EQUIVARIANT HIERARCHICALLY HYPERBOLIC STRUCTURES FOR 3–MANIFOLD GROUPS VIA QUASIMORPHISMS

MARK HAGEN, JACOB RUSSELL, ALESSANDRO SISTO, AND DAVIDE SPRIANO

Abstract. Behrstock, Hagen, and Sisto classified 3–manifold groups admitting a hierarchically hyperbolic space structure. However, these structures were not always equivariant with respect to the group. In this paper, we classify 3–manifold groups admitting equivariant hierarchically hyperbolic structures. The key component of our proof is that the admissible groups introduced by Croke and Kleiner always admit equivariant hierarchically hyperbolic structures. For non-geometric graph manifolds, this is contrary to a conjecture of Behrstock, Hagen, and Sisto and also contrasts with results about CAT(0) cubical structures on these groups. Perhaps surprisingly, our arguments involve the construction of suitable quasimorphisms on the Seifert pieces, in order to construct actions on quasi-lines.

Contents

1. Introduction

Fundamental groups of 3–manifolds are a major source of inspiration in geometric group theory, providing a great part of the motivation for the notion of Gromov-hyperbolicity and all its generalisations, the study of actions on nonpositively-curved spaces, and the increasingly important role of special cube complexes.

One notion of "coarse nonpositive curvature", inspired partly by special cube complexes, is *hierarchical hyperbolicity*. Hierarchically hyperbolic spaces and groups were introduced in [BHS17b] as a means of isolating geometric features common to mapping class groups and certain CAT(0) cubical groups. After the definition took an easier-to-verify form in [BHS19], a budding study of hierarchical hyperbolicity has emerged. This has included

- finding new examples of hierarchically hyperbolic spaces and groups [BHS19, BR20, Mil20, BHMS20, BHS17a, Ber21, BR22, Vok22, HMS21, DDLS20, HS20, RS20, Rus21, Hug22, NQ22];
- development of new tools [DHS17, DHS20, RST23, BHMS20, Spr17, Rus22];
- establishment of geometric and algebraic consequence of hierarchical hyperbolicity [BHS17a, BHS21, ANS19, HP22, Pet21, HHP20, DMS20].

Very roughly, a *hierarchically hyperbolic space structure* on a space W consists of a set $\mathfrak{S}$ indexing a collection of δ–hyperbolic space $\{\mathcal{C}(U)\}_{U\in\mathfrak{S}}$ and a collection of projection maps $\{\pi_U\colon W\to\mathcal{C}(U)\}_{U\in\mathfrak{S}}$ satisfying a collection of axioms that allow for the coarse geometry of W to be recovered from these projections; see [BHS19, Definition 1.1] for the precise definition. Often, W is a finitely-generated group G equipped with a word metric. In this case, stronger results can be achieved when G is not only a hierarchically hyperbolic space (HHS), but has a structure that is compatible with the group action. These hierarchically hyperbolic *groups* (HHG) are defined precisely in Definition 2.16, but essentially this means that G acts cofinitely on $\mathfrak{S}$, with elements $g\in G$ inducing isometries $\mathcal{C}(U)\to\mathcal{C}(gU)$ so that all of the expected diagrams involving these isometries and the projections from the definition of an HHS commute.

The difference between HHSs and HHGs is illustrated by the fact that being an HHS is a quasi-isometry invariant property, but being an HHG is not [PS23]. While considerable geometric information can be gleaned from merely knowing that G is an HHS (e.g. finiteness of the asymptotic dimension [BHS17a] or control of quasiflats [BHS21]), one gets much more from the HHG property (e.g. semihyperbolicity [HHP20, DMS20] and the Tits alternative [DHS17, DHS20], or the consequences listed in Corollary 5).

The first examples of hierarchically hyperbolic spaces beyond mapping class groups and some cube complexes were the fundamental groups of closed orientable 3–manifolds whose prime decompositions excludes Nil and Sol pieces [BHS19]. However, the hierarchically hyperbolic structures constructed for such groups in [BHS19] are in general non-equivariant. In the present paper, we use new combinatorial techniques to produce *equivariant* hierarchically hyperbolic structures for 3–manifold groups. While many of the consequences of hierarchical hyperbolicity were known previously for 3–manifold groups, we find this satisfying as a complete answer to the question of hierarchical hyperbolicity for 3–manifold groups:

Theorem 1 (Theorem 3.3)**.** *Let M be a closed oriented 3–manifold. Then $\pi_1 M$ is a hierarchically hyperbolic group if and only if M has no Nil, Sol, or non-octahedral flat manifolds in its prime decomposition.*

In light of the previous characterisation of which 3-manifold groups are HHSs, Theorem 1 says the only additional obstruction to being HHG are non-octahedral flat manifolds in the prime decomposition.

Theorem 1 disproves a conjecture of Behrstock–Hagen–Sisto that there were examples of non-geometric graph manifold groups that were hierarchically hyperbolic *spaces*, but not hierarchically hyperbolic *groups*; see [BHS19, Remark 10.2]. This is a surprising result as this conjecture had a compelling heuristic justification. We explain this heuristic justification and how we circumvent it, then discuss the outline of our proof of Theorem 1.

1.1. **Comparison with cubulations: lines vs quasi-lines.** To explain the justification for the original belief that some graph manifold groups were not HHGs, we start with the *octahedral* hypothesis in Theorem 1. This says that the flat pieces are quotients of $\mathbb{E}^3$ by crystallographic groups with point group conjugate into $O_3(\mathbb{Z})$ (see [Hag14, Definition 2.2] or [Hod20, Theorem 7.1]). For crystallographic groups in any dimension, being octahedral is equivalent to cocompact cubulation [Hag14]. Petyt–Spriano showed that this is in turn equivalent to being an HHG [PS23]. So, while every crystallographic group is an HHS via a quasi-isometry to $\mathbb{Z}^n$, many crystallographic groups, such as the $(3,3,3)$–triangle group, are not HHGs.

There is a similar obstruction to cocompactly cubulating $\pi_1 M$ when M is a non-geometric graph manifold [HP15]. Specifically, $\pi_1 M$ can be cocompactly cubulated if it is *flip* in the sense of [KL98], that is, in every Seifert piece there is a "horizontal" surface whose boundary circles are fibres in the adjacent Seifert pieces. The idea behind the obstruction to cubulation is then: if $\widetilde{T} \subset \widetilde{M}$ is an elevation of a JSJ torus to the universal cover, and $\pi_1 M$ is cubulated, then the walls in $\widetilde{M}$ cut through $\widetilde{T}$ in at least two intersecting families of parallel lines. If the "flip" condition fails, then in some $\widetilde{T}$, there will be at least three such families, and the dual cube complex will contain $\widetilde{T} \times \mathbb{R} \cong \mathbb{E}^3$, preventing cocompactness. So, the obstruction to cocompact cubulation arises from specific $\mathbb{Z}^2$ subgroups getting "over-cubulated", as is the case with crystallographic groups.

The suspicion (confirmed in [PS23]) that cocompact cubulation is equivalent to the existence of an HHG structure for virtually abelian groups, together with the restrictions on cubulating graph manifolds, motivated the now disproven belief that non-flip graph manifold groups could fail to be HHGs.

The proof of Theorem 1 shows that constructing an HHG structure needs less than is needed to cocompactly cubulate. Roughly, in a cocompact cubulation of $\pi_1 M$, the immersed walls in M cut through each Seifert piece in a collection of surfaces whose boundary circles map to fibers in adjacent blocks; for each Seifert piece B we thus need a $\pi_1 B$–action on a line where certain elements act loxodromically and specific others fix points. For an HHG structure, we only need an action of $\pi_1 B$ on a *quasi-line* such that the central $\mathbb{Z}$ acts loxodromically, but the $\mathbb{Z}$ subgroups corresponding to the fibers of the adjacent Seifert pieces act with bounded orbits. The latter constraint is satisfiable even if M is not flip. This explains the involvement of quasimorphisms in our proof. The idea of using quasimorphisms in building HHG structures originated in this project, but has already found additional applications to Artin groups [HMS21] and extensions of subgroups of mapping class groups [DDLS20].

Another simple application of these actions on quasi-lines is that central extensions of hyperbolic groups by $\mathbb{Z}$ are HHGs.

Corollary 2 (Corollary 4.3)**.** *If a group G is a central extension $\mathbb{Z} \hookrightarrow G \overset{\pi}{\twoheadrightarrow} F$ where F is a non-elementary hyperbolic group, then G is a hierarchically hyperbolic group.*

While these central extensions were known to be hierarchically hyperbolic spaces by virtue of being quasi-isometric to $\mathbb{Z} \times F$, it did not appear to be known that they are in fact HHGs.

We now discuss the proof of Theorem 1 in more detail.

1.2. **Reduction to graph manifolds and admissible groups.** Let M be a closed oriented 3–manifold. The proof of the forward direction of Theorem 1, that the existence of an HHG structure for $\pi_1 M$ implies that M has no Nil, Sol, or non-hyperoctahedral pieces in it prime decomposition, is a consequence of results in [PS23, BHS19, RST23]. The idea is that we can push the HHG structure of $\pi_1 M$ to the fundamental groups of each of M's prime pieces, implying they cannot be Nil, Sol, or non-octahedal flat.

The main part of the paper is therefore devoted to other direction of Theorem 1, namely that if the prime decomposition of M excludes Nil, Sol, and non-octahedral flat manifolds, then $\pi_1 M$ is an HHG. As the geometric cases can largely be handled by appealing to results in the literature, the main new ingredient we need is that non-geometric graph manifold groups are HHGs.

Corollary 3 (Corollary 3.2)**.** *If M is a 3-dimensional non-geometric graph manifold, then $\pi_1 M$ is a hierarchically hyperbolic group.*

With Corollary 3 in hand, we can deduce the general case of Theorem 1 using the fact that a group that is hyperbolic relative to HHGs is itself an HHG; see [BHS19].

Our proof of Corollary 3 only relies on the specific way a graph manifold group decomposes into a graph of groups. Hence, instead of working in the specific case of graph manifolds, we work in the setting of *admissible graphs of groups*. This is a class of groups introduced by Croke and Kleiner to abstract the structure of $\pi_1 M$, when M is a non-geometric graph manifold [CK02]. Roughly, an admissible graph of groups is a nontrivial finite graph of groups $\mathcal{G}$ where each edge group is $\mathbb{Z}^2$ and each vertex group G_μ has infinite cyclic center Z_μ with quotient $F_\mu = G_\mu / Z_\mu$ a non-elementary hyperbolic group. Additionally, the various edge groups need to be pairwise non-commensurable inside each vertex group. The exact definition is Definition 2.13. Hence, hierarchical hyperbolicity of $\pi_1 M$ is a special case of:

Theorem 4 (Theorem 3.1, Proposition 6.8)**.** *Let $\mathcal{G}$ be an admissible graph of groups. Then $\pi_1 \mathcal{G}$ is a hierarchically hyperbolic group. Moreover, if each quotient F_μ is a free group, then the associated hyperbolic spaces are quasi-isometric to trees.*

Recently, Nguyen and Qing showed that every admissible group that acts geometrically on a CAT(0) space is a hierarchically hyperbolic space [NQ22, Theorem A]. Their result focuses on CAT(0) geometry and does not in general produce equivariant structures. Our proof of Theorem 4 will employ a much more combinatorial framework that will ensure equivariance and avoid the need for the action on a CAT(0) space.

1.3. **Consequences.** Equivariant hierarchical hyperbolicity for fundamental groups of admissible graphs of groups has several immediate consequences for these groups.

Corollary 5. *cori:consequences Let $\mathcal{G}$ be an admissible graph of groups, and let $G = \pi_1\mathcal{G}$. Then:*

(1) G acts properly and coboundedly on an injective metric space, and is hence semihyperbolic;
(2) if G is virtually torsion-free, then G has uniform exponential growth;
(3) the action of G on the Bass–Serre tree is the largest (hence universal) acylindrical action of G on a hyperbolic space;
(4) the Morse boundary of G is an ω–cantor set. In particular, it is totally disconnected.

Proof. The first assertion follows from the fact that G is an HHG (Theorem 4) by [HHP20, Corollary 3.8, Lemma 3.10].

For the other assertions, we will need that the $\sqsubseteq$–maximal domain in the HHG structure is G–equivariantly quasi-isometric to the Bass–Serre tree T for $\mathcal{G}$. We prove this in Proposition 6.8. Because the definition of an admissible graph of groups ensures that T has infinitely many ends, [ANS19, Corollary 4.8] implies that G has uniform exponential growth. It follows from [ABD21, Theorem A] that the action of G on T is the largest acylindrical action of G on a hyperbolic space[1]. The last item on the Morse boundary follows from [Rus21, Corollary A.8] (using HHGs) or [CCS23, Theorem 1.2] (using graphs of groups). □

HHGs on quasi-trees vs cubical groups. We also note the following consequence for the question of when hierarchically hyperbolic structures are forced to arise from cubulation. Corollary 3 provides a hierarchically hyperbolic structure in which the constituent hyperbolic spaces are all quasi-isometric to trees. Such hierarchically hyperbolic structures also arise on fundamental groups of compact special cube complexes [BHS17b] and more generally, groups acting geometrically on cube complexes admitting *factor systems* [BHS17b, HS20]. However, there are many examples of graph manifolds whose fundamental groups are virtually special but not virtually *compact* special, and indeed not even virtually cocompactly cubulated [HP15]. Hence Corollary 3.2 provides examples of groups that are not cocompactly cubulated, but do admit HHG structures in which the hyperbolic spaces are all quasi-trees.

1.4. **Proof ingredients: combinatorial HHS and quasi-morphisms.** To prove admissible groups are HHGs, we employ the recent *combinatorial HHS* machinery from [BHMS20]. For a group G, this requires constructing a simplicial complex Y and a graph W, which are then combined in a graph Y^{+W}. Intuitively, the role of those spaces is as follows: the complex Y encodes the index set of a hierarchically hyperbolic structure, the complex Y and the graph Y^{+W} together encode the associated hyperbolic spaces, and W is the (equivariant) quasi-isometry model of G.

[1] Theorem A of [ABD21], as written, can be read as suggesting that 3–manifold groups are hierarchically hyperbolic *groups*, although at the time they were only known to be hierarchically hyperbolic *spaces* in the stated generality. But, as noted in [ABD21, Remark 5.3], Theorem A holds for 3–manifold groups without needing an HHG structure. Alternatively, by Theorem 1, the statement in [ABD21] holds once one excludes non-octahedral flat pieces from the prime decomposition.

In our case, the space Y is an augmented version of the Bass-Serre tree, where each vertex is "blown up" to contain a copy of the coset it represents. This technique of building combinatorial HHSs by "blowing up the vertex groups" in some naturally-occurring hyperbolic graph is quite flexible, and has analogues in a number of other contexts. For example, it is applied in the context of certain Artin groups in [HMS21], extensions of lattice Veech groups in [DDLS20], and extensions of multicurve stabilisers in [Rus21]. In [BHMS20], it is explained how to build combinatorial HHSs for right-angled Artin groups and mapping class groups by respectively blowing up the Kim–Koberda *extension graph* [KK13] and the curve graph.

In a general combinatorial HHS, Y is a simplicial complex with a G–action that has finitely many orbits of links of simplices, and W is a graph whose vertices are maximal simplices of Y, where the action of G on Y induces an isometric action of G of W. Given Y and W, the graph Y^{+W} is constructed from $Y^{(1)}$ by joining every vertex of the maximal simplex Σ to every vertex of the maximal simplex Δ by an edge whenever Σ and Δ represent adjacent vertices of W. The group G acts naturally on the resulting graph Y^{+W}.

The spaces Y, W, and Y^{+W} encode the HHS structure as follows. The elements of the index set correspond to the links $\mathrm{lk}(\Delta)$ of non-maximal simplices Δ of Y (including the empty simplex, whose link is Y). The hyperbolic space associated to $\mathrm{lk}(\Delta)$ is the subgraph $\mathrm{lk}(\Delta)^{+W}$ of Y^{+W} spanned by the vertices in $\mathrm{lk}(\Delta) \subset Y$. Accordingly, we have to choose the edges of W in a way that ensures that all of these spaces (including Y^{+W} itself) are hyperbolic, while also ensuring that the action of G on W is proper and cobounded.

Hierarchical hyperbolicity demands not only the construction of a collection of hyperbolic spaces, but also a coarse projection from W to each $\mathrm{lk}(\Delta)^{+W}$ (satisfying a list of properties [BHS19, Definition 1.1]). To arrange this, the definition of a combinatorial HHS requires the following: consider all of the simplices $\Delta' \subset Y$ with the same link as Δ, and remove their vertex sets (and incident edges) from Y^{+W} to obtain a graph Y_Δ, which contains $\mathrm{lk}(\Delta)^{+W}$. We ask that the inclusion $\mathrm{lk}(\Delta)^{+W} \hookrightarrow Y_\Delta$ is a quasi-isometric embedding, for each non-maximal simplex Δ. The exact definition of a combinatorial HHS is Definition 2.23, which involves some additional (combinatorial) conditions.

Our combinatorial HHS and the role of quasimorphisms. Given an admissible graph $\mathcal{G}$ of groups, let T be the Bass–Serre tree. The idea for constructing the simplicial complex Y for $\pi_1\mathcal{G}$ is as follows: "blow up" each vertex v of T to become the cone on a discrete set whose elements correspond to the associated coset of the vertex group. Two such cones are then graph-theoretically joined according to the edges of T, resulting in a 3–dimensional simplicial complex. The action of $\pi_1\mathcal{G}$ on T induces an action on Y, by construction.

Having constructed the simplicial complex Y, we now need to construct the graph W whose vertices are maximal simplices of Y and will serve as the geometric model for $\pi_1\mathcal{G}$. This is where quasimorphisms come in.

Specifically, for each vertex group G_μ, we construct an action of G_μ of a quasi-line L_μ so that the center Z_μ of G_μ acts loxodromically, and each cyclic subgroup conjugate to the images of the center of adjacent vertex groups acts elliptically. This is achieved by first choosing an appropriate quasimorphism and then using a result of Abbott–Balasubramanya–Osin [ABO19] to promote it to an action on a quasi-line; see Lemma 4.2.

Using the action on this quasi-line, each vertex groups in our admissible graph of group is *equivariantly* quasi-isometric to the product $L_\mu \times F_\mu$, where F_μ is the hyperbolic quotient G_μ/Z_μ. Balls in the L_μ factor therefore give us coarse "level surfaces" in this product. Moreover, if ω is adjacent to μ, the fact that the center Z_ω acts elliptically on L_μ means that Z_ω is sent into one of these coarse level surfaces by the edge maps in $\mathcal{G}$.

Now, maximal simplices of Y consist of an edge $\{u, v\}$ of T and a pair of elements s, t in the corresponding cosets of the vertex groups. Using the above product structure, each of s and t determine a level surface in each of the vertex groups corresponding to u and v, and these two level surface will intersect in uniformly bounded subsets. Roughly, we define W so that there is an edge between two vertices if these bounded diameter subsets associated to the two maximal simplices of Y are close; see see Definition 5.7 and Proposition 5.9 for details. This definition will make W an equivariant quasi-isometric model for $\pi_1\mathcal{G}$.

The definition of edges in W will ensure that the extra edges in Y^{+W} are only added between vertices of Y that are uniformly close under the collapse map $Y \to T$. Hence Y^{+W} will be quasi-isometric to the Bass–Serre tree T and hence hyperbolic. The other hyperbolic spaces coming from our combinatorial HHS are all either bounded diameter or correspond to one of the two factors of the product $L_\mu \times F_\mu$ for one of the vertex groups. One set of spaces will be quasi-isometric to the quasi-lines L_μ, while the other will be quasi-isometric to hyperbolic cone-offs of the F_μ. The quasi-isometric embedding conditions for these hyperbolic spaces are verified using a combination of closest point projection in the Bass–Serre tree T with the hyperbolic geometry of the F_μ factor of each vertex group.

1.5. **Outline.** Section 2 contains background on coarse geometry, graphs of groups and hierarchical hyperbolicity. This includes the definition of an admissible graph of groups (Section 2.2) and combinatorial HHS (Section 2.3). Section 3 presents the statements of our main results in more detail and deduces Theorem 1 from Theorem 4. The rest of the paper is devoted to the proof of Theorem 4. In Section 4, we use quasimorphism to produce actions of central $\mathbb{Z}$–extensions on quasi-lines. In Section 5, we construct the simplicial complex Y and the graph W that will comprise our combinatorial HHS for an admissible graph of groups. Section 6 contains the proof that (Y, W) is a combinatorial HHS. Section 6.1 focuses on describing the link of simplices of Y and verifying the combinatorial parts of the definition of a combinatorial HHS. Section 6.2 contains the proof that Y^{+W} and the $\mathrm{lk}(\Delta)^{+W}$ are hyperbolic. Section 6.3 is devoted to checking that the inclusions $\mathrm{lk}(\Delta)^{+W} \to Y_\Delta$ are quasi-isometric embeddings.

Acknowledgments. Hagen was partly supported by EPSRC New Investigator Award EP/R042187/1. Russell was supported by NSF grant DMS-2103191. Spriano was partly supported by the Christ Church Research Centre. We would like to thank the LabEx of the Institut Henri Poincaré (UAR 839 CNRS-Sorbonne Université) for their support during the trimester program "Groups acting on Fractals, Hyperbolicity and Self-similarity". We thank the referee for their careful reading and numerous helpful comments.

2. Preliminaries

2.1. **Coarse Geometry and Groups.** We recall some basic notions from coarse geometry and outline some techniques we will use repeatedly. For a metric space Y, we will use d_Y to denote the distance in the space Y. The metric spaces we will consider will be undirected graphs, which we always equip with the path metric coming from declaring each edge to have length 1. For a graph Y, we let $Y^{(0)}$ denote the set of vertices of Y.

Let $\kappa \geqslant 1$, $\xi \geqslant 0$ and $f\colon Y \to Q$ be a map between metric spaces. The map f is a (κ,ξ)*–quasi-isometric embedding* if for all $x, y \in Y$ we have

$$\frac{1}{\kappa} d_Q(f(x), f(y)) - \xi \leqslant d_Y(x,y) \leqslant \kappa d_Q(f(x), f(y)) + \xi.$$

The map f is ξ*–coarsely onto* (or coarsely surjective) if for all $q \in Q$, there exists $y \in Y$ so that $d_Q(q, f(y)) \leqslant \xi$. If f is a (κ,ξ)–quasi-isometric embedding that is ξ–coarsely onto, we say f is a (κ,ξ)*–quasi-isometry.* A ξ*–quasi-inverse* of f is a map $h\colon Q \to Y$ so that $d_Y(y, h(f(y)) \leqslant \xi$ for each $y \in Y$. The map f will be (κ,ξ)*–coarsely Lipschitz* if

$$d_Y(x,y) \leqslant \kappa d_Q(f(x), f(y)) + \xi$$

for all $x, y \in Y$. We often omit the constants when their specific value is not relevant. Note that the map f is a quasi-isometry if and only if f is coarsely Lipschitz and has a coarsely Lipschitz quasi-inverse (where the constants on either side of this equivalence determine the constants on the other).

A *(quasi)-geodesic* in a metric space Y is an (quasi)-isometric embedding of a closed interval $I \subseteq \mathbb{R}$ into Y. When Y is a graph, we additionally require that the endpoints of the (quasi)-geodesic are vertices of Y.

At times it will be convenient to work with coarsely defined maps. A ξ*–coarse map* from a metric space Y to a metric space Q is a function $f\colon Y \to 2^Q$ where for each $y \in Y$, $f(y)$ is a subset of Q with diameter at most ξ. By a slight abuse of notation, we still write $f\colon Y \to Q$ to denote a coarse map. We say that a coarse map is coarsely Lipschitz, coarsely onto, a quasi-isometric embedding, a quasi-inverse or a quasi-isometry if it satisfies the same inequalities as described in the previous paragraph (where the distance between two sets is the minimal distance between two elements).

For graphs, we frequently use the following criteria to determine whether a map is coarsely Lipschitz and when an inclusion is a quasi-isometric embedding. The proofs are left as straightforward exercises.

Lemma 2.1 (Locally Lipschitz is Lipschitz)**.** *For each $\xi \geqslant 0$ and $\kappa \geqslant 0$, there exists $\xi' \geqslant 0$ and $\kappa' \geqslant 1$ so that the following holds. Let Y and Q be graphs and suppose $f_0\colon Y^{(0)} \to Q$ is a ξ–coarse map. Let $f\colon Y \to Q$ be the map that extends f_0 by sending each edge e of Y to union of the images of the endpoints of e under f_0. If $d_Q(f_0(x), f_0(y)) \leqslant \kappa$ for each $x, y \in Y^{(0)}$ that are joined by an edge of Y, then f is a (κ', κ')–coarsely Lipschitz ξ'–coarse map.*

Lemma 2.2 (Coarse retracts are undistorted)**.** *Let Y and Q be graphs and assume there is an injective simplicial map $i\colon Q \to Y$. If there is a (κ,κ)–coarsely Lipschitz κ–coarse map $f\colon Y^{(0)} \to Q$ so that $f(i(Q)) = i(Q)$ and for each $q \in Q$, $d_Q(q, i^{-1} \circ f \circ i(q)) \leqslant \kappa$, then the map $i\colon Q \to Y$ is a quasi-isometric embedding with constants determined by κ.*

We will apply Lemma 2.2 exclusively in the case where Q is a connected subgraph of Y. In this case, we emphasise that the map f is coarsely Lipschitz with respect to the intrinsic path metric on Q and not the metric the Q inherits as a subset of Y. A map f satisfying the conditions of Lemma 2.2 is called a *coarse retract* of Y to Q.

We say a graph Y is *δ–hyperbolic* if for any geodesic triangle in Y, the δ–neighborhood of any two sides covers the third side. Special cases of hyperbolic graphs are *quasi-trees* and *quasi-line*, which are graphs that are quasi-isometric to a tree or line respectively. We will need to use some ideas from the theory of relatively hyperbolic groups and spaces. Given a collection of coarsely connected subsets $\mathcal{Q}$ of a graph Y, we define the *electrification of Y with respect to $\mathcal{Q}$* to be the space obtained from Y by adding an additional edge between $x, y \in Y^{(0)}$ whenever there is $Q \in \mathcal{Q}$ so that $x, y \in Q$. We denote this electrification by $\hat{Y}_{\mathcal{Q}}$. We say that Y is *hyperbolic relative to* $\mathcal{Q}$ if $\hat{Y}_Q$ is δ–hyperbolic for some $\delta \geqslant 0$ and if it satisfies the *bounded subset penetration* property; see [Sis12, Definition 3.7] for full details.

Many of the graphs we will study will be the Cayley graphs of groups.

Definition 2.3. Let G be a group and J be a symmetric generating set for G. We let $\mathrm{Cay}(G, J)$ denote the simplicial graph whose vertices are the elements of G and where two elements g, h are joined by an edge if $g^{-1}h \in J$.

Note that the generating set J does not need to be finite; in fact we will consider non-locally finite Cayley graphs throughout the paper.

Suppose a group G is acting by isometries on metric space Y. We say G acts *coboundedly* if there exists a bounded set B such that $G \cdot B = Y$. We say the action of G on Y is *metrically proper* if for any bounded diameter subset K of Y, the set $\{g \in G : g \cdot K \cap K \neq \emptyset\}$ is finite. A version of the Milnor-Schwartz lemma say that if a finitely generated group G groups act metrically proper and coboundedly on a metric space Y, then the orbit map gives a quasi-isometry $\mathrm{Cay}(G, J) \to Y$ for any finite generating set J.

A finitely generated group G is *hyperbolic* if for some (and hence any) finite generating set J, the graph $\mathrm{Cay}(G, J)$ is δ–hyperbolic for some $\delta \geqslant 0$. A finitely generated group G is *hyperbolic relative* to a finite collection of subgroups $\{Q_1, \ldots, Q_n\}$ if for some (and hence any) finite generating set J, the graph $\mathrm{Cay}(G, J)$ is hyperbolic relative to the collection of all cosets of the Q_i's. In particular, the Cayley graph $\mathrm{Cay}(G, J \cup Q_1 \ldots Q_n)$ is hyperbolic.

The next lemma is a useful tool that allows to verify that the electrification of a quasi-tree with respect to quasiconvex subsets is again a quasi-tree.

Lemma 2.4. *For all $\delta, \kappa \geqslant 1$ there exists δ' such that the following holds. Let Γ be a graph that is (δ, δ)–quasi-isometric to a tree and $\mathcal{Q}$ be a collection of κ–quasiconvex subspaces of Γ. Then the electrification $\hat{\Gamma}_{\mathcal{Q}}$ is (δ', δ')–quasi-isometric to a tree.*

Proof. We use the following consequence of Manning's bottleneck criterion [Man05, Theorem 4.6], formulated in [BBF15, Section 3.6] (see also [DDLS20, Proposition 2.3]): a space is a quasi-tree if and only if there exists ξ as follows: for any two points x, y, path p between them and point z on a geodesic between x and y, we have $d(z, p) \leqslant \xi$. Moreover, the constants of the quasi-isometry to a tree and ξ each determine the other

Let ξ be such a constant for the quasi-tree Γ and let $\hat{\Gamma} = \hat{\Gamma}_{\mathcal{Q}}$. Our goal is to find an analogous $\hat{\xi}$ for $\hat{\Gamma}$. Let x, y be two points of $\Gamma^{(0)} = \hat{\Gamma}^{(0)}$ and let $\hat{\beta}$ be a $\hat{\Gamma}$-geodesic between them. Let $\hat{z}$ be a point on $\hat{\beta}$ and $\hat{\gamma}$ be some path in $\hat{\Gamma}$ connecting x and y between them. Let β be a Γ–quasi-geodesic between x, y. By [KR14, Corollary 2.6], the Hausdorff distance in $\hat{\Gamma}$ between β and $\hat{\beta}$ is uniformly bounded by some R. Thus, there exists $z \in \beta$ such that

$d_{\hat{\Gamma}}(\hat{z}, z) \leqslant R$. Let γ be the Γ–path obtained from $\hat{\gamma}$ by replacing $\hat{\Gamma} - \Gamma$ edges with geodesics of Γ. Since Γ is a quasi-tree, there is a point $p \in \gamma$ with $d_\Gamma(p, z) \leqslant \xi$. If p is also a point of $\hat{\gamma}$ we are done. Otherwise, p is on a geodesic with endpoints on a κ-quasiconvex Q_i, we have $d_\Gamma(p, Q_i) \leqslant \kappa$. As Q_i is coned-off in $\hat{\Gamma}$ and $\hat{p}$ intersects Q_i, we obtain $d_{\hat{\Gamma}}(p, \hat{\gamma}) \leqslant \kappa + 1$. By the triangular inequality,

$$d_{\hat{\Gamma}}(\hat{z}, \hat{\gamma}) \leqslant d_{\hat{\Gamma}}(\hat{z}, z) + d_{\hat{\Gamma}}(z, p) + d_{\hat{\Gamma}}(p, \hat{\gamma}).$$

As each of the above quantities is uniformly bounded, we get the claim. □

We conclude with a lemma relating quotients and Cayley graphs with respect to infinite generators.

Lemma 2.5. *Let G be a group and $N \trianglelefteq G$ a normal subgroup. For any generating set K of G satisfying $N \subseteq K$ the quotient map $\pi\colon G \to G/N$ induces a $(2,1)$–quasi-isometry*

$$\pi\colon \mathrm{Cay}(G, K) \to \mathrm{Cay}(G/N, \pi(K)).$$

Proof. Let $\Gamma = \mathrm{Cay}(G, K)$ and $\Omega = \mathrm{Cay}(G/N, \pi(K))$. By construction, the map π gives a 1-Lipschitz map $\Gamma \to \Omega$. For each $x \in G/N$, let $\theta(x)$ be the be an element of the coset gN in G so that $\pi(gN) = x$. Given any $x_1, x_2 \in G/N$ with $x_1^{-1}x_2 \in \pi(K)$ we can find y_1 in the same coset as $\theta(x_1)$ and y_2 in the coset as $\theta(x_2)$ so that $y_1^{-1}y_2 \in K$. Since each coset gN has diameter 1 in Γ, we have $d_\Omega(\pi(x_1), \pi(x_2)) \leqslant d_\Gamma(x_1, x_2) \leqslant 2d_\Omega(\pi(x_1), \pi(x_2)) + 1$. □

2.2. **Graphs of groups.** We start with recalling some definitions and notations from Bass–Serre theory. For a comprehensive background, we refer the reader to [SW79]. Firstly, we recall that for Bass–Serre it is useful to use the language of bi-directed graphs.

Definition 2.6. A *bi-directed graph* Γ consists of sets $V(\Gamma)$, $E(\Gamma)$ and maps

$$E(\Gamma) \to V(\Gamma) \times V(\Gamma); \qquad E(\Gamma) \to E(\Gamma)$$
$$\alpha \mapsto (\alpha^+, \alpha^-) \qquad \alpha \mapsto \bar{\alpha}$$

satisfying $\bar{\bar{\alpha}} = \alpha$, $\bar{\alpha} \neq \alpha$ and $(\bar{\alpha})^- = \alpha^+$.

The elements of $V(\Gamma)$ are called *vertices*, the ones of $E(\Gamma)$ are called *edges*, the vertex α^- is the *source* of α, α^+ is the *target* and $\bar{\alpha}$ is the *reverse edge*. A bi-directed graph Γ is *finite* if both $V(\Gamma), E(\Gamma)$ are finite sets. A *subgraph* of Γ is a bi-directed graph Γ' such that $V(\Gamma') \subseteq V(\Gamma)$ and $E(\Gamma') \subseteq E(\Gamma)$. Given a bi-directed graph Γ, it is standard to associate to it a an undirected graph $|\Gamma|$, where the vertices are the elements of $V(\Gamma)$ and the edges are pairs of edges of the form $\{\alpha, \bar{\alpha}\}$. We call these pairs of edges $\{\alpha, \bar{\alpha}\}$ *undirected edges* of $\mathcal{G}$. An *orientation* on an undirected edge is choice of one of the directed edges. We say that a bi-directed graph Γ is *connected*, respectively a *tree* if $|\Gamma|$ is connected, respectively a tree. We say that a subgraph T of Γ is a *spanning tree* if $V(T) = V(\Gamma)$ and T is a tree.

The correspondence between Γ and $|\Gamma|$ gives an equivalence between undirected graphs and bi-directed graphs. The reason behind distinguishing the two classes is that the language of undirected graphs is more natural when considering graphs as metric spaces, whereas bi-directed graphs highlight combinatorial properties used to describe graphs of groups.

Definition 2.7. A *graph of groups* $\mathcal{G}$ consists of a finite connected bi-directed graph Γ, two collections of groups $\{G_\mu \mid \mu \in V(\Gamma)\}$ and $\{G_\alpha \mid \alpha \in E(\Gamma)\}$ satisfying $G_\alpha = G_{\bar{\alpha}}$, and injective homomorphisms $\tau_\alpha\colon G_\alpha \to G_{\alpha^+}$ for each $\alpha \in E(\Gamma)$.

Definition 2.8. Let $\mathcal{G} = (\Gamma, \{G_\mu\}, \{G_\alpha\}, \{\tau_\alpha\})$ be a graph of groups. We define the group $F\mathcal{G}$ as:

$$F\mathcal{G} = \left(\mathop{\ast}_{\mu \in V(\Gamma)} G_\mu \right) \ast \left(\mathop{\ast}_{\alpha \in E(\Gamma)} \langle t_\alpha \rangle \right).$$

Let $\mathrm{Sp}(\Gamma)$ be a spanning tree of Γ. The *fundamental group* of $\mathcal{G}$ with respect to $\mathrm{Sp}(\Gamma)$, denoted by $\pi_1(\mathcal{G}, \mathrm{Sp}(\Gamma))$, is the group obtained adding the following relations to $F\mathcal{G}$:

(1) $t_\alpha = t_{\bar{\alpha}}^{-1}$;
(2) $t_\alpha = 1$ if $\alpha \in E(\mathrm{Sp}(\Gamma))$;
(3) $t_\alpha \tau_\alpha(x) t_\alpha^{-1} = \tau_{\bar{\alpha}}(x)$ for all $x \in G_\alpha$.

Given a graph of groups $\mathcal{G}$ we can associate to it the Bass–Serre tree T; see e.g. [SW79, Section 4]. This is the bi-directed graph whose vertices are cosets of the vertex groups, and two cosets are joined by a directed edge if there are representatives gG_μ and hG_ω such that the vertices μ and ω are connected by an edge α with $\alpha^+ = \mu$ and $ht_\alpha = g$. For a vertex v of T we let $\check{v}$ denote the vertex μ of $\mathcal{G}$ so that $v = gG_\mu$. Similarly, given an edge e of T, we define $\check{e}$ to be the edge of $\mathcal{G}$ joining $\check{e}^+$ and $\check{e}^-$.

If the vertex $v \in T^{(0)}$ is the coset gG_μ, then stabiliser $\mathrm{Stab}_{\pi_1\mathcal{G}}(v)$ is the conjugate of the vertex group $gG_\mu g^{-1}$. Similarly, for each edge e of T, the stabiliser $\mathrm{Stab}_{\pi_1\mathcal{G}}(e)$ is $g\tau_{\bar{\alpha}}(G_\alpha)g^{-1} = gt_\alpha\tau_\alpha(G_\alpha)t_\alpha^{-1}g^{-1}$ where $\check{e} = \alpha$, and g is an element of $\pi_1\mathcal{G}$ so that $gG_{\check{e}^-}$ and $gt_\alpha G_{\check{e}^+}$ are the vertices e^- and e^+ respectively.

Even though T is a bi-directed graph, we will at times think of it as a metric space. When we do this, we are implicitly referring to $|T|$, the undirected graph obtained from T. We will use E to denote unoriented edges of T and e to denote an orientation on E.

Given a graph of groups, we want to provide a geometric model that encodes the geometry of the entire fundamental group. To achieve this we will take the cosets of the vertex groups and join them together using the information coming from the tree and the edge group. We call the resulting space the *Bass–Serre space.* In order to keep track of the geometry of the edge spaces, it is useful to introduce a combinatorial notion of edges with midpoints.

Definition 2.9. A *subdivided edge* is a (undirected) graph isomorphic to the graph with vertices v_0, v_1, v_2 and edges between v_0 and v_1, and between v_1 and v_2. The vertex v_1 is called the *middle vertex.* Two vertices x, y of a graph Γ are *connected by a subdivided edge* if there is a subgraph of Γ isomorphic to a subdivided edge with $v_0 = x$ and $v_2 = y$.

Definition 2.10 (Bass–Serre space)**.** Let $\mathcal{G}$ be a graph of finitely generated groups. For each vertex group G_μ and edge group G_α fix once and for all finite symmetric generating sets J_μ and J_α respectively, such that $J_\alpha = J_{\bar{\alpha}}$ and $\tau_\alpha(J_\alpha) \subseteq J_{\alpha^+}$. We build the *Bass–Serre space* X for the graph of groups $\mathcal{G}$ in three steps.

Step 1: vertex spaces. For each vertex $v = gG_\mu$ of T, we define X_v to be the graph with vertex set gG_μ and with an edge between $x, y \in gG_\mu$ if $x^{-1}y \in J_\mu$. We call each X_v the *vertex space* for $v \in T^{(0)}$. Because each vertex group injects into $\pi_1\mathcal{G}$, each X_v is graphically isomorphic to the Cayley graph of $G_{\check{v}}$ with respect to the generating set $J_{\check{v}}$.

Step 2: subdivided edges. Given an undirected edge E of T, pick an orientation $e \in E$ and let $\alpha = \check{e}$, $\mu = \alpha^+$, and $\omega = \alpha^-$. Fix an element $g \in \pi_1\mathcal{G}$ so that $gG_\omega = X_{e^-}$ and $gt_\alpha G_\mu = X_{e^+}$. For each $a \in G_\alpha$, add a subdivided edge between $g\tau_{\bar{\alpha}}(a) \in gG_\omega = X_{e^-}$ and $gt_\alpha\tau_\alpha(a) \in gG_\mu = X_{e^+}$. By Definition 2.8.(3), if $x = g\tau_{\bar{\alpha}}(a)$, then $xt_\alpha = g\tau_{\bar{\alpha}}(a)t_\alpha =$

$gt_\alpha\tau_\alpha(a)$. Hence, all such x and xt_α are joined by a subdivided edge and the addition of these subdivided edges does not rely on our specific choice of representative $g \in \pi_1\mathcal{G}$. Since $t_\alpha^{-1} = t_{\bar{\alpha}}$, the addition of these subdivided edges is also independent of the orientation chosen for E.

Step 3: edges spaces. Let E be an undirected edge of T with orientation e. Let $e^+ = v$, $e^- = w$, and $\check{e} = \alpha$. For each subdivided edge added between X_v and X_w there is a middle vertex. Let a, b be two of these middle vertices and x, y the vertices of X_v adjacent to them. To complete the Bass–Serre space, we add an edge between any two such a, b if $x^{-1}y \in \tau_\alpha(J_\alpha)$ in $\pi_1\mathcal{G}$. This is independent of the orientation for E because if p, q are the vertices of X_w adjacent to a and b respectively, then $p = xt_\alpha$ and $q = yt_\alpha$. Thus $p^{-1}q = t_\alpha^{-1}x^{-1}yt_\alpha$, which implies $x^{-1}y \in \tau_\alpha(J_\alpha)$ if and only if $p^{-1}q \in \tau_{\bar{\alpha}}(J_\alpha)$ by Definition 2.8.(3).

For each (directed) edge e in T, we use X_e to denote the (undirected) graph whose vertices are all of the middle vertices of the subdivided edges between X_{e^+} and X_{e^-} with the edges defined as above. We call X_e the *edge space* for e and note that $X_e = X_{\bar{e}}$. Each edge space X_e is graphically isomorphic to the Cayley graph of the edge group $G_{\check{e}}$ with generating set $J_{\check{e}}$. We let $\tau_e \colon X_e \to X_{e^+}$ denote be the map that associates to each middle vertex the only vertex of X_v adjacent to it, and we define $\tau_{\bar{e}}$ analogously. Figure 1 gives a schematic of the edge spaces and τ_e maps.

The *Bass–Serre space* X for the graph of groups $\mathcal{G}$ is the space constructed from taking all the vertex spaces in Step 1, adding in all the subdivided edges from Step 2, and then adding in all the edges of the edge spaces in Step 3. The group $\pi_1\mathcal{G}$ acts on the disjoint union of the vertex spaces by left multiplication. This action can be extended to the subdivided edges and edge spaces to give an action of $\pi_1\mathcal{G}$ of X by isometries. The edge space maps τ_e and $\tau_{\bar{e}}$ are equivariant with respect to this action.

Remark 2.11. For every $x, y \in X_e$, we have $d_X(x, \tau_e(x)) = 1$ and $d_{X_{e^+}}(\tau_e(x), \tau_e(y)) \leqslant d_{X_e}(x, y) + 2$.

While the inclusion of the vertex and edge spaces in to the Bass–Serre space are simplicial injections, their images maybe very distorted in the total metric on X. However, as there are only finitely many vertex and edge groups, we have uniform control over this distortion

Lemma 2.12. *Let $\mathcal{G}$ be a graph of groups with Bass–Serre tree T and Bass–Serre space X. There exists a monotone diverging function $h \colon [0, \infty) \to [0, \infty)$ so that for each vertex v and edge e of T we have*

$$d_{X_v}(x, y) \leqslant h\big(d_X(x, y)\big) \text{ and } d_{X_e}(x, y) \leqslant h\big(d_X(x, y)\big)$$

for any $x, y \in X_v$ or $x, y \in X_e$.

Proof. For each $v \in T^{(0)}$, define $h_v \colon [0, \infty) \to [0, \infty)$ to be

$$h_v(r) = \max_{\{x,y \in X : d_{X_v}(x,y) \leqslant r\}} \{d_{X_v}(x, y)\}.$$

Because X is locally finite and $\pi_1\mathcal{G}$ acts transitively on the vertices of the vertex and edge spaces respectively, h_v exists and is a monotone diverging function. We similarly define h_e for each edge e of T. If two vertices, v and w, or two edges, e and f, are in the same

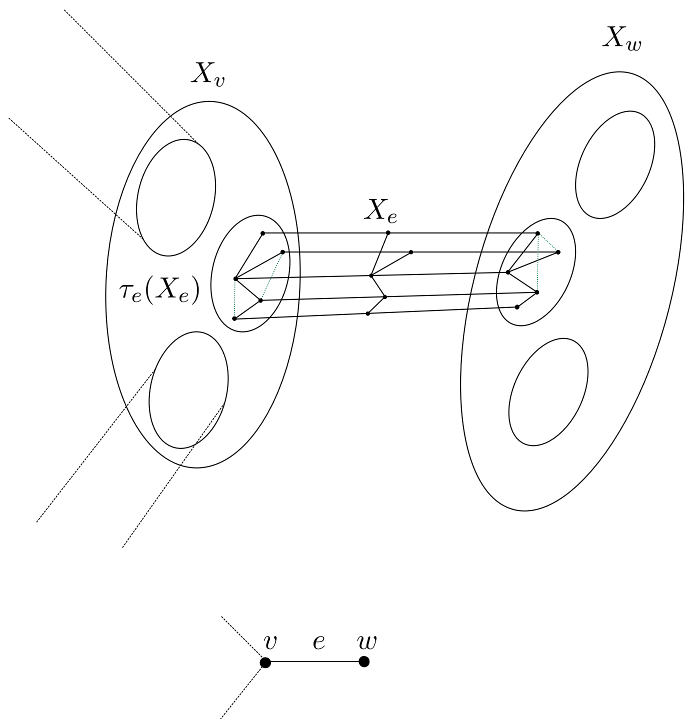

FIGURE 1. The cosets corresponding to the edge e are connected by a subdivided edge. In the picture, we assume that $e^+ = v$. To every edge of X_e corresponds an edge in X_v, X_w, but additional edges might be present.

$\pi_1\mathcal{G}$–orbit then $h_v = h_w$ and $h_e = h_f$. Since $\pi_1\mathcal{G}$ acts of T with finitely many orbits of edges and vertices, we can find the desired h by taking the minimum over all of these finitely many orbits. □

Croke and Kleiner introduce the following class of *admissible* graphs of groups to abstract the properties of the graphs of groups structure of the fundamental groups of non-geometric graph manifolds [CK02]. This will be the class of graphs of groups that we will study.

Definition 2.13. Let $\mathcal{G} = (\Gamma, \{G_\mu\}, \{G_\alpha\}, \{\tau_\alpha\})$ be a graph of groups. We say $\mathcal{G}$ is *admissible* if the following hold:

(1) Γ contains at least 1 edge.
(2) Each vertex group G_μ has center Z_μ that is a infinite cyclic group, and $G_\mu/Z_\mu = F_\mu$ is a non-elementary hyperbolic group.
(3) Each edge group G_α is isomorphic to $\mathbb{Z}^2$.
(4) If α is an edge with $\mu = \alpha^+$ and $\omega = \alpha^-$, then $\langle \tau_\alpha^{-1}(Z_\mu), \tau_{\bar{\alpha}}^{-1}(Z_\omega)\rangle$ is a finite index subgroup of $G_\alpha \cong \mathbb{Z}^2$.
(5) If α_1,α_2 are distinct edges with $\alpha_1^+ = \alpha_2^+$, then
   - for each $g \in \pi_1\mathcal{G}$, $g\tau_{\alpha_1}(G_{\alpha_1})g^{-1}$ is not commensurable with $\tau_{\alpha_2}(G_{\alpha_2})$;
   - for each $g \in \pi_1\mathcal{G} - \tau_{\alpha_1}(G_{\alpha_1})$, $g\tau_{\alpha_1}(G_{\alpha_1})g^{-1}$ is not commensurable with $\tau_{\alpha_1}(G_{\alpha_1})$.

We conclude this section with a few basic consequence of Definition 2.13. First we apply a theorem of Bowditch to obtain that the hyperbolic quotients, F_μ, are actually hyperbolic relative to the subgroups coming from the incident edge groups.

Lemma 2.14. *Let $\mathcal{G}$ be an admissible graph of groups. For each vertex μ, let π_μ be the quotient map $\pi_\mu \colon G_\mu \to F_\mu$, where F_μ is the quotient G_μ/Z_μ. The group F_μ is hyperbolic relative to the collection $\{\pi_\mu(\tau_\alpha(G_\alpha)) : \alpha$ is an edge with $\alpha^+ = \mu\}$.*

Proof. Let I_μ be the set of edges α of $\mathcal{G}$ with $\alpha^+ = \mu$ and let $A_\alpha = \tau_\alpha(G_\alpha)$ for each $\alpha \in I_\alpha$. We want to show that $\{\pi_\mu(A_\alpha) \colon \alpha \in I\mu\}$ is an almost malnormal collection of quasiconvex subgroups as this implies F_μ is relatively hyperbolic by [Bow12, Theorem 7.11].

We first establish that $\pi_\mu(A_\alpha)$ is virtually cyclic for each $\alpha \in I_\mu$. By construction, $\pi_\mu(A_\alpha)$ is the quotient of A_α by $A_\alpha \cap Z_\mu$. Since $A_\alpha \cong \mathbb{Z}^2$ and F_μ is hyperbolic, A_α must intersect Z_μ in a non-trivial subgroup. Hence $\pi_\mu(A_\alpha)$ must be virtually cyclic. Note, this implies each $\pi_\mu(A_\alpha)$ is quasiconvex in F_μ as virtually cyclic subgroups of hyperbolic groups are always quasiconvex.

We now show the set $\{\pi_\mu(A_\alpha) : \alpha \in I_\mu\}$ is an almost malnormal collection of subgroups of F_μ. Since each $\pi_\mu(A_\alpha)$ is virtually cyclic, if the collection fails to be almost malnormal, there must be $\alpha_1, \alpha_2 \in I_\mu$ so that some conjugate of $\pi_\mu(A_{\alpha_1})$ is commensurable to $\pi_\mu(A_{\alpha_2})$ in F_μ. Because $Z_\mu \cong \mathbb{Z}$ and each $A_\alpha \cong \mathbb{Z}^2$, this would imply a conjugate of A_{α_1} is commensurable to A_{α_2} in $\pi_1\mathcal{G}$. As this would contradict Definition 2.13.(5), we must have that $\{\pi_\mu(A_\alpha) : \alpha \in I_\mu\}$ is an almost malnormal. The lemma now follows by applying [Bow12, Theorem 7.11]. □

Lastly, we note that by choosing appropriate infinite generating sets for the vertex groups G_μ, we can make Cayley graphs that are quasi-isometric to the hyperbolic groups F_μ as well as the electrification of F_μ by the cyclic subgroups from the incident edge groups. Recall each vertex group is a central extension $Z_\mu \to G_\mu \to F_\mu$ where Z_μ is cyclic and F_μ is hyperbolic.

Lemma 2.15. *Let $\mathcal{G}$ be an admissible graph of groups. Let I_μ be the set of edges α of $\mathcal{G}$ with $\alpha^+ = \mu$, then let $\mathcal{E}_\mu = \bigcup_{\alpha \in I_\mu} \tau_\alpha(G_\alpha)$. For each finite generating set J_μ of G_μ we have:*

1. *The quotient map $\pi_\mu \colon G_\mu \to F_\mu$ induces a quasi-isometry*
$$\pi_\mu \colon \operatorname{Cay}(G_\mu, J_\mu \cup Z_\mu) \to \operatorname{Cay}(F_\mu, \pi_\mu(J_\mu)),$$
in particular $\operatorname{Cay}(G_\mu, J_\mu \cup Z_\mu)$ is hyperbolic and hyperbolic relative to the collection $\{g\tau_\alpha(G_\alpha) : \alpha \in I_\mu$ and $g \in G_\mu\}$.
2. *The quotient map $\pi_\mu \colon G_\mu \to F_\mu$ induces a quasi-isometry*
$$\pi_\mu \colon \operatorname{Cay}(G_\mu, J_\mu \cup \mathcal{E}_\mu) \to \operatorname{Cay}(F_\mu, \pi_\mu(J_\mu \cup \mathcal{E}_\mu)).$$
Hence, $\operatorname{Cay}(G_\mu, J_\mu \cup \mathcal{E}_\mu)$ is hyperbolic and will be a quasi-tree whenever F_μ is virtually free.

The quasi-isometry constants are independent of $\mathcal{G}$.

Proof. The fact that the map are quasi-isometries follows from Lemma 2.5. The first relative hyperbolicity follows from Lemma 2.14. For the second, since F_μ is hyperbolic relative to the subgroups $\{\pi_\mu(\tau_\alpha(G_\alpha)) : \alpha \in I_\mu\}$ (Lemma 2.14), the graph $\operatorname{Cay}(F_\mu, \pi_\mu(J_\mu \cup \mathcal{E}_\mu))$ is hyperbolic. Moreover, if F_μ is virtually free, then $\mathrm{Cay}(F_\mu, \pi_\mu(J_\mu))$ is a quasi-tree. Hence the fact that $\mathrm{Cay}(F_\mu, \pi_\mu(J_\mu \cup \mathcal{E}_\mu))$ is a quasi-tree is a consequence of Lemma 2.4. □

2.3. **Hierarchically hyperbolic groups.** As we will not directly require the full definition of a hierarchically hyperbolic space, we will only review the necessary data to define a hierarchically hyperbolic group. We direct the reader to [BHS19] or [Sis19] for complete details on the HHS axioms.

Fix $E \geqslant 1$. An *E–hierarchically hyperbolic space (HHS) structure* on a geodesic metric space $\mathcal{X}$ starts with a set $\mathfrak{S}$ indexing a collection of E–hyperbolic spaces $\{\mathcal{C}(V)\}_{V\in\mathfrak{S}}$. For each $V \in \mathfrak{S}$, there is an (E,E)–coarsely Lipschitz, E–coarsely surjective *projection map* $\varphi_V : \mathcal{X} \to \mathcal{C}(V)$. The set $\mathfrak{S}$ is also equipped with three combinatorial relations: nesting ($\sqsubseteq$), orthogonality ($\perp$), and transversality ($\pitchfork$). To be a hierarchically hyperbolic space structure for $\mathcal{X}$, the set $\mathfrak{S}$ and these relations and projections need to satisfy a number of axioms. The most relevant for us are:

- Every pair of distinct elements of $\mathfrak{S}$ is related by exactly one of $\sqsubseteq$, $\perp$, or $\pitchfork$.
- $\pitchfork$ and $\perp$ are both symmetric, while $\sqsubseteq$ is a partial order.
- If $V \perp W$ and $U \sqsubseteq V$, then $U \perp W$.
- If $V \sqsubsetneq W$ or $V \pitchfork W$, then there exists a distinguished subset $\rho^V_W \subseteq \mathcal{C}(W)$ with diameter at most E.

We use $\mathfrak{S}$ to denote the entire HHS structure (the spaces, projections, relations, and distinguished subsets) and the pair $(\mathcal{X}, \mathfrak{S})$ to denote the hierarchically hyperbolic space $\mathcal{X}$ equipped with the specific HHS structure $\mathfrak{S}$. An HHS structure can be transferred across a quasi-isometry $f : \mathcal{Y} \to \mathcal{X}$, by replacing the projection maps φ_V with $\varphi_V \circ f$. In particular, if a finitely generated group G acts metrically properly and coboundedly on an HHS $(\mathcal{X}, \mathfrak{S})$, then $\mathfrak{S}$ is also a hierarchically hyperbolic space structure for G equipped with any word metric (or equivalently any Cayley graph of G with respect to a finite generating set). However, the maps and relations defining $\mathfrak{S}$ need not be equivariant with respect to the action of G. If the HHS structure is compatible with the group action, then we can have the following stronger definition of a hierarchically hyperbolic *group*.

Definition 2.16. Suppose a finitely generated group G is acting isometrically on an E–hierarchically hyperbolic space $(\mathcal{X}, \mathfrak{S})$. We say $\mathfrak{S}$ is an *E–hierarchically hyperbolic group* structure if

1. (1) G acts metrically properly and coboundedly on $\mathcal{X}$.
2. (2) There is an $\sqsubseteq$, $\perp$, and $\pitchfork$ preserving action of G on the index set $\mathfrak{S}$ by bijections.
3. (3) $\mathfrak{S}$ has finitely many G–orbits.
4. (4) For each $V \in \mathfrak{S}$ and $g \in G$, there exists an isometry $g_V : \mathcal{C}(V) \to \mathcal{C}(gV)$ satisfying the following for all $V, U \in \mathfrak{S}$ and $g, h \in G$.
   - The map $(gh)_V : \mathcal{C}(V) \to \mathcal{C}(ghV)$ is equal to the map $g_{hV} \circ h_V : \mathcal{C}(V) \to \mathcal{C}(hV)$.
   - For each $x \in \mathcal{X}$, $g_V(\varphi_V(x)) = \varphi_{gV}(g \cdot x)$.
   - If $V \pitchfork U$ or $V \sqsubsetneq U$, then $g_V(\rho^V_U) = \rho^{gV}_{gU}$.

We say G is a *hierarchically hyperbolic group* (HHG) if there exists an HHS $(\mathcal{X}, \mathfrak{S})$ so that $\mathfrak{S}$ is an E–HHG structure for G for some $E \geqslant 1$.

Modulo the incompleteness of our description of a hierarchically hyperbolic space structure, the above definition of a hierarchically hyperbolic group is precise.

There are examples of finitely generated groups that have hierarchically hyperbolic *space* structures, but do not have any hierarchically hyperbolic *group* structures. In fact, there are groups that are not HHGs, but have finite index subgroups that are HHGs [PS23].

We will need the following proposition, originally due to Paul Plummer, to check that certain 3–manifold groups are not HHG.

Proposition 2.17 (Invariant quasiflats for virtually abelian subgroups)**.** *Let* $(G, \mathfrak{S})$ *be an HHG. Let* $A \subset G$ *be a virtually* $\mathbb{Z}^k$ *subgroup for some* $k \geqslant 1$*. Then there exists* $\ell \geqslant k$ *and* $U_1, \ldots, U_\ell \in \mathfrak{S}$ *such that the following hold:*

(1) $\{U_1, \ldots, U_\ell\}$ *is* A*–invariant.*
(2) $U_i \perp U_j$ *for* $1 \leqslant i < j \leqslant \ell$*.*
(3) *There exists* $L < \infty$ *such that* $\operatorname{diam}(\varphi_V(A)) \leqslant L$ *for* $V \in \mathfrak{S} - \{U_1, \ldots, U_\ell\}$*.*
(4) *For each* $i \leqslant \ell$*, the image* $\varphi_{U_i}(A)$ *of* A *in* $\mathcal{C}(U_i)$ *is a quasi-line.*

Hence the (A*–invariant) hierarchically quasiconvex hull* F_A *of* A *is quasi-isometric to* $\mathbb{Z}^\ell$*.*

Hierarchically quasiconvex hulls are discussed in [BHS19, Section 6].

Proof of Proposition 2.17. We adopt the standard convention that for $a, b \in G$ and $V \in \mathfrak{S}$, $d_V(a, b)$ denotes $d_V(\varphi_v(a), \varphi_V(b))$. Let $\mathbb{1}$ denote the identity in G and equip both A and G with word metrics from finite generating sets.

Apply [PS23, Theorem 5.1] to obtain a nonempty A–invariant set of elements $U_1, \ldots, U_\ell \in \mathfrak{S}$ such that

- $U_i \perp U_j$ for $1 \leqslant i < j \leqslant \ell$;
- if $W \in \mathfrak{S}$ has the property that $\varphi_W(A)$ is unbounded, then $W \sqsubseteq U_i$ for some i;
- $\varphi_{U_i}(A)$ is unbounded for each $i \leqslant \ell$.

Each $\varphi_{U_i}(A)$ is a quasiline: Since $\ell < \infty$, there is a finite-index subgroup $\ddot{A} \leqslant A$ such that $\ddot{A} \cdot U_i = U_i$ for all i. Assume, by passing to a further finite-index subgroup, that $\ddot{A} \cong \mathbb{Z}^k$. In particular, $\ddot{A}$ acts on each of the E–hyperbolic spaces $\mathcal{C}(U_i)$.

Since $\ddot{A}$ has finite index in A, and φ_{U_i} is (E, E)–coarsely lipschitz and $\ddot{A}$–equivariant, we have that $\varphi_{U_i}(A)$ and $\ddot{A} \cdot \varphi_{U_i}(\mathbb{1})$ lie at finite Hausdorff distance. In particular, the above choice of the U_i implies the orbit $\ddot{A} \cdot \varphi_{U_i}(\mathbb{1})$ is unbounded for each U_i. Proposition 3.1 of [CCMT15] therefore provides four options for the action of $\ddot{A}$ on $\mathcal{C}(U_i)$: *focal*, *general*, *horocyclic*, or *lineal*. We verify the the action must be lineal.

Since $\ddot{A}$ is abelian, it does not contain a free sub-semigroup and hence the action on $\mathcal{C}(U_i)$ action cannot be *focal* or *general*. By [DHS20, Theorem 3.1], any infinite-order element of $\ddot{A}$ is loxodromic on $\mathcal{C}(U_i)$, so the action is not *horocyclic*. Hence the action is *lineal*. In particular, the orbit $\ddot{A} \cdot \varphi_{U_i}(\mathbb{1})$, with the metric inherited from $\mathcal{C}(U_i)$, is (C, C)–quasi-isometric to $\mathbb{Z}$ and C–quasiconvex, where C depends on $\ddot{A}$ and the HHS constant E. Up to enlarging C, we can assume that $\varphi_{U_i}(A)$ is a C–quasiconvex (C, C)–quasiline. Moreover, since there are finitely many i, we can assume that the same constant C works for all i.

Bounding remaining domains: We now bound the diameter of $\varphi_V(A)$, $V \notin \{U_1, \ldots, U_\ell\}$.

Claim 2.18. *There exists* $L \geqslant 0$ *such that* $\operatorname{diam}(\varphi_V(A)) \leqslant L$ *for all* $V \in \mathfrak{S} - \{U_1, \ldots, U_\ell\}$*.*

Proof of Claim 2.18. It suffices to prove the claim for the finite-index subgroup $\ddot{A}$ of A, since the maps φ_V are all (E, E)–coarsely lipschitz. Choose $a_1, \ldots, a_k \in A$ such that $a_1, \ldots, a_k$

generate the finite-index subgroup $\ddot{A}$ isomorphic to $\mathbb{Z}^k$. For any $g \in G$, let $\mathrm{Big}(g)$ be the set $\{W \in \mathfrak{S} : \mathrm{diam}(\varphi_W(\langle g \rangle)) = \infty\}$. By [DHS17, Lemma 6.7], $\mathrm{Big}(g)$ is a finite, pairwise orthogonal subset of $\mathfrak{S}$ for any $g \in G$. Moreover, $\mathrm{Big}(g)$ is non-empty whenever g has infinite order by [DHS17, Proposition 6.4].

We claim that $\mathrm{Big}(a_i)$ is a non-empty subset of $\{U_1, \ldots, U_\ell\}$ for all i. Let $W \in \mathrm{Big}(a_i)$. By [PS23, Theorem 5.1] or [DHS17, Lemma 6.3, Proposition 6.4], there is $m \in \mathbb{N}$ so that a_i^m fixes W and has unbounded orbits on $\mathcal{C}(W)$. By the choice of the U_j, there exists j such that $W \sqsubseteq U_j$. If $W \neq U_j$, then $W \sqsubsetneq U_j$. Hence, $\rho^W_{U_j}$ is defined and is a subset of $\mathcal{C}(U_j)$ of diameter at most E. Since $\ddot{A}$ has unbounded orbits in $\mathcal{C}(U_j)$, there is $g \in \ddot{A}$ such that $d_{U_j}(\rho^W_{U_j}, g\rho^W_{U_j}) > 10^9 E$. By the definition of an HHG and the fact that $\ddot{A}$ fixes U_j, we have $g\rho^W_{U_j} = \rho^{gW}_{U_j}$, so $gW \neq W$. Now, $ga_i^m g^{-1}$ has unbounded orbits on $\mathcal{C}(gW)$, but $ga_i^m g^{-1} = a_i^m$. Hence, $W, gW \in \mathrm{Big}(a_i)$, but they are not orthogonal by [DHS17, Lemma 1.5]. This contradicts that the elements of $\mathrm{Big}(a_i)$ are pairwise orthogonal. Hence, $W = U_j$.

Since we have shown that $\mathrm{Big}(a_i) \subseteq \{U_1, \ldots, U_\ell\}$ for all i, [DHS17, Proposition 6.4] provides a constant $D(a_i)$ such that $\mathrm{diam}(\varphi_V(\langle a_i \rangle)) \leqslant D(a_i)$ for all $V \in \mathfrak{S} - \{U_1, \ldots, U_\ell\}$. Let $D = \max_{1 \leqslant i \leqslant k} D(a_i)$. For any $b \in \ddot{A}$, write $b = a_1^{n_1} \cdots a_k^{n_k}$. Since $a_1^{-n_1} V \notin \{U_1, \ldots, U_k\}$, we have

$$d_V(\mathbb{1}, b) \leqslant d_{a_1^{-n_1}V}(\mathbb{1}, a_2^{n_2} \cdots a_k^{n_k}) + d_{a_1^{-n_1}V}(\mathbb{1}, a_1^{-n_1}) \leqslant d_{a_1^{-n_1}V}(\mathbb{1}, a_2^{n_2} \cdots a_k^{n_k}) + D,$$

and we get $d_V(\mathbb{1}, b) \leqslant kD$ by induction. This bounds $\mathrm{diam}(\varphi_V(\ddot{A}))$, which proves the claim. □

This proves the enumerated statements. The distance formula in a HHG [BHS19, Theorem 4.5] now shows that the hull F_A of A is quasi-isometric to the product $\prod_{i=1}^{\ell} \varphi_{U_i}(A)$, i.e., to the product of ℓ quasilines, i.e., to $\mathbb{Z}^\ell$. Since $\ddot{A} \cong \mathbb{Z}^k$ acts properly on F_A, we must have $k \leqslant \ell$. □

2.4. **Combinatorial hierarchical hyperbolicity.** To verify that our admissible groups are hierarchically hyperbolic groups, we will employ the *combinatorial hierarchical hyperbolicity* machinery introduced in [BHMS20]. This allows us to forgo checking the axioms directly, and instead extract hierarchical hyperbolicity from an action on a well chosen simplicial complex. We recall the required definitions and theorems for this approach.

Definition 2.19 (Join, link, and star)**.** Let Y be a flag simplicial complex. If Q, Z are disjoint flag subcomplexes of Y so that every vertex of Q is joined by an edge to Z, then the *join* of Q and Z, $Q \star Z$, is the subcomplex of Y spanned by Q and Z. Given a simplex Δ of Y, the *link* of Δ, $\mathrm{lk}(\Delta)$, is the subcomplex of Y spanned by the vertices of Y that are joined by an edge to all the vertices of Δ. The *star* of Δ, $\mathrm{st}(\Delta)$, is the join $\Delta \star \mathrm{lk}(\Delta)$. We consider $\varnothing$ as a simplex of Y whose link and star are both Y.

Definition 2.20. Given a flag simplicial complex Y, a Y*–graph* is any graph W whose vertices are maximal simplices of Y. Here maximal means not contained in a larger simplex.

If W is a Y–graph for the flag simplicial complex Y, we define the W–augmented graph Y^{+W} as the graph with the same vertex set as Y and with two types of edges:

(1) (Y–edge) If two vertices $y_1, y_2 \in Y$ are joined by an edge in Y, then y_1 and y_2 are joined by an edge in Y^{+W}.
(2) (W–edge) If Δ_1 and Δ_2 are maximal simplices of Y that are joined by an edge in W, then each vertex of Δ_1 is joined by an edge to each vertex of Δ_2 in Y^{+W}.

We note that if a group G acts by simplicial automorphisms on Y that is an isometry of Y^{+W}, then there is an induced action by isometries of G on W.

Definition 2.21. Let Δ and Δ' be simplices of the flag simplicial complex Y. We write $\Delta \sim \Delta'$ if $\mathrm{lk}(\Delta) = \mathrm{lk}(\Delta')$. We define the *saturation* of Δ, $\mathrm{Sat}(\Delta)$, to be the set of vertices of Y contained in a simplex in the $\sim$–equivalence class of Δ. That is $x \in \mathrm{Sat}(\Delta)$ if and only if there exists $\Delta' \sim \Delta$ so that x is a vertex of Δ'.

Definition 2.22. Let W be a Y–graph. For each simplex Δ of Y, define Y_Δ to be the subgraph of Y^{+W} spanned by the vertices of $Y^{+W} - \mathrm{Sat}(\Delta)$.

Define $\mathcal{C}(\Delta)$ to be the subgraph of Y_Δ spanned by the vertices in $\mathrm{lk}(\Delta)$. Note, we are taking the link in Y, not in Y^{+W}, and then considering the subgraph of Y_Δ induced by those vertices. We give $\mathcal{C}(\Delta)$ its intrinsic path metric (as opposed to the metric induced as a subset of Y_Δ). By construction, we have $\mathcal{C}(\Delta) = \mathcal{C}(\Delta')$ whenever $\Delta \sim \Delta'$. Note, since $\varnothing$ is a simplex of Y with $\mathrm{lk}(\varnothing) = Y$, we have $Y_\varnothing = \mathcal{C}(\varnothing) = Y^{+W}$.

Definition 2.23. Let $\delta \geqslant 0$, Y be a flag simplicial complex and W be a Y–graph. The pair (Y, W) is a *δ–combinatorial HHS* if the following are satisfied.

(I) Any chain of the form $\mathrm{lk}(\Delta_1) \subsetneq \mathrm{lk}(\Delta_2) \subsetneq \dots$ has length at most δ.
(II) For each non-maximal simplex $\Delta \subset Y$, the space $\mathcal{C}(\Delta)$ is δ–hyperbolic.
(III) For each non-maximal simplex Δ, the inclusion $\mathcal{C}(\Delta) \to Y_\Delta$ is a (δ, δ)–quasi-isometric embedding.
(IV) Whenever Δ and Ω are non-maximal simplices of Y, there exists a (possibly empty) simplex Π of $\mathrm{lk}(\Delta)$ such that $\mathrm{lk}(\Delta \star \Pi) \subseteq \mathrm{lk}(\Omega)$ and for all non-maximal simplices Λ of Y so that $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta) \cap \mathrm{lk}(\Omega)$ either
   (a) $\mathrm{diam}(\mathcal{C}(\Lambda)) < \delta$ or;
   (b) $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta \star \Pi)$.
(V) For each non-maximal simplex $\Delta \subset Y$ and $x, y \in \mathrm{lk}(\Delta)$, if x and y are not joined by a Y–edge of Y^{+W}, but are joined by a W–edge of Y^{+W}, then there exits simplices $\Lambda_x, \Lambda_y \subseteq \mathrm{lk}(\Delta)$ so that $x \in \Lambda_x$, $y \in \Lambda_y$, and $\Delta \star \Lambda_x$ is joined by an edge of W to $\Delta \star \Lambda_y$.

Theorem 2.24 ([BHMS20, Thoerem 1.8])**.** *Let (Y, W) be a δ–combinatorial HHS.*

(1) The graph W is a connected and a hierarchically hyperbolic space.
(2) Suppose G is a finitely generated group that acts on Y by simplicial automorphism. If there are finitely many G–orbits of links of simplices of Y, and the action of G on Y induces a metrically proper and cobounded action on W, then G is a hierarchically hyperbolic group.

3. Statements of the main results

We now state the main result of the paper, summarise where the various parts of the proof are found, and then deduce our application to 3-manifolds.

Theorem 3.1. *Let $\mathcal{G}$ be an admissible graph of groups. Let $\mathcal{S}(T)$ and $W = W_{r,R}$ be the spaces from Definitions 5.1 and 5.7. For sufficiently large choices of $r \geqslant 0$ and $R \geqslant 0$, the pair $(\mathcal{S}(T), W)$ is a δ–combinatorial HHS with δ determined by $\mathcal{G}$.*

Moreover, $\pi_1\mathcal{G}$ is an HHG, because $\pi_1\mathcal{G}$ acts on $\mathcal{S}(T)$ with finitely many orbits of links of simplices, and the action on the set of maximal simplices of $\mathcal{S}(T)$ extends to a metrically proper and cobounded action on W.

Proof. Item (I) is verified in Lemma 6.4. Item (II) is verified in Proposition 6.8 and Lemma 6.2. Item (III) is immediate when $\Delta = \varnothing$ or $\mathcal{C}(\Delta)$ is bounded, while the other cases are verified in Lemmas 6.14 and Lemma 6.15 (with Corollary 6.1 guaranteeing that all cases are covered). Item (IV) is Lemma 6.6 and, finally, Item (V) is Lemma 6.5.

The statement on orbits of links is verified in Lemma 6.7, while the metrically proper and cobounded action is shown in Lemma 5.11. The conclusion that $\pi_1\mathcal{G}$ is an HHG then follows from Theorem 2.24. □

Theorem 3.1 proves that non-geometric graph manifolds are HHG.

Corollary 3.2. *If M is a non-geometric graph manifold, then $\pi_1 M$ is a hierarchically hyperbolic group, where the hyperbolic spaces in the HHG structure are all quasi-isometric to trees.*

Proof. Since $\pi_1 M$ has the structure of an admissible graph of groups, we can apply Theorem 3.1. The hyperbolic spaces in the HHS structure coming from a combinatorial HHS are the $\mathcal{C}(\Delta)$ (as stated in [BHMS20, Theorem 1.18]). These spaces are all quasi-isometric to trees in the case of $\pi_1 M$ by Proposition 6.8 (and Lemma 6.2 for the bounded $\mathcal{C}(\Delta)$). □

We can now combine Corollary 3.2 and Corollary 4.3 with results from the literature to classify when a 3-manifold group has an HHG structure in terms of the geometry of the prime pieces. We say that a flat 3-manifold is *octahedral* if it is the quotient of $\mathbb{R}^3$ by a 3–dimensional crystallographic group whose point group is conjugate in $GL_3(\mathbb{R})$ into $O_3(\mathbb{Z})$.

Theorem 3.3. *Let M be a closed oriented 3–manifold. $\pi_1 M$ is a hierarchically hyperbolic group if and only if M has no Nil, Sol, or non-octahedral flat manifolds in its prime decomposition.*

Proof. We first show that if M has no Nil, Sol, or non-octahedral flat manifolds in its prime decomposition, then $\pi_1 M$ is an HHG. Since being hyperbolic relative to HHGs will make $\pi_1 M$ an HHG [BHS19, Theorem 9.1], it suffices to prove that $\pi_1 M$ is an HHG whenever M is prime and not a Nil, Sol, or non-octahedral flat manifold.

We first analyse the possible geometric cases from the geometrisation theorem.

- $S^3, S^2 \times \mathbb{R}, \mathbb{H}^3$. In this case the fundamental group is hyperbolic, whence a hierarchically hyperbolic group.
- $\mathbb{R}^3$. The fact that the fundamental group of a manifold with geometry $\mathbb{R}^3$ is an HHG if and only if the manifold is octahedral follows from [PS23, Theorem 4.4].
- $\mathbb{H}^2 \times \mathbb{R}, \mathrm{PSL}_2(\mathbb{R})$. In these cases, the fundamental group is a central extension of $\mathbb{Z}$ by a hyperbolic surface group, so we can apply Corollary 4.3 to conclude it is an HHG (the $\mathbb{H}^2 \times \mathbb{R}$ case was previously known, see, e.g. [Hug22, Proposition 3.1]). For later purposes, note that this case also yields HHG fundamental groups when M is a $\mathbb{H}^2 \times \mathbb{R}$ manifold with toroidal boundary.

In the non-geometric case, $\pi_1 M$ is hyperbolic relative to subgroups each of which is either $\mathbb{Z}^2$ or the fundamental group of a non-geometric graph manifold (this is a consequence of [Dah03, Theorem 0.1] and is stated explicitly as [AFW15, Theorem 9.12]; see also [BW13, Corollary E]). Each peripheral is therefore an HHG, so the conclusion follows from [BHS19, Theorem 9.1].

We now assume $\pi_1 M$ has an HHG structure $\mathfrak{S}$ and show M cannot have a Nil, Sol, or non-octahedral flat manifold in its prime decomposition. If M is prime and has either Nil or Sol geometry, then $\pi_1 M$ cannot be an HHG since it would not have quadratic Dehn function, contradicting [BHS19, Corollary 7.5]. If M is prime and is a non-octahedral flat manifold, then $\pi_1 M$ is not an HHG by [PS23, Theorem 4.4].

For the non-prime case, let $M_1 \# \cdots \# M_n$ be the prime decomposition of M. Then $\pi_1 M$ is hyperbolic relative to $\pi_1 M_1, \dots, \pi_1 M_n$. As the peripheral subgroups in a relative hyperbolic group, each $\pi_1 M_i$ is strongly quasiconvex in $\pi_1 M$. Combining [RST23, Proposition 5.7] and [BHS19, Proposition 5.6], we have that restricting the projections in the HHG structure $\mathfrak{S}$ to the subgroup $\pi_1 M_i$ produces an HHS structure for $\pi_1 M_i$ (but not necessarily an HHG structure). As before, this says M_i cannot have Nil or Sol geometry.

To rule out non-octahedral flat geometry, suppose M_i is a flat manifold. Then, $\pi_1 M_i$ is virtually $\mathbb{Z}^3$ and Proposition 2.17 says there are $U_1, \dots, U_\ell \in \mathfrak{S}$ so that

- $\{U_1, \dots, U_\ell\}$ is pairwise orthogonal and $\pi_1 M_i$–invariant;
- $\operatorname{diam}(\varphi_V(\pi_1 M_i))$ is uniformly bounded for all $V \in \mathfrak{S} - \{U_1, \dots, U_\ell\}$;
- for each $i \in \{1, \dots, \ell\}$, $\varphi_{U_i}(\pi_1 M_i)$ is a quasi-line in $\mathcal{C}(U_i)$;
- the *hierarchically quasiconvex hull* of $\pi_1 M_i$ is quasi-isometric to $\mathbb{Z}^\ell$.

Since $\pi_1 M_i$ is strongly quasiconvex in $\pi_1 M$, it is undistorted and the hierarchically quasiconvex hull of $\pi_1 M_i$ is uniformly close to $\pi_1 M_i$ in $\pi_1 M$. Hence $\pi_1 M_i$ acts properly and cocompactly on its hierarchically quasiconvex hull. As $\pi_1 M_i$ it virtually $\mathbb{Z}^3$, this implies the number ℓ in the bulleted properties is 3. Hence, we can make an HHG (and not just HHS) structure for $\pi_1 M_i$ by using the three quasilines $\varphi_{U_1}(\pi_1 M_i), \varphi_{U_2}(\pi_1 M_i), \varphi_{U_3}(\pi_1 M_i)$ and a finite number of bounded diameter spaces (this is the standard HHG structure on $\mathbb{Z}^3$ with the $\varphi_{U_3}(\pi_1 M_i)$ replacing the x, y, z axes). By [PS23, Theorem 4.4], this means M_i must be an octahedral flat manifold. □

Remark 3.4 (Concrete description of octahedral flat 3–manifolds)**.** Combining [Hag14] and [Hod20], a crystallographic group is octahedral (in any dimension) if and only if it is cocompactly cubulated if and only if it is *Helly*. In [PS23], it is shown that for crystallographic groups, this is equivalent to being an HHG. However, the octahedral flat 3–manifolds can be explicitly listed, following Scott [Sco83]. Specifically, if M is a compact orientable flat 3–manifold, M is octahedral if and only if is one of the following:

- the 3–torus;
- made by gluing opposite faces of a cube with a $\frac{1}{2}$ or $\frac{1}{4}$ twist, on one pair;
- made by gluing opposite faces of a hexagonal prism with a $\frac{1}{3}$–twist of the hexagonal faces;
- the *Hantzsche–Wendt manifold*, which has point group $(\mathbb{Z}/2\mathbb{Z})^2$.

The third one is tricky to visualise as octahedral, but here is an explanation in pictures instead of matrices. Consider the tiling of $\mathbb{E}^3$ by hexagonal prisms; this is the universal cover of M and so $\pi_1 M$ acts freely with quotient the 3–manifold described above. In Figure 2, we show one of these cells, P.

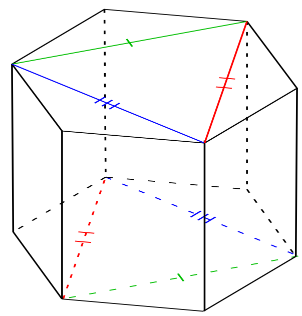

FIGURE 2. The $\frac{1}{3}$–twist prism manifold is octahedral. The set of planes through lines of the same colors are preserved by the $\frac{1}{3}$–twist.

Consider the six coloured segments in the figure, three in each of the two hexagonal faces of P. As indicated by the colours/labels, these come in three pairs of parallel segments, with each hexagonal face contributing one of the segments in each pair. Each parallel pair lies in a uniquely determined plane in $\mathbb{E}^3$. This set of three planes is invariant under an order 3 rotation of P about the central vertical line. Hence the $\pi_1 M$–orbit of this family of 3 planes is a set of planes in $\mathbb{E}^3$ falling into three parallelism classes. Cubulating the resulting wallspace (see e.g. [CN05]) therefore gives a proper cocompact action of $\pi_1 M$ on the standard tiling of $\mathbb{R}^3$ by 3–cubes, whence $\pi_1 M$ is octahedral by [Hag14] or [Hod20]. One can also directly compute a basis invariant under the point group.

According to [Sco83], there is only one more compact oriented flat 3–manifold. This is also constructed from a hexagonal prism by identifying opposite faces, but the hexagons are identified using a $\frac{1}{6}$ twist. (So, one can still cubulate $\pi_1 M$ as above, but this gives an action on $\mathbb{R}^6$, which is not cocompact.) This manifold is not octahedral since $O_3(\mathbb{Z})$ does not have an orientation-preserving element of order 6.

4. Quasi-lines from quasimorphisms

We now use quasimorphisms to construct the actions on quasi-lines. This is both an essential ingredient in our construction of a combinatorial HHS for an admissible graph of group and the key to proving that central extensions of $\mathbb{Z}$ by hyperbolic groups are HHGs.

We first build quasimorphisms for central extensions where the center is unbounded.

Lemma 4.1. *Suppose that the central extension of groups $\mathbb{Z} \xrightarrow{\iota} G \xrightarrow{\pi} F$ corresponds to a bounded element of $H^2(F,\mathbb{Z})$. Then there exists a quasimorphism $\phi\colon G \to \mathbb{Z}$ which is unbounded on $\iota(\mathbb{Z})$.*

Proof. The fact that the cohomology class associated to the central extension is bounded implies that there exists a (set-theoretic) section $s\colon F \to G$ so that there are only finitely many possible values of $s(f_1)s(f_2)s(f_1f_2)^{-1}$ as f_1, f_2 vary in F. Hence, if we define $c \in H^2(F,\mathbb{Z})$ by $c(f_1,f_2) = \iota^{-1}\left(s(f_1)s(f_2)s(f_1f_2)^{-1}\right)$, then the absolute value of $c(f_1,f_2)$ is bounded independently of f_1, f_2.

We now define ϕ. Any $x \in G$ can be written in a unique way as $s(f_x)\iota(t_x)$ for $f_x \in F$ and $t_x \in \mathbb{Z}$. Hence we can set $\phi(x) = t_x$. To show that ϕ is a quasimorphism note that, since $\iota(\mathbb{Z})$ is central and $s(f_1)s(f_2) = \iota(c(f_1, f_2))s(f_1 f_2)$, we have

$$xy = s(f_x)\iota(t_x)s(f_y)\iota(t_y) = s(f_x)s(f_y)\iota(t_x + t_y) = s(f_x f_y)\iota(c(f_x, f_y) + t_x + t_y).$$

Hence, $\phi(xy) = \phi(x) + \phi(y) + c(f_x, f_y)$, and we are done since the absolute value of $c(f_x, f_y)$ is uniformly bounded. □

We now use quasimorphisms to show that the vertex groups of an admissible graph of groups have the desired action of a quasi-line.

Lemma 4.2. *Let $\mathcal{G} = (\Gamma, \{G_\mu\}, \{G_\alpha\}, \{\tau_\alpha\})$ be an admissible graph of groups. For each edge α of $\mathcal{G}$, denote $C_\alpha = \tau_\alpha((\tau_{\bar{\alpha}})^{-1}(Z_{\alpha^-})) < G_{\alpha^+}$. Each vertex group G_μ has an infinite generating set S_μ so that the following hold.*

1. *$Cay(G_\mu, S_\mu)$ is quasi-isometric to a line,*
2. *the inclusion $Z_\mu \hookrightarrow \mathrm{Cay}(G_\mu, S_\mu)$ is a Z_μ–equivariant quasi-isometry,*
3. *for each edge α with $\alpha^+ = \mu$, C_α is bounded in $\mathrm{Cay}(G_\mu, S_\mu)$ (in fact, the bound is uniform over all α, μ since there are finitely many).*

Proof. By [ABO19, Lemma 4.15], if one can find an unbounded homogeneous quasimorphism $\bar{\phi} : G_\mu \to \mathbb{R}$, then there exists a generating set S_μ such that $\mathrm{Cay}(G_\mu, S_\mu)$ is quasi-isometric to a line and an element $g \in G$ acts loxodromically if and only $\bar{\phi}(g) \neq 0$. In particular, items (1), (2) and (3) are equivalent to the existence of a quasimorphism $\phi : G_\mu \to \mathbb{R}$ so that $\phi(Z_\mu)$ is unbounded, but $\phi(C_\alpha)$ is uniformly bounded for all edges α with $\alpha^+ = \mu$ (then $\bar{\phi}$ is the homogenization of ϕ). Note that each C_α does not intersect Z_μ in G_μ. Hence, the quotient map $\pi_\mu \colon G_\mu \to F_\mu$ is injective on C_α and we have that $\pi_\mu(C_\alpha) < \pi_\mu(\tau_\alpha(G_\alpha))$ is infinite.

We now construct certain auxiliary quasimorphisms. The first one, ϕ_μ, is just the homogenization of the quasimorphism from Lemma 4.1, which we can apply by condition Definition 2.13.(2) and the fact that every cohomology class of a hyperbolic group is bounded [Min01]. The other ones are constructed as follows. We claim that for each edge α of $\mathcal{G}$ with $\alpha^+ = \mu$, there is a homogeneous quasimorphism $\psi_\alpha \colon F_\mu \to \mathbb{R}$ so that $\psi_\alpha(\pi_\mu(c_\alpha)) = 1$, where c_α is a fixed generator of C_α, and $\psi_\alpha(c_{\alpha'}) = 0$ for all other edges α' with $\alpha'^+ = \mu$.

By Lemma 2.14, F_μ is hyperbolic relative to the subgroups $\pi_\mu(\tau_\alpha(G_\alpha))$ for edge of $\mathcal{G}$ with $\alpha^+ = \mu$. In particular the subgroups $\pi_\mu(\tau_\alpha(G_\alpha))$ are (jointly) hyperbolically embedded in F_μ. We can then appeal to [HO13, Theorem 4.2] to find the required quasimorphism. (The construction of Epstein–Fujiwara [EF97] should also be applicable to construct such quasimorphisms).

Let $\phi_\alpha = \psi_\alpha \circ \pi_\mu$ and observe that

$$\phi := \phi_\mu - \sum_{\alpha^+ = \mu} \phi_\mu(c_\alpha)\phi_\alpha$$

satisfies all the required properties. Thus, ϕ is the desired quasimorphism. □

To prove the first two bullet points of Lemma 4.2, we do not need the full definition of an admissible graph of groups. That is, if we have a central extension of groups $\mathbb{Z} \hookrightarrow G \overset{\pi}{\twoheadrightarrow} F$

corresponds to a bounded element of $H^2(F, \mathbb{Z})$, then [ABO19, Lemma 4.15] says the quasimorphism from Lemma 4.1 produces a generating set S for G so that $\mathrm{Cay}(G, S)$ is a quasi-line where the inclusion of the central $\mathbb{Z}$ is a $\mathbb{Z}$–equivariant quasi-isometry. This construction allows us to prove all such central extension are HHG.

Corollary 4.3. *If a group G is a central extension $Z \hookrightarrow G \overset{\pi}{\twoheadrightarrow} F$ where Z is an infinite cyclic group and F is a hyperbolic group, then G is a hierarchically hyperbolic group.*

Proof. Let z be the generator for Z and J be a finite symmetric generating set for G that contains z. We will identify F with the quotient G/Z and write elements of F as coset of Z. As described in the paragraph before Corollary 4.3, there is a generating set S for G so that $\mathrm{Cay}(G, S)$ is a quasi-line and the inclusion of Z into $\mathrm{Cay}(G, S)$ is a Z–equivariant quasi-isometry. Let $L = \mathrm{Cay}(G, S)$ and $H = \mathrm{Cay}(F, \pi(J))$. We will prove that the diagonal action of G on $L \times H$ is metrically proper and cobounded (where we fix, say, the ℓ_1–metric on said product). This will imply that G is an HHG as any group acting metrically properly and coboundedly on a product of hyperbolic spaces preserving the factors is an HHG; see [BHS19, Section 8.3] or [Hug22, Proposition 3.1].

To prove coboundedness, let r be large enough that every point in L is within r of an element of Z. Hence, for any vertex (k, hZ) of $L \times H$, there is a power z^n of z so that $d_L(k, z^n h) \leqslant 2r$. Thus $z^n h \cdot (1, Z) = (z^n h, hZ)$ is within $2r$ of (k, hZ) and hence the action of G of $L \times H$ is cobounded.

Moving on to metric properness, let $B^L(r)$ and $B^H(r)$ be the balls of radius $r \geqslant 0$ around the identity element in L and H respectively. Since G acts coboundedly on $L \times H$, every bounded diameter set of $L \times H$ is contained in some G–translate of $B^L(r) \times B^H(r)$ for some r. Hence it suffices to prove that the set of $g \in G$ such that $g(B^L(r) \times B^H(r)) \cap B^L(r) \times B^H(r) \neq \varnothing$ is finite.

If $g(B_L(r) \times B^H(r)) \cap (B^L(r) \times B^H(r)) \neq \varnothing$, then $gB^*(r) \cap B^*(r) \neq \varnothing$ for $* = L$ or H. The set $\{g \in G : gB^H(r) \cap B^H(r) \neq \varnothing\}$ is contained in the set $\{g \in G : d_H(gZ, Z) \leqslant 2r\}$. However, because F is finitely generated, the later is the union of finitely many cosets of Z. Now, since orbit maps of the action of Z on L are quasi-isometries, each coset of Z can only contain finitely many element g for which $gB^L(r) \cap B^L \neq \varnothing$. Together, these say that the set

$$\{g \in G : g\big(B^L(r) \times B^H(r)\big) \cap \big(B^L(r) \times B^H(r)\big) \neq \varnothing\}$$

is finite. □

We now translate the content of Lemma 4.2 into the language and notation of Bass–Serre space. Firstly, let us introduce the analogues of $\mathrm{Cay}(G_\mu, S_\mu)$.

Definition 4.4 (Space L_v)**.** Let $\mathcal{G} = (\Gamma, \{G_\mu\}, \{G_\alpha\}, \{\tau_{\alpha^\pm}\})$ be an admissible graph of groups. Let X be the Bass–Serre space associated to $\mathcal{G}$ and T be the Bass–Serre Tree of $\mathcal{G}$. For each vertex μ of $\mathcal{G}$, let S_μ be a generating set of G_μ as in Lemma 4.2. Without loss of generality we can assume $J_\mu \subseteq S_\mu$, where J_μ is the fixed finite generating set of G_μ. For a vertex $v \in T$ with $\mu = \check{v}$, let gG_μ be the corresponding coset of G_μ. Define L_v to be the graph with vertex set gG_μ and edges connecting $x, y \in gG_\mu$ if $x^{-1}y \in S_\mu$. Since L_v is obtained from X_v by adding extra edges to the same vertex set, there is a distance-non-increasing map $p_v \colon X_v \to L_v$ that is the identity on the vertices.

Proposition 4.5. *Let $\mathcal{G}$ be an admissible graph of groups with Bass–Serre tree T and Bass–Serre space X. Let e be an edge of T, with $v = e^+$ and $w = e^-$. Let $g, h \in G$ be such that $gZ_{\check{w}} \subseteq \tau_{\bar{e}}(X_w) \subset X_w$ and $hZ_{\check{v}} \subseteq X_v$. There exists $\xi \geqslant 1$, depending only on $\mathcal{G}$, so that:*

(1) $\operatorname{diam}(p_v \circ \tau_e \circ \tau_{\bar{e}}^{-1}(gZ_{\check{w}})) \leqslant \xi$.
(2) The restriction of p_v to $hZ_{\check{v}}$ (seen with the induced metric of X_v) is a (ξ, ξ)–quasi-isometry. In particular, the cosets $hZ_{\check{v}}$ are undistorted in X_v.
(3) Let $x \in X_v$. Then

$$d_{X_v}(x, \tau_e \circ \tau_{\bar{e}}^{-1}(gZ_{\check{w}})) \leqslant \xi d_{L_v}(p_v(x), p_v \circ \tau_e \circ \tau_{\bar{e}}^{-1}(gZ_{\check{w}})) + \xi$$

Proof. Item (1) is a verbatim translation of Lemma 4.2.(3) in the setting of Bass–Serre spaces. The bound is independent of e, g, h since there are finitely many orbits of vertices and edges.

For the proof of (2), fix a representative h of $hZ_{\check{v}}$. This determines a map $Z_{\check{v}} \to hZ_{\check{v}}$ defined as $z \mapsto hz$. Note that this maps is not canonical, as it depends on the choice of h, but this will not be a problem. We consider three different metrics on the set $Z_{\check{v}}$: the intrinsic word metric d_Z on $\operatorname{Cay}(Z_{\check{v}})$, the restriction of (X_v, d_{X_v}) using the inclusion $Z_{\check{v}} \to hZ_{\check{v}} \subseteq X_v$ and the restriction of (L_v, d_{L_v}) using the map p_v. In particular, by choosing an appropriate generating set on $\operatorname{Cay}(Z_{\check{v}})$, the maps $\operatorname{Cay}(Z_{\check{v}}) \to (hZ_{\check{v}}, d_v) \to (hZ_{\check{v}}, d_L)$ are all distance non-increasing. By Lemma 4.2.(2), the composition is a quasi-isometry, yielding that the $p_v|_{hZ_{\check{v}}} \colon (hZ_{\check{v}}, d_v) \to L_v$ is a quasi-isometry.

For the proof of (3), we will denote $\tau_e \circ \tau_{\bar{e}}^{-1}(gZ_{\check{w}})$ by C_e. Let $x \in \tau_e(X_e)$ and $x' \in C_e$ so that $d_{X_v}(x, C_e) = d_{X_v}(x, x')$. Now, there exist $g \in G$ so that x is contained in the coset $gZ_{\check{v}}$ in X_v. Because we have proved Item (2), there is $\kappa \geqslant 1$, depending only on $\mathcal{G}$ so that the restriction of p_v to $gZ_{\check{v}}$ is a (κ, κ)–quasi-isometry. In particular, there must be $\bar{x} \in C_e$ so that $d_{L_v}(p_v(\bar{x}), p_v(gZ_{\check{v}})) \leqslant \kappa$. Moreover, we can choose κ so that $\operatorname{diam}(p_v(C_e)) \leqslant \kappa$ as well. Using that $p_v \colon X_v \to L_v$ is distance non-increasing, we now have

$$d_{L_v}(x, x') \leqslant d_{X_v}(x, x') \leqslant d_{X_v}(x, \bar{x}) \leqslant \kappa d_{L_v}(p_v(x), p_v(\bar{x})) + \kappa \leqslant \kappa d_{L_v}(p_v(x), p_v(x')) + \kappa^2 + \kappa,$$

which implies

$$d_{L_v}(p_v(x), p_v(C_e)) \leqslant d_{X_v}(x, C_e) \leqslant \kappa d_{L_v}(p_v(x), p_v(C_e)) + \kappa^2 + \kappa.$$

The result follows by taking $\xi = \kappa^2 + \kappa$. □

We remark that the statements of Proposition 4.5 are concerned only with the metrics of the vertex spaces X_v and not on the metric on all of X, where X_v and X_e maybe distorted. In the sequel, we will often use Proposition 4.5 to establish a uniform bound on distances in X_v or X_e and then use Lemma 2.12 to translate this into a uniform bound on distances in X.

5. Defining a combinatorial HHS: a blow-up of the Bass–Serre tree

In this section, we describe how to construct the simplicial complex and graph that make a combinatorial HHS for an admissible graph of group. We prove that this construction satisfies the requirements of Theorem 2.24 in Section 6.

For the remainder of this section, let $\mathcal{G} = (\Gamma, \{G_\mu\}, \{G_\alpha\}, \{\tau_\alpha\})$ be an admissible graph of groups (Definition 2.13) and fix $G = \pi_1\mathcal{G}$. As in Section 2.2, we fix generating sets J_μ and J_α for the vertex and edge groups of $\mathcal{G}$. Let T denote the Bass–Serre tree of $\mathcal{G}$ and

X the Bass–Serre space from Definition 2.10. For vertices v and edges e of T, X_v and X_e will denote the vertex and edge spaces of X respectively. Recall that $T^{(0)}$ is the set $\{gG_\mu : g \in G, \mu \in V(\mathcal{G})\}$ and that for each $v \in T^{(0)}$, the elements of $X_v^{(0)}$ are precisely the elements of the coset $gG_\mu = v$. For an edge e of T, the maps τ_e and $\tau_{\bar{e}}$ denote the maps from the edge space X_e into the vertex spaces X_{e^+} and $X_{e^-} = X_{\bar{e}^+}$ described in Definition 2.10.

We also fix the generating sets S_μ from Lemma 4.2 for the vertex groups G_μ that produce Cayley graphs that are quasi-lines. Accordingly, for each vertex $v \in T^{(0)}$ we have the quasi-line L_v from Definition 4.4, which is the Cayley graph of the coset $gG_\mu = v$ with respect to the generating set S_μ. As described in Definition 4.4, there is a 1–Lipschitz map $p_v \colon X_v \to L_v$.

The simplicial complex for our combinatorial HHS will be the following complex $\mathcal{S}(T)$ that is a "blow-up" of the Bass–Serre tree T to include the vertices of each vertex space X_v at each vertex $v \in T$.

Definition 5.1. Let $Q = \bigsqcup_{v \in T^{(0)}} X_v^{(0)}$. Define the function $\nu \colon T^{(0)} \sqcup Q \to T^{(0)}$ as the identity on $T^{(0)}$ and as $\nu(s) = v$ if $s \in X_v$.

Let $\mathcal{S}(T)$ be the flag simplicial complex with vertex set $T^{(0)} \sqcup Q$ and the following two types of edges. First, each $s \in Q$ is connected to $\nu(s)$. Second, two vertices $s, t \in \mathcal{S}(T)$ are connected if $\nu(s), \nu(t)$ are adjacent in T. Observe that ν extends to a simplicial map $\mathcal{S}(T) \to T$ that we still denote ν.

Having constructed our simplicial complex, we now need to define a graph W whose vertices are the maximal simplices of $\mathcal{S}(T)$. We start be describing the maximal simplices of $\mathcal{S}(T)$.

Lemma 5.2 (Maximal simplicies in $\mathcal{S}(T)$)**.** *The maximal simplices of $\mathcal{S}(T)$ are exactly the simplices of the form $\{s, \nu(s), t, \nu(t)\}$, where $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$ and $\nu(s), \nu(t)$ are adjacent in T. We denote such a simplex by $\Sigma(s,t)$.*

Proof. Consider a simplex $\Sigma = \{s, \nu(s), \nu(t), t\}$ of $\mathcal{S}(T)$, where $\nu(s), \nu(t)$ are adjacent in T. Suppose that Σ is non-maximal. There then exists a vertex u of $\mathcal{S}(T)$ that is adjacent to each of $s, \nu(s), t, \nu(t)$. Since ν is simplicial, this means that $\nu(u)$ is equal to, or adjacent to, each of $\nu(s), \nu(t)$. Since T is triangle-free and $\nu(s), \nu(t)$ are adjacent, $\nu(u)$ cannot be adjacent to both $\nu(s)$ and $\nu(t)$, so, without loss of generality, $\nu(u) = \nu(s)$. Since u is different from s and $\nu(s)$, we therefore have that $\nu^{-1}(\nu(s))$ contains a 3–cycle with vertices $s, u, \nu(s)$. This contradicts the definition of $\mathcal{S}(T)$. Thus Σ is a maximal simplex.

Conversely, let Σ be a maximal simplex of $\mathcal{S}(T)$. Since $\nu : \mathcal{S}(T) \to T$ is simplicial, $\nu(\Sigma)$ is either a vertex of T or an edge of T. If $\nu(\Sigma)$ is a vertex, then T has some vertex v adjacent to $\nu(\Sigma)$ as T is a connected graph with at least two vertices. But then $v \star \Sigma$ is a simplex of $\mathcal{S}(T)$ properly containing Σ. Hence $\nu(\Sigma)$ must be an edge joining two vertices, $\nu(s)$ and $\nu(t)$, of T. So, Σ has the form $\Delta_s \star \Delta_t$, where Δ_s is a simplex projecting to $\nu(s)$ and Δ_t projects to $\nu(t)$. Maximality of Σ implies that Δ_s, Δ_t are edges, as required. □

Our goal is to define the edges in W so that G has a metrically proper and cobounded action on W, and so that we can verify the conditions of a combinatorial HHS. To accomplish the former, we want to associate to each maximal simplex $\Sigma(s,t)$ a uniformly bounded diameter subset of X and then declare two maximal simplices to be joined by any edge if their corresponding bounded diameter subsets are close in X. To facilitate this, we use the following "coarse level sets" of the map $p_v\colon X_v \to L_v$ from Definition 4.4.

Definition 5.3 ("Level surfaces")**.** Let $v \in T^{(0)}$ and $s \in X_v^{(0)}$. For $r \geqslant 0$, define $\sigma_r(s)$ to be the set

$$\sigma_r(s) := \{x \in X_v : d_{L_v}(p_v(s), p_v(x)) \leqslant r\}.$$

While we will not use this fact directly, it is helpful to think of the vertex spaces X_v as having a product structure in which the subspaces parallel to one factor are the $\sigma_r(s)$ and the subspaces parallel to the other factors are cosets of the gZ_μ. Thus, if one compares to the motivating case of a graph manifold, we can think of the $\sigma_r(s)$ as the "level surfaces" of the vertex spaces and the gZ_μ (which are quasi-isometric to the L_v) can be thought of as "lines".

The intersection of these "level surfaces" gives us a bounded diameter subset of X associated to a maximal simplex.

Definition 5.4 (Coarse points of maximal simplices)**.** Let $\mathcal{N}_c(\cdot)$ denote the c–neighborhood of a set in X. Given $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$ so that $\nu(s)$ and $\nu(t)$ are joined by an edge e of T with $e^+ = \nu(s)$, define

$$P_r(s,t) := \mathcal{N}_1(\sigma_r(s)) \cap \mathcal{N}_1(\sigma_r(t)).$$

Since $\sigma_r(s)$ and $\sigma_r(t)$ are in different vertex spaces (X_v vs X_w), $P_r(s,t)$ is precisely the set of vertices $x \in X_e$ so that $\tau_e(x) \in \sigma_t(s)$ and $\tau_{\bar{e}}(x) \in \sigma_r(t)$.

Lemma 5.5. *There exists $r_0 > 0$ such that for all $r \geqslant r_0$ there exists $\xi \geqslant 0$ so that the following holds. Let $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$ be such that $\nu(s), \nu(t)$ are joined by an edge of T. Then*

- *$P_r(s,t)$ is non-empty and has diameter at most ξ;*
- *the map $(s,t) \to P_r(s,t)$ is a (ξ,ξ)–coarsely Lipschitz, ξ–coarse map from $L_{\nu(s)} \times L_{\nu(t)}$ to X.*

Proof. Let $v = \nu(s)$ and $w = \nu(t)$, then let e be the edge of T from $e^- = w$ to $e^+ = v$. The key tool for the proof is the following quasi-isometry from X_e to $L_v \times L_w$.

Claim 5.6. *For each edge e of T, the diagonal map $\Phi_e\colon X_e \to L_{e^+} \times L_{e^-}$ given by*

$$\Phi_e(x) = (p_{e^+} \circ \tau_e(x), p_{e^-} \circ \tau_{\bar{e}}(x))$$

is a uniform quasi-isometry with $\Phi_e(g \cdot x) = g \cdot \Phi_e(x)$ for each $g \in \mathrm{Stab}_G(e)$ and $x \in X_e$.

Proof. Let $e^+ = v, e^- = w$, then let $\mu = \check{v}$, $\omega = \check{w}$, and $\alpha = \check{e}$. Equip G_α with the metric coming from $\mathrm{Cay}(G_\alpha, J_\alpha)$ and $gt_\alpha Z_\mu$, gZ_ω with the metrics as subsets of X_v. By Proposition 4.5.(2), this metric is quasi-isometric to any intrinsic metric on the cosets coming from a finite generating set of Z_μ and Z_ω.

Let $g \in G$ so that w is the coset gG_ω and v is the coset $gt_\alpha G_\mu$. If we let z_μ and z_ω be arbitrary elements of Z_μ and Z_ω respectively, define $\phi\colon gt_\alpha Z_\mu \times gZ_\omega \to G_\alpha$ by

$$(gt_\alpha z_\mu, gz_\omega) \mapsto \tau_\alpha^{-1}(z_\mu)\tau_{\bar{\alpha}}^{-1}(z_\omega).$$

Definition 2.13 says $G_\alpha \cong \mathbb{Z}^2$ and that $\langle \tau_\alpha^{-1}(Z_\mu), \tau_{\bar{\alpha}}^{-1}(Z_\omega)\rangle$ is a finite index subgroup of G_α. Hence, ϕ is a quasi-isometry.

Now define a map $\theta\colon G_\alpha \to X_e$ by $\theta(a) = \tau_{\bar{e}}^{-1}(g\tau_{\bar{\alpha}}(a))$. The construction of the Bass–Serre space tells us θ gives an isometry $\theta\colon \mathrm{Cay}(G_\alpha, J_\alpha) \to X_e$. Moreover, $\theta(a)$ also equals $\tau_e^{-1}(gt_\alpha\tau_\alpha(a)$ for each $a \in G_\alpha$. Thus, we have

$$p_v \circ \tau_e(\theta(\phi(gt_\alpha z_\mu, gz_\omega))) = p_v(gt_\alpha z_\mu) \text{ and } p_w \circ \tau_{\bar{e}}(\theta(\phi(gt_\alpha z_\mu, gz_\omega))) = p_w(gz_\omega)$$

for each $z_\mu \in Z_\mu$ and $z_\omega \in Z_\omega$. As $p_v|_{gt_\alpha Z_\mu}$ and $p_w|_{gZ_\omega}$ are uniform quasi-isometries by Proposition 4.5.(2), $p_v \times p_w\colon gt_\alpha Z \times gZ_\omega \to L_v \times L_w$ is a quasi-isometry. Hence,

$$(p_v \times p_w) \circ (\tau_e \times \tau_{\bar{e}}) = (p_v \times p_w) \circ \phi^{-1} \circ \theta^{-1}$$

is a quasi-isometry (here ϕ^{-1} is any quasi-inverse of ϕ that inverts ϕ on its image).

The constants of all these quasi-isometries can be chosen independent of α because $\mathcal{G}$ has finitely many edges. □

Let $\Phi_e\colon X_e \to L_v \times L_w$ be the quasi-isometry

$$x \to (p_v(\tau_{\bar{e}}(x)), p_w(\tau_e(x)))$$

from Claim 5.6. Because Φ_e is coarsely onto, there exists $x \in X_e$ and $r_0 > 0$ so that $p_v \circ \tau_e(x)$ is within R_0 of $p_v(s)$ in L_v and $p_w \circ \tau_{\bar{e}}(x)$ is within r_0 of $p_w(t)$ in L_w. Thus, $x \in \mathcal{N}_1(\sigma_r(s)) \cap \mathcal{N}_1(\sigma_r(t)) = P_r(s,t)$ is non-empty for each $r \geqslant r_0$.

When $P_r(s,t)$ is non-empty, then $\Phi_e(P_r(s,t)) \subseteq p_v(\sigma_r(s)) \times p_w(\sigma_r(t))$, which is a bounded diameter subset of $L_v \times L_w$. Since Φ_e is a quasi-isometry and the inclusion of X_e into X is 1-Lipschitz, $P_r(s,t)$ is then a bounded diameter subset of both X_e and X.

Finally, the map $\Psi_e\colon L_v \times L_w \to X_e$ given by $(s,t) \to P_r(s,t)$ is a quasi-inverse of Φ_e. Since the inclusion of X_e into X is 1-Lipschitz, the extension Ψ_e into X will be coarsely Lipschitz (and in fact uniformly so since there are only finitely many orbits of edges). □

We can now define the edges in our graph W, relying on the "level surfaces" $\sigma_r(s)$ from Definition 5.3 and coarse points $P_r(s,t)$ that arise as their intersections as in Definition 5.4.

Definition 5.7. (W–edges.) For $r, R > 0$, let $W = W_{r,R}$ be the graph defined as follows. The vertices of W are the maximal simplices of $\mathcal{S}(T)$. Two simplices $\Sigma(s,t)$ and $\Sigma(s',t')$ are adjacent if and only if one of the following holds.

- $\nu(s) = \nu(s')$ and $d_X(P_r(s,t), P_r(s',t')) \leqslant R$.
- $s = s'$ and $d_X(\sigma_r(t), \sigma_r(t')) \leqslant R + 2$.

Remark 5.8. In both cases, being joined by an edge of W implies that the two maximal simplices share a common vertex of T.

The first type of edge of W is needed to assure that W is quasi-isometric to the Bass–Serre space X. The second type is needed to arrange a fine combinatorial constraint in the definition of combinatorial HHS. To prove that G acts metrically properly on W we start by showing the second type of edges gives a similar bound to the first type.

Proposition 5.9. *There exists $r_1 \geqslant 0$ so that for each $r \geqslant r_1$, there exists a monotone diverging function $\Phi\colon [0,\infty) \to [0,\infty)$ so that the following holds. Consider two vertices v_1, v_2 of the Bass–Serre tree T at distance 2 from each other, with w being the vertex at distance 1 from both. Suppose that $t_1, t_2, s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ are such that $\nu(t_i) = v_i$ and $\nu(s) = w$. Then*

$$d_X(P_r(s,t_1), P_r(s,t_2)) \leqslant \Phi(d_X(\sigma_r(t_1), \sigma_r(t_2))).$$

Proof. The bulk of the technical work of the proposition is contained in the following claim.

Claim 5.10. *There exist $r_1 \geqslant 0$ so that for every $r \geqslant r_1$ there is a constant $c \geqslant 0$ and a monotone diverging function $\Phi' \colon [0,\infty) \to [0,\infty)$ so that the following holds. Let $s,t \in \mathcal{S}(T)^{(0)} - T^{(0)}$ so that $v = \nu(t)$ and $w = \nu(s)$ are joined by an edge e of T with $e^+ = v$ and $e^- = w$. Let $\mu = \check{v}$ and $\omega = \check{w}$.*

1. *If $g \in G$ is so that $gZ_\mu \subseteq X_v$, then $\sigma_r(t) \cap gZ_\mu \neq \varnothing$. Equivalently, if $gZ_\omega \subseteq X_w$, then $\sigma_r(s) \cap gZ_\omega \neq \varnothing$.*
2. *There exist $g \in G$ (depending on t), so that $gZ_\omega \subseteq \tau_{\bar{e}}(X_e)$ and*
$$\tau_e(\tau_{\bar{e}}^{-1}(gZ_\omega)) \subseteq \sigma_r(t) \cap \tau_e(X_e) \subseteq \mathcal{N}_c(\tau_e(\tau_{\bar{e}}^{-1}(gZ_\omega)).$$
3. *For each $x \in \sigma_r(t)$ there exists $x' \in \tau_e(X_e) \cap \sigma_r(t)$ with $d_X(x,x') \leqslant \Phi'(d_X(x,\tau_e(X_e)))$.*
4. *If $g_1, g_2 \in G$ so that $g_i Z_\omega \subseteq X_w$ for $i = 1,2$, then*
$$d_X(g_1 Z_\omega \cap \sigma_r(s), g_2 Z_\omega \cap \sigma_r(s)) \leqslant \Phi'(d_X(g_1 Z_\omega, g_2 Z_\omega)).$$

Proof. **Proof of** (1)**:** By Proposition 4.5.(2), the restriction of p_v to any coset $gZ_\mu \subseteq X_v$ is a uniform quasi-isometry. In particular, $p_v(gZ_\mu)$ uniformly coarsely covers L_v. Hence, there is some $r_1 \geqslant 0$ so that for all $r \geqslant r_1$, $p_v(\sigma_r(t)) \cap p_v(gZ_\mu) \neq \varnothing$, which implies $\sigma_r(t) \cap gZ_\mu \neq \varnothing$.

Proof of (2)**:** The set of cosets gZ_ω so that $gZ_\omega \subseteq \tau_{\bar{e}}(X_e)$ partition $\tau_{\bar{e}}(X_e)$. Since $\tau_{\bar{e}}$ is injective, the images of these cosets under $\tau_{\bar{e}}^{-1}$ will partition X_e. Since $p_v \circ \tau_e(X_e)$ coarsely covers L_v (Proposition 4.5.(2)), this partition implies there must be an $r_1' \geqslant 0$ and a coset $gZ_\omega \subset \tau_{\bar{e}}(X_e)$ so that $p_v(\sigma_{r_1'}(t)) \cap p_v(\tau_e \tau_{\bar{e}}(gZ_\omega)) \neq \varnothing$. By Proposition 4.5.(1), the diameter of $p_v(\tau_e \tau_{\bar{e}}(gZ_\omega))$ is uniformly bounded. Hence there is some $r_1 \geqslant r_1'$, so that whenever $r \geqslant r_1$, we have
$$\tau_e \tau_{\bar{e}}^{-1}(gZ_\omega) \subseteq \sigma_r(t) \cap \tau_e(X_e).$$

Now consider $x \in \sigma_r(t) \cap \tau_e(X_e)$. By construction, $d_{L_v}(p_v(x), p_v(\tau_e \tau_{\bar{e}}^{-1}(gZ_\omega))) \leqslant 2r$, so the fact that x is uniformly close in X_v to $\tau_e(\tau_{\bar{e}}^{-1}(gZ_\mu))$ follows from Proposition 4.5.(3). Since the inclusion $X_v \to X$ is distance non-increasing, there is $c \geqslant 0$ depending on $\mathcal{G}$ and r so that $d_X(x, \tau_e(\tau_{\bar{e}}^{-1}(gZ_\mu)) \leqslant c$.

Proof of (3)**:** Let r_1 be the lower bound from the proof of Item (1) so that for all $r \geqslant r_1$, $\sigma_r(t) \cap gZ_\mu \neq \varnothing$ when $gZ_\mu \subseteq X_v$. Fix $x \in \sigma_r(t)$, and let $\bar{x} \in \tau_e(X_e)$ be a point realizing $d_X(p, \tau_e(X_e))$. There exists some coset $gZ_\mu \in \tau_e(X_e)$ such that $\bar{x} \in gZ_\mu$ (because the cosets partition X_v). Let x' be a point of the intersection $\sigma_r(t) \cap gZ_\mu$. As $x, x' \in \sigma_r(t)$, we have
$$|d_{L_v}(p_v(\bar{x}), p_v(x)) - d_{L_v}(p_v(\bar{x}), p_v(x'))| \leqslant r.$$
By Proposition 4.5.(2), the map $p_v \colon gZ_\mu \to L_v$ is a (κ,κ)–quasi-isometry for some $\kappa \geqslant 1$ determined by $\mathcal{G}$. As $\bar{x}$ and x' both belong to gZ_μ, we have
$$\begin{aligned} d_X(\bar{x}, x') \leqslant d_{X_v}(\bar{x}, x') &\leqslant \kappa d_{L_v}(p_v(\bar{x}), p_v(x')) + \kappa \\ &\leqslant \kappa d_{L_v}(p_v(\bar{x}), p_v(x)) + \kappa r + \kappa \\ &\leqslant \kappa^2 d_{X_v}(\bar{x}, x) + \kappa^3 + \kappa r + \kappa \end{aligned}$$

By applying Lemma 2.12, the above bound on $d_X(\bar{x}, x')$ in terms of $d_{X_v}(\bar{x}, x)$ produces a diverging monotone function $\Phi' \colon [0,\infty) \to [0,\infty)$ so that $d_X(\bar{x}, x') \leqslant \Phi'(d_X(\bar{x}, x))$. The triangle inequality now yields

$$d_X(x,x') \leqslant d_X(x,\bar{x}) + d_X(\bar{x},x') \leqslant d_X(x,\bar{x}) + \Phi'(d_X(x,\bar{x})).$$

Since $d_X(x,\bar{x}) = d_X(x,\tau_e(X_e))$, this completes the proof of (3).

Proof of (4)**:** Let r_1 be the lower bound from the proof of Item (1) so that for all $r \geqslant r_1$, $\sigma_r(s) \cap gZ_\omega \neq \varnothing$ when $gZ_\omega \subseteq X_w$. Given g_1Z_ω and g_2Z_ω in X_w, let $x_i \in g_iZ_\omega$ so that $d_X(x_1,x_2) = d_X(g_1Z_\omega, g_2Z_\omega)$. Let z_1,z_2 be the elements of Z_ω so that $x_i = g_iz_i$ for $i = 1,2$.

We can assume that $x_1 \in \sigma_r(s) \cap g_1Z_\omega$ by the following argument. Suppose z is the element of Z_ω so that g_1z is a point in $\sigma_r(s) \cap g_1Z_\omega$. Because Z_ω is central in G_ω, we have

$$zz_1^{-1}x_1 = zz_1^{-1}g_1z_1 = g_1z \in \sigma_r(s) \cap g_1Z_\omega$$

and

$$zz_1^{-1}x_2 = zz_1^{-1}g_2z_2 = g_2zz_2^{-1}z_1 \in g_2Z_\omega.$$

Hence $zz_1^{-1}x_1 \in \sigma_r(s)$ and $zz_1^{-1}x_2$ are points in g_1Z_ω and g_2Z_ω that realise $d_X(g_1Z_\omega, g_2Z_\omega)$.

We can now proceed by a very similar argument as Item (3) using L_w instead of L_v. Let y_2 be a vertex in $\sigma_r(s) \cap g_2Z_\omega$. After replacing gZ_μ with g_2Z_ω and L_v with L_w in the proof of Item (3), we can repeat the same calculations with $x_1 = x$, $x_2 = \bar{x}$, and $y_2 = x'$, to produce

$$d_X(x_1,y_2) \leqslant d_X(x_1,x_2) + \Phi'(d_X(x_1,x_2)).$$

Since $x_1 \in \sigma_r(s) \cap g_1Z_\omega$, $y_2 \in \sigma_r(s) \cap g_2Z_\omega$, and $d_X(x_1,x_2) = d_X(g_1Z_\omega, g_2Z_\omega)$, this completes the proof of (4) with the same function Φ' as in the proof of (3). ☐

We now use Claim 5.10 to prove Proposition 5.9. Let $t_1, t_1, s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ be such that $\nu(t_i) = v_i$, $\nu(s) = w$ and $d_T(v_1,v_2) = 2$ with w the only vertex at distance one from both. Let e_i be the edge of T such that $e_i^+ = v_i$, $e_i^- = w$. Let $r \geqslant r_1$ where $r_1 \geqslant 0$ is the lower bound on r from Claim 5.10. Consider points $x_i \in \sigma_r(t_i)$ so that

$$d_X(x_1,x_2) = d_X(\sigma_r(t_1), \sigma_r(t_2)) := d.$$

We have to show that the $P_r(s,t_1)$ and $P_r(s,t_1)$ are at most some function of d apart in X. Because the edge spaces separate X, we have $d_X(x_i, \tau_{e_i}(X_{e_i})) \leqslant d$. By Claim 5.10.(3), there are points $x_i' \in \tau_{e_i}(X_{e_i}) \cap \sigma_r(t_i)$ such that $d_X(x_1',x_2') \leqslant d + 2\Phi'(d)$. Setting $\omega = \check{w}$, Claim 5.10.(2) gives us $c \geqslant 0$ and $g_i \in G$ so that each $g_iZ_\omega \subseteq X_w$ and

$$\tau_e(\tau_{\bar{e}}^{-1}(g_iZ_\omega)) \subseteq \sigma_r(t_i) \cap \tau_e(X_e) \subseteq \mathcal{N}_c(\tau_e(\tau_{\bar{e}}^{-1}(g_iZ_\omega)).$$

Since the map $\tau_e \circ \tau_{\bar{e}}^{-1}$ moves points distance at most 2, we have

$$d_X(g_1Z_\omega, g_2Z_\omega) \leqslant d_X(x_1',x_2') + 2(c+2) \leqslant d + 2\Phi'(d) + 2c + 4.$$

Applying Claim 5.10.(4), we have points $x_i'' \in g_iZ_\mu \cap \sigma_r(s)$ with

$$d_X(x_1'',x_2'') \leqslant \Phi'(d_X(g_1Z_\mu, g_2Z_\mu)) \leqslant \Phi'(d + 2\Phi'(d) + 2C + 4).$$

The points x_i'' are not quite in $P_r(s,t_i)$, but because $\tau_e(\tau_{\bar{e}}^{-1}(g_iZ_\omega)) \subseteq \sigma_r(t_i) \cap \tau_e(X_e)$, we have

$$\tau_{\bar{e}_i}^{-1}(x_i'') \in \tau_{e_i}^{-1}(\sigma_r(t_i)).$$

Since $x_i'' \in \sigma_r(s) \cap g_iZ_\omega$, we also have

$$\tau_{\bar{e}_i}^{-1}(x_i'') \in \tau_{\bar{e}_i}^{-1}(\sigma_r(s)).$$

Hence

$$\tau_{\bar{e}_i}^{-1}(x_i'') \in P_r(s,t_i),$$

which implies

$$d_X(P_r(s,t_1),P_r(s,t_2)) \leqslant \Phi'(d+2\Phi'(d)+2C+4)+2,$$

as desired for Proposition 5.9. □

Lemma 5.11. *Let r_0 and r_1 be as in Lemma 5.5 and Proposition 5.9 respectively. For all $r > \max\{r_0,r_1\}$ there is $R_0 = R_0(r) \geqslant 0$ so that*

- *for all $R > R_0$ the graph $W_{r,R}$ is connected;*
- *the G–action on $W_{r,R}$ by $g \cdot \Sigma(s,t) = \Sigma(gs,gt)$ is metrically proper and cobounded.*

Proof. Fix $r \geqslant \max\{r_0,r_1\}$ and let $W = W_{r,R}$ for a choice of R decided below.

W **is connected:** Because the Bass–Serre tree T is connected, given any two maximal simplices $\Sigma(s,t),\Sigma(s',t')$ of $\mathcal{S}(T)$, we can find a sequence of maximal simplices $\Sigma(s_i,t_i)$ so that $\nu(\Sigma(s_i,t_i))$ produces a path in T from $\nu(\Sigma(s,t))$ to $\nu(\Sigma(s',t'))$. Hence, it suffices to prove that two vertices $\Sigma(s,t),\Sigma(s',t') \in W$ with $\nu(s)=\nu(s')$ can connected by a path in W.

First assume $\nu(s)=\nu(s')$ and $\nu(t)=\nu(t')$. Then $P_r(s,t)$ and $P_r(s',t')$ are both subsets of X_e. Let e be the edge of T between $\nu(s)$ and $\nu(t)$ and let h be an element of G so that $\mathrm{Stab}_G(e) = hG_{\check{e}}h^{-1}$. Because $\mathrm{Stab}_G(e)$ acts transitively on the vertices of X_e, there is $k \in \mathrm{Stab}_G(e)$ so that $P_r(ks,kt) \cap P_r(s',t') \neq \varnothing$ and hence $\Sigma(ks,kt)$ is joined by a W–edge to $\Sigma(s',t')$. Because $\mathrm{Stab}_G(e)$ is generated by the finite set $hJ_{\check{e}}h^{-1}$, $\Sigma(s,t)$ will be connected to $\Sigma(ks,kt)$—and hence $\Sigma(s',t')$—if $\Sigma(gs,gt)$ is connected to $\Sigma(s,t)$ by a W–edge for each $g \in hJ_{\check{e}}h^{-1}$. There exists $R_1 \geqslant r$ depending only on $J_{\check{e}}$ so that $d_{X_e}(P_r(s,t),P_r(gs,gt)) \leqslant R_1$ for all $g \in hJ_{\check{e}}h^{-1}$. Thus, $\Sigma(gs,gt)$ is connected to $\Sigma(s,t)$ by a W–edge for each $g \in hJ_{\check{e}}h^{-1}$ provided $R \geqslant R_1$.

Now assume $\nu(s)=\nu(s')$, but $\nu(t) \neq \nu(t')$. Let e_1 be the edge of T from $\nu(t)$ to $\nu(s)$ and e_2 be the edge of T from $\nu(t')$ to $\nu(s')=\nu(s)$. Let $v=\nu(s)$ and h be an element of G so that $\mathrm{Stab}_G(v) = hG_{\check{v}}h^{-1}$. Because $\mathrm{Stab}_G(v)$ acts transitively on the vertices of X_v, there is $k \in \mathrm{Stab}_G(v)$ so that $k \cdot \tau_{e_1}(P_r(s,t)) \cap \tau_{e_2}(P_r(s',t')) \neq \varnothing$. Thus, $d_X(P_r(ks,kt),P_r(s',t')) \leqslant 2$. Hence, if $R \geqslant 2$, then $\Sigma(ks,kt)$ is joined by a W–edge to $\Sigma(s',t')$. Because $\mathrm{Stab}_G(v)$ is generated by the finite set $hJ_{\check{v}}h^{-1}$, $\Sigma(s,t)$ will be connected to $\Sigma(ks,kt)$ if $\Sigma(gs,gt)$ is connected to $\Sigma(s,t)$ by a W–edge for each $g \in hJ_{\check{v}}h^{-1}$. There exist $R_2 \geqslant r$ depending only on $J_{\check{v}}$ so that $d_X(\tau_{e_1}(P_r(s,t)),g\tau_{e_1}(P_r(s,t))) \leqslant R_2$ for all $g \in hJ_{\check{v}}h^{-1}$. Thus, $d_X(P_r(gs,gt),P_r(s,t)) \leqslant R_2+2$ for all $g \in hJ_{\check{v}}h^{-1}$ and hence $\Sigma(gs,gt)$ is connected to $\Sigma(s,t)$ be a W–edge for each $g \in hJ_{\check{v}}h^{-1}$ whenever $R \geqslant R_2+2$.

Because R_1 and R_2 depend only on the choice of finite generating set for the vertex and edge groups of $\mathcal{G}$, they can be chosen to be uniform for each vertex and edge of $\mathcal{G}$. Thus, W is connected whenever $R \geqslant R_0 = R_1+R_2+2$.

G **acts properly:** Let K_W be a bounded subset of W and let K_X be the subset of X that is the union $\bigcup_{\Sigma(s,t)\in K_W} P_r(s,t)$. We note that when we have W–adjacent maximal simplices $\Sigma(s,t),\Sigma(s',t')$ then $d_X(P_r(s,t),P_r(s',t'))$ is uniformly bounded. Indeed, if the W–edge is as in the first bullet of Definition 5.7, this is clear, and otherwise this follows from Proposition 5.9. Therefore, K_X is a bounded subset of X. Since the action of G on

X is metrically proper, the set $\{g \in G : K_X \cap gK_X \neq \emptyset\}$ is finite. Now,

$$\{g \in G : K_W \cap gK_W \neq \emptyset\} \subseteq \{g \in G : K_X \cap gK_X \neq \emptyset\}$$

because whenever $\Sigma(s,t)$ and $g\Sigma(s,t)$ are both in K_W, $P_r(s,t)$ and $gP_r(s,t)$ are both contained in K_X. As the latter set is finite, the claim follows.

G **acts coboundedly:** Since W is connected and G acts cofinitely on the edges of T, it suffices to prove that for any edge e of T, any two maximal simplices of $\mathcal{S}(T)$ that contain the edge e have $\mathrm{Stab}_G(e)$ translates that are W–adjacent. Let $\Sigma(s,t)$ and $\Sigma(s',t')$ be two maximal simplices that contain the edge e. Since $\mathrm{Stab}_G(e)$ acts transitively on the vertices of X_e, there exists $g \in \mathrm{Stab}_G(e)$ so that $P_r(gs,gt) \cap P_r(s',t') \neq \emptyset$. Thus, $d_X(P_r(gs,gt), P_r(s',t')) \leqslant R$ and $\Sigma(gs,gt)$ is W–adjacent to $\Sigma(s',t')$ as desired. □

6. Verification of combinatorial HHS axioms

We now verify that the pair $(\mathcal{S}(T), W)$ from Section 5 is a combinatorial HHS. For our admissible graph of groups $\mathcal{G}$, we fix the same notation as the beginning of Section 5 and let $\mathcal{S}(T)$ be the simplicial complex from Definition 5.1 for $\mathcal{G}$. We continue to use $\Sigma(s,t)$ to denote the maximal simplex of $\mathcal{S}(T)$ determined by $s,t \in \mathcal{S}(T)^{(0)} - T^{(0)}$ (see Lemma 5.2) and let $\sigma_r(s)$ and $P_r(s,t)$ be the sets from Definition 5.3 and 5.4 respectively.

Fix $r \geqslant 0$ and $R \geqslant 0$ large enough that Lemma 5.11 ensures the graph $W_{r,R}$ is connected and has a metrically proper and cobounded action of G. Moreover, choose R to be larger than 2ξ where $\xi = \xi(r)$ is the constant from Lemma 5.5. With these values of r and R fixed, we let $W = W_{r,R}$.

Our proof that $(\mathcal{S}(T), W)$ is a combinatorial HHS is spread over three subsections. Section 6.1 contains a description of the links of the non-maximal simplices of $\mathcal{S}(T)$ and verifies parts (I), (IV), and (V) of the definition of a combinatorial HHS (Definition 2.23). This section also includes a proof that the action of G on $\mathcal{S}(T)$ has finitely many orbits of links of simplices. Section 6.2 proves the augmented links $\mathcal{C}(\Delta)$ for simplices are hyperbolic, while Section 6.3 prove that they quasi-isometrically embed in the space Y_Δ. These are condition (II) and (III) of Definition 2.23.

6.1. **Simplices, links, and the combinatorial conditions.** We now describe the combinatorics of simplices and their links in $\mathcal{S}(T)$ and then verify three of the conditions for $(\mathcal{S}(T), W)$ to be a combinatorial HHS. In what follows, $\mathrm{lk}(\cdot)$ denotes the link in $\mathcal{S}(T)$, while $\mathrm{lk}_T(\cdot)$ denotes the link in $|T|$, the unoriented graph obtained from T by replacing each pair of oriented edges with an unoriented edge. Similarly, we use $d_T(\cdot,\cdot)$ to denote the distance in $|T|$ between two vertices of T.

A basic consequence of the description of maximal simplices of $\mathcal{S}(T)$ (Lemma 5.2), is that non-empty, non-maximal simplices come in one of the following types.

Corollary 6.1. *Every non-maximal, non-empty simplex Δ of $\mathcal{S}(T)$ is one of the following 8 types*

Type 1: $\Delta = \{v\}$ *for some* $v \in T^{(0)}$
Type 2: $\Delta = \{s\}$ *for some* $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$
Type 3: $\Delta = \{v, w\}$ *for some* $v, w \in T^{(0)}$
Type 4: $\Delta = \{s, t\}$ *for some* $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$

Type 5: $\Delta = \{s, v\}$ *for some* $v \in T^{(0)}$ *and* $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ *with* $\nu(s) \neq v$
Type 6: $\Delta = \{s, \nu(s), t\}$ *for some* $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$
Type 7: $\Delta = \{s, \nu(s)\}$ *for some* $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$
Type 8: $\Delta = \{s, \nu(s), v\}$ *for some* $v \in T^{(0)}$ *and* $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ *with* $\nu(s) \neq v$*.*

Proof. Since every non-maximal simplex can be completed to a maximal simplex by adding vertices, the above list is a consequence of Lemma 5.2 □

By examining each type of simplex, we also obtain a description of the links of each type of simplex. Figure 3 contains a schematic of each type of simplex along with its link and will be a useful reference through this section.

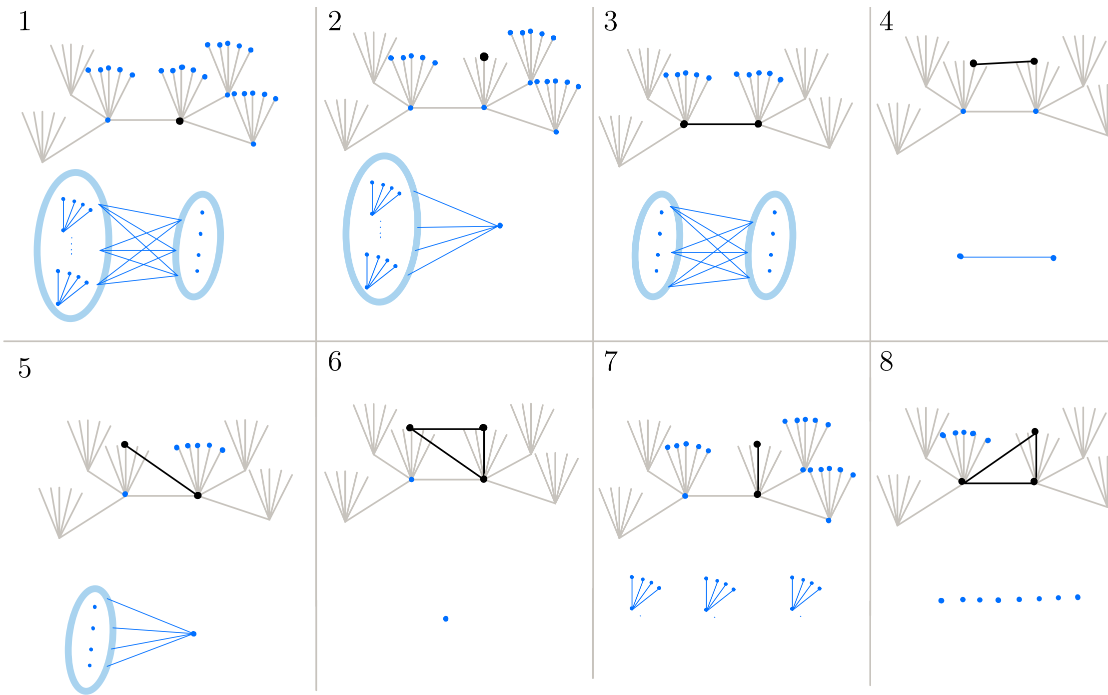

FIGURE 3. A schematic of each type of simplex and its link. The simplex is drawn in black with the vertices of the link highlighted in blue. Below, a schematic of the link is drawn in blue. To avoid clutter, most edges between vertices s, t with $\nu(s) \neq \nu(t)$ are missing, as can be seen in the links of Type 1 and Type 2.

Lemma 6.2. *Let* Δ *be a non-maximal, non-empty simplex of* $\mathcal{S}(T)$*. The link of* Δ *is determined by the type of* Δ *as follows, where* $v, w \in T^{(0)}$ *and* $s, t \in \mathcal{S}(T)^{(0)} - T^{(0)}$*:*

Type 1: if $\Delta = \{v\}$*, then* $\mathrm{lk}(\Delta)$ *is the join of* $\{s \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(s) = v\}$ *with the span of* $\{t \in \mathcal{S}(T)^{(0)} : \nu(t) \in \mathrm{lk}_T(v)\}$*.*
Type 2: if $\Delta = \{s\}$*, then* $\mathrm{lk}(\Delta)$ *is the join of* $\{\nu(s)\}$ *and the span of*

$$\{t \in \mathcal{S}(T)^{(0)} : \nu(t) \in \mathrm{lk}_T(\nu(s))\}.$$

Type 3: if $\Delta = \{v, w\}$*, then* $\mathrm{lk}(\Delta)$ *is the join of* $\{s \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(s) = v\}$ *with* $\{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(t) = w\}$.
Type 4: if $\Delta = \{s, t\}$*, then* $\mathrm{lk}(\Delta)$ *is the edge between* $\nu(s)$ *and* $\nu(t)$.
Type 5: if $\Delta = \{s, w\}$*, then* $\mathrm{lk}(\Delta)$ *is the join of* $\{\nu(s)\}$ *and* $\{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(t) = w\}$.
Type 6: if $\Delta = \{s, \nu(s), t\}$*, then* $\mathrm{lk}(\Delta)$ *is the vertex* $\nu(t)$.
Type 7: if $\Delta = \{s, \nu(s)\}$*, then* $\mathrm{lk}(\Delta)$ *is spanned by* $\{t \in \mathcal{S}(T)^{(0)} : \nu(t) \in \mathrm{lk}_T(v)\}$.
Type 8: if $\Delta = \{s, \nu(s), v\}$*, then* $\mathrm{lk}(\Delta)$ *is* $X_v^{(0)} = \{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(t) = v\}$.

In particular, if Δ *is not of Type 7 or Type 8, then* $\mathcal{C}(\Delta)$ *has diameter at most* 3*, by virtue of being a single vertex or a non-trivial join with some added edges.*

Proof. All cases are a straightforward exercise using the definitions of edges of $\mathcal{S}(T)$. The "in particular" clause follows as $\mathcal{C}(\Delta)$ is obtained from adding edges to $\mathrm{lk}(\Delta)$. □

When the link of a simplex is not a non-trivial join, we will need to understand its saturation (Definition 2.21) in order to understand the space Y_Δ.

Lemma 6.3. *Let* Δ *be a non-empty, non-maximal simplex of* $\mathcal{S}(T)$.

(1) If $\Delta = \{s, \nu(s)\}$ *is a simplex of Type 7, then*

$$\mathrm{Sat}(\Delta) = \{\nu(s)\} \cup \{s' \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(s') = \nu(s)\}.$$

(2) If $\Delta = \{s, \nu(s), v\}$ *is a simplex of Type 8, then*

$$\mathrm{Sat}(\Delta) = \{u \in T^{(0)} : d_T(v, u) \leqslant 1\} \cup \{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : d_T(v, \nu(t)) = 1\}.$$

Proof. **Case 1:** $\Delta = \{s, \nu(s)\}$ is a simplex of Type 7. First, suppose that s' is a vertex with $\nu(s) = \nu(s')$ and $s' \neq \nu(s)$, so $\Delta' = \{s', \nu(s)\}$ is a simplex of $\mathcal{S}(T)$. If $u \in \mathcal{S}(T)$ is a vertex adjacent to both s' and $\nu(s') = \nu(s)$, then $\nu(u)$ is adjacent to $\nu(s)$, which makes u adjacent to s. Hence $u \in \mathrm{lk}(\Delta)$, and we have $\mathrm{lk}(\Delta') \subseteq \mathrm{lk}(\Delta)$. By a symmetrical argument, $\mathrm{lk}(\Delta') = \mathrm{lk}(\Delta)$. Thus, every simplex of the form $\{s', \nu(s)\}$ with $\nu(s) = \nu(s')$ and $\nu(s) \neq s'$ has the same link as Δ. In particular, every such s' is in $\mathrm{Sat}(\Delta)$ and $\nu(s)$ is in $\mathrm{Sat}(\Delta)$.

Conversely, suppose that Δ' is a simplex with $\mathrm{lk}(\Delta') = \mathrm{lk}(\Delta)$. Then $\nu(\Delta') = \nu(s)$ as $\mathrm{lk}_T(\nu(\Delta)) = \mathrm{lk}_T(\nu(s))$. Now, Δ' cannot be $\nu(s) \in \mathcal{S}(T)$, because then its link would contain vertices s' with $\nu(s) = \nu(s')$, which are not in $\mathrm{lk}(\Delta)$. On the other hand, if $\Delta' = s'$ for some $s' \in \mathcal{S}(T)$ with $\nu(s') = \nu(s)$, then $\mathrm{lk}(\Delta')$ would contain $\nu(s)$, which is not in $\mathrm{lk}(\Delta)$. So Δ' must be equal to $\{s', \nu(s)\}$ for some $s' \neq \nu(s)$ with $\nu(s) = \nu(s')$. By definition, $\mathrm{Sat}(\Delta)$ is the union of these $\{s', \nu(s)\}$, which completes the proof that

$$\mathrm{Sat}(\Delta) = \bigcup_{\nu(s') = \nu(s)} \{s', \nu(s)\}.$$

Case 2: $\Delta = \{s, \nu(s), v\}$ is a simplex of Type 8. Let $u \in \mathrm{Sat}(\Delta)$. If $u \in T^{(0)}$, then u is adjacent to or equal to v in $\mathcal{S}(T)$ and hence in T. If $u \in \mathcal{S}(T)^{(0)} - T^{(0)}$, then $\nu(u)$ is adjacent to or equal to v. Conversely, suppose that $u \in T^{(0)}$ and $d_T(u, v) = 1$. Choose any vertex $t \in \nu^{-1}(u)$ with $t \neq u$. Then $\{u, t, v\}$ is a simplex with link $\mathrm{lk}(\Delta)$. Next, suppose that $u \in \mathcal{S}(T)^{(0)} - T^{(0)}$ and $d_T(v, \nu(u)) \leqslant 1$. Then $\Delta' = \{u, \nu(u), v\}$ is a simplex with $\mathrm{lk}(\Delta') = \mathrm{lk}(\Delta)$. Together, these show the

$$\mathrm{Sat}(\Delta) = \{u \in T^{(0)} : d_T(v, u) \leqslant 1\} \cup \{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : d_T(v, \nu(t)) = 1\}.$$

□

We now verify conditions (I), (IV), (V) from Definition 2.23 for $(\mathcal{S}(T), W)$ to be a combinatorial HHS.

Lemma 6.4. *If $\Delta_1, \ldots, \Delta_n$ are simplices of $\mathcal{S}(T)$ such that $\mathrm{lk}(\Delta_1) \subsetneq \cdots \subsetneq \mathrm{lk}(\Delta_n)$, then $n \leqslant 5$.*

Proof. Corollary 6.1 lists all types of non-maximal, non-empty simplices of $\mathcal{S}(T)$. Examining the links for each of these different types of simplices (Lemma 6.2) shows that if $\mathrm{lk}(\Delta) \subsetneq \mathrm{lk}(\Delta')$, then Δ must have strictly more vertices that Δ'. Thus, any chain of strictly nested links of simplices must have length at most 5 (recall, $\mathrm{lk}(\varnothing) = \mathcal{S}(T)$ by definition). □

Lemma 6.5. *Let Δ be a simplex of $\mathcal{S}(T)$ and $x, y \in \mathrm{lk}(\Delta)^{(0)}$ be vertices that are not adjacent in $\mathcal{S}(T)$, but are adjacent in $\mathcal{S}(T)^{+W}$. Then there exist two maximal simplices $\Sigma_x, \Sigma_y \subseteq \mathrm{st}(\Delta)$ that respectively contain x and y such that Σ_x and Σ_y are adjacent in W.*

Proof. Let Δ be a simplex of $\mathcal{S}(T)$ and $x, y \in \mathrm{lk}(\Delta)^{(0)}$ that are not adjacent in $\mathcal{S}(T)$, but are adjacent in $\mathcal{S}(T)^{+W}$. Let s_x, t_x and s_y, t_y be the elements of $\mathcal{S}(T) - T$ so that $x \in \Sigma(s_x, t_x)$, $y \in \Sigma(s_y, t_y)$ and $\Sigma(s_x, t_x)$ is W–adjacent to $\Sigma(s_y, t_y)$. Without loss of generality, assume x and y are respectively contained in the edges $\{t_x, \nu(t_x)\}$ and $\{t_y, \nu(t_y)\}$. It suffices to find $s \in \mathcal{S}(T) - T$ so that $\Delta \subseteq \mathrm{st}(s)$ and the simplices $\Sigma(s, t_x)$ and $\Sigma(s, t_y)$ are W–adjacent.

First assume that $\nu(x) \neq \nu(y)$. Since x and y are not joined by an $\mathcal{S}(T)$–edge, $\nu(x)$ cannot be joined to $\nu(y)$ by an edge in T. Thus, there must exist $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ so that Δ is contained in the edge $\{s, \nu(s)\}$ and $\nu(x), \nu(y) \subseteq \mathrm{lk}(\nu(s))$. The simplices $\Sigma(s, t_x)$ and $\Sigma(s, t_y)$ are therefore W–adjacent, since $\Sigma(s_x, t_x)$ and $\Sigma(s_y, t_y)$ being W–adjacent implies that $d_X(\sigma_r(t_x), \sigma_r(t_y)) \leqslant R + 2$ in both cases of edges in W.

Now assume that $\nu(x) = \nu(y)$. Since x and y are not joined by an $\mathcal{S}(T)$–edge, both x and y must be elements of $\mathcal{S}(T)^{(0)} - T^{(0)}$. This implies that Δ is contained in a 2–simplex of the form $\{s, \nu(s), \nu(x)\}$ where $s \in \mathcal{S}(T) - T$ with $\nu(s) \subseteq \mathrm{lk}(\nu(x))$. Since $x \neq y$ and $\Sigma(s_x, t_x)$ is W–adjacent to $\Sigma(s_y, t_y)$, we must have $t_x = x$, $t_y = y$, and $d_X(\sigma_r(x), \sigma_r(y)) \leqslant R + 2$ in both case of edges in W. Thus, the simplices $\Sigma(s, x)$ and $\Sigma(s, y)$ are connected by an edge in W. □

Lemma 6.6. *For any non-maximal simplices Δ and Ω of $\mathcal{S}(T)$ there exists a (possibly empty) simplex Π of $\mathrm{lk}(\Delta)$ such that $\mathrm{lk}(\Delta \star \Pi) \subseteq \mathrm{lk}(\Omega)$ and for all non-maximal simplices Λ of $\mathcal{S}(T)$ so that $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta) \cap \mathrm{lk}(\Omega)$ either*

(1) $\mathrm{lk}(\Lambda)$ *is a non-trivial join or a vertex; or*
(2) $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta \star \Pi)$.

Proof. First of all, we will implicitly assume throughout the proof that the link of the empty simplex is not contained in $\mathrm{lk}(\Delta) \cap \mathrm{lk}(\Omega)$, for otherwise we have $\Delta = \Omega = \varnothing$, and we can take Π to be empty as well.

Let Δ and Ω be as in the statement, and let U denote the union of all $\mathrm{lk}(\Lambda) \subseteq \mathrm{lk}(\Delta) \cap \mathrm{lk}(\Omega)$ such that $\mathrm{lk}(\Lambda)$ is neither a non-trivial join or a single vertex. It suffices to show $U = \mathrm{lk}(\Delta \star \Pi)$ for some simplex Π. Note that if Λ is a non-empty simplex of $\mathcal{S}(T)$ so that $\mathrm{lk}(\Lambda)$ is not a single vertex nor a non-trivial join, then Λ is either a Type 7 or Type 8 simplex.

We say that a subgraph $\mathcal{X}$ of $\mathcal{S}(T)$ satisfies property P if the following holds. For all vertices v of T, if there exist two vertices $x, y \in \mathcal{X}$ with $\nu(x)$, $\nu(y)$ at distance 1 from v in T, then we have that $\mathcal{X}$ contains the entire Type 7 link of a simplex $\{s, \nu(s) = v\}$.

We make two preliminary observations about this property. First, if two subgraphs satisfy property P, then their intersection does as well. Secondly, given a subgraph $\mathcal{X}$ satisfying property P, the (possibly empty) union of all links of Type 7 or Type 8 contained in $\mathcal{X}$ satisfies property P.

By inspection of the list of possible links of a non-empty simplex (Lemma 6.2 and Figure 3), we can check that links satisfy property P. In view of the observations above, given simplices Δ and Ω, the subgraph U of $\mathrm{lk}(\Delta)$ considered above also satisfies property P.

To conclude the proof, we go through the list of possible links one more time and we check that, given any simplex Δ and any union U of links of Type 7 or Type 8 contained in $\mathrm{lk}(\Delta)$ satisfying property P, we have $U = \mathrm{lk}(\Delta \star \Pi)$. (Note that $U = \mathrm{lk}(\Delta \star \Pi)$ is equivalent to U being a link as a subgraph of $\mathrm{lk}(\Delta)$, and note also that if U is empty then it suffices to take Π to be a maximal simplex in $\mathrm{lk}(\Delta)$.) □

We conclude this subsection by verifying that the action of G on $\mathcal{S}(T)$ has finitely many orbits of link of simplices.

Lemma 6.7. *The action of G on $\mathcal{S}(T)$ has finitely many orbits of links of simplices.*

Proof. Let Δ be a simplex of $\mathcal{S}(T)$. If Δ is maximal, then $\mathrm{lk}(\Delta) = \varnothing$, and if $\Delta = \varnothing$, then $\mathrm{lk}(\Delta) = \mathcal{S}(T)$, and we are done.

If Δ is spanned entirely by vertices of T (Type 1 or Type 3), then Δ—and hence $\mathrm{lk}(\Delta)$—belongs to one of finitely many G–orbits. Similarly, because the G–stabiliser of a vertex $v \in T^{(0)}$ acts cofinitely on the set $X_v^{(0)} = \{s \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(s) = v\}$, there are finitely many G–orbits of vertices of $\mathcal{S}(T)$ (Type 2 simplices) and simplices of Type 7, i.e., $\Delta = \{s, \nu(s)\}$ for $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$. Hence, there finitely many G–orbits of these types of simplices and their links.

If Δ is of Type 4 or Type 6, then $\mathrm{lk}(\Delta)$ is either an unoriented edge or a vertex of T (Lemma 6.2), of which there are finitely many G–orbits of both.

If $\Delta = \{s, \nu(s), v\}$ is a simplex of Type 8, then $\mathrm{lk}(\Delta)$ is

$$X_v^{(0)} = \{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(t) = v\}.$$

There are only finitely many G–orbits of these sets as there are finitely many G–orbits of vertices in T.

Finally, let $\Delta_1 = \{s_1, v_1\}$ and $\Delta_2 = \{s_2, v_2\}$ be two simplices of Type 5. For each Δ_i, there is an oriented edge e_i of T from $\nu(s_i)$ to v_i. If $g \in G$ so that $ge_1 = e_2$, then $\mathrm{lk}(g\Delta_1) = \mathrm{lk}(\Delta_2)$ (even though $g\Delta_1$ might not equal Δ_2). As there are finitely many G–orbits of edges of T, this shows there are only finitely many G–orbits of links of Type 5 simplices.

Examining the lists of types of simplices in Corollary 6.1, we see that the preceding discussion exhausts all the possibilities. □

6.2. **Hyperbolicity of non-join links.** Recall from Section 2.3 that $\mathcal{C}(\Delta)$ is the graph obtained from $\mathrm{lk}(\Delta)$ by adding an edge between every pair of of vertices x, y for which there exists maximal simplices Σ_x, Σ_y that are joined by an edge of W and contain x and y respectively. In this section, we verify that each $\mathcal{C}(\Delta)$ is hyperbolic, which is condition (II) of Definition 2.23 for $(\mathcal{S}(T), W)$ to be a combinatorial HHS. Since there are only finitely many G–orbits of the $\mathcal{C}(\Delta)$ by Lemma 6.7, the hyperbolicity constant will automatically be uniform over all simplices of $\mathcal{S}(T)$ (although this fact is independently explicit in our proof). We only need to verify $\mathcal{C}(\Delta)$ is hyperbolic for the empty simplex and simplices of Type 7 and 8 as Lemma 6.2 showed $\mathcal{C}(\Delta)$ has diameter 2 in all other cases.

Proposition 6.8 (Unbounded augmented links are hyperbolic)**.** *Let Δ be a simplex of $\mathcal{S}(T)$.*

(1) If $\Delta = \varnothing$, then $\mathcal{S}(T)^{+W} = \mathcal{C}(\varnothing)$ is (G–equivariantly) quasi-isometric to $|T|$. Hence $\mathcal{S}(T)^{+W}$ is a quasi-tree.
(2) If $\Delta = \{s, \nu(s), v\}$ is a simplex of Type 8, then the identity map on vertices gives a uniform quasi-isometry from $\mathcal{C}(\Delta)$ to the quasi-line L_v.
(3) If $\Delta = \{s, \nu(s)\}$ is a simplex of Type 7, then $\mathcal{C}(\Delta)$ is uniformly hyperbolic. Moreover, if every vertex group in $\mathcal{G}$ is virtually free, then $\mathcal{C}(\Delta)$ is quasi-isometric to a tree.

Proof. We prove the three case separately.

Proof of (1)**.** The inclusion $|T| \to \mathcal{S}(T)^{+W}$ is simplicial and hence Lipschitz, thus it suffices to find a coarsely Lipschitz quasi-inverse for the inclusion. This quasi-inverse is provided by the map $\nu : \mathcal{S}(T)^{(0)} \to T^{(0)}$, where $\mathcal{S}(T)^{(0)}$ is equipped with the metric inherited from $\mathcal{S}(T)^{+W}$. To show that the map ν is coarsely Lipschitz it suffices to prove that $d_T(\nu(x), \nu(y))$ is uniformly bounded whenever $x, y \in \mathcal{S}(T)^{(0)}$ are joined by an edge of $\mathcal{S}(T)^{+W}$.

If x and y are joined by an edge of $\mathcal{S}(T)$ then $\nu(x)$ and $\nu(y)$ are equal or joined by an edge of $\mathcal{S}(T)$ as well, hence $d_T(\nu(x), \nu(y)) = 1$. Now assume x, y are joined by a W–edge. This means that x, y respectively belong to maximal simplices $\Sigma(s, t)$ and $\Sigma(s', t')$ that are adjacent in W. The definition of edges of W (Definition 5.7), allows us to assume that $\nu(s) = \nu(s')$ without loss of generality. Hence $\nu(x)$ and $\nu(y)$ are both either equal to $\nu(s)$ or joined by an edge of T to $\nu(s)$. Hence $d_T(\nu(x), \nu(y)) \leqslant 2$ as desired.

Since ν and the inclusion are G–equivariant, the quasi-isometry is also G–equivariant. This completes the proof of (1).

Proof of (2)**.** Let $\Delta = \{s, \nu(s), v\}$ be a Type 8 simplex, and let $\mu = \check{v}$. By Lemma 6.2, the vertex set of $\mathcal{C}(\Delta)$ is exactly $\mathrm{lk}(\Delta)$, which is the set of vertices

$$X_v^{(0)} = \{t \in \mathcal{S}(T)^{(0)} - T^{(0)} : \nu(t) = v\}.$$

Recall that L_v is a copy of the vertex space X_v with extra edges between vertices that make L_v a quasi-line. Let $I \colon \mathcal{C}(\Delta)^{(0)} \to L_v^{(0)}$ be the identity on the vertex set.

We first show that I^{-1} sends edges of L_v to edges of $\mathcal{C}(\Delta)$.

Claim 6.9. *If $t_1, t_2 \in L_v^{(0)}$ are joined by an edge of L_v, then $I^{-1}(t_1)$ and $I^{-1}(t_2)$ are joined by an edge of $\mathcal{C}(\Delta)$. In particular, $\mathcal{C}(\Delta)$ is connected and I^{-1} is 1-Lipschitz.*

Proof. Let $t_1, t_2 \in L_v^{(0)}$ be joined by an edge of L_v. By Lemma 5.5, the there is a (ξ, ξ)–coarsely Lipschitz ξ–coarse map $L_{\nu(s)} \times L_v \to X$ that sends (s, t_1) to $P_r(s, t_1)$ and (s, t_2) to $P_r(s, t_2)$. Because we chose R to be greater that 2ξ, this implies $d_X(P_r(s, t_1), P_r(s, t_2)) \leqslant R$. Thus, the maximal simplices $\{s, \nu(s), t_1, v\}$ and $\{s, \nu(s), t_2, v\}$ are joined by an edge in W, which implies t_1, t_2 are joined by an edge in $\mathcal{C}(\Delta)$. □

We now prove that I is also coarsely Lipschitz. This will complete the proof of (2).

Claim 6.10. *The map I is coarsely Lipschitz.*

Proof. It suffices to show that whenever $t_1, t_2 \in \mathcal{C}(\Delta)^{(0)}$ are adjacent in $\mathcal{C}(\Delta)$, that $I(t_1)$ and $I(t_2)$ are uniformly close in L_v. Since $\mathrm{lk}(\Delta)$ contains no edges in this case, the only way t_1, t_2 can being joined by an edge in $\mathcal{C}(\Delta)$ is for them to be joined by a W–edge. Since t_1, t_2 both belong to $\mathrm{lk}(\Delta)$, Lemma 6.5 provides maximal simplices $\Sigma_{t_1} = \{s, \nu(s), v, t_1\}$ and $\Sigma_{t_2} = \{s, \nu(s), v, t_2\}$ that are joined by an edge in W. By the definition of the edges of W, we have $d_X(\sigma_r(t_1), \sigma_r(t_2)) \leqslant R + 2$. Using Lemma 2.12, there is then a constant $\kappa \geqslant 1$ (determined by r and $\mathcal{G}$) so that $d_{X_v}(\sigma_r(t_1), \sigma_r(t_2)) \leqslant \kappa$. As the map $p_v \colon X_v \to L_v$ is distance non-increasing, $I(t_i) \in p_v(\sigma_r(t_i))$, and $\mathrm{diam}(p_v(\sigma_r(t_i))) \leqslant 2r$, we have

$$d_{X_v}(\sigma_r(t_1), \sigma_r(t_2)) \leqslant \kappa \implies d_{L_v}(I(t_1), I(t_2)) \leqslant \kappa + 4r. \quad \square$$

Proof of (3)**.** Let $\Delta = \{s, \nu(s)\}$ be a simplex of Type 7, and let $v = \nu(s)$. By multiplying by an element of G, we can assume that $\mathrm{Stab}_G(v) = G_\mu$ for some vertex $\mu \in \mathcal{G}$.

Let Y be the graph obtained from $\mathrm{lk}_T(v)$ by joining distinct $x, y \in \mathrm{lk}_T(v)^{(0)}$ by an edge if and only if there exist $x', y' \in \mathrm{lk}(\Delta)^{+W}$ with $\nu(x') = x, \nu(y') = y$, and x', y' adjacent in $\mathcal{C}(\Delta)$. Note that G_μ acts on Y, and $\nu \colon \mathcal{S}(T)^{(0)} \to T^{(0)}$ induces a G_μ–equivariant simplicial map $\eta \colon \mathcal{C}(\Delta) \to Y$.

We first show Y is connected and quasi-isometric to $\mathcal{C}(\Delta)$.

Claim 6.11. *The graphs $\mathcal{C}(\Delta)$ and Y are connected.*

Proof. Because there is a simplicial surjection $\eta \colon \mathcal{C}(\Delta) \to Y$, connectedness of $\mathcal{C}(\Delta)$ will imply connectedness of Y.

Let $x, y \in \mathcal{C}(\Delta)^{(0)}$ and let Σ_x, Σ_y be maximal simplices of $\mathcal{S}(T)$ containing x and y respectively. Since $x, y \in \mathrm{lk}(\Delta)$, we can use Lemma 6.5 to assume that Σ_x, Σ_y are maximal simplices of the star of Δ. By Lemma 5.11, there is a path $\Sigma_x = \Sigma_0, \Sigma_1, \ldots, \Sigma_n = \Sigma_y$ in W, where each Σ_i is a maximal simplex of $\mathcal{S}(T)$ and Σ_i, Σ_{i+1} are joined be an edge of W for $0 \leqslant i \leqslant n-1$. We now argue by induction on n that x can be joined to y by a path in $\mathcal{C}(\Delta)$.

If $n = 0$, then both x, y are contained in Σ_x and hence are either equal or are adjacent in $\mathcal{S}(T)$ and therefore in $\mathcal{S}(T)^{+W}$. Since $x, y \in \mathrm{lk}(\Delta)$, either $x = y$ or x, y are adjacent in $\mathcal{C}(\Delta)$ by Lemma 6.5.

If $n = 1$, then Σ_0 is joined by an edge of W to Σ_1, hence $x, y \in \mathrm{lk}(\Delta)$ are either equal or adjacent in $\mathcal{S}(T)^{+W}$. Thus, by Lemma 6.5, x, y are either equal or adjacent in $\mathcal{C}(\Delta)$.

Suppose $n > 1$. Since Σ_0 and Σ_n are simplices of the star of Δ, the edges $\nu(\Sigma_0)$ and $\nu(\Sigma_n)$ contain v. Let v_{-1} be the vertex of $\nu(\Sigma_0)$ different from v, and let v_{n+1} be the vertex of $\nu(\Sigma_n)$ different from v. Since $x, y \in \mathrm{lk}(\Delta)$, we must have $\nu(x) = v_{-1}$ and $\nu(y) = v_{n+1}$. If $\nu(x) = v_{-1} = v_{n+1} = \nu(y)$, then x, y are in the link of the Type 8 simplex $\Delta' = \{v_{-1}, s, \nu(s)\}$. Claim 6.9 therefore implies x is connected to y as $\mathcal{C}(\Delta')$ has an injective simplicial inclusion into $\mathcal{C}(\Delta)$.

Now suppose $\nu(x) \neq \nu(y)$. This implies v_{-1} and v_{n+1} must lie in different components of $T - \{v\}$. The definition of edges in W (Definition 5.7) ensures that the edges $\nu(\Sigma_i)$ and $\nu(\Sigma_{i+1})$ share a vertex v_i for each all $i \in \{0, \ldots, n-1\}$. The sequence $v_{-1}, v_0, \ldots, v_n, v_{n+1}$ is then a sequence of vertices of T where consecutive vertices are either equal or adjacent in T. Because v_{-1} and v_{n+1} are in different components of $T - \{v\}$, there exists $i \in \{1, \ldots, n-1\}$

such that $\nu(\Sigma_i)$ contains v. Choose $z \in \Sigma_i^{(0)}$ such that $\nu(z) \in \nu(\Sigma_i) - \{v\}$. Then $z \in \mathrm{lk}(\Delta)$, and z is contained in the maximal simplex Σ_i. The sequence $\Sigma_0, \ldots, \Sigma_i$ is a path in W with $i < n$, and has $x \in \Sigma_0, z \in \Sigma_i$. So, by induction, x can be joined to z by a path in $\mathcal{C}(\Delta)$. Similarly, considering $\Sigma_i, \ldots, \Sigma_n$ shows that z can be joined by a path in $\mathcal{C}(\Delta)$ to y. So x, y are connected in $\mathcal{C}(\Delta)$, as required. □

We now prove that Y is quasi-isometric to $\mathcal{C}(\Delta)$.

Claim 6.12. *The map $\eta \colon \mathcal{C}(\Delta) \to Y$ induced by ν is a quasi-isometry with constants independent of Δ.*

Proof. As mentioned above, η is simplicial and hence 1-Lipschitz. Consider the composition of inclusions

$$Y^{(0)} \hookrightarrow T^{(0)} \hookrightarrow \mathcal{S}(T)^{(0)}.$$

The image of this map is in $\mathrm{lk}(\Delta)$, and the map is a quasi-inverse for η. Now, if $x, y \in Y^{(0)}$ are Y–adjacent, then let $\Sigma(s, x')$ and $\Sigma(t, y')$ be W–adjacent simplices where $\nu(x') = x$, $\nu(y') = y$, and $\nu(s) = \nu(t) = v$. Then $\nu(x') = x$ and $\nu(y') = y$ are adjacent in $\mathcal{C}(\Delta)$, so the map $Y \to \mathcal{C}(\Delta) \subseteq \mathcal{S}(T)$ induced by the above inclusions is uniformly coarsely Lipschitz. Thus η is a quasi-isometry. □

In view of Claim 6.12, it suffices to prove that Y is δ–hyperbolic (and that Y is uniformly quasi-isometric to a tree when the vertex groups are $\mathbb{Z}$–by-virtually free). For this we use the action of G_μ on Y.

Claim 6.13. *The action of G_μ on Y is cocompact.*

Proof. Because G_μ acts on Y with finitely many orbits of vertices, it suffices to prove that for each vertex $u \in Y$, there are finitely many $\mathrm{Stab}_{G_\mu}(u)$–orbits of edges of Y incident to u. Note, $\mathrm{Stab}_{G_\mu}(u) = \mathrm{Stab}_G(e_u)$ where e_u is the (oriented) edge of T from u to v.

Let y be an element of Y that is joined by an edge of Y to u. Let e_u and e_y be the (oriented) edges of T from u or y to v respectively. By construction of Y, each of u and y are contained in a maximal simplex of $\mathcal{S}(T)$ that contains v and are adjacent in W. The definition of edges in W then requires that the edge space X_{e_y} must intersect the $(R+2)$–neighborhood of X_{e_u} inside the Bass–Serre space X. Hence, each vertex y of Y that is adjacent to u in Y has a corresponding edge spaces X_{e_y} of X that is within $R + 2$ of X_{e_u}. We will argue that there are only a finite number of $\mathrm{Stab}_{G_\mu}(u)$–orbits of such edge spaces, which implies there is a finite number of $\mathrm{Stab}_{G_\mu}(u)$–orbits of vertices of Y adjacent to u.

Because $\mathrm{Stab}_{G_\mu}(u) = \mathrm{Stab}_G(e_u)$ acts cocompactly on X_{e_u}, it also acts cocompactly on the $(R+2)$–neighborhood of X_{e_u}. Since two edges spaces intersect if and only if they are equal, there can only be a finite number of $\mathrm{Stab}_{G_\mu}(u)$–orbits of vertex spaces of X that intersect the $(R+2)$–neighborhood of X_{e_u} as desired. □

Claims 6.11 and 6.13 show that Y is a connected graph (hence a length space) with a cocompact action by G_μ. Let $\{y_i\}$ be a finite set of vertices of Y containing exactly one element of each G_μ–orbit, and let $\mathcal{H}$ be the collection of stabilisers in G_μ of the vertices y_i. Given our fixed finite generating set J_μ of G_μ, Theorem 5.1 of [CC07] implies that any orbit map $G_\mu \to Y$ induces a quasi-isometry $\Gamma \to Y$ (with constants just depending on J_μ), where Γ is the Cayley graph of G_μ with respect to the infinite set $J_\mu \cup \{H\}_{H \in \mathcal{H}}$. If y is a vertex of Y, then the stabiliser in G_μ of y is exactly the stabiliser of some edge e of T with $e^+ = v$. Hence, each $H \in \mathcal{H}$ is conjugate to the image of some edge group $\tau_\alpha(G_\alpha)$ where $\alpha^+ = \mu$. Thus Γ is quasi-isometric to the Cayley graph of G_μ with respect to the generating set $J_\mu \cup \{\tau_\alpha(G_\alpha) : \alpha^+ = \mu\}$. By Lemma 2.15 the later is always hyperbolic and is a quasi-tree when F_μ is virtually free. As Γ is quasi-isometric to Y, this completes the proof of (3). □

6.3. **Quasi-isometric embedding of augmented links.** The goal of this section is to check condition (III) of Definition 2.23, that is, that each augmented link $\mathcal{C}(\Delta)$ is quasi-isometrically embedded in the corresponding space Y_Δ from Definition 2.22. Because there are finitely many orbits of links of simplices by Lemma 6.7, we will be able to choose the quasi-isometry constants uniformly over all simplices Δ.

Because the quasi-isometric embedding condition automatically holds when $\mathcal{C}(\Delta)$ is bounded, we only have to check simplices of Type 7 and 8.

6.3.1. *Type 7 links.*

Lemma 6.14. *There exists $\kappa \geqslant 1$ so that if $\Delta = \{s, \nu(s)\} \subset \mathcal{S}(T)$ is a simplex of Type 7, then $\mathcal{C}(\Delta)$ is (κ,κ)–quasi-isometrically embedded in Y_Δ.*

Proof. By Lemma 2.2, it suffices to define a coarsely Lipschitz coarse retraction $\rho\colon Y_\Delta \to \mathcal{C}(\Delta)$, with constants independent of Δ. We define ρ on the vertex set as follows: for $y \in Y_\Delta^{(0)}$, we let $\rho(y)$ be the unique vertex of T (regarded as a vertex of $\mathcal{S}(T)$) at distance 1 from $\nu(s)$ and on the geodesic in T from $\nu(y)$ to $\nu(s)$. This is well-defined because T is a tree and $y \neq \nu(s)$ if $y \in Y_\Delta^{(0)} = \mathcal{S}(T)^{(0)} - \mathrm{Sat}(\Delta)$. Moreover, because ρ coincides with ν on the vertices of $\mathcal{C}(\Delta)$, the distance between $\rho(y)$ and y is at most 1 for vertices $y \in \mathcal{C}(\Delta)$. Hence, $\mathcal{C}(\Delta)$ will be a coarse retract if ρ is coarsely Lipschitz.

If we can uniformly bound $d_{\mathcal{C}(\Delta)}(\rho(y_1), \rho(y_2))$ whenever y_1, y_2 are joined by a edge of Y_Δ, then ρ can be extended to a coarsely Lipschitz map $Y_\Delta \to \mathcal{C}(\Delta)$. We will obtain $d_{\mathcal{C}(\Delta)}(\rho(y_1), \rho(y_2)) \leqslant 3$ for such y_1, y_2.

If y_1, y_2 are joined by an $\mathcal{S}(T)$–edge, then $\nu(y_1)$ and $\nu(y_2)$ are either equal or joined by an edge of T. Since $\nu(y_i) \neq \nu(s)$ for each $i = 1, 2$, this implies $\rho(y_1) = \rho(y_2)$ because T is a tree. Hence we have $d_{\mathcal{C}(\Delta)}(\rho(y_1), \rho(y_2)) = 0$. If instead y_1, y_2 are joined by a W–edge, then $\nu(y_1)$ and $\nu(y_2)$ lie at distance at most 2 in T by Definition 5.7. If $\nu(y_1)$ and $\nu(y_2)$ are at most 1 apart in T, then $\rho(y_1) = \rho(y_2)$ as in the previous case. If the $\nu(y_1)$ and $\nu(y_1)$ are exact 2 apart in T, then there exists a unique vertex $z \in T$ at distance 1 from both $\nu(y_1)$ and $\nu(y_2)$. If $z \neq \nu(s)$, then $\nu(y_1)$ and $\nu(y_1)$ are in the same component of $T - \nu(s)$, which implies $\rho(y_1) = \rho(y_2)$. If $z = \nu(s)$, then y_1 and y_2 are in $\mathrm{lk}(\Delta)$. By Lemma 6.5, this implies y_1 and y_2 are joined by an edge in Y_Δ. Because $\rho(y_i) = \nu(y_i)$, we have $d_{\mathcal{C}(\Delta)}(\rho(y_1), \rho(y_2)) \leqslant d_{\mathcal{C}(\Delta)}(\rho(y_1), y_1) + d_{\mathcal{C}(\Delta)}(y_1, y_2) + d_{\mathcal{C}(\Delta)}(y_2, \rho(y_2)) \leqslant 3$. □

6.3.2. *Type 8 links.* We now consider simplices of the form $\Delta = \{s, \nu(s), v\}$, where $s \in \mathcal{S}(T)^{(0)} - T^{(0)}$ and $v \in T^{(0)} - \{\nu(s)\}$.

Lemma 6.15. *There exists $\kappa \geqslant 1$ with the following property. Let $\Delta = \{s, \nu(s), v\}$ be a Type 8 simplex of $\mathcal{S}(T)$. The inclusion of $\mathcal{C}(\Delta)$ into Y_Δ is a (κ,κ)–quasi-isometric embedding.*

By Proposition 6.8.(1), the Type 8 simplices are the simplices whose augmented links, $\mathcal{C}(\Delta)$, are quasi-isometric to the quasi-lines L_v. As in the previous case, we will show quasi-isometric embedding by providing a coarse retraction. However, since the identity map on vertices gives a quasi-isometry $L_v \to \mathcal{C}(\Delta)$, it suffices to build a coarsely Lipschitz coarse map $\eta\colon Y_\Delta \to L_v$, that is the the identity on the vertices $L_v^{(0)} = X_v^{(0)} \subseteq Y_\Delta^{(0)}$.

To define this map, we need to assign to each vertex space X_u of X a projection onto a hyperbolic space. Given $u \in T^{(0)}$, let $\vartheta = \check{u}$ and choose a coset representative g for gG_ϑ; recall the vertices of X_u are the elements of gG_ϑ. We now define a graph H_u as follows: the vertices of H_u are the elements of gG_ϑ and there is an edge between two elements x, y if $x^{-1}y \in J_\vartheta \cup Z_\vartheta$, where J_ϑ is our fixed finite generating set for G_ϑ and Z_ϑ is the center of G_ϑ. Since H_u is a copy of X_u with extra edges attached, there is a simplicial inclusion $\iota_u\colon X_u \to H_u$. By construction, multiplying every vertex of H_u by g^{-1} produces an isometry to the Cayley graph of G_ϑ with respect to the generating set $J_\vartheta \cup Z_\vartheta$. Thus, Lemma 2.15 implies that H_u is a hyperbolic graph.

Lemma 2.15 also shows that H_u is hyperbolic relative to the collection

$$\{\iota_u(\tau_e(X_e)) : e \text{ an edge of } T \text{ with } e^+ = u\}.$$

For an edge e with $e^+ = u$, define $\ell_e := \iota_u \circ \tau_e(X_e)$. As a peripheral subset in a relatively hyperbolic space, each ℓ_e has a coarse closest point projection $\mathfrak{p}_e\colon H_u \to \ell_e$; see, e.g., [Sis13]. This map is coarsely Lipschitz with constants independent of e or u.

The key property about the ℓ_e that we shall need is that they have a coarsely Lipschitz map onto L_{e^-}. One can show that this is in fact a quasi-isometry, but it will not be needed in the proof.

Lemma 6.16. *Let $v, u \in T^{(0)}$ and e be an edge of T with $e^+ = v$ and $e^- = u$. Let $\psi_e\colon \ell_{\bar{e}} \to L_v$ be the map given by restricting $p_v \circ \tau_e \circ \tau_{\bar{e}}^{-1} \circ \iota_u^{-1}$ to $\ell_{\bar{e}}$. Equipping $\ell_{\bar{e}}$ with the induced metric from H_u, the map $\psi_e\colon \ell_{\bar{e}} \to L_v$ is coarsely Lipschitz with constants determined by $\mathcal{G}$.*

Proof. Let $\vartheta = \check{u}$ and $\alpha = \check{e}$. Recall J_ϑ and J_α are our fixed generating sets for the vertex groups G_ϑ and the edge group G_α.

Let $g \in G$ so that the vertices of X_u (and H_u) are the elements of gG_ϑ. Since $\iota_u \circ \tau_{\bar{e}}\colon X_e \to H_u$ is a simplicial map, $\ell_{\bar{e}} = \iota_u \circ \tau_{\bar{e}}(X_e)$ is a connected subgraph of H_u. Hence it suffices to verify that whenever $x, y \in \ell_{\bar{e}}$ differ by an edge of H_u, that $d_{L_v}(\psi_e(x), \psi_e(y))$ is uniformly bounded. Let x, y be vertices of $\ell_{\bar{e}}$ that differ by an edge of H_u. Hence $x^{-1}y$ is either an element of J_ϑ or of Z_ϑ.

If $x^{-1}y \in J_\vartheta$, then x, y are elements of $\tau_{\bar{e}}(X_e)$ that are joined by an edge of X_u. Hence $x^{-1}y \in J_\vartheta \cap \tau_{\bar{\alpha}}(G_\alpha)$. Since there is a uniform bound on the number of elements of $\tau_{\bar{\alpha}}(J_\alpha)$ that are needed to write any element of $J_\vartheta \cap \tau_{\bar{\alpha}}(G_\alpha)$, there is a uniform bound on the distance between $\tau_{\bar{e}}^{-1} \circ \iota_u^{-1}(x)$ and $\tau_{\bar{e}}^{-1} \circ \iota_u^{-1}(y)$ in X_e (independence of α and ϑ comes from considering the finitely many vertices and edges of $\mathcal{G}$). Since p_v and τ_e are distance non-increasing from X_v and X_e respectively, this shows $d_{L_v}(\psi_e(x), \psi_e(y))$ is uniformly bounded.

If instead $x^{-1}y \in Z_\vartheta$, then x, y are elements of the same coset gZ_ϑ and $gZ_\vartheta \subseteq \tau_{\bar{e}}(X_e)$. Proposition 4.5.(1) provides a uniform bound on the diameter of $p_v \circ \tau_e \circ \tau_{\bar{e}}^{-1}(gZ_\vartheta)$ in L_v. Hence $d_{L_v}(\psi_e(x), \psi_e(y))$ as uniformly bounded. $\square$

We can now use the map ψ_e from Lemma 6.16 to define a map ρ_v for vertices of T that are at least distance 2 from v. We start with the case where $w \in T^{(0)}$ is exactly distance 2 from v. In this case, there is a unique vertex u at distance 1 from both v and w. If f is the oriented edge of T from w to u and e is the oriented edge from u to v, define β_e^f to be $\mathfrak{p}_{\bar{e}}(\ell_f)$. We then define

$$\rho_v(w) := \psi_e(\beta_e^f) = p_v \circ \tau_e \circ \tau_{\bar{e}}^{-1} \circ \iota_u^{-1}(\beta_e^f).$$

To define $\rho_v(w)$ when w is more than 2 away from v in T, let $\bar{w}$ be the unique vertex of T that is distance exactly 2 from v and on the geodesic in T from w to v. Define $\rho_v(w) := \rho_v(\bar{w})$.

The first thing to verify is that $\rho_v(w)$ is uniformly bounded.

Lemma 6.17. *There exist $\kappa_0 \geqslant 0$ so that for any $v \in T^{(0)}$, if $w \in T^{(0)}$ with $d_T(v,w) \geqslant 2$, then* $\operatorname{diam}(\rho_v(w)) \leqslant \kappa_0$.

Proof. By the definition of ρ_v, it suffices to verify the lemma when $d_T(v,w) = 2$. Let u be the unique vertex of T that is distance 1 from both v and w. Let e be the edge of T from u to v and f be the edge from w to u. Since $\ell_{\bar{e}}$ and ℓ_f are distinct peripheral subsets in the relatively hyperbolic space H_u, there is a uniform bound on the diameter of $\mathfrak{p}_{\bar{e}}(\ell_f) = \beta_e^f$; see, e.g., [Sis13]. Because the map ψ_e is coarsely Lipschitz (Lemma 6.16), this implies $\rho_v(w) = \psi_e(\beta_e^f)$ will be uniformly bounded in L_v. □

Next we verify that when two vertex spaces X_w and $X_{w'}$ are close in X, we have that $\rho_v(w)$ and $\rho_v(w')$ are close in L_v. This will be a key step to showing that pairs of vertices of Y_Δ that are joined by a W–edge are sent to uniformly bounded diameter set in L_v.

Lemma 6.18. *For every $q \geqslant 0$ there exists $\kappa_1 \geqslant 0$ such that the following holds for each $v \in T^{(0)}$. Let w, w' be vertices of T with $d_T(w,w') \leqslant 2$ and $d_T(w,v), d_T(w',v) \geqslant 2$. If $d_X(X_w, X_{w'}) \leqslant q$, then $d_{L_v}(\rho_v(w), \rho_v(w')) \leqslant \kappa_1$.*

Proof. Let $\bar{w}$ be the vertex of T at distance exactly 2 from v and along the geodesic from w to v, let u be the unique vertex of T at distance 1 from v and $\bar{w}$. Let f and e be the oriented edges of T from u to v and from $\bar{w}$ to u respectively. Define $\bar{w}', u', f', e'$ analogously, using w' rather than w.

If $\bar{w} = \bar{w}'$, then $\rho_v(w) = \rho_v(w')$ by definition and we are done. Otherwise, because $d_T(w,w') \leqslant 2$, we must have $w = \bar{w}$ and $w' = \bar{w}'$ and $u = u'$. This implies $e = e'$ as well.

Because each edge and vertex space of X separates X,

$$d_X(X_w, X_{w'}) \leqslant q \implies d_X(\tau_f(X_f), \tau_{f'}(X_{f'})) \leqslant q.$$

Applying Lemma 2.12 produces a $\kappa = \kappa(r, \mathcal{G}) \geqslant 1$ so that $d_{X_u}(\tau_f(X_f), \tau_{f'}(X_{f'})) \leqslant \kappa$. As the map $\iota_u \colon X_u \to H_u$ is distance non-increasing, we have $d_{H_u}(\ell_f, \ell_{f'}) \leqslant \kappa$. Because $\mathfrak{p}_{\bar{e}}$ is coarsely Lipschitz, there is now a uniform bound on the distance between β_e^f and $\beta_e^{f'}$. Since $\psi_e \colon \ell_{\bar{e}} \to L_v$ is a coarsely Lipschitz (Lemma 6.16), this implies $d_{L_v}(\rho_v(w), \rho_v(w'))$ is uniformly bounded as well. □

We now present the proof of the quasi-isometric embedding of $\mathcal{C}(\Delta)$ into Y_Δ.

Proof of Lemma 6.15. By Proposition 6.8, the identity map on vertices is a quasi-isometry $L_v \to \mathcal{C}(\Delta)$ with constants independent of Δ. Hence, the composition of this quasi-isometry with a coarsely Lipschitz coarse map $\eta\colon Y_\Delta \to L_v$ that is the the identity on the vertices $L_v^{(0)} = X_v^{(0)} \subseteq Y_\Delta^{(0)}$ will produce a coarse retraction $Y_\Delta \to \mathcal{C}(\Delta)$. By Lemma 2.2, this suffices to prove the inclusion is a quasi-isometric embedding.

By Lemma 6.3, $\mathrm{Sat}(\Delta) = \{v\} \cup \{t \in \mathcal{S}(T) : \nu(t) \in \mathrm{lk}_T(v)\}$. Since $Y_\Delta^{(0)} = \mathcal{S}(T)^{(0)} - \mathrm{Sat}(\Delta)$, we have

$$Y_\Delta^{(0)} = \left\{t \in \mathcal{S}(T)^{(0)} - \{v\} \colon \nu(t) = v \text{ or } d_T(\nu(t), v) \geqslant 2\right\}.$$

We now use the $\rho_v(w)$ from above to define the desired map $\eta\colon Y_\Delta \to L_v$. If $t \in Y_\Delta^{(0)}$ and $d_T(\nu(t), v) \geqslant 2$, then we can define $\eta(t) = \rho_v(\nu(t)) \subseteq L_v$. If instead $\nu(t) = v$, then t is a vertex of both X_v and L_v, and we define $\eta_v(t) = p_v(t) = t \in L_v$. Lemma 6.16 ensures $\mathrm{diam}(\eta(t)) \leqslant \kappa_0$ for all $t \in Y_\Delta^{(0)}$. We can extend this definition of η to a coarsely Lipschitz map on all of Y_Δ if we can show that $d_{L_v}(\eta(t_1), \eta(t_2))$ is uniformly bounded whenever t_1 and t_2 are joined by an edge of Y_Δ.

Let $t_1, t_2 \in Y_\Delta^{(0)}$ be joined by an edge. By the definition of the W–edge (Definition 5.7), this implies $d_T(\nu(t_1), \nu(t_2)) \leqslant 2$. First assume both $\nu(t_1)$ and $\nu(t_2)$ are v. Thus $t_1, t_2 \in \mathcal{C}(\Delta)$ and are joined be an edge. Since $\eta(t_1) = t_1$ and $\eta(t_2) = t_2$, the quasi-isometry between $\mathcal{C}(\Delta)$ and L_v ensures $d_{L_V}(\eta(t_1), \eta(t_2))$ is uniformly bounded.

Next suppose neither $\nu(t_1)$ or $\nu(t_2)$ equals v. If $d_T(\nu(t_1), \nu(t_2)) = 0$, then $\eta(t_1) = \eta(t_2)$ by definition. If $d_T(\nu(t_1), \nu(t_2)) = 1$, then, with out loss of generality, the geodesic in T from $\nu(t_1)$ to v must contain $\nu(t_2)$. Since each $\nu(t_i)$ are at least distance 2 from v, the definition of $\rho_v(\cdot)$ then implies $\eta(t_1) = \rho_v(\nu(t_1)) = \rho_v(\nu(t_2)) = \eta(t_2)$. Finally, if $d_T(\nu(t_1), \nu(t_2)) = 2$, then the edge between t_1 and t_2 must be a W–edge. This implies $X_{\nu(t_1)}$ and $X_{\nu(t_2)}$ are uniformly close in X. Hence the desired bound on $d_{L_v}(\eta(t_1), \eta(t_2))$ is a consequence of Lemma 6.18.

Finally consider the case where $\nu(t_1) = v$, but $\nu(t_2) \neq v$. In this case, $d_T(\nu(t_1), \nu(t_2)) = 2$, and so the edge between t_1 and t_2 must be a W–edge. Hence, $d_X(\sigma_r(t_1), \sigma_r(t_2)) \leqslant R + 2$ in either case of Definition 5.7. Let u be the vertex distance 1 from both $v = \nu(t_1)$ and $\nu(t_2)$, then let f be the oriented edge of T from $\nu(t_2)$ to u and e be the oriented edge from u to $\nu(t_1)$.

Let $\sigma_e = \tau_e^{-1}(\sigma_r(t_1))$ and $\sigma_f = \tau_{\bar f}^{-1}(\sigma_r(t_2))$. Our choice of r is large enough that Lemma 5.5 ensures σ_e and σ_f are both non-empty. Recalling that ψ_e is the restriction of $p_v \circ \tau_e \circ \tau_{\bar e}^{-1} \circ \iota_u^{-1}$, we have that

$$\psi_e(\iota_u \circ \tau_{\bar e}(\sigma_e)) = p_v(\sigma_r(t_1)). \qquad (*)$$

Claim 6.19. *There exists $\kappa' \geqslant 1$ depending only on $\mathcal{G}$ so that $d_{H_u}(\iota_u \circ \tau_{\bar e}(\sigma_e), \beta_e^f) \leqslant \kappa'$.*

Proof. Because $d_X(\sigma_r(t_1), \sigma_r(t_2)) \leqslant R + 2$, we have $d_X(\tau_{\bar e}(\sigma_e), \tau_f(\sigma_f)) \leqslant R + 6$. Applying Lemma 2.12 produces $\kappa = \kappa(R, \mathcal{G}) \geqslant 1$ so that $d_{X_u}(\tau_{\bar e}(\sigma_e), \tau_f(\sigma_f)) \leqslant \kappa$. As $\iota_u\colon X_u \to H_u$ is distance non-increasing, we have

$$d_{H_u}(\iota_u \circ \tau_{\bar e}(\sigma_e)), \iota_u \circ \tau_f(\sigma_f)) \leqslant \kappa.$$

Recall that $\beta_e^f = \mathfrak{p}_{\bar e}(\ell_f)$, that $\mathfrak{p}_{\bar e}$ is a coarse closest point projection to $\ell_{\bar e}$, which is a quasiconvex subset of a hyperbolic space. Hence, there is some κ', determined by κ and the hyperbolicity constant, so that $d_{H_u}(\iota_u \circ \tau_{\bar e}(\sigma_e), \beta_e^f) \leqslant \kappa'$. $\square$

Since ψ_e is a coarsely Lipschitz, Claim 6.19 plus $(*)$ implies that $\psi_e(\beta_e^f) = \eta(t_2)$ is uniformly close to $p_v(\sigma_r(t_1))$. Since $t_1 \in \sigma_r(t_1)$ and $\operatorname{diam}(p_v(\sigma_r(t_1))) \leqslant 2r$, this implies $\eta(t_2)$ is uniformly close to $\eta(t_1) = p_v(t_1)$ in L_v as desired. □

References

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

[ABO19] Carolyn Abbott, Sahana H. Balasubramanya, and Denis Osin. Hyperbolic structures on groups. *Algebr. Geom. Topol.*, 19(4):1747–1835, 2019.

[AFW15] Matthias Aschenbrenner, Stefan Friedl, and Henry Wilton. *3-manifold groups*. EMS Series of Lectures in Mathematics. European Mathematical Society (EMS), Zürich, 2015.

[ANS19] Carolyn Abbott, Thomas Ng, and Davide Spriano. Hierarchically hyperbolic groups and uniform exponential growth. *arXiv preprint arXiv:1909.00439*, 2019. with appendix by Radhika Gupta and Harry Petyt.

[BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publications mathématiques de l'IHÉS*, 122(1):1–64, 2015.

[Ber21] Daniel James Solomon Berlyne. *Hierarchical Hyperbolicity of Graph Products and Graph Braid Groups*. ProQuest LLC, Ann Arbor, MI, 2021. Thesis (Ph.D.)–City University of New York.

[BHMS20] Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. *arXiv preprint arXiv:2005.00567*, 2020.

[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.

[BHS19] J. Behrstock, M.F. Hagen, and A. Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299:257–338, 2019.

[BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021.

[Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[BR20] Federico Berlai and Bruno Robbio. A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 14(4):1127–1203, 2020.

[BR22] Daniel Berlyne and Jacob Russell. Hierarchical hyperbolicity of graph products. *Groups Geom. Dyn.*, 16(2):523–580, 2022.

[BW13] Hadi Bigdely and Daniel T. Wise. Quasiconvexity and relatively hyperbolic groups that split. *Michigan Math. J.*, 62(2):387–406, 2013.

[CC07] Ruth Charney and John Crisp. Relative hyperbolicity and Artin groups. *Geom. Dedicata*, 129:1–13, 2007.

[CCMT15] Pierre-Emmanuel Caprace, Yves Cornulier, Nicolas Monod, and Romain Tessera. Amenable hyperbolic groups. *J. Eur. Math. Soc. (JEMS)*, 17(11):2903–2947, 2015.

[CCS23] Ruth Charney, Matthew Cordes, and Alessandro Sisto. Complete topological descriptions of certain Morse boundaries. *Groups Geom. Dyn.*, 17(1):157–184, 2023.

[CK02] C. B. Croke and B. Kleiner. The geodesic flow of a nonpositively curved graph manifold. *Geom. Funct. Anal.*, 12(3):479–545, 2002.

[CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.

[Dah03] François Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963, 2003.

[DDLS20] Spencer Dowdall, Matthew G Durham, Christopher J Leininger, and Alessandro Sisto. Extensions of Veech groups II: Hierarchical hyperbolicity and quasi-isometric rigidity. *arXiv preprint arXiv:2111.00685*, 2020.

[DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.

[DHS20] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 24(2):1051–1073, 2020.

[DMS20] Matthew G Durham, Yair N Minsky, and Alessandro Sisto. Stable cubulations, bicombings, and barycenters. *arXiv preprint arXiv:2009.13647*, 2020.

[EF97] David B. A. Epstein and Koji Fujiwara. The second bounded cohomology of word-hyperbolic groups. *Topology*, 36(6):1275–1289, 1997.

[Hag14] Mark F. Hagen. Cocompactly cubulated crystallographic groups. *J. Lond. Math. Soc. (2)*, 90(1):140–166, 2014.

[HHP20] Thomas Haettel, Nima Hoda, and Harry Petyt. The coarse helly property, hierarchical hyperbolicity, and semihyperbolicity. *arXiv preprint arXiv:2009.14053*, 2020.

[HMS21] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type artin groups are hierarchically hyperbolic. *arXiv preprint arXiv:2109.04387*, 2021.

[HO13] Michael Hull and Denis Osin. Induced quasicocycles on groups with hyperbolically embedded subgroups. *Algebr. Geom. Topol.*, 13(5):2635–2665, 2013.

[Hod20] Nima Hoda. Crystallographic helly groups. *arXiv preprint arXiv:2010.07407*, 2020.

[HP15] Mark F. Hagen and Piotr Przytycki. Cocompactly cubulated graph manifolds. *Israel J. Math.*, 207(1):377–394, 2015.

[HP22] Mark F. Hagen and Harry Petyt. Projection complexes and quasimedian maps. *Algebr. Geom. Topol.*, 22(7):3277–3304, 2022.

[HS20] Mark F. Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *Israel J. Math.*, 236(1):45–89, 2020.

[Hug22] Sam Hughes. Lattices in a product of trees, hierarchically hyperbolic groups and virtual torsion-freeness. *Bull. Lond. Math. Soc.*, 54(4):1413–1419, 2022.

[KK13] Sang-hyun Kim and Thomas Koberda. Embedability between right-angled Artin groups. *Geom. Topol.*, 17(1):493–530, 2013.

[KL98] M. Kapovich and B. Leeb. 3-manifold groups and nonpositive curvature. *Geom. Funct. Anal.*, 8(5):841–852, 1998.

[KR14] Ilya Kapovich and Kasra Rafi. On hyperbolicity of free splitting and free factor complexes. *Groups Geom. Dyn.*, 8(2):391–414, 2014.

[Man05] Jason Manning. Geometry of pseudocharacters. *Geom. Topol.*, 9(2):1147–1185, 2005.

[Mil20] Marissa Miller. Stable subgroups of the genus two handlebody group. *arXiv preprint arXiv:2009.05067*, 2020.

[Min01] Igor Mineyev. Straightening and bounded cohomology of hyperbolic groups. *Geom. Funct. Anal.*, 11(4):807–839, 2001.

[NQ22] Hoang Thanh Nguyen and Yulan Qing. Sublinearly morse boundary of CAT(0) admissible groups. *arxiv preprint arxiv:2203.00935*, 2022.

[Pet21] Harry Petyt. Mapping class groups are quasicubical. *arXiv preprint arXiv:2112.10681*, 2021.

[PS23] Harry Petyt and Davide Spriano. Unbounded domains in hierarchically hyperbolic groups. *Groups Geom. Dyn.*, 17(2):479–500, 2023.

[RS20] Bruno Robbio and Davide Spriano. Hierarchical hyperbolicity of hyperbolic-2-decomposable groups. *arXiv preprint arXiv:2007.13383*, 2020.

[RST23] Jacob Russell, Davide Spriano, and Hung C. Tran. Convexity in hierarchically hyperbolic spaces. *Algebr. Geom. Topol.*, 23(3):1167–1248, 2023.

[Rus21] Jacob Russell. Extensions of multicurve stabilizers are hierarchically hyperbolic. *arXiv preprint arXiv:2107.14116*, 2021.

[Rus22] Jacob Russell. From hierarchical to relative hyperbolicity. *Int. Math. Res. Not. IMRN*, (1):575–624, 2022.

[Sco83] Peter Scott. The geometries of 3-manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983.

[Sis12] Alessandro Sisto. On metric relative hyperbolicity. *arXiv preprint arXiv:1210.8081*, 2012.

[Sis13] Alessandro Sisto. Projections and relative hyperbolicity. *Enseign. Math. (2)*, 59(1-2):165–181, 2013.

[Sis19] Alessandro Sisto. What is a hierarchically hyperbolic space? In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 117–148. Cambridge Univ. Press, Cambridge, 2019.

[Spr17] Davide Spriano. Hyperbolic HHS I: Factor systems and quasi-convex subgroups. *arXiv preprint arXiv:1711.10931*, 2017.

[SW79] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36, pages 137–203, 1979.

[Vok22] Kate M. Vokes. Hierarchical hyperbolicity of graphs of multicurves. *Algebr. Geom. Topol.*, 22(1):113–151, 2022.

School of Mathematics, University of Bristol, Bristol, UK
Email address: markfhagen@posteo.net

Department of Mathematics, Rice University, Houston, TX, USA
Email address: jacob.russell@rice.edu

Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, UK
Email address: a.sisto@hw.ac.uk

Mathematical Institute, University of Oxford, Oxford, UK
Email address: spriano@maths.ox.ac.uk