. Therefore $\mathsf{d}_{\mathbb{E}}(e, (g_v g_w)^m) = 2m$ for any $g_v \in G_v \smallsetminus \{e\}$ and $g_w \in G_w \smallsetminus \{e\}$, 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 4.3.11. By Lemma 4.3.9, 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 $\mathbb{E}(\Gamma_1)$ must also have bounded diameter. Since Γ_1 contains at most $n-1$ vertices, the induction hypothesis then implies 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 either splits 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 $\Omega \star \Gamma_2$. □

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 direction to escape to infinity in $\mathbb{E}(\Gamma)$.

Theorem 4.3.12. *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 $\Gamma_1 \star \Gamma_2$, 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 [BPR19, 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 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 4.3.10 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 Γ does not split as a join at all. Since the action of G_Γ on $C(\Gamma)$ by left multiplication is acylindrical (Corollary 4.2.24), Theorem 2.3.2 says 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 4.2.10 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 $\mathsf{d}_\Gamma(x,y) \leqslant \mathsf{d}_\mathbb{E}(x,y)$ 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 $\mathbb{E}(\Gamma)$ is not a quasi-line, as any two bi-infinite quasi-geodesics in a quasi-line have finite Hausdorff distance.

Now assume 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 G_Γ is not virtually cyclic. Thus we can assume Γ contains a minsquare subgraph Λ. By applying Lemma 4.3.9 iteratively, we have that Γ splits as a join $\Gamma = \Gamma_1 \star \Gamma_2$ 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 that we are assuming 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. □

Chapter 5

Non-positive curvature in graph braid groups

In this chapter we will develop an explicit HHG structure for graph braid groups by using the cubical structure described in Section 2.5. We shall then use this HHG structure to characterise when a graph braid group is hyperbolic (Theorem 5.2.1) or acylindrically hyperbolic (Theorem 5.2.4), as well as conjecturing and partially proving a characterisation of relative hyperbolicity and thickness (Conjecture 5.2.6 and Theorem 5.2.7).

Throughout this chapter we shall take Γ to be a finite, connected graph and consider the graph braid group $B_n(\Gamma)$ for $n \in \mathbb{N}$. The case where Γ is disconnected may be treated by applying Lemma 2.5.2 and using Behrstock–Hagen–Sisto's construction of HHG structures on products of HHGs [BHS19, Section 8.3]. This then reduces our analysis to the connected case.

Convention 5.0.1. In order to avoid confusion between edges of the graph Γ and edges of the cube complex $UC_n(\Gamma)$ when discussing the structure of a graph braid group $B_n(\Gamma)$, we shall adopt the convention of denoting edges of Γ by e and edges of $UC_n(\Gamma)$ by E. Moreover, e will denote a *closed* edge of Γ, unless otherwise specified.

5.1 The hierarchically hyperbolic structure on a graph braid group

Recall that in Section 2.7.4 we showed a graph braid group $B_n(\Gamma)$ has the structure of a hierarchically hyperbolic group by virtue of its cubical structure (Corollary 2.7.25). In particular, by sufficiently subdividing edges of Γ, we obtain a new graph Γ' such that $B_n(\Gamma)$ is isomorphic to the fundamental group of the unordered combinatorial configuration space $UC_n(\Gamma')$ (Theorem 2.5.5), and moreover $UC_n(\Gamma')$ is a compact special cube complex (Corollary 2.5.7).

Since $B_n(\Gamma)$ is isomorphic to $B_n(\Gamma')$, we may drop the Γ' notation entirely and simply work under the assumption that Γ satisfies the conditions of Theorem 2.5.5. We shall adopt this convention for the remainder of the chapter. Note that in particular, Γ can be assumed to be a simplicial graph.

Hierarchical hyperbolicity of $B_n(\Gamma)$ is obtained via its action on the universal cover X of $UC_n(\Gamma)$, which is a CAT(0) cube complex. Since $UC_n(\Gamma)$ is a special cube complex with finitely many hyperplanes, it follows that X has a factor system that is invariant under the action of $\pi_1(UC_n(\Gamma)) \cong B_n(\Gamma)$, by Theorem 2.7.23. We can then apply the construction in Section 2.7.4 to obtain an explicit HHG structure for $B_n(\Gamma)$.

5.1.1 The cubical structure

Recall that $UC_n(\Gamma)$ is a compact cube complex, as described in Section 2.5. Each vertex of the cube complex is a configuration of the n particles on the vertices of the graph Γ. Two vertices of $UC_n(\Gamma)$ are connected by an edge if one configuration can be obtained from the other by moving a particle along an edge of Γ to a vacant neighbouring vertex. Two adjacent edges of $UC_n(\Gamma)$ span a square if the corresponding moves can be performed independently of each other.

More concretely, we have the following construction, as described by Genevois [Gen19a].

(See Figure 5.1 for an example.)

- The vertices of $UC_n(\Gamma)$ are the subsets S of $V(\Gamma)$ with cardinality $|S| = n$.
- Two vertices S and S' of $UC_n(\Gamma)$ are connected by an edge if their symmetric difference $S \triangle S'$ is a pair of adjacent vertices of Γ. We therefore label each edge E of $UC_n(\Gamma)$ with a closed edge e of Γ. Note that $S \cap S'$ is a subset of $V(\Gamma \smallsetminus e)$ of cardinality $n-1$; that is, $S \cap S'$ is a vertex of $UC_{n-1}(\Gamma \smallsetminus e)$. Here, $\Gamma \smallsetminus e$ denotes the induced subgraph of Γ spanned by the vertices $V(\Gamma) \smallsetminus V(e)$.
- A collection of m edges of $UC_n(\Gamma)$ with a common endpoint span an m–cube if their labels are pairwise disjoint.

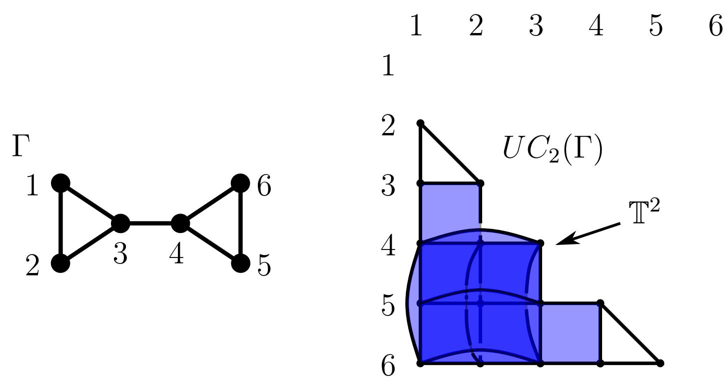

Figure 5.1: An example of an unordered combinatorial configuration space. The vertices of Γ are labelled to more easily see the construction of the cube complex. One can see that $UC_2(\Gamma) \simeq \mathbb{T}^2 \vee \mathbb{S}^1 \vee \mathbb{S}^1$, thus $B_2(\Gamma) \cong \mathbb{Z}^2 * F_2$. Note that the torus arises as the product of the two 3–cycles in Γ.

Remark 5.1.1. We make heavy use of the above description of $UC_n(\Gamma)$ in that which follows. It is important to keep in mind that when we refer to a vertex S of $UC_n(\Gamma)$, we are considering S to be a subset of $V(\Gamma)$ of cardinality n.

Genevois uses the above description of $UC_n(\Gamma)$ to give an important characterisation of the hyperplanes of $UC_n(\Gamma)$.

Lemma 5.1.2 (Hyperplane labelling; [Gen19a, Lemma 3.6])**.** *Let E_1 and E_2 be two edges of $UC_n(\Gamma)$ and denote the endpoints of E_i by S_i, S_i' for $i = 1, 2$. The edges E_1 and E_2 are dual to the same hyperplane if and only if they are labelled by the same closed edge e of Γ and $S_1 \cap S_1'$, $S_2 \cap S_2'$ are in the same connected component K_e of $UC_{n-1}(\Gamma \smallsetminus e)$.*

Remark 5.1.3 (Labelling hyperplanes and embedding combinatorial hyperplanes)**.**

(1) Lemma 5.1.2 tells us that hyperplanes H of $UC_n(\Gamma)$ can be consistently and uniquely labelled by pairs (e, K_e), where e is a (closed) edge of Γ and K_e is a connected component of $UC_{n-1}(\Gamma \smallsetminus e)$. Let E' be an edge of $UC_n(\Gamma)$. Then E' is an edge of some combinatorial hyperplane H' associated to H if and only if there exists some edge E dual to H such that E and E' span a square. Thus, the labels of E and E' are disjoint closed edges of Γ; that is, the edges of H' are labelled by edges of $\Gamma \smallsetminus e$.

(2) Crossing a hyperplane corresponds to moving a particle along the associated edge e of Γ, and travelling along a combinatorial hyperplane from S_1 to S_2 corresponds to rearranging the remaining $n-1$ particles from the configuration $S_1 \cap S_1'$ to the configuration $S_2 \cap S_2'$ without using the edge e. It therefore follows that $K_e \subseteq UC_{n-1}(\Gamma \smallsetminus e)$ has two canonical isometric embeddings in $UC_n(\Gamma)$, as the two combinatorial hyperplanes. More generally, intersections of combinatorial hyperplanes are isometric to connected components of complexes $UC_k(\Lambda)$, where $k < n$ and Λ is obtained from Γ by removing a collection of disjoint closed edges.

5.1.2 The HHG structure

The goal of this section is to describe the HHG structure on a graph braid group $B_n(\Gamma)$ afforded by Corollary 2.7.25 in terms of the graph Γ. This will enable us to characterise hyperbolicity, acylindrical hyperbolicity, relative hyperbolicity, and thickness in terms of Γ. To this end, we study a natural collection of subgroups of $B_n(\Gamma)$ that we call the *graphical* *subgroups*, constructed as follows. Again, it is important to note that in what follows, vertices of $UC_n(\Gamma)$ are considered as subsets of $V(\Gamma)$ of cardinality n.

Definition 5.1.4 (Graphical subgroup)**.** Let Γ be a finite connected graph, let $\Lambda \subseteq \Gamma$ be a subgraph with no isolated vertices, and choose a base point $T \in UC_n(\Gamma)^{(0)}$ with $|T \cap \Lambda| = k$. Then the connected component K of $UC_k(\Lambda)$ containing $S = T \cap \Lambda$ is embedded isometrically as a subcomplex of $UC_n(\Gamma)$ by fixing all particles in $T \cap (\Gamma \smallsetminus \Lambda)$ and restricting the motion of the remaining k particles to Λ. Furthermore, by [HW08, Lemma 2.11] this embedding is π_1–injective, thus induces an embedding of $B_k(\Lambda, S)$ as a subgroup of $B_n(\Gamma)$. We call this a *graphical subgroup* and denote it $\langle \Lambda, k, S \rangle$.

Remark 5.1.5. We may define an equivalence relation $\sim$ on the vertices of $UC_k(\Lambda)$ by writing $S \sim S'$ if S and S' are in the same connected component of $UC_k(\Lambda)$. Note that $\langle \Lambda, k, S \rangle$ and $\langle \Lambda, k, S' \rangle$ define the same graphical subgroup if and only if $S \sim S'$. We may therefore change the base point S to another S' in the same $\sim$–equivalence class if convenient, without affecting the graphical subgroup. We take advantage of this fact frequently.

Remark 5.1.6. Let X be the universal cover of $UC_n(\Gamma)$ and let $\langle \Lambda, k, S \rangle$ be a graphical subgroup of $B_n(\Gamma)$. Embed the connected component K of $UC_k(\Lambda)$ containing S into $UC_n(\Gamma)$ as in Definition 5.1.4, then consider its universal cover as a subcomplex $\widetilde{K}$ of X. Then $\widetilde{K}$ is quasi-isometric to $\langle \Lambda, k, S \rangle$ via the orbit map, by the Milnor–Švarc lemma.

The above remark gives us a correspondence between cosets of graphical subgroups and subcomplexes of the universal cover of $UC_n(\Gamma)$. In order to define our HHG structure for $B_n(\Gamma)$, we will need to develop a version of parallelism for our graphical subgroups.

Definition 5.1.7. (Parallelism) Let $g\langle \Lambda, k, S \rangle$ and $h\langle \Omega, l, T \rangle$ be cosets of graphical subgroups of $B_n(\Gamma)$. We say $g\langle \Lambda, k, S \rangle$ is *parallel* to $h\langle \Omega, l, T \rangle$ if $\Lambda = \Omega$, $k = l$, S and T are in the same connected component of $UC_k(\Lambda)$, and $g^{-1}h$ is in a graphical subgroup $\langle \Pi, n, S' \rangle$, where $\Pi = \Lambda \cup (\Gamma \smallsetminus \Lambda)$ and $S' \cap \Lambda = S$. This defines an equivalence relation on the collection of graphical subgroups of $B_n(\Gamma)$. We call the equivalence classes with respect to this relation the *parallelism classes* of cosets of graphical subgroups.

Remark 5.1.8. Note that:

(1) Each graphical subgroup $\langle \Pi, n, S' \rangle$ as above splits as a product

$$\langle \Pi, n, S' \rangle \cong \langle \Lambda, k, S \rangle \times \langle \Gamma \smallsetminus \Lambda, n-k, S' \cap (\Gamma \smallsetminus \Lambda) \rangle$$

by Lemma 2.5.2. Indeed, these are the maximal product subgroups with $\langle \Lambda, k, S \rangle$ as a factor.

(2) *A priori*, $\Gamma \smallsetminus \Lambda$ may contain isolated vertices, and thus $\langle \Gamma \smallsetminus \Lambda, n-k, S' \rangle$ is not well-defined as a graphical subgroup. However, we may always remove these isolated vertices from $\Gamma \smallsetminus \Lambda$ to obtain a graph $(\Gamma \smallsetminus \Lambda)'$ which defines a genuine graphical subgroup $\langle (\Gamma \smallsetminus \Lambda)', m, S' \cap (\Gamma \smallsetminus \Lambda)' \rangle$, where $m = |S' \cap (\Gamma \smallsetminus \Lambda)'|$. Moreover, the connected component of $UC_m((\Gamma \smallsetminus \Lambda)')$ containing $S' \cap (\Gamma \smallsetminus \Lambda)'$ is isometric to the connected component of $UC_{n-k}(\Gamma \smallsetminus \Lambda)$ containing S', justifying this abuse of notation.

As each graphical subgroup $\langle \Lambda, k, S \rangle$ is isomorphic to the graph braid group $B_k(\Lambda, S)$, they also admit the following useful classification of diameter due to Genevois.

Lemma 5.1.9 (Diameter of graphical subgroups; [Gen19a, Lemma 4.3])**.** *Let Γ be a finite connected graph and let $n \in \mathbb{N}$. Let $\langle \Lambda, k, S \rangle$ be a graphical subgroup of $B_n(\Gamma)$. Then $\langle \Lambda, k, S \rangle$ has infinite diameter if and only if one of the following holds; otherwise, $\langle \Lambda, k, S \rangle$ is trivial.*

(1) $k = 1$ and the connected component of Λ containing S has a cycle subgraph.

(2) $k \geqslant 2$ and either Λ has a connected component whose intersection with S has cardinality at least 1 and which contains a cycle subgraph, or Λ has a connected component whose intersection with S has cardinality at least 2 and which contains a star subgraph.

We are now ready to define our HHG structure.

Theorem 5.1.10 (HHG structure of a graph braid group)**.** *Let Γ be a finite connected graph. The graph braid group $B_n(\Gamma)$ has an HHG structure $\mathfrak{S}$ such that:*

(1) ***(Index set.)*** *The index set $\mathfrak{S}$ consists of exactly one coset of a graphical subgroup $g\langle\Lambda, k, S\rangle$ from each parallelism class. Without loss of generality, we may take $g = e$ where possible.*

(2) ***(Nesting.)*** *Given $g\langle\Lambda, k, S\rangle, h\langle\Omega, l, T\rangle \in \mathfrak{S}$, we have $g\langle\Lambda, k, S\rangle \sqsubseteq h\langle\Omega, l, T\rangle$ if $\Lambda \subseteq \Omega$, $k \leqslant l$, S and $T \cap \Lambda$ are in the same component of $UC_k(\Lambda)$, and $g\langle\Lambda, k, S\rangle$ is parallel to $h\langle\Lambda, k, S\rangle$.*

(3) ***(Orthogonality.)*** *Given $g\langle\Lambda, k, S\rangle, h\langle\Omega, l, T\rangle \in \mathfrak{S}$, we have $g\langle\Lambda, k, S\rangle \perp h\langle\Omega, l, T\rangle$ if $\Lambda \cap \Omega = \varnothing$, $k + l \leqslant n$, and there exists $a \in B_n(\Gamma)$ such that $g\langle\Lambda, k, S\rangle$ is parallel to $a\langle\Lambda, k, S\rangle$ and $h\langle\Omega, l, T\rangle$ is parallel to $a\langle\Omega, l, T\rangle$.*

Proof. Let $\mathcal{H} = \{(e, K_e) \mid e \in E(\Gamma),\ K_e \in \pi_0(UC_{n-1}(\Gamma \smallsetminus e))\}$. Remark 5.1.3(1) tells us that $\mathcal{H}$ is in one-to-one correspondence with the collection of hyperplanes of $UC_n(\Gamma)$. Let Ξ be the crossing graph of $UC_n(\Gamma)$, so that its vertices are in one-to-one correspondence with elements of $\mathcal{H}$, and let $\mathcal{R}$ be the collection of all subgraphs of Ξ.

Take $\Omega \in \mathcal{R}$, so that we have $\Omega^{(0)} \subseteq \mathcal{H}$. Given two edges E, E' of $UC_n(\Gamma)$, write $E \sim_\Omega E'$ if there is a sequence of consecutive edges $E = E_1, E_2, \dots, E_\ell = E'$ of $UC_n(\Gamma)$ such that for each i, $(e_i, K_{e_i}) \in \Omega^{(0)}$, where e_i is the label of E_i and K_{e_i} is the connected component of $UC_{n-1}(\Gamma \smallsetminus e_i)$ containing the intersection $S_i \cap S_i'$ of the endpoints of E_i. In other words, $E \sim_\Omega E'$ if there is an edge path γ in $UC_n(\Gamma)^{(1)}$ from E to E' such that each edge of γ is dual to some hyperplane of $\Omega^{(0)}$.

Let $[E]_\Omega$ denote the equivalence class of E with respect to $\sim_\Omega$, and define $UC_n^\Omega(\Gamma)$ to be the collection of induced subcomplexes of $UC_n(\Gamma)$ whose 1–skeleton is $[E]_\Omega$ for some edge E. Let Λ_Ω be the subgraph of Γ defined by taking the union of all edges $e \in E(\Gamma)$ such that $(e, K_e) \in \Omega^{(0)}$ for some K_e. Let E be an edge of $UC_n(\Gamma)$ dual to some hyperplane $(e, K_e) \in \Omega^{(0)}$, so that $S \cap \Lambda_\Omega \neq \varnothing$ and $S \cap (\Gamma \smallsetminus e) \in K_e$ for each endpoint S of E (and moreover $S \cap \Lambda_\Omega$ does not depend on the choice of endpoint). Let $k = |S \cap \Lambda_\Omega|$. Then $[E]_\Omega$ is isometric to the 1–skeleton of the connected component of $UC_k(\Lambda_\Omega)$ containing $S \cap \Lambda_\Omega$.

Theorem 2.7.23 implies that the universal cover X of $UC_n(\Gamma)$ has a factor system $\mathfrak{F}$ consisting of all lifts of subcomplexes in $\bigcup_{\Omega \in \mathcal{R}} UC_n^\Omega(\Gamma)$, where each of these subcomplexes is isometric to a connected component of some $UC_k(\Lambda_\Omega)$. Note that $\Lambda_\Omega \subseteq \Gamma$ is expressible as a union of closed edges by construction; that is, it contains no isolated vertices. Thus, each $F \in \mathfrak{F}$ is quasi-isometric to a coset of a graphical subgroup $\langle \Lambda_\Omega, k, S \cap \Lambda_\Omega \rangle$ via the orbit map, by Remark 5.1.6. Moreover, since $\mathcal{R}$ consists of all subgraphs of Ξ, all graphical subgroups can be expressed in this form.

Applying the construction in Section 2.7.4, we obtain an HHS structure on X with index set $\mathfrak{T}$ consisting of exactly one element of $\mathfrak{F}$ from each parallelism class. Furthermore, this induces an HHG structure $\mathfrak{G}$ on $\pi_1(UC_n(\Gamma)) \cong B_n(\Gamma)$ by composing the projections in $(X, \mathfrak{T})$ with an orbit map, as explained in Remark 2.7.4. Remark 5.1.6 tells us that elements of $\mathfrak{G}$ are cosets of graphical subgroups. Moreover, the characterisation of parallelism given by Lemma 2.4.7 combined with Remark 5.1.8(1) tells us that parallelism classes in $\mathfrak{F}$ correspond to parallelism classes of graphical subgroups. Thus, $\mathfrak{G}$ consists of exactly one coset of a graphical subgroup $g\langle \Lambda, k, S \rangle$ from each parallelism class. Since there are no restrictions on the elements of parallelism classes chosen for $\mathfrak{T}$, we are free to choose the elements of $\mathfrak{G}$ so that $g = e$ wherever possible. Thus, item (1) of our theorem is satisfied. Item (2) then follows immediately from the definition of nesting for $(X, \mathfrak{T})$ given in Section 2.7.4 by expressing this definition in terms of the corresponding graphical subgroups. For item (3), note that if $g\langle \Lambda, k, S \rangle, h\langle \Omega, l, T \rangle \in \mathfrak{G}$, then Lemma 2.5.2 tells us that we have a graphical subgroup $\langle \Lambda \cup \Omega, k + l, S \cup T \rangle \cong \langle \Lambda, k, S \rangle \times \langle \Omega, l, T \rangle$ if and only if $\Lambda \cap \Omega = \varnothing$ and $k + l \leqslant n$. The result then follows from the definition of orthogonality for $(X, \mathfrak{T})$ given in Section 2.7.4. $\square$

Remark 5.1.11 (The hyperbolic spaces associated to $B_n(\Gamma)$)**.** As explained in the above proof, our HHG structure $\mathfrak{S}$ on $B_n(\Gamma) = \pi_1(UC_n(\Gamma))$ is induced by an HHS structure $\mathfrak{T}$ on the universal cover X of $UC_n(\Gamma)$, by composing the projections in $\mathfrak{T}$ with an orbit map. In particular, the hyperbolic spaces of $\mathfrak{S}$ are hyperbolic spaces of $\mathfrak{T}$. That is, if $g\langle\Lambda, k, S\rangle \in \mathfrak{S}$ is quasi-isometric to a subcomplex Y of X via the aforementioned orbit map, then $C(g\langle\Lambda, k, S\rangle) = C(Y)$, where $C(Y)$ is the factored contact graph of the cube complex Y, defined in Section 2.7.4.

Note, if one graphical subgroup is nested in another, then it embeds as a subgroup.

Lemma 5.1.12. *Let $\langle\Lambda, k, S\rangle, \langle\Omega, l, T\rangle \in \mathfrak{S}$ be graphical subgroups of $B_n(\Gamma)$ and suppose $\langle\Lambda, k, S\rangle \sqsubseteq \langle\Omega, l, T\rangle$. Then $\langle\Lambda, k, S\rangle \leqslant \langle\Omega, l, T\rangle$.*

Proof. Let K_Λ denote the connected component of $UC_k(\Lambda)$ containing S, and let K_Ω denote the connected component of $UC_l(\Omega)$ containing T. Then $\langle\Lambda, k, S\rangle \cong \pi_1(K_\Lambda)$ and $\langle\Omega, l, T\rangle \cong \pi_1(K_\Omega)$. If $\langle\Lambda, k, S\rangle \sqsubseteq \langle\Omega, l, T\rangle$, then $\Lambda \subseteq \Omega$, $k \leqslant l$, and S and $T \cap \Lambda$ are in the same connected component of $UC_k(\Lambda)$. Thus, we obtain an isometric embedding of K_Λ into K_Ω by fixing the $l - k$ particles in $T \cap (\Omega \smallsetminus \Lambda)$ and restricting the motion of the remaining k particles to Λ. Furthermore, by [HW08, Lemma 2.11] this embedding is π_1–injective, thus induces an embedding of $\langle\Lambda, k, S\rangle$ as a subgroup of $\langle\Omega, l, T\rangle$. □

By applying the following theorem of Behrstock–Hagen–Sisto, we may modify the HHS structure $(B_n(\Gamma), \mathfrak{S})$ by removing any domains with finite diameter in $B_n(\Gamma)$. In the statement of the theorem below, the space F_U refers to one of the factors of the *standard product region* associated to a domain U. We shall not go into any details about these; we refer the reader to [BHS19, Section 5.2] for more information. The important point to note is that if $U \in \mathfrak{S}$ is a graphical subgroup, then F_U is quasi-isometric to U itself, by [BHS17b, Remark 13.5].

Theorem 5.1.13 ([BHS17a, Proposition 2.4]). *Let $(X, \mathfrak{S})$ be an HHS, and let $\mathfrak{U} \subseteq \mathfrak{S}$ be closed under nesting. Suppose there exists $D > 0$ such that $\operatorname{diam}(F_U) \leqslant D$ for each $U \in \mathfrak{U}$. Then $(X, \mathfrak{S} \smallsetminus \mathfrak{U})$ is an HHS, where the associated $C(*), \pi_*, \rho^*_*, \sqsubseteq, \perp, \pitchfork$ are the same as in the original structure.*

Corollary 5.1.14. *Let $\mathfrak{U} \subseteq \mathfrak{S}$ be the collection of domains $g\langle \Lambda, k, S\rangle$ with finite diameter in $B_n(\Gamma)$. Then $(B_n(\Gamma), \mathfrak{S} \smallsetminus \mathfrak{U})$ is an HHS.*

Proof. Let $g\langle \Lambda, k, S\rangle \in \mathfrak{U}$ and suppose $h\langle \Omega, l, T\rangle \sqsubseteq g\langle \Lambda, k, S\rangle$. By definition of nesting, we have $\Omega \subseteq \Lambda$, $l \leqslant k$, and T and $S \cap \Omega$ are in the same connected component of $UC_l(\Omega)$. Thus, $\langle \Omega, l, T\rangle$ embeds as a subgroup of $\langle \Lambda, k, S\rangle$ by Lemma 5.1.12. Since $g\langle \Lambda, k, S\rangle$ has finite diameter, it therefore follows that $h\langle \Omega, l, T\rangle$ has finite diameter too, and so $h\langle \Omega, l, T\rangle \in \mathfrak{U}$. Thus, $\mathfrak{U}$ is closed under nesting. Furthermore, the bound on the diameter of graphical subgroups in $\mathfrak{U}$ is uniform by Lemma 5.1.9 (in fact, it is 0). This in turn uniformly bounds the diameter of the spaces F_U for $U \in \mathfrak{U}$, as F_U is quasi-isometric to U by [BHS17b, Remark 13.5]. We can therefore apply Theorem 5.1.13 to conclude that $(B_n(\Gamma), \mathfrak{S} \smallsetminus \mathfrak{U})$ is an HHS. □

In Section 5.2.2 we shall take advantage of this result to assume that every domain $g\langle \Lambda, k, S\rangle \in \mathfrak{S}$ has infinite diameter in $B_n(\Gamma)$.

5.2 Detecting other forms of hyperbolicity in a graph braid group

In this section, we classify when a graph braid group is hyperbolic or acylindrically hyperbolic, and provide a conjectural classification of relative hyperbolicity and thickness. This builds upon results of Genevois, who obtained a classification of hyperbolicity, acylindrical hyperbolicity, and toral relative hyperbolicity [Gen19a]. In the hyperbolic case, we use the bounded orthogonality criterion (Theorem 2.7.12) to recover a version of Genevois' theorem.

In the acylindrically hyperbolic case, we use Behrstock–Hagen–Sisto's criteria for acylindrical hyperbolicity in HHGs (Theorem 2.7.13). In the relatively hyperbolic case, we adapt techniques developed by Levcovitz in his classification of relative hyperbolicity and thickness for right-angled Coxeter groups [Lev20]. In particular, we introduce a sequence of hypergraphs which encode collections of mutually orthogonal domains arising in the HHG structure of a graph braid group $B_n(\Gamma)$. By analysing connectedness properties of these hypergraphs and applying Russell's isolated orthogonality criterion (Theorem 2.7.15), we conjecture a characterisation of when the graph braid group is relatively hyperbolic. By construction, our hypergraphs show that any graph braid group which does not satisfy the isolated orthogonality criterion is in fact strongly thick, and moreover we obtain an upper bound on the order of thickness.

5.2.1 Hyperbolicity and acylindrical hyperbolicity

The HHG structure $(B_n(\Gamma), \mathfrak{S})$ allows us to easily detect when the graph braid group is hyperbolic or acylindrically hyperbolic. Indeed, the bounded orthogonality criterion (Theorem 2.7.12) allows us to obtain a classification of hyperbolicity of graph braid groups, giving an alternate proof of a theorem of Genevois.

Theorem 5.2.1 (Characterisation of hyperbolicity; [Gen19a, Theorem 4.1])**.** *Let Γ be a finite connected graph and let $n \in \mathbb{N}$. The graph braid group $B_n(\Gamma)$ is hyperbolic if and only if one of the following holds.*

(1) $n = 1$.

(2) $n = 2$ *and* Γ *does not contain two disjoint cycle subgraphs.*

(3) $n = 3$ *and* Γ *does not contain two disjoint cycle subgraphs, nor does it contain a disjoint star subgraph and cycle subgraph.*

(4) $n \geqslant 4$ *and* Γ *does not contain two disjoint subgraphs, each of which is a star or a cycle.*

In order to prove this theorem, we modify the HHG structure on $B_n(\Gamma)$ slightly by choosing a smaller factor system for the universal cover X of $UC_n(\Gamma)$, obtained by closing the set of subcomplexes parallel to combinatorial hyperplanes under large projections. Using our characterisation of combinatorial hyperplanes (Remark 5.1.3), we see that this gives us a smaller index set $\mathfrak{S}' \subseteq \mathfrak{S}$ for $B_n(\Gamma)$, consisting of cosets of graphical subgroups of the form $\langle \Gamma \smallsetminus (e_1 \cup \cdots \cup e_m), k, S \rangle$ for some (possible empty) set of disjoint closed edges $e_1, \ldots, e_m$ of Γ. Recall that $\Gamma \smallsetminus e$ denotes the induced subgraph of Γ spanned by the vertices $V(\Gamma) \smallsetminus V(e)$. Using this new HHG structure $(B_n(\Gamma), \mathfrak{S}')$, we have the following result.

Lemma 5.2.2. *Let $\langle \Lambda, k, S \rangle \in \mathfrak{S}'$. The hyperbolic space $C(\langle \Lambda, k, S \rangle)$ is unbounded if and only if Λ is connected and $\langle \Lambda, k, S \rangle$ has infinite diameter.*

Proof. Suppose $\langle \Lambda, k, S \rangle$ has infinite diameter and suppose Λ is connected. We claim that the universal cover Y of $UC_k(\Lambda)$ does not split as a product of subcomplexes.

Claim 5.2.3. Suppose Λ is connected. Then the universal cover Y of $UC_k(\Lambda)$ does not split as a product of subcomplexes.

Proof of claim. Suppose Y splits as a direct product $Y = Y_1 \times Y_2$. Recall that each edge of $UC_k(\Lambda)$ is labelled by a closed edge of Λ, two adjacent edges of $UC_k(\Lambda)$ span a square if and only if they are labelled by disjoint edges of Λ, and opposite edges of a square are labelled by the same edge of Λ. Thus, we may also label edges of Y by closed edges of Λ by projecting to $UC_k(\Lambda)$, and moreover if two edges of Y span a square then they are labelled by disjoint edges of Λ. Furthermore, since every edge of $Y = Y_1 \times Y_2$ has the form $E_1 \times \{y_2\}$ or $\{y_1\} \times E_2$ where E_i is an edge of Y_i and y_i is a vertex of Y_i, the labelling of edges of Y induces a labelling of edges of Y_i for $i = 1, 2$ (this induced labelling is well-defined since opposite edges of a square in Y have the same labels). Let Λ_i be the subgraph of Λ spanned by the edge labels of Y_i, for $i = 1, 2$. Then $\Lambda_1 \cup \Lambda_2 = \Lambda$. Moreover, every edge of Λ_1 must be disjoint from every edge of Λ_2, since every edge of Y_1 spans a square in Y with every edge of Y_2. This contradicts connectedness of Λ. Thus, Y does not split as a product of subcomplexes. □

A result of Hagen now tells us that the contact graph $C_0(Y)$ must be unbounded [Hag12, Theorem 6.3.6]. (Note, we can ensure that Y is leafless by applying Caprace–Sageev's pruning procedure, which will not affect $\langle\Lambda, k, S\rangle$ [CS11].) Furthermore, $C(\langle\Lambda, k, S\rangle)$ is the factored contact graph $C(Y)$ by Remark 5.1.11, which is quasi-isometric to the contact graph $C_0(Y)$ by [BHS17b, Remark 8.18]. Thus, $C(\langle\Lambda, k, S\rangle)$ is unbounded.

Conversely, suppose either $\langle\Lambda, k, S\rangle$ has finite diameter or Λ is disconnected. If $\langle\Lambda, k, S\rangle$ has finite diameter then it must be trivial by Lemma 5.1.9. Thus, the cube complex Y has finitely many hyperplanes and hence $C(Y) = C(\langle\Lambda, k, S\rangle)$ is bounded. If Λ is disconnected then Y splits as a product by Lemma 2.5.2. Thus, the factored contact graph $C(Y) = C(\langle\Lambda, k, S\rangle)$ is bounded by [Hag12, Theorem 6.2.3]. □

Proof of Theorem 5.2.1. We wish to use the bounded orthogonality criterion (Theorem 2.7.12) on $\mathfrak{S}'$ to classify hyperbolicity of $B_n(\Gamma)$. Lemma 5.2.2 tells us that $C(\langle\Lambda, k, S\rangle)$ is unbounded if and only if $\langle\Lambda, k, S\rangle \in \mathfrak{S}'$ is an infinite-diameter graphical subgroup of $B_n(\Gamma)$ with Λ connected. Moreover, there are finitely many subgraphs $\Lambda \subseteq \Gamma$, finitely many $k \leqslant n$, and finitely many base points $S \in UC_k^{(0)}(\Lambda)$, therefore the bounded hyperbolic spaces can be bounded uniformly. The bounded orthogonality criterion therefore says that $B_n(\Gamma)$ is hyperbolic if and only if there do not exist two infinite-diameter graphical subgroups $\langle\Lambda_1, k_1, S_1\rangle, \langle\Lambda_2, k_2, S_2\rangle \in \mathfrak{S}'$ with Λ_1, Λ_2 connected, $\Lambda_1 \cap \Lambda_2 = \varnothing$, and $k_1 + k_2 \leqslant n$. We classify these by applying the characterisation of diameter of graphical subgroups (Lemma 5.1.9).

Lemma 5.1.9 tells us that if Λ is connected, then $\langle\Lambda, k, S\rangle$ has infinite diameter if and only if either $k = 1$ and Λ contains a cycle subgraph, or $k \geqslant 2$ and Λ contains a cycle or star subgraph. Note that if there exist two disjoint cycle or star subgraphs Ω_1, Ω_2 of Γ, then by subdividing edges of Γ sufficiently, we can always find two connected disjoint subgraphs

Λ_1, Λ_2 of Γ of the form $\Gamma \smallsetminus (e_1 \cup \cdots \cup e_m)$ such that $\Omega_i \subseteq \Lambda_i$ for $i = 1, 2$. That is, we can find $\langle \Lambda_1, k_1, S_1 \rangle, \langle \Lambda_2, k_2, S_2 \rangle \in \mathfrak{S}'$ with Λ_i connected, disjoint, and $\Omega_i \subseteq \Lambda_i$ for $i = 1, 2$. The desired characterisation of hyperbolicity of $B_n(\Gamma)$ therefore follows from Lemma 5.1.9 by analysing when $k_1 + k_2 \leqslant n$. □

The criteria for acylindrical hyperbolicity for HHGs (Theorem 2.7.13) allow us to recover another theorem of Genevois regarding acylindrical hyperbolicity.

Theorem 5.2.4 ([Gen19a, Theorem 4.10])**.** *Let Γ be a finite connected graph and let $n \in \mathbb{N}$. The graph braid group $B_n(\Gamma)$ is either cyclic or acylindrically hyperbolic.*

Proof. Suppose $B_n(\Gamma)$ is non-trivial; by Lemma 5.1.9 it must have infinite diameter. Let $(B_n(\Gamma), \mathfrak{S}')$ be the HHG structure given above Lemma 5.2.2, and let X be the universal cover of $UC_n(\Gamma)$. The $\sqsubseteq$–maximal element $S \in \mathfrak{S}'$ is the entire graph braid group $S = B_n(\Gamma)$. By Lemma 5.2.2, $C(S)$ must be unbounded since Γ is connected and $B_n(\Gamma)$ has infinite diameter. Furthermore, as $B_n(\Gamma)$ is torsion-free by [Abr00, Corollary 3.7], it is virtually cyclic if and only if it is cyclic. The criteria for acylindrical hyperbolicity for HHGs (Theorem 2.7.13) therefore imply that $B_n(\Gamma)$ is either cyclic or acylindrically hyperbolic. □

5.2.2 Relative hyperbolicity and thickness

We wish to classify when a graph braid group $B_n(\Gamma)$ is relatively hyperbolic or thick. We do this by introducing a sequence of hypergraphs $\mathcal{O}_k$ called the *orthogonality graphs*, which encode collections of mutually orthogonal domains arising in the HHS structure of $B_n(\Gamma)$. These mimic the hypergraphs employed by Levcovitz in characterising relative hyperbolicity and thickness of right-angled Coxeter groups [Lev19, Lev20].

Note that we may assume that all graphical subgroups in $\mathfrak{S}$ have infinite diameter, as discussed at the end of Section 5.1 (see Corollary 5.1.14). The graphical subgroups given in the definition below can therefore all be assumed to have infinite diameter.

Definition 5.2.5 (Orthogonality graph, hypergraph index)**.** Let Γ be a finite graph and let $n \in \mathbb{N}$. The *orthogonality graphs* $\mathcal{O}_i = \mathcal{O}_i(\Gamma, n)$ of the graph braid group $B_n(\Gamma)$ are hypergraphs, defined inductively as follows.

(1) $\mathcal{O}_0$ is the hypergraph whose vertices are proper graphical subgroups $\langle \Lambda, k, S\rangle \in \mathfrak{S}$, and whose hyperedges are maximal collections $\{\langle \Lambda_1, k_1, S_1\rangle, \ldots, \langle \Lambda_m, k_m, S_m\rangle\}$ of pairwise orthogonal domains. Given a hyperedge $E = \{\langle \Lambda_1, k_1, S_1\rangle, \ldots, \langle \Lambda_m, k_m, S_m\rangle\}$, let $\langle E\rangle$ be the subgroup generated by $\langle \Lambda_1, k_1, S_1\rangle, \ldots, \langle \Lambda_m, k_m, S_m\rangle$.

(2) Define an equivalence relation $\equiv_i$ on the set of hyperedges $\mathcal{E}(\mathcal{O}_i)$ by setting $E \equiv_i E'$ if there exists a sequence $E = E_1, \ldots, E_m = E'$ of hyperedges in $\mathcal{E}(\mathcal{O}_i)$ such that for each $1 \leqslant j < m$, there exists $\langle \Lambda, k, S\rangle \in \mathfrak{S}$ and $\langle \Omega_j, l_j, T_j\rangle \in E_j$, $\langle \Omega_{j+1}, l_{j+1}, T_{j+1}\rangle \in E_{j+1}$ with $\langle \Lambda, k, S\rangle \sqsubseteq \langle \Omega_j, l_j, T_j\rangle$ and $\langle \Lambda, k, S\rangle \sqsubseteq \langle \Omega_{j+1}, l_{j+1}, T_{j+1}\rangle$.

(3) For each $i \geqslant 0$, define $V(\mathcal{O}_{i+1}) = V(\mathcal{O}_0)$ and define $E \subseteq V(\mathcal{O}_0)$ to be a hyperedge of $\mathcal{O}_{i+1}$ if and only if $E = E_1 \cup \cdots \cup E_m$ for some maximal collection $\{E_1, \ldots, E_m\}$ of $\equiv_i$–equivalent hyperedges of $\mathcal{O}_i$. Let $\langle E\rangle$ be the subgroup generated by $\langle E_1\rangle, \ldots, \langle E_m\rangle$.

We define the *hypergraph index* of $B_n(\Gamma)$ to be the smallest integer i such that $\mathcal{O}_i$ contains a hyperedge E with $\langle E\rangle = B_n(\Gamma)$. If no such i exists, then the we define the hypergraph index to be ∞.

Recall that the characterisation of diameter of graphical subgroups (Lemma 5.1.9) tells us that a graphical subgroup $\langle \Lambda, k, S\rangle$ has infinite diameter if and only if Λ contains a cycle or star subgraph in one of its components and S has at least one particle in that component (in the cycle case) or at least two particles in that component (in the star case). Since we are considering all domains of $\mathfrak{S}$ to have infinite diameter, one may therefore interpret the hypergraph index as a measure of complexity of intersection patterns of cycle and star subgraphs occurring in Γ. For example, if $B_n(\Gamma)$ has hypergraph index 0, then there exists a collection $\{\Lambda_1, \ldots, \Lambda_m\}$ of disjoint subgraphs of Γ such that each cycle or star subgraph of

Γ is contained in some Λ_j. If $B_n(\Gamma)$ has hypergraph index 1, then there exists a collection $\{\Lambda_1, \ldots, \Lambda_m\}$ of subgraphs of Γ such that: each cycle or star subgraph of Γ is contained in some Λ_j; for each j there is a collection of disjoint subgraphs of Λ_j containing all cycle and star subgraphs of Λ_j; and for any $j \neq j'$ there exists a sequence $\Lambda_j = \Lambda_{j_1}, \ldots, \Lambda_{j_r} = \Lambda_{j'}$ such that $\Lambda_{j_i} \cap \Lambda_{j_{i+1}}$ contains a cycle or star subgraph for each i. One must also be careful to keep track of the number of particles in each subgraph.

We wish to classify relative hyperbolicity and thickness of $B_n(\Gamma)$ in terms of this complexity. We claim that if the hypergraph index of $B_n(\Gamma)$ is $k < \infty$, then $B_n(\Gamma)$ is strongly thick of order k, and if the hypergraph index is ∞, then $B_n(\Gamma)$ is relatively hyperbolic. To prove the latter, we claim that if the hypergraph index is ∞, then there exists some i such that the hyperedges $E \in \mathcal{E}(\mathcal{O}_i)$ isolate orthogonality of $\mathfrak{S}$ in the sense of Definition 2.7.14. However, *a priori*, the subgroups $\langle E \rangle$ may not themselves be graphical subgroups, and thus may not be domains in $\mathfrak{S}$. We conjecture that the subgroups $\langle E \rangle$ are indeed graphical subgroups.

Conjecture 5.2.6. Let Γ be a finite connected graph and let $n \in \mathbb{N}$. For each $i \geqslant 0$ and each hyperedge $E \in \mathcal{E}(\mathcal{O}_i)$, the subgroup $\langle E \rangle$ of $B_n(\Gamma)$ is a graphical subgroup. Furthermore, if $E = \{\langle \Lambda_1, k_1, S_1 \rangle, \ldots, \langle \Lambda_m, k_m, S_m \rangle\}$, then $\langle \Lambda_j, k_j, S_j \rangle \sqsubseteq \langle E \rangle$ for each $1 \leqslant j \leqslant m$.

Suppose $E \in \mathcal{E}(\mathcal{O}_i)$ has vertex set $\{\langle \Lambda_1, k_1, S_1 \rangle, \ldots, \langle \Lambda_m, k_m, S_m \rangle\}$. The naive approach to this conjecture would be to show that $\langle E \rangle$ is a graphical subgroup of the form $\langle \bigcup_j \Lambda_j, k, S \rangle$ for some $k \geqslant \max\{k_1, \ldots, k_m\}$ and some initial configuration S such that $S \cap \Lambda_j$ is equivalent to S_j in the sense of Remark 5.1.5 for all j. The latter two conditions ensure that $\langle \Lambda_j, k_j, S_j \rangle \sqsubseteq \langle \bigcup_j \Lambda_j, k, S \rangle$ for all j. However, we already see that this fails for hyperedges of $\mathcal{O}_0$. Indeed, suppose Γ contains m disjoint cycles $\Lambda_1, \ldots, \Lambda_m$, where $m > n \geqslant 2$. Then the graphical subgroups $\langle \Lambda_j, 1, S_j \rangle$ (where S_j is any vertex of Λ_j) are pairwise orthogonal and thus are contained in a common hyperedge $E \in \mathcal{E}(\mathcal{O}_0)$. Suppose $\langle E \rangle$ is of the form $\langle \bigcup_j \Lambda_j, k, S \rangle$ described above. Since we only have $n < m$ particles in $B_n(\Gamma)$, there must be some $1 \leqslant N \leqslant m$ such that the initial configuration S does not have a particle in Λ_N. But then $S \cap \Lambda_N = \varnothing$, contradicting our assumption that $S \cap \Lambda_N$ is equivalent to S_N. In particular, Lemma 2.5.2 implies that $\langle \bigcup_j \Lambda_j, k, S \rangle \cong \langle \bigcup_{j=1}^{N-1} \Lambda_j \cup \bigcup_{j=N+1}^{m} \Lambda_j, k, S \rangle$. However, $\langle \Lambda_N, 1, S_N \rangle$ is not a subgroup of this, and therefore $\langle E \rangle$ cannot have this form.

In order to solve this problem, one may try to define a subgraph Λ of Γ by 'connecting up' the subgraphs Λ_j so that particles may move freely between them, reducing the number of particles required for the graphical subgroup $\langle \Lambda, k, S \rangle$ to contain all subgroups $\langle \Lambda_j, k_j, S_j \rangle$. Note that since Γ is connected, we can always connect two subgraphs Λ_j with a path in Γ. We claim that $\langle E \rangle$ is a graphical subgroup of the form $\langle \Lambda, k, S \rangle$, where Λ is obtained by 'connecting up' some of the subgraphs Λ_j in a minimal way, $k \geqslant \max\{k_1, \ldots, k_m\}$, and $S \cap \Lambda_j$ is equivalent to S_j in the sense of Remark 5.1.5 for all j. Again, the latter two conditions ensure that $\langle \Lambda_j, k_j, S_j \rangle \sqsubseteq \langle \Lambda, k, S \rangle$ for all j.

Assuming Conjecture 5.2.6 is true, we are able to prove our classification theorem.

Theorem 5.2.7. *Let Γ be a finite graph and let $n \geqslant 1$, $k \geqslant 0$ be integers. Suppose Conjecture 5.2.6 is true.*

(1) If $B_n(\Gamma)$ has hypergraph index k, then $B_n(\Gamma)$ is strongly thick of order at most k. In particular, $B_n(\Gamma)$ is not relatively hyperbolic.

(2) If $B_n(\Gamma)$ has hypergraph index ∞, then $B_n(\Gamma)$ is relatively hyperbolic.

Proof. Suppose Γ is disconnected. Then Lemma 2.5.2 tell us that $B_n(\Gamma, S)$ splits as a direct product

$$B_n(\Gamma, S) \cong B_{n_1}(\Lambda_1) \times \cdots \times B_{n_d}(\Lambda_d),$$

where $\Lambda_1, \ldots, \Lambda_d$ are the connected components of Γ and $n_i = |S \cap \Lambda_i|$. Furthermore, by Lemma 5.1.9 $B_{n_i}(\Lambda_i)$ is either infinite or trivial for each i. If there exist $i \neq j$ such that $B_{n_i}(\Lambda_i)$ and $B_{n_j}(\Lambda_j)$ have infinite diameter, then there exists $E \in \mathcal{E}(\mathcal{O}_0)$ with $\langle E \rangle = B_n(\Gamma)$; take E to be the collection of graphical subgroups given by those $B_{n_i}(\Lambda_i)$ with infinite diameter. Thus, $B_n(\Gamma)$ has hypergraph index 0. Moreover, $B_n(\Gamma)$ splits as a product with infinite factors and is therefore strongly algebraically thick of order 0. On the other hand, if there exists precisely one $1 \leqslant i \leqslant d$ such that $B_{n_i}(\Lambda_i)$ has infinite diameter, then $B_n(\Gamma, S) \cong B_{n_i}(\Lambda_i)$ and so the proof reduces to the connected case. If $B_{n_i}(\Lambda_i)$ is trivial for all i, then $B_n(\Gamma, S)$ is trivial. We may therefore assume henceforth that Γ is connected.

Proof of (1). Suppose $B_n(\Gamma)$ has hypergraph index k. Then there exists a hyperedge $E_{max} \in \mathcal{E}(\mathcal{O}_k)$ such that $\langle E_{max} \rangle = B_n(\Gamma)$. Let $0 \leqslant i \leqslant k$. We prove by induction on i that for each $E \in \mathcal{E}(\mathcal{O}_i)$, the subgroup $\langle E \rangle$ is strongly algebraically thick of order at most i. In particular, this implies that $B_n(\Gamma)$ is strongly algebraically thick of order at most k.

In the base case of $E \in \mathcal{E}(\mathcal{O}_0)$, we have $E = \{\langle \Lambda_1, k_1, S_1 \rangle, \ldots, \langle \Lambda_m, k_m, S_m \rangle\}$ for some maximal collection of pairwise orthogonal domains in $\mathfrak{S}$, which are assumed to have infinite diameter by Corollary 5.1.14. Moreover, $\langle E \rangle$ is the subgroup of $B_n(\Gamma)$ generated by the graphical subgroups $\langle \Lambda_j, k_j, S_j \rangle$. By definition of orthogonality, $\Lambda_j \cap \Lambda_r = \varnothing$ and $k_j + k_r \leqslant n$ for all $j \neq r$. Since Γ is connected, there exists some Λ_r and some connected subgraph $\Omega \subseteq \Gamma$ such that $\Lambda_r \cap \Omega = \varnothing$ and $\Lambda_j \subseteq \Omega$ for all $j \neq r$. Let $l = \max_{j \neq r}\{k_j\}$ and let T be any configuration of l particles in Ω. Then we have $\langle \Lambda_j, k_j, S_j \rangle \sqsubseteq \langle \Omega, l, T \rangle$ for all $j \neq r$, and hence $\langle \Lambda_j, k_j, S_j \rangle \leqslant \langle \Omega, l, T \rangle$ for all $j \neq r$ by Lemma 5.1.12. Moreover, since $k_r + l \leqslant n$, we have a graphical subgroup $\langle \Lambda_r \cup \Omega, k_r + l, S_r \cup T \rangle$ which splits as the direct product $\langle \Lambda_r, k_r, S_r \rangle \times \langle \Omega, l, T \rangle$. We therefore have $\langle E \rangle \leqslant \langle \Lambda_r, k_r, S_r \rangle \times \langle \Omega, l, T \rangle$, so $\langle E \rangle$ splits as a direct product with infinite factors, and hence is strongly algebraically thick of order 0.

Now let $E \in \mathcal{E}(\mathcal{O}_i)$ for $i > 0$, and suppose that for all $E' \in \mathcal{E}(\mathcal{O}_{i-1})$, the subgroup $\langle E' \rangle$ is strongly algebraically thick of order $i - 1$. By definition of $\mathcal{E}(\mathcal{O}_i)$, there is a finite collection of $\equiv_{i-1}$–equivalent hyperedges $\{E_\alpha\}_{\alpha \in I} \subseteq \mathcal{E}(\mathcal{O}_{i-1})$ such that the subgroups $\langle E_\alpha \rangle$ generate $\langle E \rangle$. These subgroups therefore satisfy the coarse covering condition in the definition of strong algebraic thickness (Definition 2.2.7), and are themselves strongly algebraically thick of order $i - 1$ by the inductive hypothesis. Furthermore, each subgroup $\langle E_\alpha \rangle$ is quasi-convex in the sense of Definition 2.1.4. Indeed, each $\langle E_\alpha \rangle$ is a graphical subgroup by Conjecture

5.2.6, and thus is a domain in the HHS structure $\mathfrak{S}$. That is, $\langle E_\alpha \rangle$ is quasi-isometric to an element F of a factor system for the universal cover X of $UC_n(\Gamma)$, via the orbit map. Since elements of factor systems are convex subcomplexes, it follows that $\langle E_\alpha \rangle$ is quasi-convex.

Now consider two subgroups $\langle E_\alpha \rangle$ and $\langle E_{\alpha'} \rangle$ for $\alpha, \alpha' \in I$. Since $E_\alpha \equiv_{i-1} E_{\alpha'}$, item (2) of the definition of the orthogonality graph $\mathcal{O}_{i-1}$ tells us there exists a sequence $E_\alpha = E_1, \ldots, E_m = E_{\alpha'}$ of hyperedges in $\mathcal{E}(\mathcal{O}_{i-1})$ such that for each $1 \leqslant j < m$, there exists $\langle \Lambda, r, S \rangle \in \mathfrak{S}$ and $\langle \Omega_j, l_j, T_j \rangle \in E_j$, $\langle \Omega_{j+1}, l_{j+1}, T_{j+1} \rangle \in E_{j+1}$ with $\langle \Lambda, r, S \rangle \sqsubseteq \langle \Omega_j, l_j, T_j \rangle$ and $\langle \Lambda, r, S \rangle \sqsubseteq \langle \Omega_{j+1}, l_{j+1}, T_{j+1} \rangle$. By Lemma 5.1.12, we therefore have $\langle \Lambda, r, S \rangle \leqslant \langle \Omega_j, l_j, T_j \rangle \leqslant \langle E_j \rangle$ and $\langle \Lambda, r, S \rangle \leqslant \langle \Omega_{j+1}, l_{j+1}, T_{j+1} \rangle \leqslant \langle E_{j+1} \rangle$ for each j. Since we are working under the assumption that all graphical subgroups in $\mathfrak{S}$ have infinite diameter (see Corollary 5.1.14), this means $\langle E_j \rangle \cap \langle E_{j+1} \rangle$ has infinite diameter for each j. Thus, the thick chaining condition in the definition of strong algebraic thickness is satisfied, and hence $\langle E \rangle$ is strongly algebraically thick of order at most i.

We conclude that $B_n(\Gamma) = \langle E_{max} \rangle$ is strongly algebraically thick of order at most k, and hence $B_n(\Gamma)$ is strongly thick of order at most k by [BDM09, Proposition 7.6]. In particular, [BDM09, Corollary 7.9] tells us $B_n(\Gamma)$ is not relatively hyperbolic. □

Proof of (2). Suppose $B_n(\Gamma)$ has hypergraph index ∞. We wish to show that $\mathfrak{S}$ satisfies the isolated orthogonality conditions (Definition 2.7.14), implying relative hyperbolicity of $B_n(\Gamma)$. In particular, we claim that there exists some $i \geqslant 0$ such that the collection

$$\mathcal{I}_i = \{\langle E \rangle \in \mathfrak{S} \mid E \in \mathcal{E}(\mathcal{O}_i)\}$$

satisfies the isolated orthogonality conditions.

Every pair of orthogonal domains of $\mathfrak{S}$ is contained in some hyperedge $E \in \mathcal{E}(\mathcal{O}_0)$ by definition, thus Conjecture 5.2.6 implies every pair of orthogonal domains is nested into some domain $\langle E \rangle$ for $E \in \mathcal{E}(\mathcal{O}_0)$. Since each hyperedge of $\mathcal{O}_{i-1}$ is contained in a hyperedge of $\mathcal{O}_i$

for each $i \geqslant 1$ by Definition 5.2.5(3), it follows that every pair of orthogonal domains of $\mathfrak{S}$ is nested into some domain $\langle E \rangle$ for $E \in \mathcal{E}(\mathcal{O}_i)$ for each i. Thus, $\mathcal{I}_i$ satisfies condition (1) of isolated orthogonality for all i.

Suppose $E_1 \not\equiv_i E_2$ for all pairs of hyperedges $E_1, E_2 \in \mathcal{E}(\mathcal{O}_i)$. Then it follows from Definition 5.2.5(2) that no domain of $\mathfrak{S}$ is nested into two domains of $\mathcal{I}_i$. Thus, $\mathcal{I}_i$ also satisfies condition (2) of isolated orthogonality. On the other hand, if there exist non-trivial collections of $\equiv_i$–equivalent hyperedges of $\mathcal{O}_i$, then there exists a hyperedge of $\mathcal{O}_{i+1}$ which strictly contains some hyperedge of $\mathcal{O}_i$. Note that each $\mathcal{O}_i$ has the same finite number of vertices: $V(\mathcal{O}_i)$ consists of all proper infinite-diameter graphical subgroups of the form $\langle \Lambda, k, S \rangle$, for which there are finitely many choices of subgraphs $\Lambda \subseteq \Gamma$, integers $1 \leqslant k \leqslant n$, and base points $S \in UC_k(\Lambda)^{(0)}$. It therefore follows by induction that either there exists some i such that $\mathcal{I}_i$ satisfies the isolated orthogonality conditions or there exists i such that $\mathcal{O}_i$ has a hyperedge E with $E = V(\mathcal{O}_i)$, and hence $\langle E \rangle = B_n(\Gamma)$. Since the hypergraph index of $B_n(\Gamma)$ is ∞, the former must be true. □

This concludes the proof of Theorem 5.2.7. □

As an immediate corollary, we have that graph braid groups form another example of a class of HHGs which satisfy the dichotomy between thickness and relative hyperbolicity.

Corollary 5.2.8. *Let Γ be a finite graph and suppose Conjecture 5.2.6 is true. The graph braid group $B_n(\Gamma)$ is strongly thick if and only if it is not relatively hyperbolic.*

We also conjecture that a stronger version of Theorem 5.2.7 is true, analogous to Levcovitz's characterisation of thickness in right-angled Coxeter groups [Lev20, Theorem A], which is stated in Section 2.4.2 as Theorem 2.4.26. Levcovitz obtains a lower bound on the order of strong thickness by studying the divergence of right-angled Coxeter groups and using Behrstock–Druţu's relationship between divergence and strong thickness (see Theorem 2.2.9 or [BD14, Corollary 4.17]). In particular, Levcovitz makes use of disc diagram techniques in order to construct geodesics with polynomial divergence of degree $k+1$ in right-angled Coxeter groups with hypergraph index k. We conjecture that similar disc diagram techniques may be developed for the universal cover of the cube complex $UC_n(\Gamma)$, giving the following result.

Conjecture 5.2.9. Let Γ be a finite connected graph and let $n \geqslant 1$, $k \geqslant 0$ be integers.

- If $B_n(\Gamma)$ has hypergraph index k, then $B_n(\Gamma)$ is strongly thick of order k and has polynomial divergence of degree $k+1$.
- If Γ has hypergraph index ∞, then $B_n(\Gamma)$ is relatively hyperbolic. Moreover, $B_n(\Gamma)$ has exponential divergence if it is one-ended, and infinite divergence otherwise.

Bibliography

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

[Abr00] Aaron Abrams. Configuration spaces and braid groups of graphs. *PhD thesis, University of California, Berkeley*, 2000.

[Art26] Emil Artin. Theorie der Zöpfe. *Abh. Math. Sem. Hamburg*, 4:47–72, 1926.

[BC12] Jason Behrstock and Ruth Charney. Divergence and quasi-morphisms of right-angled Artin groups. *Math. Ann.*, 352(2):339–356, 2012.

[BD14] Jason Behrstock and Cornelia Druţu. Divergence, thick groups, and short conjugators. *Illinois J. Math.*, 58(4):939–980, 2014.

[BDM09] Jason Behrstock, Cornelia Druţu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. *Math. Ann.*, 344(3):543–595, 2009.

[Beh06] Jason A. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geom. Topol.*, 10:1523–1578, 2006.

[BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)*, 114(5):890–926, 2017.

[BHS17b] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[BHS17c] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *arXiv:1704.04271*, 2017.

[BHS17d] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Thickness, relative hyperbolicity, and randomness in Coxeter groups. *Algebr. Geom. Topol.*, 17(2):705–740, 2017.

[BHS19] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299:257–338, 2019.

[Bir69] Joan Birman. On braid groups. *Communications on Pure and Applied Mathematics*, 22:41–72, 1969.

[BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.

[BM08] Jeffrey Brock and Howard Masur. Coarse and synthetic Weil-Petersson geometry: quasi-flats, geodesics and relative hyperbolicity. *Geom. Topol.*, 12(4):2453–2495, 2008.

[BMW94] Hans-Jürgen Bandelt, Henry Mulder, and Elke Wilkeit. Quasi-median graphs and algebras. *Journal of Graph Theory*, 18(7):681–703, 1994.

[Bow08] Brian Bowditch. Tight geodesics in the curve complex. *Invent. Math.*, 171(2):281–300, 2008.

[Bow12] Brian Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[BPR19] Hyungryul Baik, Bram Petri, and Jean Raimbault. Subgroup growth of virtually cyclic right-angled Coxeter groups and their free products. *Combinatorica*, 39(4):779–811, 2019.

[BR18] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *arXiv:1810.06476*, 2018.

[BR20] Daniel Berlyne and Jacob Russell. Hierarchical hyperbolicity of graph products. *arXiv:2006.03085*, 2020.

[Bro03] Jeffrey Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, 16(3):495–535, 2003.

[Che00] Victor Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21, 2011.

[Dav08] Michael Davis. *The geometry and topology of Coxeter groups*, volume 32 of *London Mathematical Society Monographs Series.* Princeton University Press, 2008.

[DHS17] Matthew G. Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.

[DHS19] Matthew G. Durham, Mark F. Hagen, and Alessandro Sisto. Corrigendum to: Boundaries and automorphisms of hierarchically hyperbolic spaces. `https://www.wescac.net/undistorted_cyclic.pdf`, 2019.

[DJ00] Michael Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *Journal of Pure and Applied Algebra*, 153:229–235, 2000.

[Dru09] Cornelia Druţu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. *Comment. Math. Helv.*, 84(3):503–546, 2009.

[DS05] Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis V. Osin and Sapir.

[Dur16] Matthew G. Durham. The augmented marking complex of a surface. *Journal of the London Mathematical Society*, 94(3):933, 2016.

[EMR17] Alex Eskin, Howard A. Masur, and Kasra Rafi. Large-scale rank of Teichmüller space. *Duke Math. J.*, 166(8):1517–1572, 2017.

[Far98] Benson Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.

[Fox62] Ralph Fox. The braid groups. *Math. Scand.*, 10:119–126, 1962.

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

[Gen18] Anthony Genevois. Negative curvature of automorphism groups of graph products with applications to right-angled Artin groups. *arXiv:1810.10240*, 2018.

[Gen19a] Anthony Genevois. Negative curvature in graph braid groups. *arXiv:1912.10674*, 2019.

[Gen19b] Anthony Genevois. Quasi-isometrically rigid subgroups in right-angled Coxeter groups. *arXiv:1909.04318*, 2019.

[Gen21] Anthony Genevois. Algebraic characterisation of relatively hyperbolic special groups. *Israel J. Math*, 2021.

[GM18] Anthony Genevois and Alexandre Martin. Automorphisms of graph products of groups from a geometric perspective. *arXiv:1809.08091*, 2018.

[Gre90] Elisabeth Green. Graph products of groups. *PhD thesis, University of Warwick*, 1990.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[Hag12] Mark F. Hagen. Geometry and combinatorics of cube complexes. *PhD thesis, McGill University*, 2012.

[Hag14] Mark F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.*, 7(2):385–418, 2014.

[Hag19] Mark F. Hagen. A remark on thickness of free-by-cyclic groups. *Illinois J. Math.*, 63(4):633–643, 2019.

[Hru10] Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[HW08] Frédéric Haglund and Daniel Wise. Special cube complexes. *Geom. funct. anal.*, 17:1551–1620, 2008.

[HW10] Frédéric Haglund and Daniel Wise. Coxeter groups are virtually special. *Advances in Mathematics*, 224(5):1890–1903, 2010.

[KK14] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.*, 24(2):121–169, 2014.

[Lev19] Ivan Levcovitz. A quasi-isometry invariant and thickness bounds for right-angled Coxeter groups. *Groups Geom. Dyn.*, 13(1):349–378, 2019.

[Lev20] Ivan Levcovitz. Characterizing divergence and thickness in right-angled Coxeter groups. *arXiv:2007.13796*, 2020.

[Man05] Jason Fox Manning. Geometry of pseudocharacters. *Geom. Topol.*, 9(2):1147–1185, 2005.

[Mei96] John Meier. When is the graph product of hyperbolic groups hyperbolic? *Geom Dedicata*, 61(1):29–41, 1996.

[MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.

[MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[Mul80] Henry Mulder. *The interval function of a graph*, volume 181 of *Mathematical Centre Tracts*. Mathematisch Centrum, 1980.

[Ol'92] Alexander Ol'shanskii. Almost every group is hyperbolic. *International Journal of Algebra and Computation*, 2(1):1–17, 1992.

[Osi06] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.

[Osi16] Denis V. Osin. Acylindrically hyperbolic groups. *Trans. Amer. Math. Soc.*, 368(2):851–888, 2016.

[PS14] Paul Prue and Travis Scrimshaw. Abrams's stable equivalence for graph braid groups. *Topology and its Applications*, 178:136–145, 2014.

[Raf07] Kasra Rafi. A combinatorial model for the Teichmüller metric. *Geom. Funct. Anal.*, 17(3):936–959, 2007.

[Rus20] Jacob Russell. From hierarchical to relative hyperbolicity. *International Mathematics Research Notices*, 2020. rnaa141.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.*, 71(3):585–617, 1995.

[Sel97] Zlil Sela. Acylindrical accessibility for groups. *Invent. Math.*, 129(3):527–565, 1997.

[Sis12] Alessandro Sisto. On metric relative hyperbolicity. *arXiv:1210.8081*, 2012.