Finite Covers of Graphs and Cube Complexes

Sam Shepherd

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2021

Acknowledgements

Firstly, I thank my supervisor Martin Bridson for all that he has taught me and advised me over the past four years. He's always enthusiastic to explain some mathematical topic or idea to me, but he's also happy to get stuck in to the details of my own work, giving me lots of feedback (on both writing style and mathematics). Plus I'm grateful for the journal club that Martin hosts in his office (or virtually as it has been recently), which fosters much friendly discussion and keeps us all in the habit of reading papers beyond those directly related to our own research.

I have had the pleasure of collaborating with both Martin and Daniel Woodhouse. Working with Daniel has been great fun, and has arguably produced my best work so far, despite our awkward first meeting when we discovered that we'd proven the same theorem! Thanks also to my current collaborators Alex Margolis and Emily Stark.

Further thanks go to Ric Wade, David Hume, Mehdi Yazdi, Panos Papasoglu, Cornelia Drutu, Marc Lackenby, Elia Fioravanti, Jingyin Huang and Davide Spriano for fruitful discussions and/or for lecturing graduate courses that have shaped my mathematical development. And I thank my fellow DPhil students for the great times we've had together, from reading groups in the Quillen Room to eating cake in the Common Room; these wonderful people include Nici Heuer, Jan Steinebrunner, Luci Bonatto, Christoph Weis, Alice Kerr, Joe Scull, Ido Grayevsky Adele Jackson, Dario Ascari and Yonatan Fruchter.

Finally, huge thanks to my mum, dad and sister, and the rest of my family, for all their love and support over the past four years (and the twenty-two years that came before!).

Abstract

In this thesis we study finite covers of graphs, cube complexes and related spaces, and explore applications to rigidity of groups.

Leighton's Theorem states that two finite graphs with a common universal cover must have a common finite cover. We prove three generalisations of this. The first restricts how balls of a given size in the universal cover can map down to the two finite graphs when factoring through the common finite cover. The second generalises to covers of graphs of spaces that restrict to isomorphisms between vertex spaces. And thirdly, if the two finite graphs admit regular coverings by the same quasitree, then we arrange for the common finite cover to also be covered by this quasitree. In addition, we provide counter-examples to show that the assumptions in these generalisations cannot be relaxed.

Central to Haglund and Wise's theory of special cube complexes is the construction of the canonical completion and retraction, which enables one to build finite covers of special cube complexes in a highly controlled manner. We give a new approach to this construction with the idea of imitator covers. This idea provides greater insight into the construction and leads to powerful generalisations. It enables us to prove various results about finite covers of special cube complexes - most of which generalise existing theorems of Haglund and Wise to the non-hyperbolic setting. In particular, we prove a convex version of omnipotence for virtually special cubulated groups. Continuing the theme of special cube complexes, we give a novel account of Agol's proof that hyperbolic cubulated groups are virtually special; we retain the underlying ideas and constructions of Agol, but substantially change or add to many parts of the argument to give a more transparent and detailed account.

Finally, we apply Leighton-type arguments similar to those in the aforementioned results to obtain a quasi-isometric rigidity theorem. To be more precise, let G be a group that is one-ended, hyperbolic relative to virtually abelian subgroups, and has JSJ decomposition over two-ended subgroups containing only virtually free vertex groups that aren't quadratically hanging; we prove that any group quasi-isometric to G is abstractly commensurable to G. In particular, this theorem applies to certain "generic" HNN extensions of a free group over cyclic subgroups.

Statement of Originality

Chapter 2 is an exposition of background material written by the author but drawing on the sources cited within. Chapter 4 is written by the author in collaboration with Martin Bridson. Chapter 6 is written by the author but based on the proof of Agol [1] - the similarities and differences are explained within the chapter. Chapter 7 is written by the author in collaboration with Daniel Woodhouse. Chapters 3, 5 and 8 are entirely the author's own work.

Contents

Chapter 1

Introduction

The focus of this thesis is on the existence and construction of finite covers of graphs, cube complexes and related spaces. The study of such covers is at the heart of many of the spectacular advances in geometric group theory in recent years, which have had broad ramifications in adjacent areas of mathematics, particularly low-dimensional topology and geometry. Some examples include Haglund and Wise's theory of special cube complexes [44], which has its origins in Stallings' study of finite covers of graphs, Behrstock and Neumann's work on quasi-isometric classification of 3-manifold groups [9, 10], and Agol's proof of the virtual specialness of hyperbolic cubulated groups [1], leading to his resolution of Thurston's virtual fibering conjecture.

The problem of finding finite covers with prescribed properties of a space X can often be translated, via the Galois correspondence, into a problem about the existence of a certain type of finite-index subgroup of the fundamental group of X. So group theoretic properties of $\pi_1 X$, such as residual finiteness, subgroup separability or omnipotence, can be formulated in terms of lifts and elevations to finite covers of X. These three properties in particular have been central in the study of special cube complexes referred to above. Another circle of developments regarding finite-index subgroups is the study of rigidity phenomena, such as profinite rigidity and quasi-isometric rigidity. Some examples here include the work of Bridson, Reid, Wilton [16] and Wilkes [93, 94] on profinite rigidity of 3-manifold groups, and the result of Schwarz on the quasi-isometric rigidity of non-uniform lattices in rank one Lie groups [79].

The main results of this thesis concern the existence of finite covers with certain properties, the existence of common finite covers for pairs of spaces, and the consequences of such theorems, particularly with regard to quasi-isometric rigidity. The starting point of our investigations is the classical theorem of Leighton concerning the existence of a common finite cover for pairs of graphs that are covered by the same tree.

1.1 Graph-like spaces

The following theorem due to Leighton [57] motivates much of this thesis.

Theorem 1.1.1. *(Leighton)*
If two finite graphs have a common universal cover then they have a common finite cover.

It is natural to ask in what other contexts analogous theorems hold. More precisely, if Ω is a class of spaces then we will say that *Leighton's Theorem holds for* Ω if any two compact spaces in Ω with a common universal cover have a common finite cover. To fix ideas, we describe two simple examples of classes of spaces where we can show that Leighton's Theorem holds.

Example 1.1.2. Define a *graph of polygons* to be a space consisting of solid regular polygons with some edges joining vertices of the polygons. It was not previously known that Leighton's Theorem holds for such spaces.

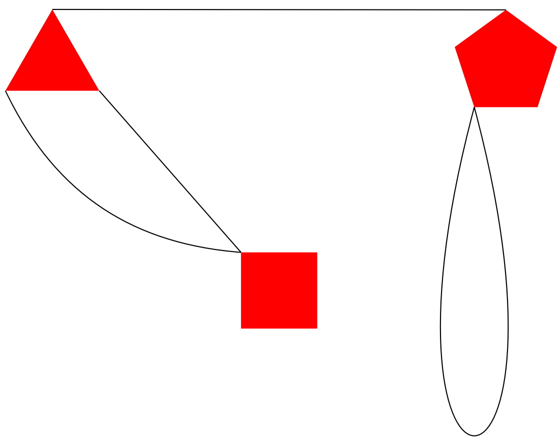

Figure 1.1: Example of a graph of polygons.

Example 1.1.3. Let X be a graph and let Δ be a collection of combinatorial immersions $\gamma : S \to X$, where each S is a (subdivided) circle or a bi-infinite line. We define the corresponding *graph with fins* to be the non-positively curved square complex obtained by taking the mapping cylinder of the following immersion:

$$\cup_\Delta \gamma : \sqcup_\Delta S \to X$$

Leighton's Theorem for graphs with fins was proved by Woodhouse [101], but we give an alternative proof here.

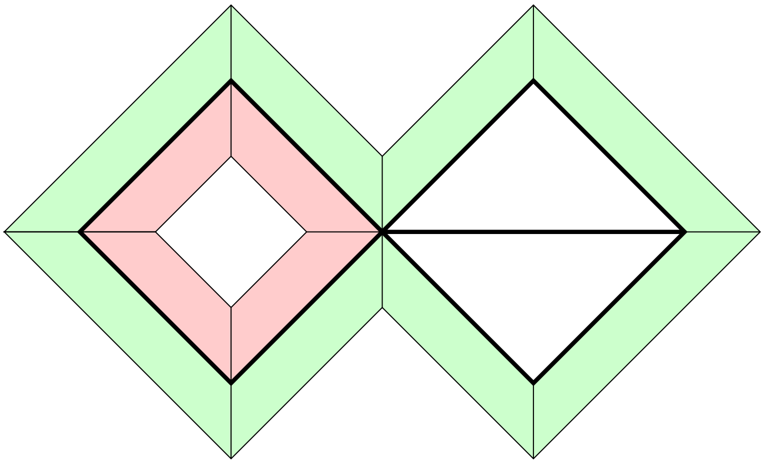

Figure 1.2: An example of a graph with fins. The underlying graph is bold, one fin is red and the other is green.

If two finite graphs X_1 and X_2 have a common universal cover T, with covering maps $p_i : T \to X_i$, then standard covering space theory tells us that any common finite cover $\hat{X}$ fits into a commutative diagram of covering maps as follows

$$\begin{array}{ccccc} T & & \xrightarrow{\quad g \quad} & & T \\ & \searrow^{q_1} & & \swarrow^{q_2} & \\ \downarrow{\scriptstyle p_1} & & \hat{X} & & \downarrow{\scriptstyle p_2} \\ & \swarrow_{\mu_1} & & \searrow_{\mu_2} & \\ X_1 & & & & X_2 \end{array} \tag{1.1.1}$$

where g is an automorphism of T.

Let the deck transformation group of $p_i : T \to X_i$ be denoted $\Gamma_i < \mathrm{Aut}(T)$. With the usual (compact-open) topology on $\mathrm{Aut}(T)$, we have that the Γ_i are free uniform lattices in $\mathrm{Aut}(T)$. Diagram (1.1.1) implies that the deck transformation group of $q_2 : T \to \hat{X}$ is contained in $\Gamma_1^g \cap \Gamma_2$. In particular, $\Gamma_1^g \cap \Gamma_2$ has finite index in Γ_1^g and Γ_2 - and we say that Γ_1^g is *commensurable* to Γ_2 in $\mathrm{Aut}(T)$. In fact Leighton's Theorem for graphs is equivalent to the statement that, for T a locally finite tree and for any two free uniform lattices $\Gamma_1, \Gamma_2 < \mathrm{Aut}(T)$, there exists $g \in \mathrm{Aut}(T)$ such that Γ_1^g is commensurable to Γ_2.

This viewpoint also works for other spaces, such as those from Examples 1.1.2 and 1.1.3. With this in mind, a possible strategy for solving Leighton's Theorem for other spaces is as follows. Given two spaces with universal cover $\tilde{X}$, find a "nice" map $\tilde{X} \to T$ to a tree T that induces a homomorphism $\mathrm{Aut}(\tilde{X}) \to \mathrm{Aut}(T)$. Given two free uniform lattices in $\mathrm{Aut}(\tilde{X})$ coming from the deck transformation groups, push them forward to $\mathrm{Aut}(T)$, then conjugate them to make them commensurable, then pull them back to $\mathrm{Aut}(\tilde{X})$. The problem with this approach is that it may not be possible to pull back to $\mathrm{Aut}(\tilde{X})$ if the homomorphism $\mathrm{Aut}(\tilde{X}) \to \mathrm{Aut}(T)$ is not surjective. To fix this problem, we need a way of keeping the automorphism g from diagram (1.1.1) within a given subgroup of $\mathrm{Aut}(T)$. This motivates the following definition and theorem.

Definition 1.1.4. Let T be a tree, let $H < \mathrm{Aut}(T)$, and let R be an integer. For a vertex $x \in VT$, we denote the R-ball centred at x by $B_R(x)$. We define the *R-symmetry restricted closure* of H to be:

$$\mathscr{S}_R(H) := \{g \in \mathrm{Aut}(T) \mid \forall x \in VT, \exists h \in H \text{ s.t. } h \text{ agrees with } g \text{ on } B_R(x)\}$$

It is easy to check that $\mathscr{S}_R(H) = \bigcap_{x \in VX} H\,\mathrm{Fix}(B_R(x))$, where $\mathrm{Fix}(B_R(x))$ is the pointwise stabiliser of $B_R(x)$ in $\mathrm{Aut}(T)$.

Theorem 1.1.5. *(Symmetry-restricted Leighton's Theorem)*
Let T be a tree, and $H < \mathrm{Aut}(T)$, and let $\Gamma_1, \Gamma_2 < H$ be free uniform lattices in $\mathrm{Aut}(T)$. Then for all $R \in \mathbb{N}$ there exists $g \in \mathscr{S}_R(H)$ such that Γ_1^g is commensurable to Γ_2 in $\mathrm{Aut}(T)$.

This theorem was proved by the author, but also independently proved by Gardam and Woodhouse - and the particular way we have stated the theorem is due to Gardam and Woodhouse. Both proofs are given in [84], we include the proof of the author in Section 3.3, which actually deduces the theorem from a more general result (see Theorem 1.1.6 below). Leighton's Theorem for Examples 1.1.2 and 1.1.3 can be deduced from this theorem as we sketched above, the details are in Section 3.3. This theorem also answers a question of Neumann from [65]. Versions of Leighton's Theorem, including Theorem 1.1.5, have been used to study commensurability and rigidity within certain classes of groups, which we discuss in Section 1.3.

Another form of Leighton's Theorem that it is natural to explore is the following: if two finite graphs of spaces are covered by the same tree of spaces, then is there always a finite graph of spaces that also covers them? One difficulty here is that a general covering between graphs of spaces induces coverings between the vertex spaces rather than homeomorphisms. In fact one can exploit this difficulty to construct a pair of finite graphs of spaces with a common universal cover but no common finite cover; this can be done with Baumslag-Solitar groups or with a graph of graphs due to Wise, which are discussed further in Section 3.2. Another serious difficulty that arises is when the symmetry group of an edge space is infinite. A way of solving this for many examples is by working in a different category of spaces, such as the category of finite simplicial complexes, where the automorphism groups of edge spaces are guaranteed to be finite.

To deal with both problems we define a *graph of objects* to be a graph of spaces with respect to a given category of spaces, and we define *coverings of graphs of objects* to be coverings of graphs of spaces that restrict to isomorphisms between edge and vertex spaces. We also define a bijective covering from a graph of objects X to itself to be an *automorphism* of X, and we denote the group of automorphisms by $\mathrm{Aut}(X)$. If X is a graph of objects and X_e is an edge space, then we have a homomorphism from the stabiliser $\mathrm{Aut}(X)_e$ of the edge e to $\mathrm{Aut}(X_e)$, the automorphism group of X_e with respect to the given category, and we call the image the *isotropy group of e in* $\mathrm{Aut}(X)$. Precise definitions for all these notions are given in Section 3.2. Our version of Leighton's Theorem for graphs of objects is then as follows.

Theorem 1.1.6. *(Graph of Objects Leighton's Theorem)*
Let X^1 and X^2 be finite graphs of objects covered by a tree of objects X. If $\mathrm{Aut}(X)$ has finite edge isotropy groups, then X^1 and X^2 have a common finite cover.

We actually prove a stronger version of this theorem which incorporates a notion of symmetry-restriction for graphs of objects (Theorem 3.2.10), and we deduce Theorem 1.1.5 from this version. Leighton's Theorem for Examples 1.1.2 and 1.1.3 can also be deduced straight from Theorem 1.1.6, as can the following example.

Example 1.1.7. Consider a closed 3-manifold that is cut into vertex spaces by a collection of disjoint embedded tori. Suppose also that the interior of each vertex space supports a complete finite-volume hyperbolic metric (which is unique by Mostow rigidity). Viewing the vertex space as a quotient of $\mathbb{H}^3$, each torus (edge space) adjacent to it will be a cusp cross-section, i.e. the image of a horosphere in $\mathbb{H}^3$, and as such will inherit a Euclidean metric that is unique up to scaling; we make it unique by requiring all the tori in our decomposition to have area a suitable constant. We make the vertex spaces compact by truncating the cusps with these horospherical tori. Note that each torus will have two Euclidean metrics, one induced from each of its adjacent vertex spaces - we do not require these to agree.

Let M_1, M_2 and $\tilde{M}$ be 3-manifolds with decompositions as above, such that the underlying graphs for M_1 and M_2 are finite and the underlying graph for $\tilde{M}$ is a tree. Suppose that there are coverings $\tilde{M} \to M_i$ that respect the decompositions and restrict to isometries between vertex spaces. Then there is another 3-manifold $\hat{M}$, with such a decomposition over a finite underlying graph, that also covers M_1 and M_2 in a way that restricts to isometries between vertex spaces.

A different way to extend Leighton's Theorem is to remain in the category of graphs but consider common covers other than the universal cover. The following theorem is from the joint paper with Bridson [17].

Theorem 1.1.8. *(Quasitree Leighton's Theorem)*
Given finite graphs X_1 and X_2, if there is a quasitree X and regular coverings $X \to X_i$, then there is a covering $X \to \hat{X}$ of a finite graph $\hat{X}$ such that $\hat{X}$ covers both X_1 and X_2.

We actually prove a slightly stronger theorem which incorporates symmetry-restriction (Theorem 4.2.5), and we deduce this from the Symmetry-restricted Leighton's Theorem via an argument similar to that described before Definition 1.1.4. We also get the following corollary.

Corollary 1.1.9. *Leighton's Theorem holds for simplicial complexes with virtually free fundamental group.*

It is natural to ask if the assumptions of Theorem 1.1.8 can be relaxed. The following result, also from [17], shows that we cannot take X to be an arbitrary regular cover of the X_i.

Theorem 1.1.10. *There exist connected graphs X, X_1, X_2 and regular covering maps $X \to X_1$ and $X \to X_2$, such that X_1 and X_2 are finite but there does not exist a covering $X \to \hat{X}$ of any finite graph $\hat{X}$ such that $\hat{X}$ covers both X_1 and X_2.*

Our proof of this theorem encodes structures from another class of spaces where Leighton's Theorem fails, namely that of compact, non-positively curved squared complexes. Key examples in this setting were constructed by Wise [100] and Burger-Mozes [19]. Our final result of the section, again from [17], shows that the regularity of the covers in Theorem 1.1.8 is necessary as well.

Theorem 1.1.11. *There exists a quasitree X, finite graphs X_1 and X_2, and covering maps $X \to X_1$ and $X \to X_2$, with $X \to X_2$ regular, such that there does not exist a covering $X \to \hat{X}$ of any finite graph $\hat{X}$ such that $\hat{X}$ covers both X_1 and X_2.*

1.2 Cube complexes

Gromov introduced the notion of CAT(0) cube complex into group theory [36], laying the foundation for an active subfield of Geometric Group Theory, driven by Wise, Sageev and others, in which groups are studied via their actions on such cube complexes. Many groups have been cubulated (sometimes referred to as cocompactly cubulated), including small cancellation groups, finite-volume hyperbolic 3-manifold groups, limit groups, many Coxeter groups, and hyperbolic free-by-cyclic groups - see for example [98]. Consequences of cubulating a group include that the group is CAT(0), bi-automatic [88] and that it satisfies the Tits Alternative [78]. Haglund and Wise introduced an important subclass of cube complexes, called special cube complexes, which grant groups even stronger properties [44]. Any group which is virtually the fundamental group of a finite special cube complex, referred to as a *virtually special group*, is residually finite, $\mathbb{Z}$-linear, conjugacy separable and has many separable subgroups [44, 63]. Agol proved that hyperbolic cubulated groups are virtually special, and used this to resolve the long-standing

Virtual Haken and Virtual Fibering Conjectures in 3-manifold theory [1]. In fact, being hyperbolic and cubulated is equivalent to being hyperbolic and virtually special [11].

Theorem 1.2.1. *(Agol)*
Let G be a hyperbolic group acting properly and cocompactly on a CAT(0) cube complex X. Then G has a finite index subgroup G' that acts freely on X such that the quotient X/G' is special.

In Chapter 6, we give a proof of Theorem 1.2.1. We retain the underlying ideas and constructions from Agol's proof, but substantially change or add to many parts of the argument to give a more transparent and detailed account. Our version of the proof has recently been used by Oregón-Reyes to generalise Theorem 1.2.1 to the relatively hyperbolic setting [67].

At the heart of the theory of special cube complexes is Haglund and Wise's construction of the canonical completion and retraction for a local isometry $\phi : Y \to X$ of (directly) special cube complexes [44], which enables one to build finite covers of special cube complexes in a highly controlled manner. In Chapter 5 we give a new interpretation of this construction: we imagine two people wandering around in these cube complexes, the *walker* wanders around X while the *imitator* wanders around Y, and the imitator tries to copy the walker by crossing over the same hyperplanes (pulled back along ϕ). We give a precise description in Section 5.1. Our interpretation not only provides greater insight into the Haglund–Wise construction, but also facilitates new applications. Furthermore, our construction admits natural generalisations. In Section 5.5 we generalise to multiple imitators, and in Section 5.7 we consider imitators who try to copy other imitators (forming a *hierarchy of imitators*). Throughout Chapter 5 we use these constructions to prove various results about virtually special cube complexes and virtually special groups, most of which generalise existing results of Haglund and Wise to the non-hyperbolic setting. We now describe our main results. First we need a definition.

Definition 1.2.2. (Commanding group elements)
A group G *commands* a set of elements $\{g_1, ..., g_n\} \subset G$ if there exists an integer $N > 0$ such that for any integers $r_1, ..., r_n > 0$ there exists a homomorphism to a finite group $G \to \bar{G}, g \mapsto \bar{g}$ such that the order of $\bar{g}_i$ is Nr_i. If this can always be done with $\langle \bar{g}_i \rangle \cap \langle \bar{g}_j \rangle = \{1\}$ for all $i \neq j$ then we say that G *strongly commands* $\{g_1, ..., g_n\}$.

Clearly a group can't command arbitrary sets of elements, for instance conjugate elements must have the same order in finite quotients. The best one can hope for is that a group commands any *independent* set of elements $\{g_1, ..., g_n\}$ (meaning the g_i have infinite order and no non-zero power of g_i is conjugate to a non-zero power of g_j for $i \neq j$) - such a group is called *omnipotent*. Omnipotence was defined by Wise in [97], where he proved that free groups are omnipotent. Subsequently omnipotence has been proven for surface groups [4], Fuchsian groups [95], and virtually special hyperbolic groups [96]. This third example naturally leads to the question of whether all virtually special groups are omnipotent. This however is far from true, for example $\mathbb{Z}^2$ does not command the set of elements $\{(1,0),(0,1),(1,1)\}$ (note this set is independent but not linearly independent). Our big insight is that virtually special groups do command independent sets of elements if we add the additional assumption that the elements are convex with respect to the given cubulation (see Definition 2.4.25). For the standard cubulation of $\mathbb{Z}^2$ by a square tiling of the plane, $(1,0)$ and $(0,1)$ are convex but not $(1,1)$. In general, convexity of a group element depends on the choice of cubulation, however it follows from the Cubical Flat Torus Theorem [99] that an infinite order element will be convex with respect to every cubulation if it has no non-trivial power contained in a $\mathbb{Z}^2$ subgroup. The following theorem is proved in Section 5.6.4. See Section 2.4.6 for the precise definitions involved. As an application, we use it to deduce that right-angled Artin groups command random sets of elements (Theorem 5.6.9).

Theorem 1.2.3. *Every virtually special cubulated group $G \curvearrowright X$ strongly commands every independent set of convex elements.*

We say that a group is *strongly omnipotent* if it strongly commands any independent set of elements. This notion was considered by Bridson and Wilton in [18], where they proved that virtually free groups are strongly omnipotent - this was a key step in their proof of the undecidability of the triviality problem for profinite completions. As a corollary of Theorem 1.2.3, we obtain strong omnipotence for virtually special hyperbolic groups (such a group has a virtually special cubulation $G \curvearrowright X$ by [1, 11], and the hyperbolicity ensures that all infinite order elements are convex [44, Proposition 7.2]), thus generalising the theorems of Wise and Bridson–Wilton. Note that the proof of Theorem 1.2.3 does not rely on any previous omnipotence results, so can also be viewed as a new proof for these results.

Corollary 1.2.4. *All virtually special hyperbolic groups are strongly omnipotent.*

The notion of commanding a collection of elements extends naturally to a notion of commanding a collection of subgroups (Definition 5.4.2). This provides a powerful tool for constructing finite covers of graphs of groups if the vertex groups command their incident edge groups (Proposition 5.4.8). Although the terminology of commanding is new to this thesis, there are similar ideas in the literature which have been used to study graphs of groups, which we discuss in Section 5.4.2. In particular, we have the following deep theorem as a consequence of Wise's Malnormal Special Quotient Theorem. This theorem was explained to the author by Woodhouse before the author formulated the notion of commanding. Note that a subgroup of a hyperbolic cubulated group $G \curvearrowright X$ is convex if and only if it is quasiconvex [44, Proposition 7.2].

Theorem 1.2.5. *Every virtually special hyperbolic group commands every almost malnormal collection of quasiconvex subgroups.*

We prove this result in Section 5.4.1, where we also extend it to the relatively hyperbolic setting (Theorem 5.4.6). We conjecture that this extends to the non-hyperbolic setting as well.

Conjecture 1.2.6. *Every virtually special cubulated group $G \curvearrowright X$ commands every almost malnormal collection of convex subgroups.*

We fall short of proving Conjecture 1.2.6, although we do prove some weaker versions (Theorems 5.6.1 and 5.6.4). Also, much of the groundwork in this thesis for proving Theorem 1.2.3 is done in the context of arbitrary convex subgroups, such as the following key theorem.

Theorem 1.2.7. *(Elevating to trivial wall projections)*
Let $Y_1, Y_2 \to X$ be local isometries of finite virtually special cube complexes, and let $K_1, K_2 < \pi_1 X$ be the corresponding subgroups (well-defined up to conjugacy). Suppose that K_1 has trivial intersection with every conjugate of K_2. Then there is a finite directly special cover $\hat{X} \to X$ such that all elevations of Y_1 and Y_2 are embedded, and each elevation of Y_1 has trivial wall projection onto each elevation of Y_2.

This partially generalises a theorem of Haglund and Wise [45, Theorem 4.25] to the non-hyperbolic setting. Trivial wall projection is defined in Definition 5.3.5; it is a concept due to Haglund and Wise that has a natural connection to imitator covers. Roughly speaking, the utility of Theorem 1.2.7 is that it allows us to control the elevations of Y_1 and Y_2 independently when passing to further finite covers of $\hat{X}$. This was a key ingredient in the proof of Haglund and Wise's combination theorem for special cube complexes [45]. The reason Theorem 1.2.7 is only a partial generalisation of [45, Theorem 4.25] is because it does not say anything about wall projections between different elevations of Y_1 (or Y_2); we believe this is the main hurdle to proving Conjecture 1.2.6, and also to potentially proving a non-hyperbolic combination theorem for special cube complexes.

The proof of Theorem 1.2.7 uses the separability of triple cosets of convex subgroups in virtually special cubulated groups - another new result of this thesis. The separability of convex subgroups themselves is due to [44, Corollary 7.9], while the separability of double cosets of convex subgroups is due to [67, Theorem A.1]. We give proofs of all three statements in Section 5.2, which all make direct use of imitator covers.

In Section 5.7 we construct covers using hierarchies of imitators, and prove the following theorem. This generalises [45, Theorem 4.25] to the non-hyperbolic setting.

Theorem 1.2.8. *(Virtually connected intersections)*
For $i = 1, ..., n$ let $(Z_i, z_i) \to (X, x)$ be based local isometries of finite virtually special cube complexes. Then there is a finite cover $(\dot{X}, \dot{x}) \to (X, x)$ such that the based elevations $\dot{Z}_i$ of Z_i are embedded in $\dot{X}$ and have connected intersections $\cap_{i \in E} \dot{Z}_i$ for any $\emptyset \neq E \subset \{1, ..., n\}$. Moreover, if the Z_i are embedded in X and do not inter-osculate with hyperplanes of X, then we may assume that the full intersection $\cap_{i=1}^n Z_i$ is isomorphic to its based elevation.

The notion of a subcomplex inter-osculating with a hyperplane generalises the notion of two hyperplanes inter-osculating, which is one of the pathologies excluded in the definition of special cube complex. Subcomplex-hyperplane inter-osculation first appeared in [67, Remark A.9], and we define it in Definition 2.4.17; the exclusion of such inter-osculations is a key criterion used throughout Chapter 5.

All of the results in this section are useful for controlling finite covers of special cube complexes, which lends credence to the following conjecture of Haglund [41, Problem 2.4].

Conjecture 1.2.9. *(Haglund)*
Leighton's Theorem holds for special cube complexes.

We note that Leighton's Theorem does not hold for arbitrary non-positively curved cube complexes, with counter-examples coming from work of Wise [100] and Burger–Mozes [19]. There are several cases where Conjecture 1.2.9 is known to hold, and the author has plans to tackle various other cases - we discuss this further in Chapter 8.

1.3 Applications to rigidity of groups

A major theme in geometry and group theory is the study of rigidity properties, and since Gromov's seminal essay [36] quasi-isometric rigidity of groups has been of particular importance. A finitely generated group is viewed as a geometric object by considering its word metric, and a map between groups is called a *quasi-isometry* if it distorts distances by at most a given linear amount and has a coarse inverse (see Section 2.7.1 for precise definitions). Quasi-isometry is an equivalence relation, so it divides the universe of groups into *quasi-isometry classes*; and a group is said to be *quasi-isometrically rigid* if it is equivalent in some strong algebraic sense (such as abstract commensurability or virtual isomorphism) to all groups in its quasi-isometry class. Groups known to be quasi-isometrically rigid include abelian groups [14, 70], free groups [30], surface groups [24, 34], non-uniform lattices in semisimple Lie groups [32, 79], and mapping class groups [8, 47] - see [55] for a survey. On the other hand, many groups are known to not be quasi-isometrically rigid, such as uniform lattices in rank-1 symmetric spaces, free products of surface groups [73, 92] and simple surface amalgams [27, 60, 86].

Leighton-type theorems are natural tools for proving quasi-isometric rigidity results: if one can promote a quasi-isometry of groups $G_1 \sim G_2$ into proper cocompact actions on a common space X, then the existence of a common finite cover of the orbi-spaces X/G_1 and X/G_2 would imply that G_1 and G_2 are abstractly commensurable. Indeed, this sort of argument was used by Behrstock–Neumann to prove a quasi-isometric rigidity result for certain non-geometric 3-manifold groups [10] - a result which the author is in the process of strengthening with Woodhouse.

Even if one cannot get control of quasi-isometries, Leighton-type theorems are still very useful for proving that certain groups are abstractly commensurable. For example Levitt used Leighton's Theorem to solve the commensurability problem for certain generalised Baumslag–Solitar groups [59]. And more recently, Stark and Woodhouse used Theorem 1.1.5 to prove action rigidity for free products of closed hyperbolic surface groups [87]. We discuss action rigidity further in Chapter 8, where we also describe an ongoing project joint with Margolis, Stark and Woodhouse to prove action rigidity for many more graphs of groups.

The main focus of this section is to describe joint work with Woodhouse on quasi-isometric rigidity for certain graphs of virtually free groups [85], the proofs for which are in Chapter 7. An important special case of our work is the following theorem.

Theorem 1.3.1. *Let G be a cyclic HNN extension or amalgamation of finite rank free groups of either of the following forms:*

$$\mathbb{F}_n *_{\mathbb{Z}} = \langle \mathbb{F}_n \mid t w_1 t^{-1} = w_2 \rangle \quad \textit{or} \quad \mathbb{F}_m *_{\mathbb{Z}} \mathbb{F}_n = \langle \mathbb{F}_m, \mathbb{F}_n \mid w_1 = w_2 \rangle$$

where $n, m \geq 2$ and $w_1, w_2 \in \mathbb{F}_m \cup \mathbb{F}_n$ are suitably random/generic elements that are not proper powers. If a finitely generated group G' is quasi-isometric to G, then G' is abstractly commensurable to G.

For the HNN extension, if $g^{-1} w_1 g = w_2$ or w_2^{-1} for $g \in \mathbb{F}_n$ then G is hyperbolic relative to $\langle w_1, gt \rangle$ - which is isomorphic to either $\mathbb{Z}^2$ or the Klein bottle group (and the latter has an index two $\mathbb{Z}^2$ subgroup). Otherwise G is hyperbolic; and the amalgamation is always hyperbolic. When we say that w_1 and w_2 are suitably random, what we really mean is that the induced line patterns on the vertex groups $\mathbb{F}_n$ and $\mathbb{F}_m$ are rigid, see Section 7.1 for more about rigid line patterns, and Remark 7.1.11 for a simple sufficient condition (which justifies the use of the word random). Theorem 1.3.1 follows from our more general Theorem 1.3.3 and Example 7.1.15.

We would like to work with groups that split as finite graphs of groups with virtually free vertex groups and two-ended edge groups. Denote this family of groups by $\mathscr{C}$. The subfamily of torsion-free groups $\mathscr{C}_{tf} \subseteq \mathscr{C}$ has finitely generated free vertex groups and infinite cyclic edge groups. These are very wide families of groups containing many surface groups, Baumslag–Solitar groups, and one-relator groups. Many of these groups will not be one-ended, or Gromov hyperbolic, or relatively hyperbolic, or residually finite, or quasi-isometrically rigid.

In [97], Wise showed that subgroup separability of a group $G \in \mathscr{C}_{tf}$ is equivalent to the non-existence of a non-Euclidean Baumslag–Solitar subgroup, or the non-existence of $1 \neq g, t \in G$ such that $tg^pt^{-1} = g^q$, where $|p| > |q|$. In Theorem 1.3.2 we give a similar criterion for all groups in $\mathscr{C}$, which we prove in Section 7.2. Furthermore, in [51] Wise and Hsu showed that all subgroup separable $G \in \mathscr{C}_{tf}$ are cubulated, and moreover in the hyperbolic case are virtually special. We show that all subgroup separable $G \in \mathscr{C}$ are virtually special.

Theorem 1.3.2. *Let $G \in \mathscr{C}$. The following are equivalent.*

(1) G is subgroup separable,

(2) G is balanced (in the sense of Definition 7.2.1, and with respect to any finite graph of groups decomposition of G with virtually free vertex groups and two-ended edge groups),

(3) G is hyperbolic relative to peripheral subgroups that are virtually $\mathbb{Z} \times \mathbb{F}_n$ $(n \geq 0)$.

(4) G is virtually special.

As a consequence of Theorem 1.3.2 we can deduce that subgroup separability is an invariant up to quasi-isometry within $\mathscr{C}$ (an application of [29, Theorem 1.6] to (3), and of [61] the quasi-isometric rigidity of $\mathbb{Z} \times \mathbb{F}_n$).

For our quasi-isometric rigidity theorem we require even stronger properties. This involves considering JSJ decompositions and quadratically hanging (QH) vertex groups, which we define in Section 2.6.2.

Theorem 1.3.3. *Let G be a one-ended group, with JSJ decomposition over two-ended subgroups containing only virtually free vertex groups and no QH vertex groups. If G is hyperbolic relative to virtually abelian subgroups, then any group quasi-isometric to G is abstractly commensurable to G.*

We prove this theorem in Sections 7.4 and 7.5. If T is a JSJ tree for G, then being hyperbolic relative to virtually abelian subgroups is equivalent to the stabilisers of the cylinders in T being virtually abelian (in fact virtually $\mathbb{Z}$ or $\mathbb{Z}^2$) - see Proposition 2.8.8 and Remark 7.2.10.

Closely related to Theorem 1.3.3 are the recent results of Taam–Touikan [89]. They prove the corresponding result in the hyperbolic setting with rigid vertex groups that are hyperbolic (closed) surface groups instead of rigid free vertex groups.

The rough strategy for proving Theorem 1.3.3 is as follows. Given a group G' quasi-isometric to G, we use Papasoglu's quasi-isometric invariance of JSJ splittings [72] to build graphs of spaces for G and G' (after passing to finite-index subgroups) in which the vertex spaces are graphs and the edge spaces are circles. We attach fins to each vertex space corresponding to the incident edge spaces. We then use work of Cashen–Macura on rigid line patterns in free groups [21] to ensure that the universal covers of the vertex spaces, viewed as trees with fins, are isomorphic for G and G'. These trees with fins also exhibit rigidity with respect to fin-preserving quasi-isometries. Finally, we apply a two-step Leighton-type argument to construct a common finite cover of the graphs of spaces for G and G'; first constructing common finite covers for pairs of vertex spaces, then gluing them together along fins. This Leighton-type argument contains a number of subtleties that we do not go into here, so it accounts for most of Chapter 7.

To close this section, we give some justification why the assumptions in Theorem 1.3.3 cannot be weakened. In Section 7.6 we show that the relative hyperbolicity assumption is necessary by giving an explicit counter-example that is hyperbolic relative to $\mathbb{Z} \times \mathbb{F}_2$. Quasi-isometries between infinite-ended groups are very flexible as shown by the work of Papasoglu and Whyte [73], so infinite-ended groups are typically not quasi-isometrically rigid; indeed combining this work with [92] shows that free products of surface groups are not quasi-isometrically rigid. So one-endedness is also a natural assumption for our theorem. Finally, simple surface amalgams are shown to not be quasi-isometrically rigid in [27, 60, 86], so it is necessary to exclude QH vertex groups in our theorem.

1.4 Gluing Equations

Many of the constructions in this thesis use Gluing Equations, which we now describe.

We have two spaces X_1 and X_2, which might be graphs or cube complexes or something similar, and we wish to build a common finite cover $\hat{X}$. Suppose for a moment that we already have $\hat{X}$, and suppose it admits a decomposition as a graph of spaces. For example, if $\hat{X}$ is just a graph then we can form such a decomposition by cutting every edge, so the edge spaces are just the midpoints of edges and each vertex space consists of a vertex of $\hat{X}$ together with a bunch of half edges joining to it. Then each vertex space is equipped with local isometries to X_1 and X_2 formed by restricting the coverings from $\hat{X}$ to X_1 and X_2. The strategy to build $\hat{X}$ is to reverse engineer this, start by constructing vertex spaces equipped with local isometries to X_1 and X_2 (which we will refer to as *vertex space triples*), and then try to glue them together along edge spaces in a way that respects the local isometries. For this gluing to be possible we must have an appropriate collection of vertex space triples. It is usually easy to figure out what all the possible vertex space triples are, the tricky part is determining how many copies of each triple to include in our collection. This task can be formulated by a set of Gluing Equations, where each variable tells us how many copies of a particular vertex space triple to include, and we have one equation for each type of edge space triple.

The above strategy was effectively how Leighton proved Theorem 1.1.1, although he didn't formulate it in these terms. In this case the graph of spaces decomposition is very simple and there are numerous approaches one could take to solving the Gluing Equations. The approach that Leighton took was particular to graphs, so to generalise this construction to other types of spaces we must find other approaches.

The proof of the Graph of Objects Leighton's Theorem (Theorem 1.1.6) solves the Gluing Equations by defining a groupoid that consists of maps between the vertex spaces. In general, groupoids are very natural objects that arise in a wide range of mathematical contexts, and they are especially powerful in topology for stitching together local information in a basepoint free manner into a global structure. We also include a proof of Theorem 1.1.1 using the groupoid technique.

Theorem 1.2.1 is a little different; instead of constructing a common finite cover we must construct a finite cube complex which itself is covered by the given CAT(0) cube complex X, and crucially this complex must be special. The proof does still follow the above template but with the definition of vertex space triple kind of reversed: instead of having two local isometries mapping from the vertex space we have one local isometry mapping from a subspace of X to the vertex space. The other difference with this construction is that we must do the gluing in several stages, corresponding to a hierarchy of hyperplanes; and at some stages we may even need to pass to finite covers of the vertex spaces before their edge spaces can be glued together. Nevertheless, all of these various stages of gluing can be encapsulated in a single set of Gluing Equations, and these are solved by constructing an invariant measure on a space of hyperplane colourings, which involves a compactness argument.

In the proof of Theorem 1.3.3 we again want to build a common finite cover of two spaces X_1 and X_2. In this setting X_1 and X_2 are graphs of spaces, and the vertex spaces are graphs with fins. There are two stages of gluing: first we build common finite covers for pairs of vertex spaces, then we glue these together along fins. For the first stage we consider the tree with fins that covers the pair of vertex spaces, and we solve the Gluing Equations by utilising the Haar measure on the automorphism group of the tree with fins. The second stage of gluing has a separate set of Gluing Equations, and the solution is more involved, in particular it relies on doing the first stage of gluing in a highly symmetrical manner. It's also necessary to pass to finite covers of the vertex spaces in between the two stages of gluing.

Chapter 2

Preliminaries

2.1 Graphs, links and covers

Graphs, links and covers are basic concepts used throughout this thesis.

Definition 2.1.1. (Graphs)
A *graph* G is defined by the following data:

- A vertex set VG.
- An edge set EG.
- Maps $\iota, \tau : EG \to VG$ to denote the initial and terminal vertex of each edge.
- An involution $EG \to EG$, $e \mapsto \bar{e}$, which denotes the inversion of an edge, such that e and $\bar{e}$ are always distinct, and such that $\iota(e) = \tau(\bar{e})$ and $\iota(\bar{e}) = \tau(e)$ for any $e \in EG$.

Note that ι is redundant if τ and edge inversion have already been defined. For $A \subset EG$ we will use the notation $\bar{A} := \{\bar{e} \mid e \in A\}$.

A *graph morphism* $\alpha : G_1 \to G_2$ is given by maps $\alpha : VG_1 \to VG_2$ and $\alpha : EG_1 \to EG_2$ that preserve the graph structure given by ι, τ and edge inversion. Note that it is enough to check that τ and edge inversion are preserved. A graph morphism $G \to G$ that is bijective on edge and vertex sets is called an *automorphism*, and the group of automorphisms of a graph G is denoted $\mathrm{Aut}(G)$.

Definition 2.1.2. (Links and covers)
Let G be a graph. For $v \in VG$ define the *link* of v by

$$\mathrm{lk}(v) = \{e \in EG \mid \tau e = v\}.$$

A graph morphism $\alpha : G_1 \to G_2$ is a *covering* if it is surjective and the induced maps $\mathrm{lk}(v) \to \mathrm{lk}(\alpha(v))$ are bijections. In this case we say that G_1 is a *cover* of G_2.

2.2 Groupoids

Groupoids are natural objects that arise in a wide range of mathematical contexts, and they are especially powerful in topology for stitching together local information in a basepoint free manner into a global structure. We will use them to build finite covers in Chapter 3. See [15, Chapter III.$\mathcal{G}$] for further background on groupoids.

Definition 2.2.1. (Groupoids)
A *groupoid* $\mathcal{G}$ is a small category in which all morphisms are invertible. We will use $\mathrm{Ob}(\mathcal{G})$ to denote the set of objects, and when referring to a morphism g we will simply write $g \in \mathcal{G}$. For $g \in \mathcal{G}$, we will denote the initial and terminal objects by $i(g)$ and $t(g)$. For $x \in \mathrm{Ob}(\mathcal{G})$ we write 1_x for the identity morphism of x.

For $x, y \in \mathrm{Ob}(\mathcal{G})$ it will helpful to have the following additional notation.

$$\begin{aligned} \mathcal{G}(x,y) &:= \{g \in \mathcal{G} \mid i(g) = x,\ t(g) = y\} \\ \mathcal{G}(x,-) &:= \{g \in \mathcal{G} \mid i(g) = x\} \\ \mathcal{G}(-,y) &:= \{g \in \mathcal{G} \mid t(g) = y\} \end{aligned}$$

A *subgroupoid* of $\mathcal{G}$ is a subcategory in which all morphisms are invertible.

When piecing together the finite cover in our proof of Leighton's Theorem, we will make use of groupoid actions. These are a direct analogue to group actions, and they also give rise to notions of orbit and stabiliser. The definition of groupoid action given below is from [15, III.$\mathcal{G}$.2.8(3)].

Definition 2.2.2. (Groupoid actions)
An *action* of a groupoid $\mathcal{G}$ on a set A consists of a map $\varepsilon : A \to \mathrm{Ob}(\mathcal{G})$ and a map

$$\begin{aligned} \{(g,a) \in \mathcal{G} \times A \mid i(g) = \varepsilon(a)\} &\to A \\ (g,a) &\mapsto g \cdot a, \end{aligned}$$

such that

(1) $\varepsilon(g \cdot a) = t(g)$,

(2) $(g'g) \cdot a = g' \cdot (g \cdot a)$,

(3) $1_{\varepsilon(a)} \cdot a = a$,

for any $a \in A$ and $g, g' \in \mathcal{G}$ satisfying $i(g) = \varepsilon(a)$ and $i(g') = t(g)$.

Definition 2.2.3. (Orbits and stabilisers of groupoid actions)
If a groupoid $\mathcal{G}$ acts on a set A and $\mathcal{H} \subset \mathcal{G}$, define the *$\mathcal{H}$-orbit of* $a \in A$ by

$$\mathcal{H} \cdot a := \{h \cdot a \mid h \in \mathcal{H},\ i(h) = \varepsilon(a)\}.$$

Similarly, for $B \subset A$ write $\mathcal{H} \cdot B := \cup_{b \in B} \mathcal{H} \cdot b$. Define the *stabiliser of* $a \in A$ by

$$\mathrm{Stab}_{\mathcal{G}}(a) := \{g \in \mathcal{G} \mid i(g) = \varepsilon(a),\, g \cdot a = a\}.$$

When building the finite cover in Leighton's Theorem we must find appropriate matchings between the pieces we wish to stitch together. The following lemma will help us achieve this.

Lemma 2.2.4. *(Groupoid Orbit-Stabiliser Theorem)*
Let $\mathcal{G}$ be a groupoid acting on a set A, and fix $a \in A$ with $\varepsilon(a) = x$. Then the fibres of the map

$$\begin{aligned} \phi_a : \mathcal{G}(x,-) &\to \mathcal{G} \cdot a \\ g &\mapsto g \cdot a \end{aligned}$$

are cosets $g\,\mathrm{Stab}_{\mathcal{G}}(a) := \{gg' \,|\, g' \in \mathrm{Stab}_{\mathcal{G}}(a)\}$ for $g \in \mathcal{G}(x,-)$. If $\mathcal{G}(x,-)$ is finite, we deduce that

$$|\mathcal{G}(x,-)| = |\,\mathrm{Stab}_{\mathcal{G}}(a)||\mathcal{G} \cdot a|.$$

Proof. If $g, h \in \mathcal{G}(x,-)$ satisfy $g \cdot a = h \cdot a$, then $g^{-1}h \in \mathrm{Stab}_{\mathcal{G}}(a)$, so $h \in g\mathrm{Stab}_{\mathcal{G}}(a)$. □

2.3 CAT(0) spaces

Alexandrov's definition of CAT(0) space generalises non-positive curvature of Riemannian manifolds to arbitrary geodesic spaces. In the past forty years CAT(0) spaces have played an important role in geometric group theory, in large part due to the influence of Gromov. Our use of CAT(0) spaces is largely via CAT(0) cube complexes, which appear in Section 2.4. This section is based on the book of Bridson–Haefliger [15].

Definition 2.3.1. (Comparison triangles and Alexandrov angles)
Let X be a metric space. A *comparison triangle* for a triple of points (p, q, r) in X is a triangle in the Euclidean plane with vertices $\bar{p}, \bar{q}, \bar{r}$ such that $d(p, q) = d(\bar{p}, \bar{q})$, $d(q, r) = d(\bar{q}, \bar{r})$ and $d(r, p) = d(\bar{r}, \bar{p})$. Such a triangle is unique up to isometry. The interior angle of the comparison triangle at $\bar{p}$ is denoted $\bar{\angle}_p(q, r)$.

Let $c : [0, a] \to X$ and $c' : [0, a'] \to X$ be two geodesic paths with $c(0) = c'(0)$. The *Alexandrov angle* between is c and c' is defined by

$$\angle(c, c') := \limsup_{t,t' \to 0} \bar{\angle}_{c(0)}(c(t), c'(t')).$$

This coincides with the usual notion of angle in Euclidean space.

There are several equivalent definitions for CAT(0) space [15, Proposition II.1.7], here is one of them.

Definition 2.3.2. (CAT(0) spaces)
A geodesic space X is *CAT(0)* if for any geodesic triangle Δ in X, with vertices (p, q, r), the Alexandrov angle of Δ at p is no greater than $\bar{\angle}_p(q, r)$.

CAT(0) spaces enjoy many nice properties. For example they are contractible, any two points are connected by a unique geodesic, and any local geodesic is a geodesic. CAT(0) spaces also exhibit a number of rigidity phenomena concerning the existence of flat subspaces, such as the following theorem [15, Theorem II.2.13].

Theorem 2.3.3. *(Flat Strip Theorem)*
Let γ, γ' be bi-infinite geodesics in a CAT(0) space that lie within bounded neighbourhoods of each other. Then the convex hull of $\gamma \cup \gamma'$ is isometric to a flat strip $\mathbb{R} \times [0, D] \subset \mathbb{E}^2$.

We now consider isometries.

Definition 2.3.4. Let g be an isometry of a metric space X. The *translation length* of g is the number

$$|g| := \inf_{x \in X} d(gx, x).$$

We say that g is *semi-simple* if this infimum is attained. Such a g is called *elliptic* if $|g| = 0$ and *hyperbolic* if $|g| > 0$.

The following general proposition about isometries of metric spaces appears as [15, Proposition II.6.10(2)] (the assumption that X is proper is missing from [15]).

Proposition 2.3.5. *Suppose that a group G acts properly (i.e. with finite point stabilisers and discrete orbits) and cocompactly on a proper metric space X. Then every element of G is semi-simple.*

Isometries of CAT(0) spaces are especially well-behaved, as exemplified by the following theorem [15, II.6.7, II.6.8].

Theorem 2.3.6. *Let g be an isometry of a complete CAT(0) space X.*

(1) g is elliptic if and only if it has bounded orbits. In particular, any finite-order isometry has a fixed point.

(2) g is hyperbolic if and only if there is a bi-infinite geodesic $c : \mathbb{R} \to X$ which is translated non-trivially by g, namely $g \cdot c(t) = c(t + a)$ for some $a > 0$. The set $c(\mathbb{R})$ is called an axis *of g, and the number a is actually equal to $|g|$.*

Corollary 2.3.7. *Let G be a group acting properly and cocompactly on a proper CAT(0) space X. Suppose $g \in G$ has infinite order and suppose $hg^kh^{-1} = g^l$ for $h \in G$ and $0 \neq k, l \in \mathbb{Z}$. Then $k = \pm l$.*

Proof. By Proposition 2.3.5 and Theorem 2.3.6, g translates along an axis $\gamma \subset X$ with translation length $|g| > 0$. So g^k and g^l also translate along γ, with translation lengths $|k||g|$ and $|l||g|$ respectively. Theorem 2.3.6(2) tells us that $|g^k| = |k||g|$ and $|g^l| = |l||g|$; but g^k and g^l are conjugate, so $|g^k| = |g^l|$ and $|k| = |l|$. □

We close the section with a proposition about projection maps.

Proposition 2.3.8. *Let X be a CAT(0) space and let C be a convex subset which is complete in the induced metric. Then for each $x \in X$ there exists a unique closest point $p(x) \in C$. We call $p : X \to C$ the* projection to C. *Moreover, p is distance non-increasing and commutes with any isometry of X that preserves C.*

2.4 Cube complexes

Cube complexes have been important objects of study in geometric group theory, particularly in the last twenty years, as explained in Section 1.2. They are the focus of Chapters 5 and 6, although they appear in the other chapters as well. The definitions in this section are mostly based on those of Haglund–Wise [44].

2.4.1 Cube complexes and hyperplanes

Definition 2.4.1. (Cube complex and metric)
A *cube complex* is a metric polyhedral complex in which all polyhedra are unit Euclidean cubes. We will denote the metric by d - and also use d for the distance between subsets of a cube complex, $d(A, B) = \inf\{d(x, y) | \, x \in A, \, y \in B\}$. Any finite dimensional cube complex is a complete geodesic space [15, I.7.33].

Another metric which is commonly used on the vertex set of a cube complex, referred to as the ℓ_1 or *combinatorial metric*, defines the distance between two vertices to be the length of a shortest edge path between them. We will not use this metric directly, although we will often consider shortest edge paths.

Notation 2.4.2. Let X be a cube complex. We let X^n denote its n-skeleton. We refer to the 0-cubes, 1-cubes and 2-cubes of X as *vertices*, *edges* and *squares* respectively. Note that, in this thesis, edges of graphs are oriented, and denoted $e \in EX$ (Definition 2.1.1), whereas edges in cube complexes are unoriented, and denoted $e \in X^1$. By a *path* in X we will always mean an edge path, and a *loop* is an edge path that starts and finishes at the same vertex. We let link(x) denote the *link* of a vertex x; this combinatorial complex is the space of directions at x. In this thesis all cube complexes and their covers will be connected unless otherwise stated.

Definition 2.4.3. (Hyperplanes)
An n-cube $C = [-1, 1]^n \subset \mathbb{R}^n$ has n midcubes, each obtained by setting one coordinate to zero. The face of a midcube of C is naturally the midcube of a face of C, and so the collection of midcubes of a cube complex X can be given the structure of a cube complex. We call this the *hyperplane complex*; it has a natural immersion $q : \mathcal{H} \to X$ induced by inclusions of the midcubes - note this immersion is neither combinatorial nor an embedding, but it can be made combinatorial by passing to the cubical subdivisions of $\mathcal{H}$ and X. A component of the hyperplane complex is called a *hyperplane*, but when talking about a hyperplane we will usually be referring to its image under q. A picture of a hyperplane in a cube complex is given in Figure 2.1.

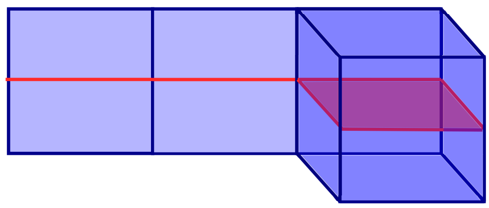

Figure 2.1: A hyperplane in a cube complex.

For an edge e in a cube complex X, write $H(e)$ for the unique hyperplane that it intersects, and say that e is *dual* to $H(e)$. If a group G acts on X then it also acts on the set of hyperplanes, and clearly $gH(e) = H(ge)$ for $g \in G$, $e \in X^1$. For a hyperplane H, let G_H denote its stabiliser. Given an edge e dual to H, $g \in G_H$ if and only if $H(ge) = H$ - thus G_H is also the stabiliser of the set of edges dual to H.

Remark 2.4.4. If the action of G on X is proper and cocompact then so is the action of G_H on H. Properness is immediate, and cocompactness is a consequence of the following easy argument. Let $\{e_1, ..., e_k\}$ be edges dual to H which are representatives of those G-orbits of edges in X that include edges dual to H. For an edge e dual to H, there exists $g \in G$ with $ge = e_j$ some $1 \leq j \leq k$, but then $g \in G_H$. Therefore the union of midcubes that intersect $\{e_1, ..., e_k\}$ is a compact subset of H with G_H-translates that cover H.

2.4.2 Directly special cube complexes

Definition 2.4.5. (NPC and CAT(0) cube complexes)
A cube complex is *nonpositively curved* (we will use the shorthand NPC) if the link of each vertex is a flag complex, and an NPC cube complex is *CAT(0)* if it is simply

connected. This is equivalent to a cube complex being CAT(0) in the sense of Definition 2.3.2 [56, appendix].

Definition 2.4.6. (Parallelism)
Let X be a cube complex. Two edges $e_1, e_2 \in X^1$ are *elementary parallel* if they appear as opposite edges of some square of X. The relation of elementary parallelism generates the equivalence relation of *parallelism*. We write $e_1 \parallel e_2$ if edges e_1 and e_2 are parallel. We define $H(e_1)$ to be the hyperplane dual to an edge e_1 - note that $e_1 \parallel e_2$ is equivalent to $H(e_1) = H(e_2)$.

Definition 2.4.7. (Intersecting and osculating hyperplanes)
Let X be a cube complex. Suppose distinct edges e_1 and e_2 of X are incident at a vertex x.

(1) If e_1 and e_2 form the corner of a square at x, then we say that the hyperplanes $H(e_1)$ and $H(e_2)$ *intersect at* $(x; e_1, e_2)$. If in addition $H(e_1) = H(e_2)$, then we say that $H(e_1)$ *self-intersects at* $(x; e_1, e_2)$. Note that a pair of hyperplanes intersect (as subsets of X) if and only if they are equal or they intersect at some $(x; e_1, e_2)$.

(2) If e_1 and e_2 do not form the corner of a square at x, then we say that the hyperplanes $H(e_1)$ and $H(e_2)$ *osculate at* $(x; e_1, e_2)$. If in addition $H(e_1) = H(e_2)$, then we say that $H(e_1)$ *self-osculates at* $(x; e_1, e_2)$. Alternatively, if e_1 has both its ends incident at x, then we say that $H(e_1)$ *self-osculates at* $(x; e_1)$. We say that a pair of hyperplanes *osculate* if they osculate at some $(x; e_1, e_2)$.

(3) We say that distinct hyperplanes H_1 and H_2 *inter-osculate* if they both intersect and osculate.

We will sometimes just say that a pair of hyperplanes intersect or osculate at a vertex x if we do not wish to specify edges e_1 and e_2. The notation from [44] is slightly different as they work with oriented edges and distinguish between direct and indirect self-osculations.

Definition 2.4.8. (Two-sided hyperplanes)
A hyperplane H is *two-sided* if the map $H \to X$ extends to a combinatorial map $H \times [-1, 1] \to X$ (where we consider H with its induced cube complex structure). This map is unique up to a reflection of $[-1, 1]$. We define sets of vertices $H^+, H^- \subset X^0$, where $H^\pm$ is the image of the vertices $H^0 \times \{\pm 1\}$ in X (so there is a choice of which set is H^+ and which is H^-).

Remark 2.4.9. For any hyperplane H the map $H \to X$ extends to a combinatorial map of a bundle over H with fibre $[-1, 1]$, H being two-sided means that this bundle is trivial.

Definition 2.4.10. (Directly special and virtually special cube complexes)
We say that an NPC cube complex X is *directly special* if every hyperplane is two-sided, no hyperplane self-intersects or self-osculates, and no pair of hyperplanes inter-osculate. These forbidden behaviours are shown in Figure 2.2. An NPC cube complex X is *virtually special* if it has a finite-sheeted directly special cover.

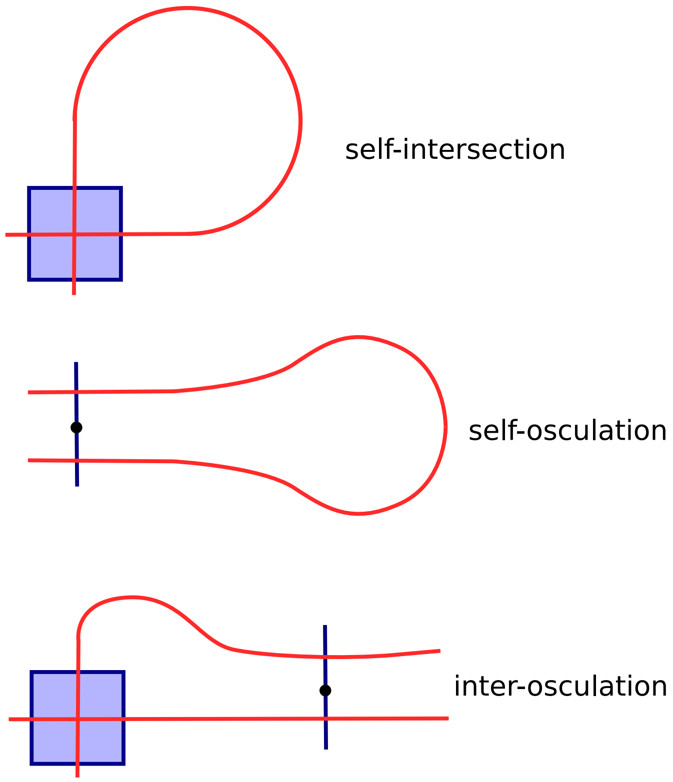

Figure 2.2: Forbidden hyperplane behaviours. The hyperplanes are red, the edges of the cube complex are blue.

Remark 2.4.11. The definition of directly special cube complex appears in [52]; it is not the same as the original notion of specialness [44], but among finite NPC cube complexes these notions are equivalent up to finite covers [44, Proposition 3.10].

Proposition 2.4.12. *If $X_1, ..., X_n$ are directly special cube complexes and $x_i \in X_i$ are choices of base vertices, then the cube complex obtained from $\sqcup_i X_i$ by identifying $x_1, ..., x_n$ is also directly special.*

Proof. This follows immediately from the definition of directly special cube complex because each hyperplane will be contained in a single X_i factor. □

2.4.3 Local isometries

Definition 2.4.13. (Local isometries)
A *local isometry* $\phi : Y \to X$ of NPC cube complexes is a combinatorial map such that

each induced map $\text{link}(y) \to \text{link}(\phi(y))$ is an embedding with image a full subcomplex of $\text{link}(\phi(y))$ (a subcomplex C of a simplicial complex D is *full* if any simplex of D whose vertices are in C is in fact entirely contained in C). Note that ϕ will also be a local isometry with respect to the metrics of X and Y, in the sense that small balls in Y will be isometrically embedded in X. A subcomplex $Y \subset X$ is *locally convex* if the inclusion $Y \hookrightarrow X$ is a local isometry.

Remark 2.4.14. Suppose $\phi : Y \to X$ is a local isometry of NPC cube complexes which contain no edges that are loops (such as directly special cube complexes), and suppose f_1, f_2 are edges incident at $y \in Y^0$. Then f_1, f_2 form the corner of a square at y if and only if $\phi(f_1), \phi(f_2)$ form the corner of a square at $\phi(y)$ (this is true without the no-edge-loop assumption if we work with oriented edges).

Remark 2.4.15. If $\phi : Y \to X$ is a local isometry and X is directly special, then Y is directly special. Covering maps are local isometries, so in particular any cover of a directly special cube complex is directly special.

Notation 2.4.16. In a directly special cube complex we write $e \in \text{link}(x)$ if e is an edge incident at a vertex x. This notation makes sense because only one end of e is incident at x ($H(e)$ doesn't self-osculate), and so e defines a unique vertex of $\text{link}(x)$.

2.4.4 More intersections and osculations

The notion of two hyperplanes in a cube complex X intersecting or osculating generalises to the notion of a hyperplane and a complex $Y \to X$ intersecting or osculating as follows. This generalisation first appeared in [67, Remark A.9].

Definition 2.4.17. (Intersections and osculations of hyperplanes with complexes)
Let $\phi : Y \to X$ be a local isometry of NPC cube complexes. If H is a hyperplane of Y, we denote by $\phi[H]$ the unique hyperplane of X that contains the image of H (equivalently, if $H = H(f)$ then $\phi[H] = H(\phi(f))$). Now let H be a hyperplane in X.

(1) We say that H and Y *intersect* if H intersects $\phi(Y)$, or equivalently if $H = \phi[H']$ for some hyperplane H' in Y.

(2) We say that H and Y *osculate at* $(y; e)$ if $y \in Y^0$, $e \in X^1$ is an edge incident to $\phi(y)$, $H = H(e)$, and no edge $f \in Y^1$ incident to y has $\phi(f) = e$.

(3) We say that H and Y *inter-osculate* if they both intersect and osculate.

If $Y \subset X$ is a locally convex subcomplex, then we can talk about Y intersecting/osculating/inter-osculating with a hyperplane of X by applying the above definitions to the inclusion $Y \hookrightarrow X$.

Definition 2.4.18. (Hyperplane carriers)
Let H be a hyperplane in a directly special cube complex X. The *carrier of* H, denoted $N(H)$, is the smallest subcomplex of X containing H. If X is not directly special, one can still define the carrier of H as a certain immersion $N(H) \to X$ as in Remark 2.4.9, but this is more delicate and will not concern us.

Remark 2.4.19. Let H be a hyperplane in a directly special cube complex X. Because H does not self-intersect or self-osculate, the map $H \times [-1, 1] \to X$ from Definition 2.4.8 is an embedding with image $N(H)$, and it maps $H \times \{0\}$ identically onto H. In particular, the inclusion $H \hookrightarrow N(H)$ is a homotopy equivalence.

Remark 2.4.20. In a directly special cube complex X, a pair of hyperplanes H_1 and H_2 intersect (resp. osculate) if and only if H_1 and $N(H_2) \to X$ intersect (resp. osculate) - in fact this remains true in arbitrary cube complexes with the more general definition of hyperplane carrier. Also, the carriers $N(H_1)$ and $N(H_2)$ are disjoint if and only if H_1 and H_2 neither intersect nor osculate.

2.4.5 Elevations

Definition 2.4.21. (Elevations)
Let $\phi : Y \to X$ be a local isometry and $\mu : \hat{X} \to X$ a finite cover. We say that a map $\hat{\phi} : \hat{Y} \to \hat{X}$ is an *elevation* of ϕ to $\hat{X}$ if there exists a covering $\nu : \hat{Y} \to Y$ fitting into the commutative diagram

$$\begin{array}{ccc} \hat{Y} & \xrightarrow{\hat{\phi}} & \hat{X} \\ \downarrow{\scriptstyle \nu} & & \downarrow{\scriptstyle \mu} \\ Y & \xrightarrow{\phi} & X, \end{array} \tag{2.4.1}$$

and such that a path in $\hat{Y}$ closes up as a loop if and only if its projections to Y and $\hat{X}$ both close up as loops. The map $\hat{\phi}$ is necessarily a local isometry. Equivalently, $\hat{Y}$ is a component of the pullback of ϕ and μ. If (2.4.1) is a diagram of based spaces, then we call it a *based elevation*.

Remark 2.4.22. Based elevations correspond to intersecting subgroups. Indeed, if (2.4.1) is a diagram of based spaces with respect to basepoints $y \in Y$, $\hat{y} \in \hat{Y}$, $x \in X$ and $\hat{x} \in \hat{X}$, then

$$(\phi\nu)_*\pi_1(\hat{Y}, \hat{y}) = \phi_*\pi_1(Y, y) \cap \mu_*\pi_1(\hat{X}, \hat{x}).$$

And choosing a different elevation corresponds to conjugating. Indeed if $\hat{\phi}(\hat{y}) \neq \hat{x}$ but instead $\hat{\gamma}$ is a path in $\hat{X}$ from $\hat{x}$ to $\hat{\phi}(\hat{y})$, then

$$(\phi\nu)_*\pi_1(\hat{Y}, \hat{y}) = \gamma\phi_*\pi_1(Y, y)\gamma^{-1} \cap \mu_*\pi_1(\hat{X}, \hat{x}),$$

where $\gamma = \mu\hat{\gamma}$.

2.4.6 Cubulated groups and convex subgroups

Definition 2.4.23. (Cubulated and virtually special groups)
A *cubulated group* $G \curvearrowright X$ is a finitely generated group G together with a geometric action on a CAT(0) cube complex X by cubical automorphisms. We say that $G \curvearrowright X$ is *virtually special* if there is a finite-index subgroup $\hat{G} < G$ acting freely on X with special quotient $X/\hat{G}$.

Remark 2.4.24. In this thesis we will usually consider a finite NPC cube complex X with fundamental group G, in this case $G \curvearrowright \tilde{X}$ is a cubulated group, where $\tilde{X}$ is the universal cover of X and G acts on $\tilde{X}$ by deck transformations. Note that G is torsion-free since any torsion element would fix a point in $\tilde{X}$ by Theorem 2.3.6(1).

Definition 2.4.25. (Convex subgroups and elements)
Let $G \curvearrowright X$ be a cubulated group. A subgroup $K < G$ is *convex* if it stabilises a convex subcomplex $Y \subset X$ with finite quotient Y/K. An element $g \in G$ is *convex* if $\langle g \rangle < G$ is convex.

Remark 2.4.26. If $G = \pi_1(X, x)$ is the fundamental group of a finite NPC cube complex X, then $K < G$ being convex is equivalent to K being the image of a homomorphism $\phi_* : \pi_1(Y, y) \hookrightarrow G$, defined by a local isometry $\phi : Y \to X$, with Y finite, and (a homotopy class of) a path γ in X from x to $\phi(y)$. We recover the first definition with respect to a based universal cover $(\tilde{X}, \tilde{x}) \to (X, x)$ by letting $\tilde{\gamma}$ be a lift of γ from $\tilde{x}$ to $\tilde{y}$, and setting $\tilde{Y}$ to be the based elevation of $\phi : Y \to X$ with respect to basepoints $y, \phi(y)$ and $\tilde{y}$ (note that $\tilde{Y} \to \tilde{X}$ is an embedding because local isometries of CAT(0) cube complexes are embeddings). The other elevations of Y are stabilised by conjugates of K.

One can eliminate the need for the path γ in Remark 2.4.26 by the following lemma.

Lemma 2.4.27. *Let X be a finite NPC cube complex and let $K < G := \pi_1(X, x)$ be a convex subgroup. Then there is a based local isometry of finite cube complexes $\phi : (Y, y) \to (X, x)$ with $K = \phi_* \pi_1(Y, y)$.*

Proof. Let $(\tilde{X}, \tilde{x}) \to (X, x)$ be the universal cover. Then G acts on $\tilde{X}$ by deck transformations, and there is a convex subcomplex $\tilde{Y} \subset \tilde{X}$ stabilised by K with finite quotient $\tilde{Y}/K$. By replacing $\tilde{Y}$ with a cubical thickening as in [45, Section 4], we can assume that $\tilde{x} \in \tilde{Y}$. Then we can put $(Y, y) := (\tilde{Y}, \tilde{x})/K$, and define ϕ as the quotient map to $(X, x) = (\tilde{X}, \tilde{x})/G$. □

2.4.7 CAT(0) cube complexes

For any cube complex X, the map $q : \mathcal{H} \to X$ from Definition 2.4.3 is a local isometry in the metric sense. If X is CAT(0) we have that, for any hyperplane H, $q : H \to X$ is an embedding with convex image (because in a CAT(0) space, local geodesics are geodesics and geodesics are unique), and so by identifying H with its image we can view H as a closed convex subspace of X. In particular, each cube of X will have at most one midcube belonging to H.

Proposition 2.4.28. *Hyperplanes do not self-osculate in a CAT(0) cube complex.*

Proof. Let e_1, e_2 be edges incident at a vertex x in a CAT(0) cube complex X, and suppose $H = H(e_1) = H(e_2)$. H is a closed convex subspace of X, so there is a well-defined closest point projection $p : X \to H$. If y_1 is the midpoint of e_1, then $d(x, y_1) = 1/2$, and no other point of H can be closer to x (the open ball of radius 1/2 about x is contained in the cubes incident at x and doesn't touch any hyperplane), so $p(x) = y_1$. But similarly if y_2 is the midpoint of e_2 then $p(x) = y_2$; hence $y_1 = y_2$ and $e_1 = e_2$. □

Again using CAT(0) geometry, one can show that pairs of hyperplanes do not inter-osculate in a CAT(0) cube complex. It then follows from the following proposition that CAT(0) cube complexes are directly special.

Proposition 2.4.29. *In a CAT(0) cube complex X, each hyperplane is two-sided and separates X into two connected components.*

Proof. Let H be a hyperplane and let γ be an edge loop. X is CAT(0), so in particular it is simply connected, thus γ can be homotoped down to a constant loop by a sequence of moves that add/remove backtracks or push a subpath of γ across a square in X as shown in Figure 2.3 (see [15, I.8A.4]). The parity of the number of times γ crosses H is preserved by these moves, so γ must originally have crossed H an even number of times.

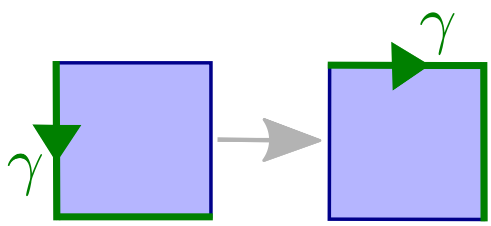

Figure 2.3: Pushing a subpath across a square.

It now follows from Remark 2.4.9 that H is two-sided. Since H does not self-intersect or self-osculate, we deduce that the carrier splits as a product $N(H) \cong H \times [-1, 1]$, with H mapping identically to $H \times \{0\}$.

Now consider the equivalence relation $\sim$ on X^0, where $x \sim y$ if x can be joined to y by an edge path that never crosses H. Write $[x]$ for the equivalence class of x. Let

e be an edge with endpoints x, y that is dual to H. Note that $[x] \neq [y]$ because loops containing e cross H an even number of times. Clearly any vertex in X is equivalent to one in $N(H)$, so we deduce from the above product splitting of $N(H)$ that $[x]$ and $[y]$ are the only classes.

Let $X(x) := \{z \in X - H \mid d(z, [x]) < d(z, [y])\}$ - this is the union of all cubes containing only vertices in $[x]$ plus, for every cube with a midcube in H, the open half cube that has a vertex in $[x]$. Define $X(y)$ similarly. By construction $X(x)$ and $X(y)$ are open, disjoint and their union is $X - H$. □

The closures of $X(x)$ and $X(y)$ above are called *half-spaces*. Each half-space is a convex subcomplex of the cubical subdivision $\dot{X}$, because a geodesic joining two points in the same half-space but crossing into the other half-space would create a geodesic between points in H that leaves H, contradicting the convexity of $H \subset X$.

Proposition 2.4.30. *Let x, y be vertices in a CAT(0) cube complex X. An edge path between x and y will be of minimal length if and only if it only crosses hyperplanes that separate x and y and it crosses each of these once.*

Proof. An edge path between x and y must cross every hyperplane that separates x and y, so any edge path from x to y that only crosses these hyperplanes, and crosses each of them once, is necessarily of minimal length.

For the converse implication, let γ be a shortest edge path from x to y and suppose for contradiction that it crosses some hyperplane twice. Say $e_1, ..., e_n$ are the edges of a subpath of γ with $H(e_1), ..., H(e_{n-1})$ distinct and $H = H(e_1) = H(e_n)$.

Suppose first that $e_2, ..., e_{n-1}$ all lie in $N(H)$. Then there must be a square containing e_2 that intersects H; and if e_1' is the edge of this square that crosses H and meets e_2 at the same vertex as e_1, then Proposition 2.4.28 implies that $e_1 = e_1'$. But then we can replace e_1, e_2 with edges f_1, f_2 going the other way round the square. Thus we have pushed the first crossing of H further along the subpath, whilst preserving the length of γ. Repeating this we can replace the subpath with $f_1, ..., f_{n-1}, e_n$, where $H = H(f_{n-1})$. But f_{n-1}, e_n must be a backtrack by Proposition 2.4.28, contradicting the minimality of γ.

Suppose now that e_i is the first edge in the subpath that leaves $N(H)$. Note that $N(H)$ is the $\frac{1}{2}$-neighbourhood of H, so it is convex. The first half of e_i is a geodesic η between $N(H)$ and $H(e_i)$; if $N(H)$ and $H(e_i)$ intersect then we can form a geodesic triangle between a point of the intersection and η, but the Alexandrov angles at either end of η are at least $\pi/2$, contradicting X being CAT(0) (Definition 2.3.2). Thus $N(H) \cap H(e_i) = \emptyset$ and $H(e_i)$ is disjoint from H. But then our subpath must recross $H(e_i)$ before it can get back to H, which is a final contradiction. □

We close this section by recalling a result about intersecting hyperplanes, which appears as [44, Lemma 13.13].

Proposition 2.4.31. *(Helly's Theorem for hyperplanes of CAT(0) cube complexes)* *If $H_1, ..., H_n$ are pairwise intersecting hyperplanes in a CAT(0) cube complex X, then there is a cube C with $C \cap H_1 \cap ... \cap H_n \neq \emptyset$.*

2.5 Subgroup separability

Separability of subgroups has been an idea of interest to both group theorists and topologists since the work of Wall, Mal'cev and Scott, and it is a key concept used throughout Chapters 5–7. In particular, we will explore its interactions with cubulated groups (Section 5.2) and graphs of groups (Sections 5.4.2 and 7.2).

Definition 2.5.1. (Separable subgroups, residual finiteness and subgroup separability) Let G be a group. A subgroup $K < G$ is *separable* if it is an intersection of finite-index subgroups in G. Equivalently, for every $g \in G - K$ there is a homomorphism $\rho : G \to \bar{G}$ to a finite group such that $\rho(g) \notin \rho(K)$. G is *residually finite* if the trivial subgroup is separable in G. G is *subgroup separable* if every finitely generated subgroup is separable.

Remark 2.5.2. If G is subgroup separable and $H < G$, then H is subgroup separable. If $\hat{G} < G$ is finite index and $\hat{G}$ is subgroup separable, then G is subgroup separable.

Proposition 2.5.3. *Let G be a subgroup separable group acting on a tree T. Suppose $U \subset VT$ is a finite set of vertices, and for each $u \in U$ let $\dot{G}_u$ be a finite index subgroup of G_u. Then G contains a finite index normal subgroup $\hat{G}$ such that $\hat{G}_u < \dot{G}_u$ for all $u \in U$. This also implies that $\hat{G}_{gu} = g\hat{G}_u g^{-1} < g\dot{G}_u g^{-1} < G_{gu}$ for all $u \in U$ and $g \in G$.*

Proof. We know that $\dot{G}_u < G$ is separable, so for any $g \in G - \dot{G}_u$ there is a homomorphism $\rho : G \to \bar{G}$ to a finite group such that $\rho(g) \notin \rho(\dot{G}_u)$. By taking products of these homomorphisms, we can produce a homomorphism $\rho : G \to \bar{G}$ to a finite group such that $\rho(g_i) \notin \rho(\dot{G}_u)$ for $\{g_i\}$ a set of representatives for the left cosets of $\dot{G}_u$ in G_u that are not equal to $\dot{G}_u$. This implies that $\ker \rho \cap G_u < \dot{G}_u$. The proposition then follows by taking products of these homomorphisms for all of the vertices in U, and setting $\hat{G}$ equal to the kernel. □

Definition 2.5.4. (Profinite topology) Let G be a group. The collection of all cosets of finite-index subgroups of G is a basis for the *profinite topology* on G. A subset $A \subset G$ is *separable* if it is closed in the profinite topology (note that this is equivalent to Definition 2.5.1 for subgroups).

2.6 Group Splittings

The modern study of group splittings is largely due to the work of Bass and Serre, and it has become a central concept in geometric group theory, with further influences from 3-manifold theory and the study of quasi-isometries. Group splittings play an important role in Chapters 6 and 7, although they also appear in other chapters.

2.6.1 Graphs of groups and spaces

The basic definitions for graphs of groups and graphs of spaces are as follows. We refer to [5, 80, 81] for further background.

Definition 2.6.1. (Graphs of groups)
A *graph of groups* is a finite graph Γ together with the following data:

(1) a *vertex group* G_v for each vertex $v \in V\Gamma$,

(2) an *edge group* G_e for each edge $e \in E\Gamma$ such that $G_{\bar{e}} \cong G_e$,

(3) an injective homomorphism $\zeta_e : G_e \to G_{\tau(e)}$ for each $e \in E\Gamma$.

Given a vertex group G_v, we refer to the subgroups $\zeta_e(G_e)$ for $e \in \mathrm{lk}(v)$ as its *incident edge groups*.

The *fundamental group* is defined by the following presentation, where $\Delta \subset \Gamma$ is a spanning tree:

$$G = \left\langle \mathop{*}_{v \in V\Gamma} G_v \mathop{*}_{e \in E\Gamma} \langle e \rangle \;\middle|\; \begin{matrix} \bar{e} = e^{-1},\, e \in E\Gamma;\, e = 1,\, e \in E\Delta; \\ e\zeta_e(g)e^{-1} = \zeta_{\bar{e}}(g),\, e \in E\Gamma, g \in G_e \end{matrix} \right\rangle$$

G is independent of the choice of Δ. This presentation induces natural homomorphisms $G_v, G_e \to G$, which are injective, so we can view the vertex and edge groups as subgroups of G. We write (G, Γ) to refer to the graph of groups together with the fundamental group G. A given group G could be isomorphic to the fundamental group of many different graphs of groups, so we refer to (G, Γ) as a *graph of groups decomposition* or a *splitting* for the group G.

Definition 2.6.2. (Graphs of spaces)
Given a graph of groups decomposition (G, Γ), an *associated graph of spaces* (X, Γ) is a topological space X together with the graph Γ and the following data:

(1) a *vertex space* X_v for each $v \in V\Gamma$ such that $\pi_1(X_v) \cong G_v$,

(2) an *edge space* X_e for each $e \in E\Gamma$ such that $\pi_1(X_e) \cong G_e$ and $X_{\bar{e}} \cong X_e$,

(3) a map $\phi_e : X_e \to X_{\tau(v)}$ for each $e \in E\Gamma$ that induces the homomorphism ζ_e on the fundamental groups.

Then we construct X as the following quotient space:

$$X = \bigsqcup_{v \in V\Gamma} X_v \bigsqcup_{e \in E\Gamma} X_e \times [0,1] / \sim$$

where $\sim$ is the relation that identifies $X_e \times \{0\}$ with $X_{\bar{e}} \times \{0\}$ via the identification of X_e with $X_{\bar{e}}$, and the point $(x, 1) \in X_e \times [0,1]$ with $\phi_e(x)$. One can verify that the fundamental group of X as a topological space, $\pi_1 X$, is isomorphic to G.

Definition 2.6.3. (Bass-Serre trees)
Given a graph of groups decomposition (G, Γ) and an associated graph of spaces (X, Γ), the universal cover $\widetilde{X} \to X$ has vertex and edge spaces corresponding to the components of the preimages of the vertex and edge spaces of X, and these are joined together in a tree formation. The underlying tree T is called the *Bass-Serre tree*. $G \cong \pi_1 X$ acts on $\widetilde{X}$ by deck transformations, so also acts on T. For vertices and edges v, e in T, we denote the G-stabilisers by G_v and G_e. This notation is justified by the fact that the vertex and edge stabilisers correspond precisely to the conjugates of the vertex and edge groups from Definition 2.6.1. Similarly, for a vertex stabiliser G_v, we refer to the subgroups G_e for $e \in \mathrm{lk}(v)$ as its *incident edge groups*.

This construction can also be run in reverse, given an action of a group G on a tree T without edge inversions we obtain an associated graph of groups decomposition (G, Γ), with the underlying graph Γ given by the quotient T/G.

2.6.2 JSJ decompositions of finitely presented groups

Inspired by the JSJ decomposition of 3-manifolds there has been extensive work generalizing such results to finitely generated groups. We refer to Guirardel and Levitt [38] for the most modern and comprehensive overview of this field. Our definitions in this section come from [38], but specialised to splittings of one-ended groups over two-ended, and frequently just infinite cyclic, subgroups. The *number of ends* of a geodesic space X is the supremum of the number of unbounded components of $X - K$, where K ranges over all compact subsets, and the *number of ends* of a group is the number of ends of a Cayley graph (which is the same for all Cayley graphs).

Definition 2.6.4. (JSJ trees)
Let G be a one-ended group. Consider trees with a *minimal* G-action (i.e. with no G-invariant subtree), *without inversion* (i.e. if $g \in G$ stabilises an edge e then it fixes e pointwise), and such that all edge stabilisers are two-ended. Such a tree is *universally elliptic* if its edge stabilisers are elliptic in every such tree. A tree T is a *JSJ tree* for G if it is universally elliptic and its vertex stabilisers are elliptic in every other universally elliptic tree. A vertex group G_v, where v is a vertex in a JSJ tree, is *rigid* if it is elliptic in every splitting over two-ended edge groups, and *flexible* otherwise.

The idea of a JSJ decomposition over two-ended subgroups is that we wish to find a "maximal" splitting over two-ended subgroups and claim that it is essentially canonical. Certain vertex groups in the decomposition, the *rigid* vertex groups, are unambiguously going to belong to any maximal splitting as they cannot be split any further. There is potential for ambiguity when there may be many ways to split a vertex group over two-ended groups. Consider, for example, a vertex group coming from a hyperbolic surface with boundary, such that the incident edge groups correspond to the boundary components. The multitude of pants decompositions of the surface give many ways to split the vertex group relative to its boundary subgroups; but the edge groups in one such splitting will not be elliptic in other such splittings, hence splitting this vertex group would not give a *universally elliptic* tree. This would be an example of a *flexible* vertex group. Definition 2.6.4 thus gives us a splitting that is canonical in the sense that the collection of non-two-ended vertex stabilisers is the same for every JSJ tree (easy exercise). The great success of JSJ theory is that hyperbolic surface vertex groups as described above are in some sense the only examples of flexible vertex groups. We make this precise with the following definition.

Definition 2.6.5. (Quadratically hanging vertex groups)
Let G_v be a vertex group for a splitting of a one-ended group G over two-ended subgroups. We say that G_v is *quadratically hanging* (QH) if it surjects to a compact hyperbolic 2-orbifold group $\pi_1\Sigma$, with finite kernel, such that images of incident edge groups in $\pi_1\Sigma$ are contained in boundary subgroups. If G is torsion-free, this reduces to saying that G_v is the fundamental group of a compact hyperbolic surface with boundary, such that incident edge groups are contained in boundary subgroups.

For Theorem 1.3.3 we will be interested in the following family of groups (plus the assumption of hyperbolicity relative to virtually abelian subgroups). This family lies inside the family $\mathscr{C}$ from the introduction and satisfies the conclusions of Theorem 1.3.2.

Definition 2.6.6. Let $\mathscr{C}^\bullet$ denote the family of one-ended, subgroup separable groups that have JSJ decompositions over two-ended subgroups containing only virtually free vertex groups and no QH vertex groups. Let $\mathscr{C}^\bullet_{tf}$ denote the subfamily of torsion-free groups.

Remark 2.6.7. For a one-ended group $G \in \mathscr{C}$, it follows from [38, Theorem 6.2] that flexible vertex groups are always QH, thus vertex groups of $G \in \mathscr{C}^\bullet$ are always rigid. In fact there are only finitely many possibilities for rigid QH vertex groups [38, 5.12, 5.16 and 5.18], and if G is torsion-free then the pair of pants is the only possibility - so the assumption that there are no QH vertex groups is very close to assuming that all vertex groups are rigid.

Lemma 2.6.8. *The definition of $\mathscr{C}^\bullet$ is independent of the choice of JSJ decomposition: if $G \in \mathscr{C}^\bullet$, then the vertex stabilisers of any JSJ tree T for G are virtually free and not QH.*

Proof. The vertex stabilisers not being QH follows from [38, Proposition 5]. As for the virtual freeness, $G \in \mathscr{C}$, so it has some splitting over two-ended subgroups by acting on a tree T_0 with virtually free vertex stabilisers. By Remark 2.6.7, the vertex stabilisers of the JSJ tree T are rigid, hence they are elliptic in T_0. Then each vertex stabiliser of T is contained in a vertex stabiliser of T_0, so the former must be virtually free. □

2.7 Quasi-isometries

Quasi-isometries are another central concept in geometric group theory, and again it was Gromov who laid the foundations for much of the theory. Of particular interest is quasi-isometric rigidity, which we discuss in Section 1.3. Quasi-isometries play a crucial role in Chapter 7.

2.7.1 Quasi-isometries of metric spaces and groups

Definition 2.7.1. (Quasi-isometries)
A (Q, ϵ)-*quasi-isometry* between two metric spaces $f : (X, d_X) \to (Y, d_Y)$ is a function that satisfies the following two conditions:

(1) For all $x, x' \in X$ the following inequality is satisfied:

$$\frac{1}{Q} d_X(x, x') - \epsilon \leq d_Y(f(x), f(x')) \leq Q d_X(x, x') + \epsilon.$$

(2) For all $y \in Y$, there exists $x \in X$ such that $d_Y(f(x), y) < \epsilon$.

Definition 2.7.2. (Hausdorff equivalence and quasi-isometry groups)
Two quasi-isometries $f, h : X \to Y$ are said to be *Hausdorff equivalent* if there exists some $0 < B < \infty$ such that $d_Y(f(x), h(x)) < B$ for all $x \in X$ - we will write this as $f \approx h$, and denote the equivalence class of f by $[f]$.

For any quasi-isometry $f : X \to Y$ there exists a quasi-isometry $g : Y \to X$, called the *quasi-inverse*, with $gf \approx \mathrm{id}_X$ and $fg \approx \mathrm{id}_Y$. If $h : Y \to Z$ is another quasi-isometry, then the composition $hf : X \to Z$ is also a quasi-isometry, and the class $[hf]$ only depends on the classes $[f]$ and $[h]$. In particular, quasi-isometry between metric spaces is an equivalence relation on every set of metric spaces.

The set of quasi-isometries of a space X up to Hausdorff equivalence forms a group $\mathcal{QI}(X)$ called the *quasi-isometry group* of X. Note that a quasi-isometry $\psi : X \to Y$ induces an isomorphism of the quasi-isometry groups given by $[f] \mapsto [\psi f \psi^{-1}]$ (where ψ^{-1} is any quasi-inverse to ψ).

Definition 2.7.3. (Hausdorff distance)
If A and B are subsets of a metric space (X, d), we define the *Hausdorff distance* between A and B to be

$$d_H(A, B) := \sup\{d(a, B), d(A, b) \mid a \in A\, b \in B\}.$$

We will also write $A \sim B$ for $d_H(A, B) < \infty$ and $A \sim_D B$ for $d_H(A, B) \leq D$.

Definition 2.7.4. (Quasi-isometries of groups)
If G is a finitely generated group with finite generating set S, then the induced path metric on the associated Cayley graph $\mathrm{Cay}(G, S)$ is the *word metric* d_S on G. Different finite generating sets S, S' give distinct word metrics, but the identity map $\mathrm{id}_G : (G, d_S) \to (G, d_{S'})$ is a quasi-isometry. When working with quasi-isometries of G we will implicitly pick a finite generating set and work with the corresponding word metric. We will use $\mathscr{G} := \mathcal{QI}(G)$ to denote the quasi-isometry group of G.

2.7.2 Trees of cylinders

In general there is no entirely canonical JSJ decomposition (see Section 2.6.2). An alternative is the tree of cylinders, which is canonical, and has the added advantage of being preserved by quasi-isometries.

Definition 2.7.5. (Trees of cylinders)
Subgroups $A, A' < G$ are *commensurable* (in G) if $A \cap A'$ is a finite index subgroup of A and A'. Commensurability is an equivalence relation. Let G act on a tree T with two-ended edge stabilisers. We say that two edges $e_1, e_2 \in ET$ are *equivalent* if G_{e_1} and G_{e_2} are commensurable. The union of all edges in an equivalence class gives a subtree (that is to say a connected subcomplex of T), which is called a *cylinder*. The *tree of cylinders* T_c is the bipartite tree with vertex set $V_0T_c \sqcup V_1T_c$, where V_0T_c are the vertices of T which lie in at least two cylinders, and V_1T_c is the set of cylinders. The edges of T_c are of the form (v, Y) where v is a vertex in T that lies in the cylinder $Y \subset T$.

Notation 2.7.6. Let a finitely generated group G act minimally on a tree T without edge inversions, and let $\{v_1, ..., v_n\} \subset VT$, $\{e_1, ..., e_m\} \subset ET$ be orbit representatives for the vertices and edges, with $\{e_1, ..., e_m\}$ closed under edge inversions. For $v \in VT$ in the same orbit as v_i, let $G(v)$ denote the left coset of G_{v_i} consisting of $g \in G$ with $g(v_i) = v$. Similarly, define cosets $G(e)$ for $e \in ET$, noting that $G(e) = G(\bar{e})$. In the rest of this section we always pick orbit representatives for vertices and edges so that we can define the cosets $G(v)$ and $G(e)$; the choice of such representatives will not affect the results of this section, only the size of the constants contained within them.

Remark 2.7.7. We always have $G(v) \sim G_v$ and $G(e) \sim G_e$, just not with a uniform constant (where $\sim$ denotes finite Hausdorff distance, as in Definition 2.7.3).

The following theorem of Papasoglu [72, Theorem 7.1] says that quasi-isometries coarsely preserve vertex stabilisers of JSJ trees.

Theorem 2.7.8. *(Papasoglu)*
Let $\psi : G \to G'$ be a quasi-isometry of finitely presented one-ended groups, with JSJ trees T and T'. Then there exists a constant $D > 0$, such that for each $v \in VT$ there exists $v' \in VT'$ with $\psi(G(v)) \sim_D G'(v')$, and for each $e \in ET$ there exists $e' \in ET'$ with $\psi(G(e)) \sim_D G'(e')$. Moreover, the type of vertex stabiliser is preserved, so G_v is QH if and only if $G'_{v'}$ is QH.

Note that one needs to use cosets rather than the vertex stabilisers themselves in order to obtain the uniform constant D (this was not stated correctly in [72]). In the following theorem we deduce from Papasoglu's theorem that quasi-isometries also preserve the tree of cylinders decomposition - in fact this time we can say something even stronger than in Papasoglu's theorem, namely that a quasi-isometry induces an isomorphism between trees of cylinders. In this sense we can think of trees of cylinders as being canonical. This theorem appears as [23, Theorem 2.8], but with the same mistake regarding vertex stabilisers versus cosets mentioned above. Margolis [62, Theorem 2.9] avoids this confusion in his statement; since we state the theorem in slightly different terms, we include a brief explanation of how we deduce our version from Margolis' version.

Theorem 2.7.9. *Let $\psi : G \to G'$ be a quasi-isometry of finitely presented one-ended groups, with JSJ trees T and T', and let T_c and T'_c be the corresponding trees of cylinders. Then there is a unique isomorphism $\hat{\psi} : T_c \to T'_c$ such that:*

(1) $\hat{\psi}(V_0T_c) = V_0T'_c$ and $\hat{\psi}(V_1T_c) = V_1T'_c$.

(2) There is a constant $K > 0$, such that $\psi(G(v)) \sim_K G'(\hat{\psi}(v))$ for $v \in VT_c$, and $\psi(G(e)) \sim_K G'(\hat{\psi}(e))$ for $e \in ET_c$.

(3) $\psi(G_v) \sim G'_{\hat{\psi}(v)}$ for $v \in VT_c$ and $\psi(G_e) \sim G'_{\hat{\psi}(e)}$ for $e \in ET_c$. Moreover, the restrictions $\psi : G_v \to G'_{\hat{\psi}(v)}$ and $\psi : G_e \to G'_{\hat{\psi}(e)}$ are quasi-isometries with respect to the intrinsic metrics of the vertex and edge stabilisers (these restrictions are well-defined up to Hausdorff-equivalence by (2) and Remark 2.7.7).

Proof. [62, Theorem 2.9] gives a unique isomorphism $\hat{\psi} : T_c \to T'_c$ that satisfies (1) and satisfies (2) for the vertex spaces in the trees of spaces corresponding to (G, T_c) and (G', T'_c). But [62, Proposition 2.3] allows us to transfer between the edge and vertex spaces in these trees of spaces and the corresponding cosets in G and G' as described in Notation 2.7.6. Hence we get a constant $K > 0$, such that $\psi(G(v)) \sim_K G'(\hat{\psi}(v))$ for $v \in VT_c$. Because of the tree structure of the trees of spaces, each edge space is $\sim$-equivalent to the intersection of the R-neighbourhoods of the adjacent vertex spaces for all $R \geq 1$; so we deduce that ψ coarsely maps edge spaces to edge spaces, and that $\psi(G(e)) \sim_K G'(\hat{\psi}(e))$ for $e \in ET_c$ (possibly increasing K). Lastly, (3) follows from (2) and Remark 2.7.7, and the moreover part follows from [33, Lemma 2.1]. □

Corollary 2.7.10. *If G is a finitely presented one-ended group with JSJ tree T, then the group $\mathscr{G}$ of quasi-isometries of G acts on the tree of cylinders T_c by*

$$\begin{aligned} \mathscr{G} &\to \mathrm{Aut}(T_c) \\ [f] &\mapsto \hat{f}. \end{aligned}$$

Proof. Properties (1)–(3) of Theorem 2.7.9 remain true if we perturb f by a bounded amount, hence $\hat{f}$ is determined by the Hausdorff class $[f]$. If $[f_1], [f_2] \in \mathscr{G}$, then $\hat{f}_1 \circ \hat{f}_2$ clearly satisfies properties (1)–(3) of Theorem 2.7.9 with respect to the quasi-isometry $f_1 \circ f_2$, therefore $\widehat{f_1 \circ f_2} = \hat{f}_1 \circ \hat{f}_2$ by the uniqueness in Theorem 2.7.9. □

Remark 2.7.11. G acts on itself by left translations, which induces a homomorphism $G \to \mathscr{G}$, and restricting the action of $\mathscr{G}$ on T_c to G on T_c recovers the original action of G on T_c.

In the case that G acts on T with hyperbolic vertex stabilisers, we get that the edge stabilisers of T_c are two-ended. This holds for example if $G \in \mathscr{C}^\bullet$ and T is a JSJ tree.

Lemma 2.7.12. *Let G act on a tree T with two-ended edge stabilisers and hyperbolic vertex stabilisers. Let H be the stabiliser of an edge $(v, Y) \in ET_c$. Then for every $e \in EY \subset ET$ incident at v, G_e is a finite index subgroup of H - so in particular H is two-ended. H is also equal to its normaliser in G_v.*

Proof. H consists of those elements $h \in G_v$ such that $h(e) \in EY$, or equivalently that G_e is commensurable to hG_eh^{-1} in G, or equivalently that $G_e \sim hG_e$. Since G_e is two-ended, the cosets hG_e are all quasi-geodesics with uniform constants, so they must all be at uniform Hausdorff distance from G_e by the Morse Lemma. This implies that there are only finitely many cosets hG_e with $h \in H$, so G_e has finite index in H. If $g \in G_v$ normalises H, then G_e and gG_eg^{-1} will both be finite index subgroups of H, so they will be commensurable and $g \in H$; hence H is equal to its normaliser in G_v. □

2.8 Hyperbolicity and relative hyperbolicity

Negative curvature of manifolds has long been an important concept in differential geometry. It was one of the great insights of Gromov to generalise this notion to arbitrary geodesic metric spaces, and to groups, with the following definition [36].

Definition 2.8.1. (Hyperbolic spaces and groups)
Let $\delta \geq 0$. We say that a geodesic triangle in a geodesic space is *δ-slim* if each side is contained in the δ-neighbourhood of the other two sides. We say that a geodesic space is *δ-hyperbolic* if every geodesic triangle is δ-slim, and *hyperbolic* if it is δ-hyperbolic for some δ. A group is *hyperbolic* if it is hyperbolic with respect to the word metric of some finite generating set.

Remark 2.8.2. Hyperbolicity is preserved by quasi-isometries [15, III.H.1.9], so a hyperbolic group is hyperbolic with respect to the word metric of any finite generating set.

Remark 2.8.3. For a geodesic n-gon in a δ-hyperbolic space X, each side is within the $(n-2)\delta$ neighbourhood of the union of the other sides (proof: subdivide the n-gon into triangles).

Hyperbolic groups have been well studied in the literature. Among many other properties, they have solvable word and conjugacy problems, and possess a well-behaved notion of boundary [15, Chapter III.H].

In this thesis we will also be interested in a generalisation of hyperbolicity called relative hyperbolicity. The following definition is due to Bowditch [13], but there are many equivalent definitions - see for example [29, 68].

Definition 2.8.4. (Relatively hyperbolic groups)
A group G is *hyperbolic relative to subgroups* $\{P_1, ..., P_n\}$ if G acts on a connected graph X with the following properties:

(1) X is hyperbolic, and each edge of X is contained in only finitely many circuits of length k for any given integer k.

(2) There are finitely many G-orbits of edges, and the edge stabilisers are finite.

(3) The conjugates of the P_i are precisely the infinite vertex stabilisers of X, and the P_i are pairwise non-conjugate.

(4) Each P_i is finitely generated.

The P_i are called *peripheral subgroups.* Note that G is hyperbolic in the case $n = 0$.

It follows easily from this definition that the collection $\{P_1, ..., P_n\}$ is almost malnormal, defined as follows.

Definition 2.8.5. (Almost malnormal)
A collection of subgroups $\{P_1, ..., P_n\}$ in a group G is *almost malnormal* if $gP_ig^{-1} \cap P_j$ is finite whenever $g \notin P_i$ or $i \neq j$.

If G itself is hyperbolic, then Bowditch proves that almost malnormality is actually sufficient for relative hyperbolicity [13, Theorem 7.11] (peripheral subgroups are always quasiconvex [29, Lemma 4.15]).

Theorem 2.8.6. *(Bowditch)*
Let G be a hyperbolic group. Then G is hyperbolic relative to a collection $\{P_1, ..., P_n\}$ of infinite subgroups if and only if $\{P_1, ..., P_n\}$ is an almost malnormal collection of quasi-convex subgroups.

There are theorems that allow one to deduce that a graph of groups is hyperbolic given certain conditions on the vertex and edge groups. For example, the following result is (a special case of) the Bestvina–Feighn Combination Theorem [12, first corollary in §7], albeit stated in slightly more modern language here (see also [54]).

Theorem 2.8.7. *(Bestvina–Feighn Combination Theorem)*
A graph of groups has hyperbolic fundamental group if it satisfies the following:

- *The vertex groups are hyperbolic.*
- *The inclusions of edge groups into vertex groups are quasi-isometric embeddings.*
- *For each vertex group, the collection of incident edge groups is an almost malnormal family.*

There is also a combination theorem for relatively hyperbolic groups due to Dahmani [26]. We will be interested in the following application, which concerns the tree of cylinders from Section 2.7.2.

Proposition 2.8.8. *Let G be a group acting on a tree T with two-ended edge stabilisers and hyperbolic vertex stabilisers. Then G is hyperbolic relative to (conjugacy representatives of) its cylinder stabilisers.*

Proof. Let T_c be the tree of cylinders corresponding to T, and let (G, Γ) be the quotient graph of groups for the action of G on T_c. The partition $VT_c = V_0T_c \sqcup V_1T_c$ induces a partition $V\Gamma = V_0\Gamma \sqcup V_1\Gamma$. We wish to show that G is hyperbolic relative to its vertex groups G_v for $v \in V_1\Gamma$ - which we call its cylinder vertex groups. For the original tree T, two stabilisers of edges in different cylinders will have finite intersection, so for $u \in V_0\Gamma$ Lemma 2.7.12 implies that different G_u-conjugates of edge groups incident at G_u also have finite intersection. Then by Theorem 2.8.6 and Lemma 2.7.12, G_u is hyperbolic relative to its incident edge groups. Next, for each $u \in V_0\Gamma$, let (G^u, Γ^u) be the graph of groups obtained by amalgamating G_u with its neighbouring cylinder vertex groups in (G, Γ). By [26, Theorem 0.1(2)], G^u is hyperbolic relative to its cylinder vertex groups in (G^u, Γ^u). We can then join together the graphs of groups (G^u, Γ^u) via a sequence of amalgamations and HNN extensions to recover the graph of groups (G, Γ), and this will be hyperbolic relative to its cylinder vertex groups by [26, Theorem 0.1(3)+(3')]. □

Another useful result is the following, due to Drutu–Sapir [29, Corollary 1.14]. In particular, given the conclusion of Proposition 2.8.8, we can remove any virtually cyclic cylinder stabilisers from the family of peripheral subgroups.

Theorem 2.8.9. *(Drutu–Sapir)*
If G is hyperbolic relative to $\{P_1, ..., P_n\}$ and each P_i is hyperbolic relative to $\{P_i^1, ..., P_i^{n_i}\}$, then G is hyperbolic relative to $\{P_i^j \mid 1 \leq i \leq n, 1 \leq j \leq n_i\}$.

Chapter 3

Two generalisations of Leighton's Theorem

In this chapter we present a new groupoid proof of Leighton's Theorem (Theorem 1.1.1), and use this as a model to prove two generalisations. The first works with graphs of spaces, and more general objects, and considers covering maps that restrict to isomorphisms between vertex spaces (Theorem 1.1.6). The second, which we deduce from the first, restricts how balls of a given size in the universal cover can map down to the two finite graphs when factoring through the common finite cover (Theorem 1.1.5).

3.1 Original Leighton's Theorem

We now present our new groupoid proof of Leighton's Theorem.

Theorem 1.1.1. *(Leighton's Theorem)*
Let G_1 and G_2 be finite connected graphs with a common universal cover. Then they have a common finite cover.

Proof. We will define a finite groupoid $\mathcal{S}$, consisting of maps between links in G_1 and G_2, and we'll use this to label the vertices of a finite graph G. We'll then consider an action of $\mathcal{S}$ on the edges of G_1 and G_2, and use this to connect up the vertices of G with edges in a way that makes it a cover of G_1 and G_2.

We divide the proof into four steps. We define $\mathcal{S}$ in the first step, then set up the action of $\mathcal{S}$ in the second step. In the third step we build the graph G, using the action to connect up the vertices, and finally we construct the covering maps to G_1 and G_2 in the fourth step.

Step 1: We will have $\mathrm{Ob}(\mathcal{S}) = VG_1 \sqcup VG_2$, and each $s \in \mathcal{S}$ will be a bijection

$$s : \mathrm{lk}(i(s)) \to \mathrm{lk}(t(s)). \tag{3.1.1}$$

Composition of groupoid elements is just composition of bijections. The set of all such bijections forms a groupoid. Let $p_1 : T \to G_1$ and $p_2 : T \to G_2$ be coverings of G_1 and G_2 by a tree T. We define $\mathcal{S}$ to be the subgroupoid consisting of bijections

$$s := p_i \circ g \circ (p_j|_{\mathrm{lk}(x)})^{-1} : \mathrm{lk}(p_j(x)) \to \mathrm{lk}(p_i g(x)), \tag{3.1.2}$$

for $i, j \in \{1, 2\}$, $x \in VT$ and $g \in \mathrm{Aut}(T)$. Think of s as a map that lifts a link in G_1 to a link in T, maps it across to another link in T, and then projects it down to a link in G_2. $\mathcal{S}$ is closed under composition because if $x, y \in VT$ with $p_j(x) = p_j(y)$ then

$$(p_j|_{\mathrm{lk}(x)})^{-1} \circ p_j|_{\mathrm{lk}(y)} = g|_{\mathrm{lk}(y)},$$

for some $g \in \mathrm{Aut}(T)$ a deck transformation of $p_j : T \to G_j$. $\mathcal{S}$ also contains inverses because $\mathrm{Aut}(T)$ contains inverses.

Step 2: There is a natural action of $\mathcal{S}$ on $EG_1 \sqcup EG_2$ defined by $s \cdot e = s(e)$ for $\tau(e) = i(s)$ - the associated map $\varepsilon : EG_1 \sqcup EG_2 \to VG_1 \sqcup VG_2$ is given by $\varepsilon(e) := \tau(e)$. The following claim will be vital in the next step of the proof.

<u>Claim:</u> $\mathcal{S} \cdot \bar{e} = \overline{\mathcal{S} \cdot e}$

<u>Proof:</u> It suffices to show the inclusion $\overline{\mathcal{S} \cdot e} \subset \mathcal{S} \cdot \bar{e}$, because replacing e by $\bar{e}$ gives $\overline{\mathcal{S} \cdot \bar{e}} \subset \mathcal{S} \cdot e$, which implies $\mathcal{S} \cdot \bar{e} \subset \overline{\mathcal{S} \cdot e}$. So let s be as in (3.1.2) and take $\tau(e) = i(s) = p_j(x)$. We will show there is $s' \in \mathcal{S}$ with $s' \cdot \bar{e} = \overline{s \cdot e}$. Let $\hat{e} \in \mathrm{lk}(x)$ with $p_j(\hat{e}) = e$, and consider $y := \iota(\hat{e})$. Put $f := p_i g(\hat{e}) = s \cdot e$, and define $s' \in \mathcal{S}$ by

$$s' = p_i \circ g \circ (p_j|\mathrm{lk}(y))^{-1}.$$

We have $i(s') = p_j(y) = \iota(e)$, so $s' \cdot \bar{e}$ is defined; and $s' \cdot \bar{e} = p_i g(\bar{\hat{e}}) = \bar{f} = \overline{s \cdot e}$. ■

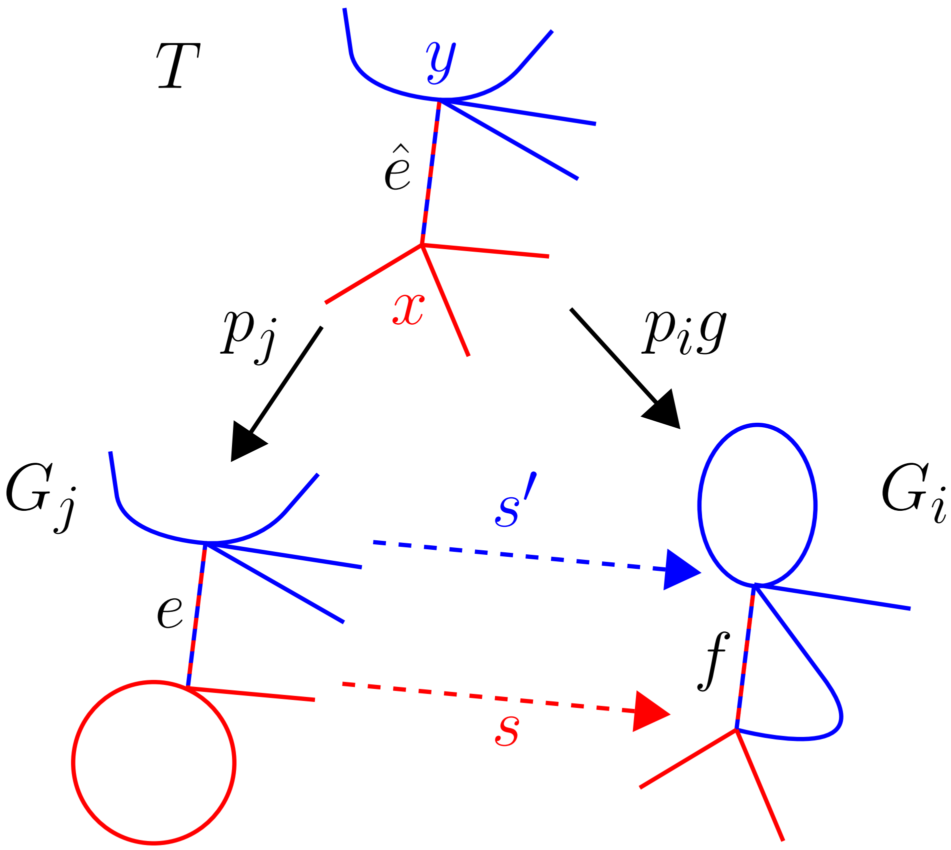

Figure 3.1: Diagram of s and s'.

Step 3: We are now ready to define a finite graph G that covers G_1 and G_2. Define the vertex set by

$$VG := \left\{ (s,l) \mid s \in \mathcal{S}(x,y),\ x \in VG_1,\ y \in VG_2,\ 1 \leq l \leq \frac{N}{|\mathcal{S}(x,-)|} \right\},$$

and the edge set by

$$EG := \left\{ (e,f,k) \mid f \in \mathcal{S} \cdot e,\ e \in EG_1,\ f \in EG_2,\ 1 \leq k \leq \frac{N}{|\mathcal{S} \cdot e|} \right\},$$

where N is a fixed positive integer that is a common multiple of all the integers $|\mathcal{S}(x,-)|$ and $|\mathcal{S} \cdot e|$. We define edge inversion in G by $\overline{(e,f,k)} := (\bar{e}, \bar{f}, k)$, and this is well-defined by step 2.

To define G, it remains to define the map $\tau : EG \to VG$. Fix $e \in EG_1$ and $f \in EG_2$ such that $f \in \mathcal{S} \cdot e$, and say $\tau(e) = x$. For reasons that will become clear in the next step, each edge $(e,f,k) \in EG$ must satisfy

$$\tau(e,f,k) = (s,l), \qquad \text{for some } s \in \mathcal{S}(x,-) \text{ such that } s \cdot e = f. \tag{3.1.3}$$

We define such a map τ by choosing an arbitrary matching between such vertices (s,l) and integers $1 \leq k \leq N/|\mathcal{S} \cdot e|$ - to verify that this is valid, we must check that we have equal numbers of each. Indeed, Lemma 2.2.4 tells us that

$$|\{s \in \mathcal{S}(x,-) \mid \mathcal{S} \cdot e = f\}| = |\mathrm{Stab}_{\mathcal{S}}(e)|,$$

and so

$$\begin{aligned}\left|\left\{(s,l) \mid s \in \mathcal{S}(x,-),\ s \cdot e = f,\ 1 \leq l \leq \frac{N}{|\mathcal{S}(x,-)|}\right\}\right| &= |\mathrm{Stab}_{\mathcal{S}}(e)| \frac{N}{|\mathcal{S}(x,-)|} \\ &= \frac{N}{|\mathcal{S} \cdot e|}\end{aligned}$$

again by Lemma 2.2.4.

Step 4: In this last step we define the covering maps from G down to G_1 and G_2.

$$G_1 \xleftarrow{\ \mu_1\ } G \xrightarrow{\ \mu_2\ } G_2$$

These are defined on the edge and vertex sets by

$$\begin{aligned}\mu_1(s,l) &:= i(s), & \mu_1(e,f,k) &:= e, \\ \mu_2(s,l) &:= t(s), & \mu_2(e,f,k) &:= f.\end{aligned}$$

These maps clearly preserve edge inversion, and (3.1.3) ensures that they are well-defined graph morphisms, because then $\tau(e,f,k) = (s,l)$ with $i(s) = x$ implies

$$\begin{aligned}\tau\mu_1(e,f,k) &= \tau(e) & \tau\mu_2(e,f,k) &= \tau(f) \\ &= x & &= \tau(s \cdot e) \\ &= i(s) & &= t(s) \\ &= \mu_1(s,l) & &= \mu_2(s,l) \\ &= \mu_1\tau(e,f,k), & &= \mu_2\tau(e,f,k).\end{aligned}$$

By construction, the link of a vertex (s,l) in G, with $s \in \mathcal{S}(x,y)$, takes the form

$$\mathrm{lk}(s,l) = \{(e, s \cdot e, k_e) \mid e \in \mathrm{lk}(x)\}$$

where each k_e is some integer associated to e. Now $\mu_1(e, s \cdot e, k_e) = e$ and $\mu_2(e, s \cdot e, k_e) = s \cdot e$, so μ_1 and μ_2 induce bijections from $\mathrm{lk}(s,l)$ to $\mathrm{lk}(x)$ and $\mathrm{lk}(y)$ respectively. We conclude that μ_1 and μ_2 are coverings, which completes the proof of the theorem.

□

Remark 3.1.1. The finite cover G constructed in the proof above may not be connected, but of course we can obtain a connected cover by choosing a component of G.

3.2 Graph of objects version

In this section we generalise Leighton's Theorem to graphs of objects, which can be thought of as a "rigid" type of graph of spaces.

3.2.1 Definitions

One limitation to generalising Leighton's Theorem to graphs of spaces is illustrated by the following example.

Example 3.2.1. The Baumslag Solitar groups BS(1,3) and BS(2,2) both arise as fundamental groups of graphs of spaces with a single circular vertex space and a single circular edge space. For BS(1,3) the maps from edge space to vertex space will be coverings of degree 1 and 3, whereas for BS(2,2) they will both be coverings of degree 2. In both cases the sum of these degrees equals 4, and so both graphs of spaces are covered by a 4-regular tree of spaces in which all edge and vertex spaces are copies of the real line and all edge maps are homeomorphisms. However, there is no common finite cover for these graphs of spaces because BS(1,3) and BS(2,2) are not commensurable (BS(1,3) is solvable whereas BS(2,2) contains a non-abelian free subgroup).

This example would no longer work if we endowed the vertex and edge spaces with metrics and required the edge maps and coverings to respect these metrics, as then the tree of spaces would have two very different metrics induced by BS(1,3) and BS(2,2). This shows that the category of spaces we use to define our graphs of spaces matters. However, there are other more restrictive categories of spaces where Leighton's Theorem still fails. For example, in the category of graphs, Wise constructs a finite graph of spaces with non-residually finite fundamental group that is covered by a 4-regular tree of spaces in which all edge and vertex spaces are 6-regular trees and edge maps are isomorphisms [100]. Such a tree of spaces also covers a finite graph of spaces whose fundamental group is a product of free groups, hence not commensurable to the group that Wise constructed.

It is clear then that one must impose some strong conditions in order to get a version of Leighton's Theorem for graphs of spaces. We do this by working with graphs of spaces with respect to a given category of spaces, and by restricting to coverings between graphs of spaces that induce isomorphisms between vertex spaces rather than just coverings. To emphasise that this is not the general setting we call them *graphs of objects* instead of graphs of spaces. The full definitions of graphs of objects and their coverings are given below. Note that these definitions actually work with any category, not just categories of spaces - see Example 3.2.5 for some suggestions of categories that could be used.

Definition 3.2.2. (Graph of Objects)
Let $\mathcal{C}$ be a category and $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \mathrm{Hom}(\mathcal{C})$ such that $(\mathrm{Ob}(\mathcal{C}), \mathcal{M}_1)$ and $(\mathrm{Ob}(\mathcal{C}), \mathcal{M}_2)$

are subcategories. Suppose that $\mathcal{M}_2\mathcal{M}_1\mathcal{M}_2 = \mathcal{M}_2$ (if $\mathcal{M}, \mathcal{N}$ are classes of morphisms within some category, we will always write $\mathcal{M}\mathcal{N}$ to denote the class of all morphisms of the form mn with $m \in \mathcal{M}$ and $n \in \mathcal{N}$). Further suppose that $(\mathrm{Ob}(\mathcal{C}), \mathcal{M}_1)$ is a groupoid. The next two definitions will mirror those for graphs of spaces.

A *graph of objects* X with respect to $(\mathcal{C}, \mathcal{M}_1, \mathcal{M}_2)$ consists of

- a graph $G = G_X$,
- objects $X_v \in \mathrm{Ob}(\mathcal{C})$ for $v \in VG$, called *vertex objects*,
- objects $X_e \in \mathrm{Ob}(\mathcal{C})$ for $e \in EG$, called *edge objects*, with $X_e = X_{\bar{e}}$,
- and morphisms in $\mathcal{M}_2$

$$\phi_e : X_e \to X_{\tau(e)},$$

  for $e \in EG$, called *edge morphisms*.

We say that a graph of objects X is *finite* if G_X is finite.

Definition 3.2.3. (Morphisms between graphs of objects)
A *morphism* $f : X \to Y$ between graphs of objects with respect to $(\mathcal{C}, \mathcal{M}_1, \mathcal{M}_2)$ consists of

- a graph morphism $\hat{f} : G_X \to G_Y$,
- morphisms $f_v : X_v \to X_{\hat{f}(v)}$ in $\mathcal{M}_1$ for $v \in VG_X$,
- and morphisms $f_e : X_e \to X_{\hat{f}(e)}$ in $\mathcal{M}_1$ for $e \in EG_X$,

such that $f_e = f_{\bar{e}}$ and

$$\begin{array}{ccc} X_e & \xrightarrow{f_e} & Y_{\hat{f}(e)} \\ \downarrow{\scriptstyle \phi_e} & & \downarrow{\scriptstyle \phi_{\hat{f}(e)}} \\ X_u & \xrightarrow{f_u} & Y_{\hat{f}(u)} \end{array} \tag{3.2.1}$$

commute whenever $u \in VG_X$ and $e \in \mathrm{lk}(u)$.

Definition 3.2.4. (Coverings of graphs of objects)
We say that a morphism $f : X \to Y$ between graphs of objects is a *covering* if $\hat{f}$ is a covering of graphs. We say that X is a *cover* of Y. Similarly, we say that $f : X \to X$ is an *automorphism* if $\hat{f}$ is a graph automorphism. Let $\mathrm{Aut}(X)$ denote the group of automorphisms of X.

Example 3.2.5. In the following table we give some examples of triples $(\mathcal{C}, \mathcal{M}_1, \mathcal{M}_2)$ that can be used to define graphs of objects.

Name	$\mathcal{C}$	$\mathcal{M}_1$	$\mathcal{M}_2$
Graph of topological spaces	topological spaces	homeomorphisms	continuous maps
Graph of finite simplicial complexes	finite simplicial complexes	simplicial isomorphisms	simplicial maps
Graph of finite cube complexes	finite non-positively curved cube complexes	cubical isomorphisms	locally isometric cubical immersions
Surface amalgams	compact surfaces with boundary	homeomorphisms up to isotopy	immersions up to isotopy
Graph of groups	groups	group isomorphisms	group monomorphisms

Observe that the $\mathcal{M}_1$ morphisms are all isomorphisms in the appropriate categories, as was required in Definition 3.2.2. Comparing with Definition 3.2.2, we see that graphs of objects corresponding to the first four rows are just examples of graphs of spaces with respect to different categories. As discussed earlier, the difference in definitions comes when we consider morphisms and coverings. A morphism of graphs of spaces is usually defined in a similar way to Definition 3.2.3, except that the morphisms f_v and f_e between vertex and edge spaces are not required to lie in $\mathcal{M}_1$. Similarly, a covering of graphs of spaces usually requires that the maps f_v and f_e are coverings of vertex and edge spaces, but this is still weaker than having them in $\mathcal{M}_1$, and the map $\hat{f}$ of underlying graphs is not required to be a covering of graphs. We can think of coverings of graphs of objects as "rigid" examples of coverings of graphs of spaces.

The last entry in the table, graphs of groups, is not an example of a graph of spaces, but it still gives a valid example of a graph of objects. However the usual notions of morphism and covering between graphs of groups are again weaker than those for graphs of objects, because we require the maps f_v and f_e to be group isomorphisms for graphs of objects.

To prove Leighton's Theorem for graphs of objects it turns out that we need a certain finiteness condition on the edge spaces, and so we'll need the following definition.

Definition 3.2.6. (Isotropy groups)
Let X be a graph of objects and $H < \mathrm{Aut}(X)$. For each $e \in EG_X$ and $v \in VG_X$ define the *isotropy groups* of e and v in H as

$$H(e) := \{h_e \mid h \in H,\, \hat{h}(e) = e\} < \mathrm{Aut}_{\mathcal{M}_1}(X_e),$$
$$H(v) := \{h_v \mid h \in H,\, \hat{h}(v) = v\} < \mathrm{Aut}_{\mathcal{M}_1}(X_v).$$

These are different from the stabilisers H_e and H_v; there are homomorphisms $H_e \to H(e)$ and $H_v \to H(v)$ sending h to h_e or h_v respectively, these are surjective but not necessarily injective.

3.2.2 Theorems

We can now state our version of Leighton's Theorem for graphs of objects. See Example 3.2.14 for why the assumption of finite edge isotropy groups is necessary.

Theorem 1.1.6. *(Graph of Objects Leighton's Theorem)*
Let X^1 and X^2 be finite graphs of objects covered by a tree of objects X. If $\mathrm{Aut}(X)$ has finite edge isotropy groups, then X^1 and X^2 have a common finite cover.

Remark 3.2.7. The assumption that $\mathrm{Aut}(X)$ has finite edge isotropy groups will be automatically satisfied if the edge objects have finite $\mathcal{M}_1$-automorphism groups. This is the case for graphs of finite simplicial or cube complexes from Example 3.2.5, and also for surface amalgams if we assume that the edge objects are annuli. For the latter example the vertex objects might have infinite $\mathcal{M}_1$-automorphism groups (in this case each vertex object is a surface and the $\mathcal{M}_1$-automorphism group is the mapping class group), but this is not a problem because Theorem 1.1.6 does not require $\mathrm{Aut}(X)$ to have finite vertex isotropy groups.

We will actually prove a stronger version of Theorem 1.1.6 which incorporates a notion of symmetry-restricted closure analogous to that of Theorem 1.1.5. For this we need two more definitions.

Definition 3.2.8. (Symmetry-restricted closure)
Let X be a graph of objects and $H < \mathrm{Aut}(X)$. Define the *symmetry-restricted closure of H* to be the subgroup $\mathscr{S}(H) < \mathrm{Aut}(X)$ consisting of automorphisms g such that:

(1) For all $e \in EG_X$ there exists $h \in H$ with $\hat{g}(e) = \hat{h}(e)$ and $g_e = h_e$.

(2) For all $v \in VG_X$ there exists $h \in H$ with $\hat{g}(v) = \hat{h}(v)$ and $g_v = h_v$.

Definition 3.2.9. (Deck transformations)
If $f : \tilde{X} \to X$ is a covering of graphs of objects such that $T := G_{\tilde{X}}$ is a tree, then $\pi_1 G_X$ acts on T as the group of deck transformations of the cover $\hat{f} : T \to G_X$. We also get an action of $\pi_1 G_X$ on $\tilde{X}$ defined by the homomorphism $\rho : \pi_1 G_X \to \mathrm{Aut}(\tilde{X})$, where $\widehat{\rho(g)}$ is given by the aforementioned action of $\pi_1 G_X$ on T, and the morphisms $\rho(g)_v, \rho(g)_e \in \mathcal{M}_1$ are uniquely determined by the equations $f_v = f_{g(v)}\rho(g)_v$ and $f_e = f_{g(e)}\rho(g)_e$ (remember $\mathcal{M}_1$ forms a groupoid). It is easy to check that $\rho(g)$ satisfies the commutative square (3.2.1) and that ρ is a homomorphism. Clearly $f = f\rho(g)$ for all $g \in \pi_1 G_X$, and so we call the image of ρ the *group of deck transformations of the cover $f : \tilde{X} \to X$*.

Our symmetry-restricted version of Leighton's Theorem for graphs of objects is the following.

Theorem 3.2.10. *Let*

$$f^1 : \tilde{X} \to X^1$$
$$f^2 : \tilde{X} \to X^2$$

be coverings of graphs of objects, with $G_1 := G_{X^1}$ *and* $G_2 := G_{X^2}$ *both finite, and* $T := G_{\tilde{X}}$ *a tree. Let* Γ_1 *and* Γ_2 *be the groups of deck transformations for* f^1 *and* f^2*, and suppose that* $\Gamma_1, \Gamma_2 < H < \mathrm{Aut}(\tilde{X})$*. Suppose also that* H *has finite edge isotropy groups. Then* X^1 *and* X^2 *have a common finite cover* X*, and there exists* $g \in \mathscr{S}(H)$ *that fits into the following commutative diagram of coverings.*

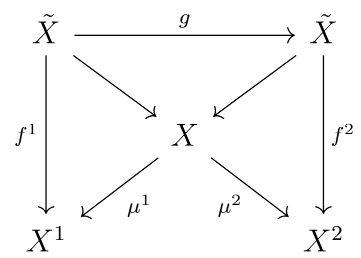

(3.2.2)

This theorem can be stated more concisely by working entirely in the tree of objects $\tilde{X}$ as follows.

Theorem 3.2.11. *Let* X *be a tree of objects with* $G_X = T$*, and let* $H < \mathrm{Aut}(X)$ *be a subgroup with finite edge isotropy groups. Suppose that* $\Gamma_1, \Gamma_2 < H$ *act freely cocompactly on* T*. Then there exists* $g \in \mathscr{S}(H)$ *such that* Γ_1^g *is commensurable with* Γ_2 *in* $\mathrm{Aut}(X)$*.*

In Theorems 3.2.10 and 3.2.11 it would of course suffice to define H as the subgroup of $\mathrm{Aut}(X)$ generated by Γ_1 and Γ_2, we state the theorems for a general H simply because there can be other subgroups H that arise naturally in examples, such as in Section 3.3. The equivalence of Theorems 3.2.10 and 3.2.11 follows from the Galois correspondence for coverings of graphs of objects, which we describe in the following remark.

Remark 3.2.12. Let $\tilde{X}$ be a tree of objects with underlying tree T, and let $f : \tilde{X} \to X$ be a covering of a finite graph of objects with deck transformation group Γ. The usual covering theory of graphs gives us a correspondence between subgroups of Γ and intermediate covers of $T \to G_X$; and given an intermediate cover of $T \to G_X$ there is a unique way of assigning edge and vertex objects (up to $\mathcal{M}_1$-isomorphism) and morphisms that make diagram (3.2.1) commute, as we explain below. Thus we have a correspondence between subgroups of Γ and intermediate covers of $f : \tilde{X} \to X$ (up to isomorphism).

An intermediate cover of graphs $T \xrightarrow{\hat{g}} G_Y \xrightarrow{\hat{h}} G_X$ induces an intermediate cover of graphs of objects $\tilde{X} \xrightarrow{g} Y \xrightarrow{h} X$ as follows (with $hg = f$). Define vertex and edge objects for Y by $Y_u := X_{\hat{h}(u)}$ and $Y_e := X_{\hat{h}(e)}$, define the morphisms h_u and h_e by the identity, and define the edge morphisms in Y so that diagram (3.2.1) commutes for the map $h : Y \to X$. For $u \in VT$ and $e \in ET$ we define $g_u := h^{-1}_{\hat{g}(u)} f_u : \tilde{X}_u \to Y_{\hat{g}(u)}$ and $g_e := h^{-1}_{\hat{g}(e)} f_e : \tilde{X}_e \to Y_{\hat{g}(e)}$, the map $g : \tilde{X} \to Y$ satisfies (3.2.1) because h and f do. The uniqueness of this construction up to $\mathcal{M}_1$-isomorphism essentially follows because each stage of the construction was forced by diagram (3.2.1).

Proof of Theorem 3.2.10.

As for the original Leighton's Theorem, we will build a common finite cover by first constructing a finite groupoid $\mathcal{S}$ consisting of "maps between links" in $\tilde{X}$. But now each link is not just a set of edges meeting at a common vertex, as each link is endowed with the extra data of edge objects and edge morphisms. Thus these "maps between links" in $\mathcal{S}$ must have the additional data of morphisms between edge objects, and these morphisms must act naturally with respect to the edge morphisms. Once we have defined $\mathcal{S}$, the proof will follow that of Theorem 1.1.1 quite closely - but with an extra step at the end to verify that we get a commutative diagram as in (3.2.2).

Step 1: Before constructing $\mathcal{S}$, we will define a general notion of "link map".

Given graphs of objects X, Y and $u \in VG_X$, $v \in VG_Y$, a *link map* from u to v is given by the data $s = (\hat{s},\, s_e : e \in \text{lk}(u))$, where:

(a) $\hat{s} : \text{lk}(u) \to \text{lk}(v)$ is a bijection.

(b) $s_e : X_e \to Y_{\hat{s}(e)}$ is a morphism in $\mathcal{M}_1$. There must also exist a morphism $s_u : X_u \to X_v$ in $\mathcal{M}_1$ such that $s_u \phi_e = \phi_{\hat{s}(e)} s_e$ for all $e \in \text{lk}(u)$ - but s_u is not part of the data of s (for a given s there could be many choices of s_u, whenever we write s_u we refer to some arbitrary choice).

If Z is another graph of objects and $w \in VG_Z$ and t is a link map from v to w, then we can compose s with t to produce a link map ts from u to w with $\widehat{ts} = \hat{t}\hat{s}$, $(ts)_e = t_{\hat{s}(e)} s_e$ and $(ts)_u = t_v s_u$. There is a natural notion of identity link map at a vertex u in which the morphisms s_e will be identity morphisms, and any link map will have an inverse by replacing $\hat{s}$ and the morphisms s_e with their inverses. Therefore the class of all link maps forms a category with inverses, where the objects are vertices in graphs of objects, and a link map s from u to v has $i(s) = u$ and $t(s) = v$.

If $f : X \to Y$ is a covering of graphs of objects, then for each $u \in VG_X$ there is a link map f^u from u to $\hat{f}(u)$ in which $\hat{f}^u$ is the restriction of $\hat{f}$ to $\text{lk}(u)$, $f^u_e := f_e$ and $f^u_u := f_u$. And if $g : Y \to Z$ is another covering, then it is easy to check that $(gf)^u = g^{\hat{f}(u)} f^u$.

Our groupoid $\mathcal{S}$ will be the subcategory of link maps, with $\text{Ob}(\mathcal{S}) = VG_1 \sqcup VG_2$, and $\text{Hom}(\mathcal{S})$ consisting of link maps

$$s = (f^i)^{\hat{h}(z)} h^z ((f^j)^z)^{-1} \tag{3.2.3}$$

for $z \in VT$, $h \in H$ and $i, j \in \{1, 2\}$. This is a link map from $u := \hat{f}^j(z)$ to $v := \hat{f}^i \hat{h}(z)$. Intuitively, think of s as lifting the link of u up to the link of z, mapping across to the link of $\hat{h}(z)$ by h, and then projecting down to the link of v - a cartoon of this is given in Figure 3.2.

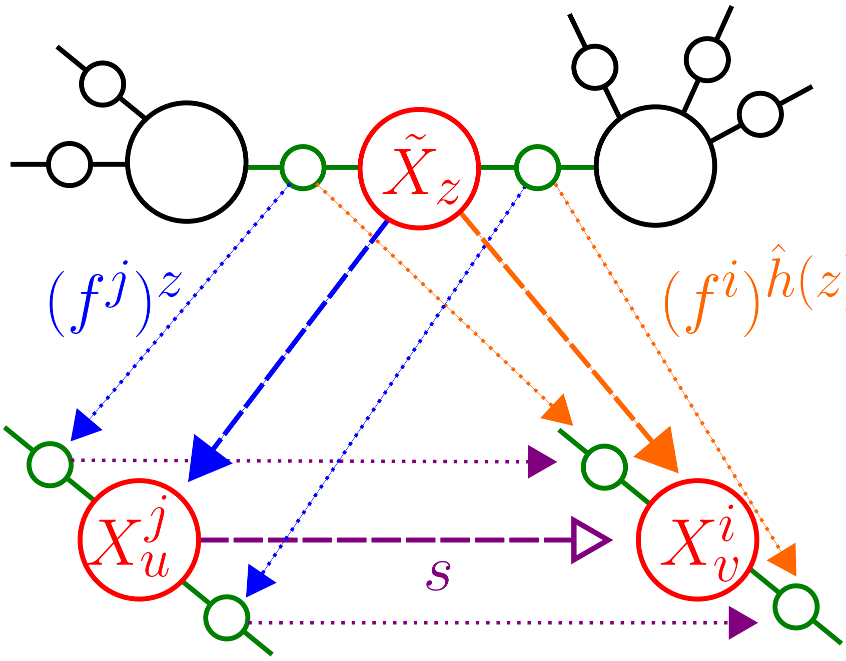

Figure 3.2: Diagram of the link map s from (3.2.3).

Unlike for the original Leighton's theorem, it is not obvious that $\mathcal{S}$ is finite, so we prove this now.

Claim: $\mathcal{S}$ is finite.

Proof: It is enough to show that each $\mathcal{S}(u, v)$ is finite. But $\mathcal{S}(u, v)$ is a coset of $\mathcal{S}(u, u)$, so it suffices to show that each $\mathcal{S}(u, u)$ is finite. There is a homomorphism $\theta : \mathcal{S}(u, u) \to \text{Aut}(\text{lk}(u))$ with finite image, so it is enough to show that $\ker\theta$ is finite.

Let $s \in \ker\theta$ and suppose $u \in VG_1$. Fix $z_0 \in VT$ with $\hat{f}^1(z_0) = u$. By (3.2.3) we can write

$$s = (f^1)^{z'} h^z ((f^1)^z)^{-1}$$

for some $z, z' \in VT$ and $h \in H$ with $\hat{h}(z) = z'$. But then $\hat{f}^1(z) = \hat{f}^1(z') = u$, thus there exist $g^1, g^2 \in \Gamma_1$ with $\hat{g}^1(z_0) = z$ and $\hat{g}^2(z') = z_0$, and we get

$$s = (f^1)^{z_0}(g^2)^{z'}h^z(g^1)^{z_0}((f^1)^{z_0})^{-1}.$$

Moreover, $s \in \ker\theta$ so $\widehat{g^2hg^1}$ must fix $\mathrm{lk}(z_0)$, hence $(g^2hg^1)_e \in H(e)$ for each $e \in \mathrm{lk}(z_0)$. By assumption the groups $H(e)$ are finite, hence there are only finitely many possibilities for $s \in \ker\theta$, as required. ∎

Step 2: As for the original Leighton's Theorem, we now define an action of $\mathcal{S}$ on a set of edge-related things, and show that it respects some notion of edge inversion. Define a finite set

$$A := \{(e, \hat{a}(e), a_e) \mid e \in EG_1 \sqcup EG_2,\ a \in \mathcal{S}(\tau(e), -)\}.$$

The observant will note that A can also be given the structure of a groupoid, but we won't need this here.

We define an action of $\mathcal{S}$ on A by $s \cdot (e, \hat{a}(e), a_e) := (e, \widehat{sa}(e), (sa)_e)$ for $s \in \mathcal{S}(t(a), -)$, with associated map $\varepsilon : A \to VG_1 \sqcup VG_2$ given by $\varepsilon(e, \hat{a}(e), a_e) = \tau\hat{a}(e) = t(a)$.

We want to define an involution $A \to A$ given by $(e, \hat{a}(e), a_e) \mapsto (\bar{e}, \overline{\hat{a}(e)}, a_e)$. This is well-defined by the following claim.

<u>Claim:</u> For any $(e, \hat{a}(e), a_e) \in A$ there exists $a' \in \mathcal{S}(\tau(\bar{e}), -)$ with $\hat{a'}(\bar{e}) = \overline{\hat{a}(e)}$ and $a'_{\bar{e}} = a_e$.

<u>Proof:</u> As in (3.2.3), write

$$a = (f^i)^{\hat{h}(z)}h^z((f^j)^z)^{-1}$$

for $z \in VT$, $h \in H$ and $i, j \in \{1, 2\}$, with $\hat{f}^j(z) = \tau(e)$. Let $\hat{e} \in \mathrm{lk}(z)$ with $\hat{f}^j(\hat{e}) = e$ and put $z' := \iota(\hat{e})$. Then define $a' \in \mathcal{S}(\tau(\bar{e}), -)$ by

$$a = (f^i)^{\hat{h}(z')}h^{z'}((f^j)^{z'})^{-1}.$$

We have $i(a') = \hat{f}^j(z') = \hat{f}^j(\iota(\hat{e})) = \iota(e) = \tau(\bar{e})$ as required. We also have $\hat{f}^j(\bar{\hat{e}}) = \bar{e}$, so $\hat{a'}(\bar{e}) = \hat{f}^i\hat{h}(\bar{\hat{e}}) = \overline{\hat{f}^i\hat{h}(\hat{e})} = \overline{\hat{a}(e)}$. And finally we have

$$\begin{aligned} a'_{\bar{e}} &= (f^i h)_{\bar{\hat{e}}}(f^j_{\bar{e}})^{-1} \\ &= (f^i h)_{\hat{e}}(f^j_e)^{-1} \\ &= a_e. \end{aligned}$$

■

A can be partitioned into sets

$$A(e) := \{(e, \hat{a}(e), a_e) \mid a \in \mathcal{S}(\tau(e), -)\},$$

for $e \in EG_1 \sqcup EG_2$. These sets are related to the action of $\mathcal{S}$ by the following claim.

Claim: $\mathcal{S} \cdot (e, \hat{a}(e), a_e) = A(e)$

Proof: The inclusion $\subset$ is clear from the definitions. The inclusion $\supset$ is also easy, because for $(e, \hat{s}(e), s_e) \in A(e)$ we have $sa^{-1} \cdot (e, \hat{a}(e), a_e) = (e, \hat{s}(e), s_e)$. ■

Step 3: We can now construct our common finite cover X of X^1 and X^2. The underlying graph $G := G_X$ will have vertex set given by

$$VG := \left\{(s, l) \mid s \in \mathcal{S}(u_1, u_2),\ u_1 \in VG_1,\ u_2 \in VG_2,\ 1 \le l \le \frac{N}{|\mathcal{S}(u_1, -)|}\right\},$$

and edge set given by

$$EG := \left\{(e_1, e_2, m, k) \mid e_1 \in EG_1,\ e_2 \in EG_2,\ (e_1, e_2, m) \in A,\ 1 \le k \le \frac{N}{|A(e_1)|}\right\},$$

where N is a fixed positive integer that is a common multiple of all the integers $|\mathcal{S}(u_1, -)|$ and $|A(e_1)|$.

A admits an involution $(e, \hat{a}(e), a_e) \mapsto (\bar{e}, \overline{\hat{a}(e)}, a_e)$, as in Step 2, which induces bijections $A(e) \to A(\bar{e})$. Hence we can define edge inversion in G by $\overline{(e_1, e_2, m, k)} := (\bar{e}_1, \bar{e}_2, m, k)$ (note that $a_{\bar{e}} = a_e$).

Vertex and edge objects in X will be given by

$$X_{(s,l)} := X^1_{i(s)}, \qquad\qquad X_{(e_1,e_2,m,k)} := X^1_{e_1}.$$

To complete the construction of G, we must define the map $\tau : EG \to VG$, and the edge morphisms in X. Fix $e_1 \in EG_1$, $e_2 \in EG_2$ and $(e_1, e_2, m) \in A$, and say $\tau(e_q) = u_q$ $(q = 1, 2)$. We would like each edge $(e_1, e_2, m, k) \in EG$ to satisfy

$$\tau(e_1, e_2, m, k) = (s, l) \tag{3.2.4}$$

for some $s \in \mathcal{S}(u_1, u_2)$ such that $\hat{s}(e_1) = e_2$ and $s_{e_1} = m$. Note this is equivalent to $s \cdot (e_1, e_1, 1_{e_1}) = (e_1, e_2, m)$, where 1_{e_1} is the identity morphism $X^1_{e_1} \to X^1_{e_1}$. We arrange this by choosing an arbitrary matching between such vertices (s, l) and integers $1 \le k \le N/|A(e_1)|$ - to verify that this is valid, we must check that we have equal numbers of each. Indeed, Lemma 2.2.4 tells us that

$$|\{s \in \mathcal{S}(u_1, -) \mid s \cdot (e_1, e_1, 1_{e_1}) = (e_1, e_2, m)\}| = |\mathrm{Stab}_{\mathcal{S}}(e_1, e_1, 1_{e_1})|,$$

and so

$$\begin{aligned} \left| \left\{ (s,l) \;\middle|\; \begin{array}{r} s \in \mathcal{S}(u_1, -),\ s \cdot (e_1, e_1, 1_{e_1}) = (e_1, e_2, m), \\ 1 \le l \le \dfrac{N}{|\mathcal{S}(u_1, -)|} \end{array} \right\} \right| &= |\mathrm{Stab}_{\mathcal{S}}(e_1, e_1, 1_{e_1})| \frac{N}{|\mathcal{S}(u_1, -)|} \\ &= \frac{N}{|\mathcal{S} \cdot (e_1, e_1, 1_{e_1})|} \\ &= \frac{N}{|A(e_1)|}. \end{aligned}$$

The second equality is by Lemma 2.2.4 while the third equality follows from Step 2.

Finally, we define the edge morphisms in X using the edge morphisms in X^1:

$$\phi_{(e_1, e_2, m, k)} := \phi_{e_1} : X^1_{e_1} \to X^1_{u_1}.$$

Step 4: In this last step, we define coverings from X down to X^1 and X^2.

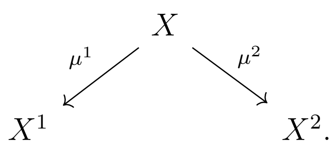

We define the maps $\hat{\mu}^1$ and $\hat{\mu}^2$ by

$$\begin{aligned} \hat{\mu}^1(s,l) &:= i(s), & \hat{\mu}^1(e_1, e_2, m, k) &:= e_1, \\ \hat{\mu}^2(s,l) &:= t(s), & \hat{\mu}^2(e_1, e_2, m, k) &:= e_2. \end{aligned}$$

These clearly preserve edge inversion, and $\tau(e_1, e_2, m, k) = (s, l)$ as in (3.2.4) implies

$$\begin{aligned} \tau\hat{\mu}^1(e_1, e_2, m, k) &= \tau(e_1) & \tau\hat{\mu}^2(e_1, e_2, m, k) &= \tau(e_2) \\ &= u_1 & &= u_2 \\ &= i(s) & &= t(s) \\ &= \hat{\mu}^1(s,l) & &= \hat{\mu}^2(s,l) \\ &= \hat{\mu}^1\tau(e_1, e_2, m, k), & &= \hat{\mu}^2\tau(e_1, e_2, m, k), \end{aligned}$$

thus $\hat{\mu}^1$ and $\hat{\mu}^2$ are well-defined graph morphisms.

We can then define μ^1 and μ^2 by the morphisms

$$\begin{aligned} \mu^1_{(s,l)} &:= 1_{i(s)}, & \mu^1_{(e_1, e_2, m, k)} &:= 1_{e_1}, \\ \mu^2_{(s,l)} &:= s_{i(s)}, & \mu^2_{(e_1, e_2, m, k)} &:= m, \end{aligned}$$

where $1_{i(s)}$ is the identity morphism $X^1_{i(s)} \to X^1_{i(s)}$, and $s_{i(s)}$ is not uniquely determined by s (see Step 1, part (b) in the definition of link map), we will just make some arbitrary choice for each s. Note that the above morphisms do go between the appropriate vertex and edge objects as specified by the maps $\hat{\mu}^1$ and $\hat{\mu}^2$.

Again with $\tau(e_1, e_2, m, k) = (s, l)$ as in (3.2.4), we get the following commutative squares, demonstrating that μ^1 and μ^2 are well-defined morphisms of graphs of objects. The bottom left square commutes precisely because s is a link map.

$$\begin{array}{ccc} X^1_{e_1} & = & X^1_{e_1} \\ \downarrow{\scriptstyle \phi_{e_1}} & & \downarrow{\scriptstyle \phi_{e_1}} \\ X^1_{u_1} & = & X^1_{u_1} \end{array} \quad \Rightarrow \quad \begin{array}{ccc} X_{(e_1,e_2,m,k)} & \xrightarrow{\mu^1_{(e_1,e_2,m,k)}} & X^1_{e_1} \\ \downarrow{\scriptstyle \phi_{(e_1,e_2,m,k)}} & & \downarrow{\scriptstyle \phi_{e_1}} \\ X_{(s,l)} & \xrightarrow{\mu^1_{(s,l)}} & X^1_{u_1} \end{array}$$

$$\begin{array}{ccc} X^1_{e_1} & \xrightarrow{m = s_{e_1}} & X^2_{e_2} \\ \downarrow{\scriptstyle \phi_{e_1}} & & \downarrow{\scriptstyle \phi_{e_2}} \\ X^1_{u_1} & \xrightarrow{s_{u_1}} & X^2_{u_2} \end{array} \quad \Rightarrow \quad \begin{array}{ccc} X_{(e_1,e_2,m,k)} & \xrightarrow{\mu^2_{(e_1,e_2,m,k)}} & X^2_{e_2} \\ \downarrow{\scriptstyle \phi_{(e_1,e_2,m,k)}} & & \downarrow{\scriptstyle \phi_{e_2}} \\ X_{(s,l)} & \xrightarrow{\mu^2_{(s,l)}} & X^2_{u_2} \end{array}$$

By construction, the link of a vertex (s, l) in G, with $s \in \mathcal{S}(u_1, u_2)$, takes the form

$$\mathrm{lk}(s, l) = \{(e, \hat{s}(e), s_e, k_e) \mid e \in \mathrm{lk}(u_1)\}$$

where each k_e is some integer associated to e. Now $\hat{\mu}^1(e, \hat{s}(e), s_e, k_e) = e$ and $\hat{\mu}^2(e, \hat{s}(e), s_e, k_e) = \hat{s}(e)$, so $\hat{\mu}^1$ and $\hat{\mu}^2$ induce bijections from $\mathrm{lk}(s, l)$ to $\mathrm{lk}(u_1)$ and $\mathrm{lk}(u_2)$ respectively. We conclude that $\hat{\mu}^1$ and $\hat{\mu}^2$ are graph coverings, which makes μ^1 and μ^2 coverings of graphs of objects.

Step 5: Finally we must construct diagram (3.2.2) with $g \in \mathscr{S}(H)$. We can assume that the graph G is connected, as otherwise we just restrict to a component. The usual covering space theory of graphs allows us to draw the following commutative diagram of graph coverings.

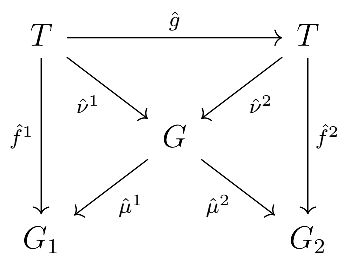

(3.2.5)

As in Remark 3.2.12, there is then a unique way of upgrading $\hat{\nu}^1$, $\hat{\nu}^2$ and $\hat{g}$ to coverings of graphs of objects ν^1, ν^2 and g. So we now have the following commutative diagram of graphs of objects, with $g \in \mathrm{Aut}(\tilde{X})$.

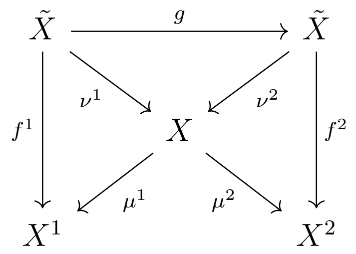

(3.2.6)

It remains to prove that $g \in \mathscr{S}(H)$. Consider a link map $s \in \mathcal{S}$ with $i(s) = u_1 \in VG_1$ and $t(s) = u_2 \in VG_2$. Suppose that s takes the form

$$s = (f^2)^{z_2} h^{z_1} ((f^1)^{z_1})^{-1}$$

from (3.2.3), with $z_1, z_2 \in VT$, $\hat{f}^i(z_i) = u_i$, $h \in H$, and $\hat{h}(z_1) = z_2$. We can then choose the morphism s_{u_1} to fit into the following commutative diagram.

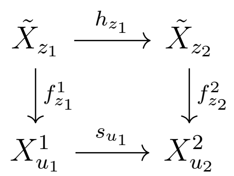

(3.2.7)

If $X_{(s,l)}$ is a vertex object of X, and $v_1, v_2 \in VT$ are such that $\hat{\nu}^i(v_i) = (s, l)$ (so $\hat{f}^i(v_i) = u_i$) and $\hat{g}(v_1) = v_2$, then we get a larger commutative diagram as follows.

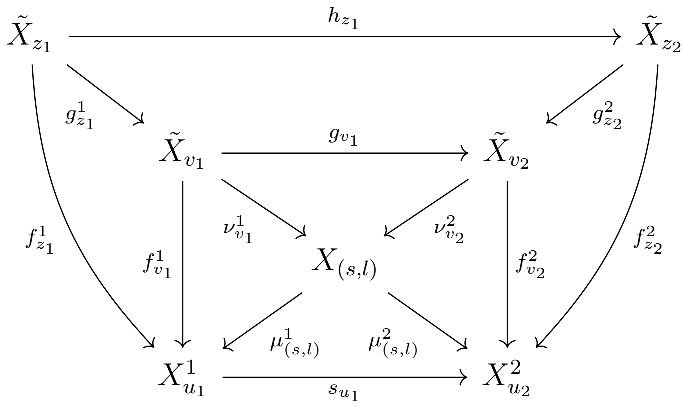

(3.2.8)

Here $g^i \in \Gamma_i$ is the element of the deck transformation group with $\hat{g}^i(z_i) = v_i$. Since $\Gamma_1, \Gamma_2 < H$, the top square of (3.2.8) implies that $g_{v_1} = h'_{v_1}$ for some $h' \in H$.

A very similar argument can be run for edge objects, and so we conclude that $g \in \mathscr{S}(H)$.

□

Remark 3.2.13. Given $v \in VT$, we can choose the automorphism $g \in \mathrm{Aut}(\tilde{X})$ from Theorem 3.2.10 such that $\hat{g}(v) = v$ and $g_v = 1_{\tilde{X}_v}$ (and similarly for $e \in ET$). This follows by examining Step 5 of the proof. Indeed, let $s \in \mathcal{S}$ be defined by $s = (f^2)^v((f^1)^v)^{-1}$, and when restricting to a component of G at the beginning of Step 5 make sure that the vertex $(s, 1)$ is included. Then choose the coverings $\hat{\nu}^1$, $\hat{\nu}^2$ and $\hat{g}$ so that $\hat{\nu}^i(v) = (s, 1)$ and $\hat{g}(v) = v$. If $\hat{f}^i(v) = u_i$, then we can draw diagram (3.2.8) with $z_i = v_i = v$ and $g^i = h = 1 \in \mathrm{Aut}(\tilde{X})$. It follows that $g_v = 1_{\tilde{X}_v}$ as required.

3.2.3 Counter-example for infinite edge isotropy groups

To close the section, we give an example to show that the assumption of finite edge isotropy groups in Theorem 3.2.10 is necessary.

Example 3.2.14. We work with graphs of topological spaces as in Example 3.2.5. We can exploit the infinite symmetry of the circle $S^1 \subset \mathbb{C}$ to build finite graphs of objects X^1, X^2, with a common cover $\tilde{X}$, but no common finite cover.

- Let X^1 have a single vertex object X^1_v and a single edge object X^1_e, both equal to S^1, and let the edge maps $\phi_e, \phi_{\bar{e}}$ both be the identity.
- Let X^2 also have a single vertex object X^2_v and a single edge object X^2_e, both equal to S^1, and let ϕ_e be the identity, but this time take $\phi_{\bar{e}}$ to be the rotation $r : z \mapsto e^i z$. The important feature of this rotation is that it has infinite order in the homeomorphism group of S^1.
- Let $\tilde{X}$ have underlying graph consisting of an infinite chain of edges $(e_i)_{i\in\mathbb{Z}}$ and vertices $(v_i)_{i\in\mathbb{Z}}$ with $\iota(e_i) = v_i$ and $\tau(e_i) = v_{i+1}$. Let all the edge and vertex objects of $\tilde{X}$ equal S^1 and let all edge maps be the identity.

  There are covers

$$f^1 : \tilde{X} \to X^1$$
$$f^2 : \tilde{X} \to X^2$$

  defined by

  (1) $\hat{f}^1(v_i) = v$, $\hat{f}^1(e_i) = e$, and $f^1_{v_i} = f^1_{e_i} = \mathrm{id}_{S^1}$ for all $i \in \mathbb{Z}$.

  (2) $\hat{f}^2(v_i) = v$, $\hat{f}^2(e_i) = e$, and $f^2_{v_i} = f^2_{e_i} = r^i$ for all $i \in \mathbb{Z}$.

Why do X^1 and X^2 have no common finite cover? Well any finite cover $g^1 : X \to X^1$ must be a circuit of copies of S^1, more precisely it must take the following form (up to isomorphism of X).

- $VX = \{v_1, ..., v_n\}$ and $EX = \{e_1, ..., e_n\}$ for some $n \in \mathbb{N}$.

- $\iota(e_i) = v_i$ $(1 \leq i \leq n)$, $\tau(e_i) = v_{i+1}$ $(1 \leq i \leq n-1)$ and $\tau(e_n) = v_1$.
- $\hat{g}^1(v_i) = v$, $\hat{g}^1(e_i) = e$ for all i.
- $g^1_{v_i} = g^1_{e_i} = \mathrm{id}_{S^1}$ for all i.

If there was a covering $g^2 : X \to X^2$, there would be two possibilities for $\hat{g}^2$ corresponding to $\hat{g}^2(e_1) = e$ or $\bar{e}$. Suppose we're in the first case (the second will lead to a contradiction similarly), then $\hat{g}^2 = \hat{g}^1$. Put $a = g^2_{v_1}$. The commutative square (3.2.1) then forces $g^2_{e_1} = a, g^2_{v_2} = ra, g^2_{e_2} = ra, g^2_{v_3} = r^2a, g^2_{e_3} = r^2a$ and so on. But taking this right round the circuit we deduce that $g^2_{v_1} = r^n a$, which is a contradiction because $r^n \neq \mathrm{id}_{S^1}$.

3.3 Symmetry-restricted version

We now show how to deduce the Symmetry-restricted Leighton's Theorem from its graph of objects counterpart, Theorem 3.2.11.

3.3.1 The theorem

Recall the definition of symmetry-restricted closure, for a tree T, $H < \mathrm{Aut}(T)$, and R an integer, we define

$$\mathscr{S}_R(H) := \{g \in \mathrm{Aut}(T) \mid \forall x \in VT, \exists h \in H \text{ s.t. } h \text{ agrees with } g \text{ on } B_R(x)\}.$$

Theorem 1.1.5. *(Symmetry-restricted Leighton's Theorem)*
Let T be a tree, and $H < \mathrm{Aut}(T)$, and let $\Gamma_1, \Gamma_2 < H$ be free uniform lattices in $\mathrm{Aut}(T)$. Then for all $R \in \mathbb{N}$ there exists $g \in \mathscr{S}_R(H)$ such that Γ_1^g is commensurable to Γ_2 in $\mathrm{Aut}(T)$.

Proof. We will turn T into a tree of objects X (ie. $G_X = T$). This will be with respect to $(\mathcal{C}, \mathcal{M}_1, \mathcal{M}_2)$, where $\mathcal{C}$ is the category of pairs (Y, U) for Y a finite tree and $U \subset VY$, and a morphism in $\mathcal{C}$ from (Y, U) to (Y', U') is a tree embedding $Y \hookrightarrow Y'$ such that U' is contained in the image of U. A morphism is in $\mathcal{M}_1$ if $Y \to Y'$ is an isomorphism and U' equals the image of U, and all morphisms are in $\mathcal{M}_2$. We then define the vertex objects for X as based R-balls $X_v := (B_R(v), \{v\})$ for $v \in VT$, and the edge objects as $X_e := (N_{R-1}(e), \{\iota(e), \tau(e)\})$ for $e \in ET$, where $N_{R-1}(e)$ is the $(R-1)$-neighbourhood of e. The morphisms $\phi_e : X_e \to X_{\tau(e)}$ are given by the inclusions $N_{R-1}(e) \hookrightarrow B_R(\tau(e))$.

For $g \in \mathrm{Aut}(T)$ let g_v be the restriction of g to $B_R(v)$ and let g_e be the restriction of g to $N_{R-1}(e)$. We then have a homomorphism $\psi : \mathrm{Aut}(T) \to \mathrm{Aut}(X)$ defined by $\widehat{\psi(g)} := g$ and

$$\psi(g)_v := g_v : B_R(v) \to B_R(gv),$$
$$\psi(g)_e := g_e : N_{R-1}(e) \to N_{R-1}(ge).$$

It is easy to check that ψ is a well-defined homomorphism. The key point is that ψ is actually an isomorphism, as we will now show.

Claim: ψ is an isomorphism.

Proof: The homomorphism ψ admits a retraction $r : \mathrm{Aut}(X) \to \mathrm{Aut}(T)$ given by $g \mapsto \hat{g}$, so we must show that r has trivial kernel.

Consider $g \in \ker(r)$ and pick $v \in VT$. Take a vertex $u \in B_R(v)$. We will show that $g_v(u) = u$ by induction on the distance $d(u, v)$. We know that $g_v(v) = v$ by definition of morphisms in $\mathcal{C}$, so we may assume that $u \neq v$. Suppose that the segment $[v, u]$ in T starts with the edge e such that $\iota(e) = v$ and $\tau(e) = w$. As g is a morphism of graphs of objects, we have the following commutative diagram.

$$\begin{array}{ccccc} B_R(v) & \hookleftarrow & N_{R-1}(e) & \hookrightarrow & B_R(w) \\ \downarrow g_v & & \downarrow g_e & & \downarrow g_w \\ B_R(v) & \hookleftarrow & N_{R-1}(e) & \hookrightarrow & B_R(w) \end{array} \tag{3.3.1}$$

We know that $d(u, w) < d(u, v)$, so $u \in B_R(w)$, and by induction we have $g_w(u) = u$. Diagram (3.3.1) then implies that g_v also fixes u. This holds for all $u \in B_R(v)$, so g_v is the identity map on $B_R(v)$, and diagram (3.3.1) implies that g_e is the identity map on $N_{R-1}(e)$. Therefore g is the identity on all edge and vertex objects, so $g = 1 \in \mathrm{Aut}(X)$ as required. ■

As a consequence of the claim we know that $\hat{g}_v = g_v$ and $\hat{g}_e = g_e$ for any $g \in \mathrm{Aut}(X)$, $v \in VT$ and $e \in ET$. It follows easily that, for $H < \mathrm{Aut}(T)$, the R-symmetry-restricted closure from Definition 1.1.4 corresponds to the symmetry-restricted closure from Definition 3.2.8:

$$\psi(\mathscr{S}_R(H)) = \mathscr{S}(\psi(H))$$

We are then done by Theorem 3.2.11. □

Remark 3.3.1. It follows from Remark 3.2.13 that given $v \in VT$ we can choose the conjugating automorphism $g \in \mathrm{Aut}(T)$ to restrict to the identity on the ball $B_R(v)$.

3.3.2 Graphs of polygons

Recall from Example 1.1.2 that a *graph of polygons* is a space consisting of solid regular polygons with some edges joining vertices of the polygons. And recall that a *covering* of graphs of polygons is a topological covering that restricts to isometries between polygons. We now give two proofs of why Leighton's Theorem holds for graphs of polygons, the first proof views graphs of polygons as graphs of objects while the second proof uses the Symmetry-restricted Leighton's Theorem.

Proposition 3.3.2. *Leighton's Theorem holds for graphs of polygons: if two finite graphs of polygons (ie. with finite underlying graphs) are covered by the same tree of polygons, then they have a common finite cover.*

Proof 1. Graphs of polygons are graphs of objects with respect to

$$(\mathcal{C}, \mathcal{M}_1, \mathcal{M}_2) = (\text{metric spaces, isometries, isometric embeddings}).$$

The vertex objects are polygons and the edge objects are points. The proposition then follows from Theorem 1.1.6. □

Proof 2. Let X be a graph of polygons. Each polygon contains a *polygon star* consisting of the centre of the polygon and edges joining the centre to each vertex. Let X^* be the graph obtained from X by retracting each polygon to its polygon star. This induces an $\mathrm{Aut}(X)$-invariant retraction $X \to X^*$, which in turn induces an injective homomorphism $\mathrm{Aut}(X) \to \mathrm{Aut}(X^*)$. An example is illustrated in Figure 3.3.

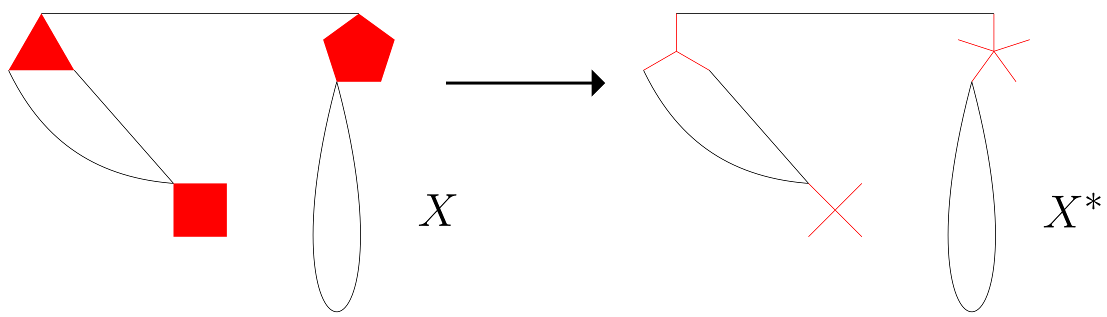

Figure 3.3: The retraction $X \to X^*$.

If two finite graphs of polygons X_1 and X_2 are covered by a tree of polygons T, with covering maps $p_i : T \to X_i$, then standard covering space theory tells us that any common finite cover $\hat{X}$ fits into a commutative diagram of covering maps as follows, where g is an automorphism of the tree of polygons T.

$$
\begin{array}{ccccc}
T & & \xrightarrow{\quad g \quad} & & T \\
 & \searrow^{\nu_1} & & \swarrow^{\nu_2} & \\
\Big\downarrow p_1 & & \hat{X} & & \Big\downarrow p_2 \\
 & \swarrow_{\mu_1} & & \searrow_{\mu_2} & \\
X_1 & & & & X_2
\end{array}
\tag{3.3.2}
$$

Let the deck transformation group of $p_i : T \to X_i$ be denoted $\Gamma_i < \mathrm{Aut}(T)$. As described for graphs in the introduction, the existence of $\hat{X}$ is equivalent to the existence of $g \in \mathrm{Aut}(T)$ such that Γ_1^g is commensurable to Γ_2, so let's now prove the latter.

As explained above, we have a retraction $T \to T^*$ that induces an injective homomorphism $\mathrm{Aut}(T) \to \mathrm{Aut}(T^*)$. Let $H < \mathrm{Aut}(T^*)$ be the image. The edges in each polygon star have a cyclic ordering given by the cyclic ordering on the vertices of the polygon, and an automorphism of T^* extends to an automorphism of T if and only if it maps polygon stars to polygon stars in a way that preserves the cyclic orderings. In the notation of Definition 1.1.4, this means that $H = \mathscr{S}_1(H)$. We may then apply Theorem 1.1.5 to obtain $g \in H$ that conjugates the image of Γ_1 in H onto a subgroup commensurable with the image of Γ_2. Pulling this back to $\mathrm{Aut}(T)$ gives $g \in \mathrm{Aut}(T)$ such that Γ_1^g is commensurable to Γ_2, as required. □

Chapter 4

Leighton's Theorem for quasitrees

In this chapter we prove the Quasitree Leighton's Theorem (Theorem 1.1.8): given finite graphs X_1 and X_2, if there is a quasitree X and regular coverings $X \to X_i$, then there is a covering $X \to \hat{X}$ of a finite graph $\hat{X}$ such that $\hat{X}$ covers both X_1 and X_2. We also construct examples demonstrating the necessity of the assumptions that X is a quasitree (Theorem 1.1.10) and that both the coverings $X \to X_i$ are regular (Theorem 1.1.11).

4.1 The need to control the geometry of the common cover

The aim of this section is to prove the following theorem.

Theorem 1.1.10. *There exist connected graphs X, X_1, X_2 and regular covering maps $X \to X_1$ and $X \to X_2$, such that X_1 and X_2 are finite but there does not exist a covering $X \to \hat{X}$ of any finite graph $\hat{X}$ such that $\hat{X}$ covers both X_1 and X_2.*

4.1.1 Square completion and #-subdivision

Definition 4.1.1. (Square completion)
A *reduced circuit* in a graph X is the loop determined by a finite cyclically-ordered set of edges $(e_1, \dots, e_n)$ with $\tau(e_i) = \iota(e_{i+1})$ and $e_{i+1} \neq \bar{e}_i$ for $i = 1, \dots, n$, indices mod n. (We specify *cyclic* ordering so that $(e_1, \dots, e_n) = (e_2, \dots, e_n, e_1)$ etc.) If the vertices $\tau(e_i)$ are all distinct, then the subgraph consisting of these vertices and the unoriented edges corresponding to the e_i is called an *n-cycle*.

The *square completion* $\square(X)$ of a graph X is the combinatorial 2-complex obtained from X by attaching a 2-cell to each reduced circuit of length 4.

Definition 4.1.2. (#-subdivision)
Let K be a squared 2-complex, i.e. a combinatorial 2-complex such that the attaching

map of each 2-cell is a reduced circuit of length 4. We define $K_\#$ to be the squared 2-complex obtained from K by introducing new vertices and edges so as to divide each edge of K into a path of combinatorial length 5 and each 2-cell (square) into 25 squares in the obvious 5-by-5 pattern.

The integer 5 is used in this definition because it restricts the nature of short loops in $K_\#$; the following lemma would fail if we used $2, 3$ or 4. For example if we used 4, then any 1-cycle in K would give rise to a reduced circuit of length 4 in $K_\#$, and $\square(K_\#^{(1)})$ would contain a 2-cell attached along this circuit, but such a 2-cell does not exist in $K_\#$. We use the standard notation $K^{(1)}$ for the 1-skeleton of a combinatorial complex K.

Lemma 4.1.3. *If a squared 2-complex K is non-positively curved, then $K_\# = \square(K_\#^{(1)})$.*

Proof. For an arbitrary squared 2-complex, there are three types of reduced circuits of length 4 in $K_\#^{(1)}$, illustrated in Figure 4.1. The first type consists of the boundaries of the 2-cells in $K_\#$. The second and third types occur when the link of a vertex in K contains a 1-cycle or 2-cycle respectively. If K is non-positively curved then only circuits of the first type are possible. $\square$

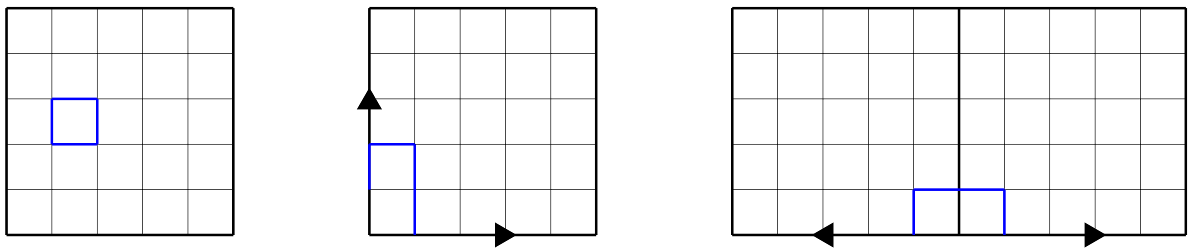

Figure 4.1: The three types of reduced circuits of length 4 in $K_\#^{(1)}$, highlighted in blue. The edges of K are in bold, and identifications between them are indicated by arrows.

4.1.2 Proof of Theorem 1.1.10

We make use of the following group constructed by Wise [100].

Example 4.1.4. The standard 2-complex K corresponding to the following presentation has universal cover $\tilde{K}$ a product of a regular tree of valence 4 and a regular tree of valence 6. Π does not contain a subgroup of finite index that splits as a direct product of free groups.

$$\Pi = \langle a, b, x, y, z \mid aya^{-1}x^{-1}, byb^{-1}x^{-1}, azb^{-1}z^{-1}, axb^{-1}y^{-1}, bxa^{-1}z^{-1}, bza^{-1}y^{-1} \rangle$$

We set $X := \tilde{K}_\#^{(1)}$ and $X_1 := K_\#^{(1)}$. The covering $\tilde{K} \to K$ induces a covering $X \to X_1$. The vertices of X_1 can be divided into three types: the original vertex of K has valence 10; the new vertices introduced in the interior of 1-cells have valence 6 or 8; and the new vertices introduced in the interior of 2-cells have valence 4. By Lemma 4.1.3, the only reduced circuits of length 4 in X_1 are the boundary cycles of the 2-cells of $K_\#$. It follows that the number of such cycles passing through each vertex is uniquely determined by the valence of that vertex, and this number is equal to the number of 4-cycles passing through a vertex of the same valence in X.

Let $\hat{X}$ be a graph and let $\pi : \hat{X} \to X_1$ be a covering. As X_1 contains no reduced circuits of length less than 4, neither does $\hat{X}$. So the number of 4-cycles passing through a vertex $v \in \hat{X}$ is no greater than the number passing through $\pi(v)$. (It will be strictly less if some loop of length 4 based at $\pi(v)$ does not lift to a loop based at v.) Similarly, if $q : X \to \hat{X}$ is a covering map then the number of reduced circuits of length 4 based at $w \in VX$ will be no greater than the number passing through $q(w)$. In our setting, the number of such loops at w is the same as the number at $\pi \circ q(w)$, because these vertices have the same valence. Thus q and π both induce bijections on the set of reduced circuits of length 4 passing through each vertex. These circuits are the attaching maps for the 2-cells of the square-completion of each graph. Therefore the coverings $X \to \hat{X} \to X$ extend to coverings of combinatorial 2-complexes $\tilde{K}_\# = \square(X) \to \square(\hat{X}) \to \square(X_1) = K_\#$.

To complete the proof, it suffices to take X_2 to be the 1-skeleton of the quotient Q of $\tilde{K}_\#$ by a direct product of free groups $F_2 \times F_3$, acting freely and transitively on the vertex set of $\tilde{K}$. Lemma 4.1.3 tells us that $Q = \square(X_2)$, so in particular $\pi_1(\square(X_2)) = F_2 \times F_3$. Arguing as in the previous paragraph, we see that any coverings $X \to \hat{X} \to X_2$ will extend to coverings of combinatorial 2-complexes $\tilde{K}_\# = \square(X) \to \square(\hat{X}) \to \square(X_2) = Q$. Thus if there were a finite graph $\hat{X}$ that admitted coverings $X \to \hat{X}$ and $\hat{X} \to X_1, X_2$, then $\pi_1(\square(\hat{X}))$ would be a subgroup of finite index in both $\Pi = \pi_1 K_\#$ and $F_2 \times F_3 = \pi_1 Q$. Passing to a subgroup of finite index in $\pi_1(\square(\hat{X}))$ would then yield a direct product of free groups that had finite index in Π, which is a contradiction. $\square$

4.2 Leighton's Theorem for regular coverings by quasitrees

In this section we prove Theorem 1.1.8, in fact we prove a more general symmetry-restricted version (Theorem 4.2.5).

4.2.1 Reducing to trees

As we explained in the introduction, our strategy for proving Theorem 1.1.8 is to reduce it to the case where X is a tree by means of a general construction that promotes actions on quasitrees to actions on trees. This is a variation on a result of Mosher, Sageev and Whyte [64] (see Remark 4.2.3). In order to understand the proof that we shall present, the reader should be familiar with the duality between CAT(0) cube complexes and spaces with walls (see for instance [25, 43, 66, 76, 77]).

Our use of these ideas is typical in the subject. In brief, we identify a natural notion of a *wall*, which separates our underlying space (the vertex set VX of a quasitree) into two disjoint subspaces (half-spaces); the wall is said to *separate* $x, y \in VX$ if x and y lie in opposite halfspaces. It is required that there be only finitely many walls separating each pair of points x, y. In our situation, every pair of distinct points is separated by at least one wall. An *ultrafilter* ω is a collection of half-spaces that gives a coherent choice of side across the collection of all walls: if the pair of half-spaces $\{\mathfrak{h}, \mathfrak{h}'\}$ is a wall, then exactly one of $\mathfrak{h}, \mathfrak{h}'$ lies in ω, and if $\{\mathfrak{h}_1, \mathfrak{h}_1'\}$ and $\{\mathfrak{h}_2, \mathfrak{h}_2'\}$ are walls with $\mathfrak{h}_1 \subset \mathfrak{h}_2$ and $\mathfrak{h}_1 \in \omega$ then $\mathfrak{h}_2 \in \omega$. Each $x \in X$ defines an ultrafilter ω_x that picks out the half-spaces containing x, and the vertex set of the *Sageev cube complex* dual to the wall structure is the set of ultrafilters at finite distance from these ω_x, where the distance between two ultrafilters is the number of walls for which the ultrafilters make a different choice of half-space. The natural map $x \mapsto \omega_x$ embeds VX in this vertex set. The process of completing this vertex set to a cube complex depends on the pattern of intersections of the half-spaces associated to walls and is described in detail in each of the above references.

Theorem 4.2.1. *There is a canonical process that, given a locally finite quasitree X on which $G = \mathrm{Aut}(X)$ acts cocompactly, will construct a locally finite tree T with an action of G on T and a continuous G-equivariant quasi-isometry $f : X \to T$ that restricts to an injection $VX \hookrightarrow VT$.*

Proof. The idea of the proof is to define walls on X that enable us to embed it in a CAT(0) cube complex Ψ, then use the panel-collapse procedure of Hagen and Touikan [40] to canonically retract Ψ onto a tree.

To this end, we define a constant C_X by considering all continuous quasi-isometries $h : X \to Y$ from X to locally finite trees Y that send vertices to vertices, and set

$$C_X := 3 + \inf_{h:X\to Y} \sup_{e\in EY} \mathrm{diam}(h^{-1}(e))$$

For each set of edges $S \subset EX$, we let $\hat{S}$ denote the set of midpoints, and we equip the vertex set VX with the structure of a wall space by defining $\mathcal{W}$ to be the set of all $\hat{S}$ such that $X \smallsetminus \hat{S}$ has exactly two connected components and the diameter of the union of edges in S is less than C_X. (More formally, the wall $\hat{S}$ is the partition of VX into the intersections of VX with the two connected components of $X \smallsetminus \hat{S}$.)

Given vertices x_1 and x_2, fix a shortest edge path P from x_1 to x_2, and note that $\hat{S}$ can only separate x_1 and x_2 if S includes an edge in P. Since X is locally finite and cocompact, there is a uniform bound, N say, on the number of sets $S \subset EX$ of diameter less than C_X that contain any given edge $e \in EX$. Applying this to each edge in P, we see that at most $N\, d(x_1, x_2)$ walls separate x_1 from x_2.

Let Ψ be the cube complex dual to this wall-space; its vertex set is defined in terms of ultrafilters as above.

Claim: The canonical map $\theta : VX \to \Psi$ is injective.

Proof: To verify this basic property, we must argue that each pair of distinct vertices $x_1, x_2 \in VX$ is separated by some $\hat{S} \in \mathcal{W}$. The set S of edges incident at x_1 has diameter less than C_X and separates x_1 from x_2. If $S' \subseteq S$ is a minimal subset such that $\hat{S}'$ separates x_1 from x_2, then $\hat{S}' \in \mathcal{W}$. ∎

The canonical nature of the construction ensures that the action of $G = \mathrm{Aut}(X)$ on VX extends to an action on Ψ making θ equivariant. Each $\hat{S} \in \mathcal{W}$ can intersect only a uniformly bounded number of other walls in $\mathcal{W}$, so the hyperplanes of Ψ are finite and of bounded diameter. The cocompactness of the action of G on X implies that there are only finitely many G-orbits of walls, so G acts cocompactly on Ψ.

Claim: $\theta : VX \hookrightarrow \Psi$ is a quasi-isometry with respect to the combinatorial metric on Ψ.

Proof: As G acts cocompactly on Ψ, every point is within a bounded distance of the image of θ. By definition, for $x_1, x_2 \in VX$, the combinatorial distance between $\theta(x_1)$ and $\theta(x_2)$ is the number of walls $\hat{S} \in \mathcal{W}$ that separate x_1 and x_2, so we must find upper and lower bounds for this number of walls that are linear functions of $d(x_1, x_2)$. We already have the upper bound of $N\, d(x_1, x_2)$, as explained earlier, so it remains to compute a lower bound.

To this end, we fix a tree Y and a continuous quasi-isometry $h : X \to Y$ such that $2 + \mathrm{diam}(h^{-1}(e)) < C_X$ for all $e \in EY$ and

$$\tfrac{1}{K} d(x_1, x_2) - K \leq d(h(x_1), h(x_2)) \leq K d(x_1, x_2) + K$$

for all $x_1, x_2 \in VX$, where $K > 1$ is constant. For each $e \in EY$, the union of the edges in $S(e) := \{e' \in EX \mid h(e') \cap e \neq \emptyset\}$ has diameter less than C_X (if non-empty). Moreover, if $X \smallsetminus \widehat{S(e)}$ has more than one component, then $\widehat{S'(e)} \in \mathcal{W}$ for some $S'(e) \subseteq S(e)$. There are at least $\frac{1}{K} d(x_1, x_2) - K - 2$ edges on the geodesic in Y joining $h(x_1)$ to $h(x_2)$, and for each such edge $\widehat{S(e)}$ will separate x_1 and x_2. Since $\mathrm{diam}(h(e)) \leq 2K$ for all edges $e \in EX$, the sets $S(e_1)$ and $S(e_2)$ are disjoint for $e_1, e_2 \in EY$ with $d(e_1, e_2) > 2K$. Thus we have a family of at least $(\frac{1}{K} d(x_1, x_2) - K - 2)/(2K + 2)$ disjoint sets $\widehat{S'(e)} \in \mathcal{W}$ separating x_1 from x_2. ∎

Claim: Ψ is locally finite.

Proof: The hyperplanes in Ψ are finite and each intersects only finitely many other hyperplanes, so if Ψ were not locally finite, there would be an infinite family of edges e_i incident at some vertex $v \in \Psi$ with the dual hyperplanes H_i all disjoint. Let $\hat{S}_i \in \mathcal{W}$ be the wall corresponding to H_i and fix $e'_i \in S_i$. The map θ sends the endpoints of e'_i to different sides of the hyperplane H_i; let x_i be the endpoint such that H_i separates $\theta(x_i)$ from v. There are only finitely many G-orbits of pairs $(\hat{S}, x) \in \mathcal{W} \times VX$ such that x is an endpoint of an edge in S, so there is a uniform bound on the diameter of $H_i \cup \{\theta(x_i)\}$. And since the H_i are all adjacent to v, this bounds the diameter of the set $\{\theta(x_i)\}$. But the H_i are disjoint, so the $\theta(x_i)$ are all distinct vertices. This provides us with the contradiction that we seek, because the previous claim shows that if the diameter of the set $\{\theta(x_i)\}$ is bounded then so is the diameter of $\{x_i\}$, contradicting the local finiteness of X. ■

By passing to the first cubical subdivision, we can assume that G acts on Ψ without inversions in hyperplanes. Since Ψ has finite hyperplanes and a cocompact G-action, we can apply the panel collapse procedure of Hagen and Touikan [40] to obtain a locally finite tree T with a G-equivariant embedding $T \hookrightarrow \Psi$ that is a quasi-isometry and induces a bijection between the vertex sets. (The edges of T will not in general map to edges of Ψ, they will just map to CAT(0) geodesic segments between the appropriate vertices.) The fact that $T \hookrightarrow \Psi$ is a quasi-isometry inducing a bijection between the vertex sets is not explicitly stated in the theorem of [40], but it is obvious from the proof. By composing the inverse of this bijection with $\theta : VX \hookrightarrow \Psi$ we get a G-equivariant quasi-isometry $f : VX \hookrightarrow T$, and by extending linearly along edges we obtain a continuous G-equivariant quasi-isometry $f : X \to T$ that proves the theorem. □

Remark 4.2.2. The tree T constructed above will not be a minimal G-tree in general, and the action of G on the geodesic core of T will not be faithful, even though the action of G on T is faithful. For example, given any finite group Ω, one can manufacture a quasi-line X with a cocompact action of the wreath product $G = \Omega \wr \mathbb{Z}$ by stringing together copies of the suspension of the Cayley graph of Ω in a linear fashion. The construction described in Theorem 4.2.1 will produce a G-equivariant map $X \to T$ where the geodesic core of T is a simplicial line on which Ω will act trivially. The full tree T is obtained from the core by attaching finite subtrees of uniformly bounded diameter, and it is the action of G on these finite subtrees that makes T a faithful G-tree. In the case where X is not a quasi-line, a similar phenomenon still occurs with finite subgroups $\Omega < \mathrm{Aut}(X)$ that have finite support.

Remark 4.2.3. Theorem 4.2.1 is reminiscent of the work of Mosher, Sageev and Whyte [64] who were concerned with promoting quasi-actions on bushy quasitrees to genuine actions on trees. Indeed, for quasitrees with infinitely many ends, one can craft a (less-canonical) alternative to Theorem 4.2.1 by appealing to their work, as we now describe.

Recall that a tree is termed *bushy* if it has bounded valence and each vertex is a uniformly bounded distance from a vertex having at least three unbounded complementary components. Given a cocompact quasitree X with infinitely many ends and $h : X \to Y$ a quasi-isometry to a tree, we modify Y to make it bushy as follows. (The choice of $h : X \to Y$ is what makes this argument less canonical.) Because X is cocompact, we can cover X with translates of a finite subgraph $U \subset X$ that has at least three unbounded complementary components. For suitable $R > 0$, the R-neighbourhood of $h(g.U)$ has at least three unbounded complementary components in Y. It follows that every vertex of Y is a uniformly bounded distance from one with at least three unbounded complementary components. Each such vertex lies in the core $Y' \subset Y$, which is the union of the geodesic lines. Composing h with the nearest-point retraction $Y \to Y'$, we obtain a quasi-isometry from X to a bushy tree, and [64, Theorem 1] promotes the resulting quasi-action of $G = \mathrm{Aut}(X)$ on Y' to a cocompact isometric action on a locally finite tree T. This action is coarsely G-equivariant in the sense that there is a quasi-isometry $f : X \to T$ such that $d(f(gx), gf(x))$ is uniformly bounded independent of $g \in \mathrm{Aut}(X)$ and $x \in X$.

Modifying f by a bounded map, we can assume that it sends vertices to vertices. Then, to make it G-equivariant, one decomposes VX into G-orbits, and for each orbit representative v_i one defines $f'(g.v_i)$ to be the centre of the bounded set $gG_i.f(v_i)$, where $G_i < G$ is the stabiliser of v_i. (It may be necessary to subdivide T to ensure that this centre is a vertex.) Extend f' linearly across edges to get a G-equivariant map $f' : X \to T$.

The final thing to arrange is that f' should be injective on VX. This can be done by adding various decorations. For example, one might add a leaf labelled e_x to $f'(x)$ for each $x \in VX$, and perturb f' equivariantly so that it sends x to the new vertex at the end of e_x.

4.2.2 Symmetry-restricted version

Before we can state the symmetry-restricted version of the Quasitree Leighton's Theorem, we must generalise the notion of symmetry restricted closure to arbitrary graphs. In what follows $B_R(x)$ will denote the R-ball centred on a vertex x of a graph.

Definition 4.2.4. Let X be a graph, let $H < \mathrm{Aut}(X)$, and let R be an integer. We define the *R-symmetry restricted closure* of H to be:

$$\mathscr{S}_R(H, X) := \{g \in \mathrm{Aut}(X) \mid \forall x \in VX, \exists h \in H \text{ s.t. } h \text{ agrees with } g \text{ on } B_R(x)\}$$

It is easy to check that $\mathscr{S}_R(H,X) = \bigcap_{x\in VX} H\,\mathrm{Fix}(B_R(x))$, where $\mathrm{Fix}(B_R(x))$ is the pointwise stabiliser of $B_R(x)$ in $\mathrm{Aut}(X)$. If X is locally finite then the topological group $\mathrm{Aut}(X)$ has a basis of open sets consisting of cosets of pointwise stabilisers of balls, and $\mathscr{S}_R(H,X)$ is a closed subgroup. A discrete subgroup $\Gamma < \mathrm{Aut}(X)$ is a *uniform lattice* if it acts properly and cocompactly on X.

The theorem we wish to prove is as follows. The case where X is a tree is precisely Theorem 1.1.5, and we rely on this basic case. We do not need to assume that the groups acting are free because any group that acts properly and cocompactly on a quasitree is virtually free. Theorem 1.1.8 follows from the case $H = \mathrm{Aut}(X)$ and $R = 0$ (see the discussion following (1.1.1)).

Theorem 4.2.5. *Let X be a locally finite quasitree graph, let $H < \mathrm{Aut}(X)$, and let $\Gamma_1, \Gamma_2 < H$ be uniform lattices in $\mathrm{Aut}(X)$. Then for all $R \in \mathbb{N}$ there exists $g \in \mathscr{S}_R(H,X)$ such that Γ_1^g is commensurable to Γ_2 in $\mathrm{Aut}(X)$.*

Proof. By subdividing we may assume that X is simplicial. By Theorem 4.2.1 there exists a locally finite tree T, an injective homomorphism $\iota : \mathrm{Aut}(X) \to \mathrm{Aut}(T)$ and a continuous ι-equivariant quasi-isometry $f : X \to T$ which is injective on VX. As Γ_1 and Γ_2 act properly and cocompactly on X, the groups $\iota(\Gamma_1)$ and $\iota(\Gamma_2)$ act properly and cocompactly on T, i.e. they are uniform lattices in $\mathrm{Aut}(T)$. Let M be large enough so that for all $x \in VX$

$$f(B_R(x)) \subset B_M(f(x)). \tag{4.2.1}$$

We then have a homomorphism

$$\alpha : \mathscr{S}_M(\iota(H), T) \to \mathscr{S}_R(H, X),$$

satisfying $\alpha \circ \iota = \mathrm{id}$ on restriction to $\mathscr{S}_R(H,X)$, defined as follows. Take $\phi \in \mathscr{S}_M(\iota(H),T)$. For $x \in VX$ there exists $h \in H$ such that $\iota(h)$ agrees with ϕ on $B_M(f(x))$. By ι-equivariance of f and (4.2.1), we get the following commutative diagram.

$$\begin{array}{ccc} B_R(x) & \xrightarrow{\;h\;} & X \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f} \\ B_M(f(x)) & \underset{\phi}{\xrightarrow{\;\iota(h)\;}} & T \end{array} \tag{4.2.2}$$

We define the restriction of $\alpha(\phi)$ to $B_R(x)$ to agree with h. This gives a well-defined description of an automorphism $\alpha(\phi) \in \mathrm{Aut}(X)$, because for each vertex $y \in B_R(x)$ the diagram (4.2.2) implies that $f(\alpha(\phi)(y)) = \phi f(y)$; this formula, which is independent of x and h, uniquely determines $\alpha(\phi)(y)$ since f is injective on VX. And since $\alpha(\phi)$ agrees on R-balls with elements of H, we get that $\alpha(\phi) \in \mathscr{S}_R(H,X)$. It is also clear that $\alpha \circ \iota = \mathrm{id}$ on restriction to $\mathscr{S}_R(H,X)$.

We know from Theorem 1.1.5 that the theorem is true when $X = T$, so there exists $g \in \mathscr{S}_M(\iota(H), T)$ such that $\iota(\Gamma_1)^g$ is commensurable to $\iota(\Gamma_2)$ in $\mathrm{Aut}(T)$. Therefore $\alpha \circ \iota(\Gamma_1)^{\alpha(g)} = \Gamma_1^{\alpha(g)}$ is commensurable to $\alpha \circ \iota(\Gamma_2) = \Gamma_2$ in $\mathrm{Aut}(X)$. □

As a consequence of Theorem 4.2.5, we obtain a version of Leighton's Theorem for simplicial complexes with virtually free fundamental groups.

Corollary 1.1.9. *Let X_1 and X_2 be finite simplicial complexes with virtually free fundamental groups and a common universal cover. Then X_1 and X_2 have a common finite cover.*

Proof. Let $\tilde{X}$ be the universal cover of X_1 and X_2, and let $\Gamma_1, \Gamma_2 < \mathrm{Aut}(\tilde{X})$ be the corresponding deck-groups. To obtain a common finite cover of X_1 and X_2, we must find $g \in \mathrm{Aut}(\tilde{X})$ such that Γ_1^g is commensurable to Γ_2. Since Γ_1 and Γ_2 are virtually free groups, we see that $\tilde{X}$ is a quasitree. Let Y be the quasitree graph obtained from the 1-skeleton of the barycentric subdivision of $\tilde{X}$ by attaching a leaf to each vertex of $\tilde{X}$ (to distinguish them from the new vertices of the barycentric subdivision). It is easy to check that there is a natural isomorphism $\mathrm{Aut}(Y) \cong \mathrm{Aut}(\tilde{X})$ of topological groups, and so the result follows by applying Theorem 4.2.5 to Y. □

4.3 Leighton's Theorem fails for quasitrees if one covering is irregular

In this section we prove Theorem 1.1.11. We repeat the statement below for the reader's convenience. The non-existence of $\hat{X}$ comes down to a counting argument (Section 4.3.4) that draws inspiration from the following elementary example based on [6, Example 4.12(1)].

Example 4.3.1. Let T be the simplicial tree whose vertex set VT is divided into three infinite sets V_2, V_3, V_4 such that each $v \in V_i$ has valence i, no pair of adjacent vertices lies in the same V_i, each $v \in V_2$ is adjacent to one vertex in V_3 and one in V_4, and each vertex in $V_3 \cup V_4$ is adjacent to a unique vertex in V_2. One can show that $\mathrm{Aut}(T)$ acts cocompactly on T with 3-orbits of edges and three orbits of vertices (namely V_2, V_3, V_4). But T does not cover any finite graph, for if Y were such a finite graph and n_2, n_3, n_4 were the number of vertices of valence $2, 3, 4$ respectively, then by counting neighbours of vertices of valence 2 we would have $n_2 = n_3 = n_4$, whereas by counting valence-3 neighbours of valence-4 vertices we would have $2n_3 = 3n_4$.

Theorem 1.1.11. *There exists a locally finite quasitree X, finite graphs X_1 and X_2, and covering maps $X \to X_1$ and $X \to X_2$, with $X \to X_2$ regular, for which there is no finite graph $\hat{X}$ fitting into the following diagram of covers.*

$$\begin{array}{ccc} & X & \\ & \downarrow & \\ & \hat{X} & \\ \swarrow & & \searrow \\ X_1 & & X_2 \end{array} \tag{4.3.1}$$

4.3.1 The graphs X_1, X_2 and X

We will build X by assembling infinitely many copies of a finite graph Y in a tree-like pattern. We shall imagine these copies of Y being made out of solid edges, with the edges joining them being dashed, and all graph morphisms we consider must map solid edges to solid edges and dashed to dashed. Formally speaking we subdivide the dashed edges to distinguish them from the solid edges, but the arguments (and diagrams) run more smoothly if we think of them as being dashed. We consider a certain covering map $p_1 : Y \to Y_1$, and X_1 and X_2 will be obtained from the graphs Y_1 and Y respectively by adding dashed edges in a particular way. The covering map $p_1 : Y \to Y_1$ is portrayed in the following figure, where we label vertices to encode the map.

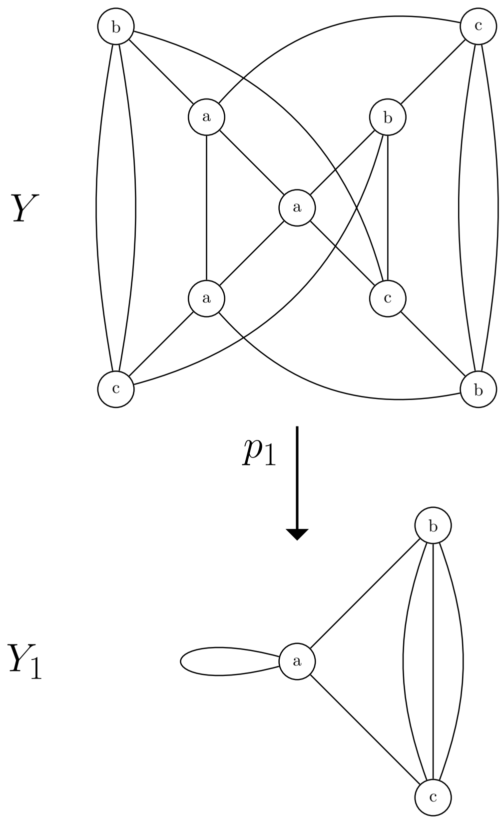

Figure 4.2: The covering $p_1 : Y \to Y_1$.

Let $\mathcal{P}$ be the collection of all covering maps $Y \to Y_1$ (not just those two that respect the labelling of Figure 4.2). For each vertex $v \in VY$ we want to determine the set $\{p(v) \mid p \in \mathcal{P}\}$. For any $p \in \mathcal{P}$, the loop of length 1 in Y_1 must lift to a collection of cycles in Y of total length 3. Since there are no 1-cycles or 2-cycles in Y, it must lift to a 3-cycle, and there are only two of those, namely the two triangles containing the central vertex, which we shall henceforth refer to as the *left triangle* and *right triangle*. Put another way, every covering map $Y \to Y_1$ maps all of the vertices of either the left or right triangle onto a. Therefore the only vertices in Y that can map to the vertex a in Y_1 are the five central vertices, and the central vertex must map to vertex a. It turns out that these are the only restrictions; in the following diagram we label the vertices $v \in Y$ with the sets $\{p(v) \mid p \in \mathcal{P}\}$.

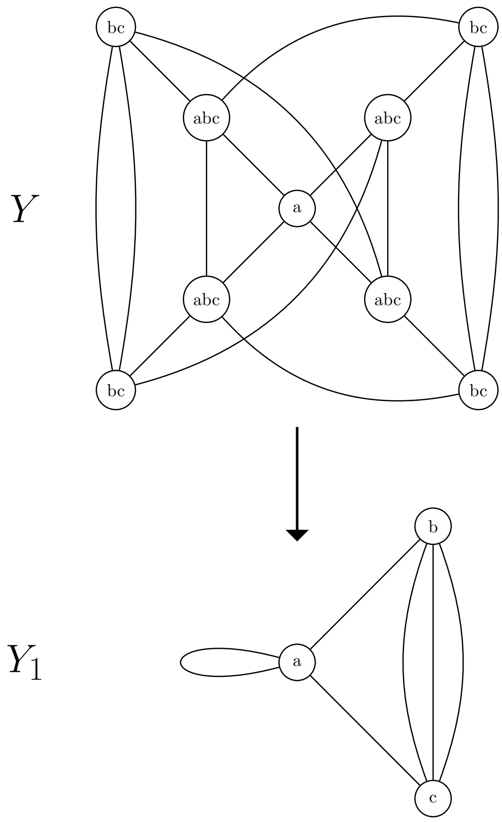

Figure 4.3: Vertices $v \in Y$ labelled by the sets $\{p(v) \mid p \in \mathcal{P}\}$.

To see that each possibility for p can be obtained, it is enough to compose p_1 with symmetries of Y: the automorphism group of Y acts transitively on the set of four outside vertices and on the set of vertices labelled *abc*.

The graphs X_1 and X_2 are obtained from Y_1 and Y by attaching dashed edges as follows (the dashed edges are blue in the figures for extra clarity). Throughout this section, we will implicitly make the identifications $VX_1 = VY_1$ and $VX_2 = VY$.

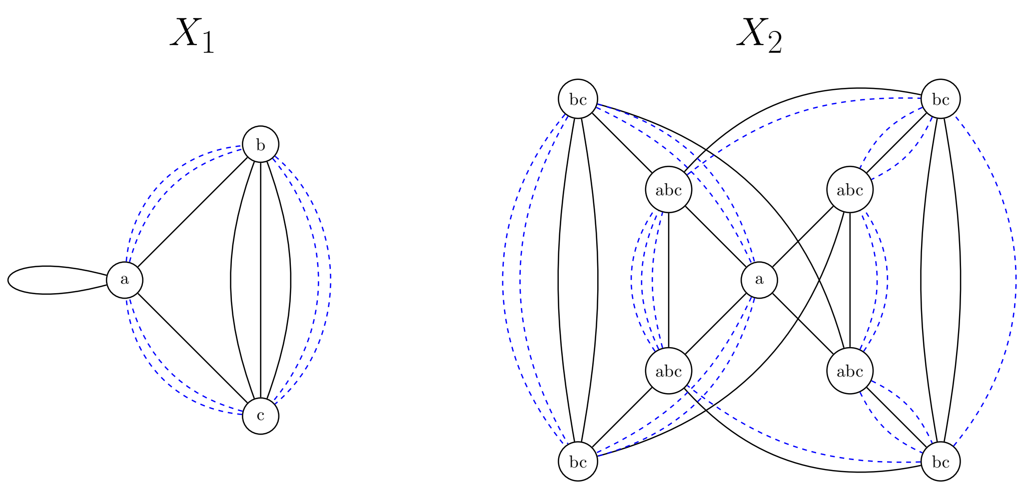

Figure 4.4: The graphs X_1 and X_2.

We define the graph X by taking the regular covering $q_2 : X \to X_2$ corresponding to the normal subgroup $\langle\langle \pi_1 Y \rangle\rangle$ of $\pi_1 X_2$; the Galois group Γ of this covering is free of rank 18. Note that $\Gamma < \mathrm{Aut}(X)$ is a uniform lattice. X is (up to isomorphism) the unique regular cover of X_2 that has no cycle containing a dashed edge and is such that each connected component of the solid subgraph is isomorphic to Y. In particular, X can be regarded as infinitely many copies of Y assembled in a tree-like template. We will use this template to build maps by working outwards from a basepoint in a radial manner.

4.3.2 Extendible edges and maps

We want to prove that there is a covering map $q_1 : X \to X_1$. This covering map will restrict to a covering $p : Y \to Y_1$ on each copy of Y in X. We have already considered limitations on where p can send vertices, and we have to work with these constraints as we determine where the dashed edges of X map: each dashed edge in X must map to a dashed edge in X_1 in such a way that we can extend continuously to coverings $Y \to Y_1$ on adjacent copies of Y. This motivates the following definition, which is made more awkward by the need to exclude an exceptional set of edges $\mathcal{E}$.

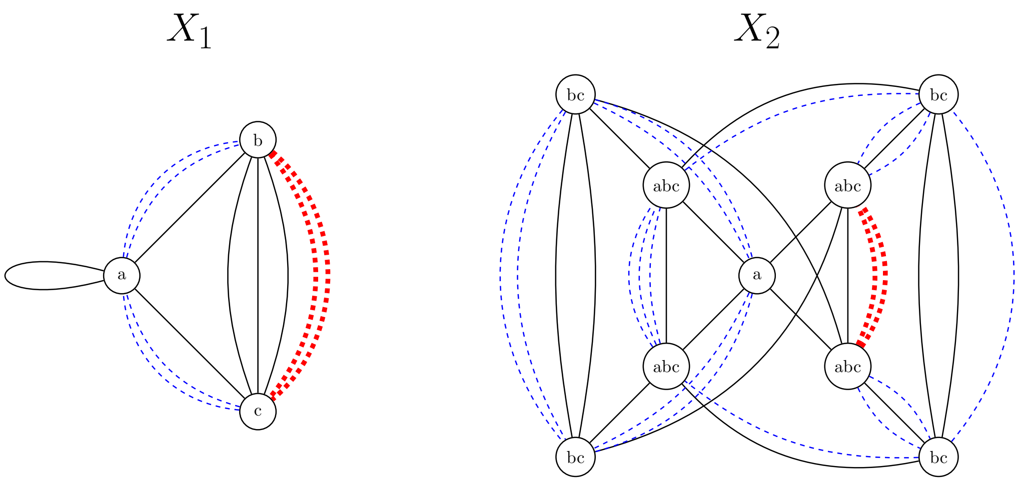

Figure 4.5: The *exceptional set* $\mathcal{E}$ consists of pairs of red thick-dashed edges (with any orientations).

We define a pair of oriented dashed edges $(e_1, e_2) \in EX_1 \times EX_2$ to be *extendible* if $(e_1, e_2) \notin \mathcal{E}$ and there exist $p, p' \in \mathcal{P}$ such that $p(\iota(e_2)) = \iota(e_1)$ and $p'(\tau(e_2)) = \tau(e_1)$. Extendible pairs can easily be read off from Figure 4.5, because the condition just requires that the label of $\iota(e_1)$ is an element of the set of labels of $\iota(e_2)$ and that the label of $\tau(e_1)$ is an element of the set of labels at $\tau(e_2)$. Note that (e_1, e_2) is extendible if and only if the pair with reversed orientations $(\bar{e}_1, \bar{e}_2)$ is extendible.

Given a vertex v let $\mathrm{lk_d}(v)$ denote the set of oriented dashed edges with terminal vertex v. If $v \in VX_2$ and $p \in \mathcal{P}$, we say that a bijection $\sigma : \mathrm{lk_d}(v) \to \mathrm{lk_d}(p(v))$ is *extendible* if every pair $(\sigma(e), e)$ is extendible.

Lemma 4.3.2. *(1) For any $p \in \mathcal{P}$ and vertex $v \in VX_2$, there exists an extendible bijection $\sigma : \mathrm{lk_d}(v) \to \mathrm{lk_d}(p(v))$.*

(2) For any $p \in \mathcal{P}$, vertex $v \in VX_2$ and extendible pair $(e_1, e_2) \in \mathrm{lk_d}(p(v)) \times \mathrm{lk_d}(v)$, there exists an extendible bijection $\sigma : \mathrm{lk_d}(v) \to \mathrm{lk_d}(p(v))$ such that $\sigma(e_2) = e_1$.

Proof. To prove these claims one must consider every pair of vertices $(p(v), v)$ with $v \in VX_2$ and $p \in \mathcal{P}$, and inspect the labellings on the dashed edges in the links of these vertices. We will do this explicitly for the particular pair of vertices highlighted in Figure 4.6 below; the arguments for other pairs are similar, but some remarks will be needed concerning the exceptional set $\mathcal{E}$. We have coloured various edges in Figure 4.6 so that we can refer to them accurately without getting overrun by notation.

To begin, note that each yellow edge in $\mathrm{lk_d}(v)$ ends at a vertex labelled abc, so an extendible σ can map these edges to either the green or orange edges. The red edge goes to a vertex labelled bc, so to be extendible σ must map this red edge to a green one; these constraints can obviously be satisfied, so claim (1), the existence of extendible maps σ, is established in this case. Assertion (2) is also clear in this case: if we specify that the red edge maps to one of the green edges then any extension of σ to the yellow edges will be extendible, whilst if we specify where one of the yellow edges goes then this leaves at least one green edge available for the red edge to map to; so in either case, we can extend the initial assignment to an extendible σ.

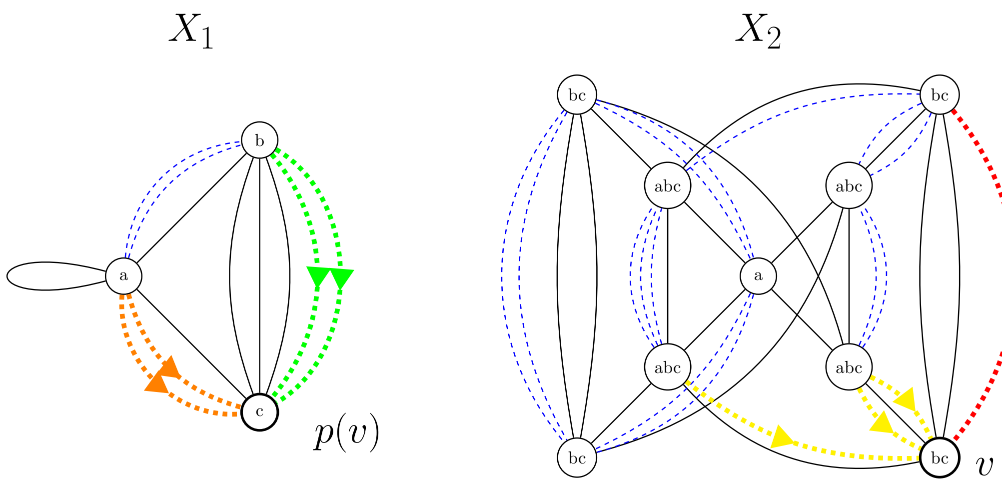

Figure 4.6: Analysis of maps $\sigma : \mathrm{lk_d}(v) \to \mathrm{lk_d}(p(v))$ for a particular pair $(p(v), v)$.

For the other pairs $(p(v), v)$ the proof is equally straightforward, having excluded the troublesome edges from Figure 4.5. If v and $p(v)$ are vertices incident at the thick-dashed red edges, then any pair of these red edges would be extendible if we had not explicitly excluded the set $\mathcal{E}$, but (2) would fail because there is no extendible bijection

$\sigma : \mathrm{lk_d}(v) \to \mathrm{lk_d}(p(v))$ mapping a red edge to a red edge, since then one of the blue edges in $\mathrm{lk_d}(v)$ would have to map to a blue edge in $\mathrm{lk_d}(p(v))$. We removed this problem by fiat. □

4.3.3 The covering $X \to X_1$

For Z a subgraph of X, we say that an immersion (locally injective graph morphism) $q : Z \to X_1$ is *extendible* if the pair $(q(e), q_2(e))$ is extendible for every dashed edge e in Z. We construct the cover $q_1 : X \to X_1$ by working outwards from a base copy of Y inductively using the following lemma. Note that defining q_1 on the base copy of Y and the dashed edges that meet it also follows from the lemma by setting $Z = \emptyset$.

Lemma 4.3.3. *Suppose Z is a finite connected subgraph of X consisting of a number of (solid) copies of Y and all of the dashed edges that meet them, and suppose $q : Z \to X_1$ is extendible. If $Y' \subset X$ is a copy of Y adjacent to Z, then we can extend q to an extendible map $q' : Z' \to X_1$, where Z' is the union of Z with Y' and any dashed edges that meet it.*

Proof. Y' is identified with Y via the covering $q_2 : X \to X_2$. We are forced to put $q' = q$ on Z, so it remains to define q on Y' and the dashed edges that meet it. We do this by setting $q' = pq_2$ on Y' for some $p \in \mathcal{P}$, and for each $v \in VY'$ we set $q' = \sigma_v q_2$ on $\mathrm{lk_d}(v)$ for some extendible bijection $\sigma_v : \mathrm{lk_d}(q_2(v)) \to \mathrm{lk_d}(pq_2(v))$. The fact that the σ_v are extendible ensures that q' is extendible. However, q' must agree with q on the unique dashed edge $\tilde{e} \in EZ$ with $\tilde{v} := \tau(\tilde{e}) \in VY'$. This gives us the following two constraints:

(1) $pq_2(\tilde{v}) = q(\tilde{v})$

(2) $\sigma_{\tilde{v}} q_2(\tilde{e}) = q(\tilde{e})$

The idea is then to meet these constraints by using the fact that q is extendible. Indeed, extendibility of q implies that the pair $(q(\tilde{e}), q_2(\tilde{e}))$ is extendible, so there exists $p \in \mathcal{P}$ satisfying (1). And Lemma 4.3.2(2) ensures that there exists an extendible bijection $\sigma_{\tilde{v}} : \mathrm{lk_d}(q_2(\tilde{v})) \to \mathrm{lk_d}(pq_2(\tilde{v}))$ satisfying (2). Finally, the existence of extendible bijections σ_v for the other vertices v follows from Lemma 4.3.2(1). □

4.3.4 Final step in the proof of Theorem 1.1.11

To complete the proof of the theorem, we must prove that there does not exist a finite graph $\hat{X}$ fitting into the diagram (4.3.1). Let $\hat{p}_i : \hat{X} \to X_i$ denote the coverings of X_1 and X_2. Each connected component $\hat{Y}$ of the solid subgraph of $\hat{X}$ would fit into the following diagram of covers:

$$\begin{array}{ccc} & Y & \\ & \downarrow & \\ & \hat{Y} & \\ {\scriptstyle \hat{p}_1} \swarrow & & \searrow {\scriptstyle \hat{p}_2} \\ Y_1 & & Y \end{array} \tag{4.3.2}$$

From this it is immediate that $\hat{Y} \cong Y$, so $\hat{X}$ is a finite graph of copies of Y joined together by dashed edges. We call the copies of Y in $\hat{X}$ (i.e. the connected components of the solid subgraph) *pieces*. Each piece is equipped with a covering $p = \hat{p}_1 \circ \hat{p}_2^{-1} : Y \to Y_1$ induced from the above diagram. Recall that Y contains two triangles, called the left and right triangles (highlighted in bold below). We saw from the discussion after Figure 4.2 that p will map either the left or the right triangle three times around the edge loop based at the vertex a in Y_1; we call the piece a *left piece* or a *right piece* correspondingly.

For any dashed edge e in $\hat{X}$, the existence of the maps $\hat{p}_1$ and $\hat{p}_2$ on the adjacent pieces tells us that the pair $(\hat{p}_1(e), \hat{p}_2(e))$ is extendible. The way we will arrive at a contradiction is to use lifts of the dashed edges highlighted green and red below to compute two different ratios between the number of left and right pieces in $\hat{X}$ (cf. Example 4.3.1). Let A_g [resp. A_r] denote the set of lifts of the green [resp. red] edges to $\hat{X}$ that map down to one of the orange edges in X_1.

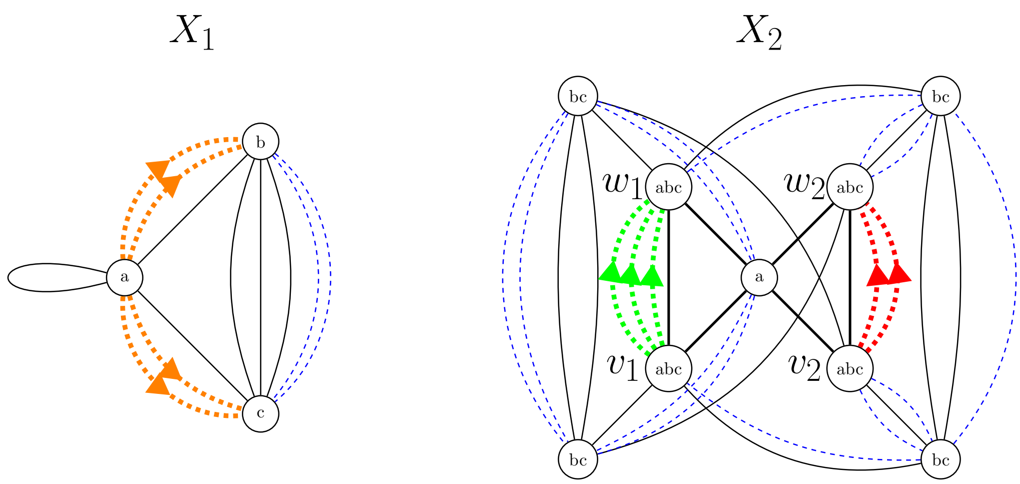

Figure 4.7: Counting left and right pieces.

Each lift of v_1 to a left piece in $\hat{X}$ will map down to a in X_1, so it has all three of its outgoing green edges in A_g. Each lift of w_1 to a right piece in $\hat{X}$ will map down to either b or c in X_1, and exactly two of its incoming green edges will be in A_g since its other incident dashed edge must map to an edge in X_1 connecting b to c. Each edge in A_g has its initial vertex (the v_1 end) in a left piece and its terminus (the w_1 end) in a right piece, so we deduce the following equation.

$$3(\#\text{ left pieces}) = |A_g| = 2(\#\text{ right pieces})$$

Meanwhile, each lift of v_2 to a right piece in $\hat{X}$ will map down to a in X_1, so both its outgoing red edges will be in A_r. And each lift of w_2 to a left piece will map down to either b or c in X_1, and both its incoming red edges will be in A_r since its other incident dashed edges must map to edges connecting b to c in X_1. Each edge in A_r has its origin (the v_2 end) in a right piece and its terminus (the w_2 end) in a left piece, so we deduce the following equation.

$$2(\#\text{ left pieces}) = |A_r| = 2(\#\text{ right pieces})$$

These two equations are inconsistent, so we get our desired contradiction, proving that $\hat{X}$ does not exist. □

Remark 4.3.4. We believe that our example for Theorem 1.1.11 is probably minimal among examples where X_2 decomposes into solid and dashed edges, with the solid edges forming a graph Y, and where X is the regular cover of X_2 corresponding to the normal subgroup $\langle\langle \pi_1 Y \rangle\rangle$ of $\pi_1 X_2$. A clue as to why it might be minimal is the fact that the covering $Y \to Y_1$ that we use is the smallest example of a non-regular covering between non-vertex-transitive graphs. We do not give a proof that both these properties on $Y \to Y_1$ are necessary, but we can share an insight as to why choosing $Y = Y_1$ will not work. Indeed, if $Y = Y_1$ then we could take the cone

$$Y \times [0,1]/[(y_1, 0) \sim (y_2, 0)],$$

and glue it to X_1 by identifying $Y \subset X_1$ with $Y \times \{1\}$ to produce a simplicial complex Z_1; similarly we could glue the cone to each copy of Y that appears in X, to produce a simply connected simplicial complex Z. The covering $X \to X_1$ would then extend to a covering $Z \to Z_1$, and this would be a regular covering since Z is simply connected. We could then deduce that $X \to X_1$ was a regular covering, and find a finite graph $\hat{X}$ fitting into diagram (4.3.1) by applying Theorem 1.1.8.

Chapter 5

Imitator covers of special cube complexes

In this chapter we describe imitator covers of special cube complexes, which provides a new interpretation of Haglund and Wise's canonical completion and retraction, and leads to powerful generalisations. We use this to prove various results about finite covers of special cube complexes (including Theorems 1.2.3, 1.2.7 and 1.2.8) - most of which generalise existing theorems of Haglund and Wise to the non-hyperbolic setting.

5.1 A single imitator

In this section we define the basic version of imitator covers, in which there is only one imitator. Firstly, we must describe what the walker and imitator are.

5.1.1 Walker and imitator

Construction 5.1.1. (Walker and imitator)
Let $\phi : Y \to X$ be a local isometry of directly special cube complexes. We consider two people wandering around the 1-skeleta of X and Y: the *walker* wanders around X^1 while the *imitator* wanders around Y^1. The walker wanders freely around X^1, while the imitator tries to "copy" the walker in the following way: if the imitator and walker are at vertices $(y, x) \in Y^0 \times X^0$ and the walker traverses the edge $e \in \text{link}(x)$, then the imitator traverses the edge $f \in \text{link}(y)$ with $\phi(f) \parallel e$, if such an edge exists, otherwise they remain at y. Note that the edge f, if it exists, will be unique because $H(e)$ doesn't self-osculate at $\phi(y)$. If the walker immediately returns along e, then the imitator returns along f in the case that f exists, otherwise they remain at y again, so the whole process is reversible.

Iterating the process, if the walker travels along a path γ, starting at x, then the imitator travels along a path that we denote $\delta = \delta(\gamma, y)$, starting at y. Since the process is reversible, if the walker immediately backtracks along the path γ then the imitator will backtrack along the path δ.

Lemma 5.1.2. *If γ is a path going once round the boundary of a square in X, starting and finishing at $x \in X^0$, then $\delta(\gamma, y)$ closes up as a loop for any $y \in Y^0$.*

Proof. We have the imitator and walker starting at vertices (y, x), the walker travels along γ, and the imitator travels along $\delta = \delta(\gamma, y)$. Say that γ consists of edges e_1, e_2, e_1', e_2', so $H_1 := H(e_1) = H(e_1')$ and $H_2 := H(e_2) = H(e_2')$. Note that H_1 and H_2 intersect at $(x; e_1, e_2')$, so $H_1 \neq H_2$. The proof splits into four cases:

- If no $f \in \text{link}(y)$ has $\phi(f)$ parallel to e_1 or e_2 then δ is the constant path at y.

- Suppose $f_1 \in \text{link}(y)$ has $\phi(f_1) \parallel e_1$, but no $f_2 \in \text{link}(y)$ has $\phi(f_2) \parallel e_2$. Then the imitator traverses f_1 as the walker traverses e_1. Let $y' \in Y^0$ be the vertex at the other end of f_1.

  We now argue that no $f_2 \in \text{link}(y')$ has $\phi(f_2) \parallel e_2$. Indeed if such f_2 did exist then H_1 and H_2 would intersect at $(\phi(y'); \phi(f_1), \phi(f_2))$ (they can't osculate here because they already intersect at $(x; e_1, e_2')$, and inter-osculating is forbidden), and so $\phi(f_1), \phi(f_2)$ would form the corner of a square at $\phi(y')$; but then local injectivity of ϕ implies f_1, f_2 would form the corner of a square S at y', and the edge $f_2' \in \text{link}(y)$ opposite f_2 in S would have $\phi(f_2') \parallel e_2$, a contradiction.

  As a result, we see that the imitator stays at y' as the walker traverses e_2, returns along f_1 as the walker traverses e_1', and stays at y as the walker traverses e_2'. So ultimately the imitator returns to y as required.

- If no $f_1 \in \text{link}(y)$ has $\phi(f_1) \parallel e_1$ but some $f_2 \in \text{link}(y)$ has $\phi(f_2) \parallel e_2$, then we can deduce that δ goes back and forth along f_2 as in the previous case.

- Suppose $f_1 \in \text{link}(y)$ has $\phi(f_1) \parallel e_1$ and $f_2 \in \text{link}(y)$ has $\phi(f_2) \parallel e_2$. Then H_1 and H_2 intersect at $(\phi(y); \phi(f_1), \phi(f_2))$, and local injectivity of ϕ implies that f_1, f_2 form the corner of a square S at y. It follows that the imitator circumnavigates S, returning to y, as the walker travels the length of γ.

□

Remark 5.1.3. Lemma 5.1.2 serves as the basic homotopy move that will allow us to lift homotopies of paths in X to homotopies of paths in the imitator cover - which we now go on to define.

5.1.2 Imitator covers

The following construction appears as the canonical completion and retraction in [44, Proposition 6.5]. Formally speaking our version is exactly the same, the difference is that we describe it using the cartoon picture of walker and imitator from Construction 5.1.1. This not only helps in understanding the construction, but also facilitates the generalisations which we make in subsequent sections. Note that the canonical completion is defined slightly differently elsewhere in the literature, such as in [98], by using a certain fibre product over a Salvetti complex of a RAAG - but the two constructions coincide for local isometries of directly special cube complexes.

Construction 5.1.4. (Imitator cover)
Given the setup of Construction 5.1.1, define a cube complex $\mathsf{C}(Y, X)$ as follows (same notation as [44]). The 0-skeleton $\mathsf{C}(Y, X)^0 := Y^0 \times X^0$ is the set of possible positions of imitator and walker on vertices. The 1-skeleton describes all possible movements of walker and imitator: given $(y, x) \in \mathsf{C}(Y, X)^0$ and $e \in \text{link}(x)$, suppose the walker traverses e from x to x' while the imitator moves from y to y' (either traversing an edge f from y to y' or staying put at $y = y'$), then $\mathsf{C}(Y, X)$ has an edge $\mathsf{e}(e; y, x)$ joining (y, x) to (y', x'). Running this in reverse, if the walker traverses e from x' to x then the imitator would move from y' to y, and this defines the edge $\mathsf{e}(e; y', x')$ of $\mathsf{C}(Y, X)$, but in this case we define $\mathsf{e}(e; y, x) = \mathsf{e}(e; y', x')$ (one could think of these as two orientations of the same edge, but as we only work with 1-cells in this thesis we will just consider them to be the same edge).

By projecting to the positions of imitator and walker, we have combinatorial maps $\mathsf{C}(Y, X)^1 \to Y^1$ and $\mathsf{C}(Y, X)^1 \to X^1$, and the map $\mathsf{C}(Y, X)^1 \to X^1$ is clearly a covering. By Lemma 5.1.2, the boundary of each square in X lifts to a collection of 4-cycles in $\mathsf{C}(Y, X)^1$, so we can form $\mathsf{C}(Y, X)^2$ by attaching squares with boundary maps equal to these 4-cycles. This gives a covering $\mathsf{C}(Y, X)^2 \to X^2$, which can then be completed to a covering $\mathsf{C}(Y, X) \to X$ by adding higher dimensional cubes. Moreover, this will be a finite-sheeted covering if Y is finite.

The map $\mathsf{C}(Y, X)^1 \to Y^1$ extends to a cellular map $r : \mathsf{C}(Y, X) \to Y$, called the *canonical retraction*, as we now describe. Suppose C is an n-cube in $\mathsf{C}(Y, X)$ with corner (y, x), and say the edges of C incident to (y, x) project to edges $e_1, ..., e_n \in \text{link}(x)$. Suppose $f_1, ..., f_m \in \text{link}(y)$ satisfy $\phi(f_i) \parallel e_i$, and suppose that no $f \in \text{link}(y)$ has $\phi(f) \parallel e_i$ for $m < i \leq n$. Then by arguments similar to those in Lemma 5.1.2, we know that $f_1, ..., f_m$ form the corner of a cube C_Y in Y, and the projection $\mathsf{C}(Y, X)^1 \to Y$ maps $C^1 \to C_Y^1$, sending $\mathsf{e}(e_i; y, x) \mapsto f_i$ for $1 \leq i \leq m$ and $\mathsf{e}(e_i; y, x) \mapsto y$ for $m < i \leq n$. So this extends to a map $C \to C_Y$ that collapses the dimensions of C corresponding to $m < i \leq n$. These maps then fit together to give a cellular map $r : \mathsf{C}(Y, X) \to Y$.

The map $Y^0 \hookrightarrow \mathsf{C}(Y,X), y \mapsto (y, \phi(y))$ extends to a combinatorial embedding $j : Y \hookrightarrow \mathsf{C}(Y,X)$ called the *canonical completion of* $\phi : Y \to X$. Note that the composition $r \circ j$ is the identity on Y, and we also have a commutative diagram:

$$\begin{array}{ccc} Y & \xrightarrow{\ j\ } & \mathsf{C}(Y,X) \\ & {\scriptstyle\phi}\searrow & \downarrow \\ & & X \end{array} \tag{5.1.1}$$

In general $\mathsf{C}(Y,X)$ will not be connected, so we will usually choose a component. We let $\mathsf{C}(Y,X;y,x)$ denote the component of $\mathsf{C}(Y,X)$ containing (y,x), and we call it the *imitator cover of* $\phi : Y \to X$ *based at* (y,x). This encapsulates the data of all paths the imitator and walker can take starting at the position (y,x).

5.2 Applications of a single imitator

In this section we apply the construction of imitator covers to prove various propositions and lemmas, some new and others already known, culminating with Proposition 5.2.10 about separability of triple cosets.

5.2.1 Subgroup separability

In [44] it was shown that fundamental groups of finite special cube complexes are virtually subgroups of right angled Coxeter groups, hence they are $\mathbb{Z}$-linear and residually finite. In fact residual finiteness can be deduced rather more directly using imitator covers.

Proposition 5.2.1. *Fundamental groups of finite virtually special cube complexes are residually finite.*

Proof. Residual finiteness is a commensurability invariant, so it suffices to work with a finite directly special cube complex X. Let $x \in X^0$ be a basepoint, and let γ be an essential loop based at x. Our task is to find a finite based cover of X such that γ lifts to a path with distinct endpoints. Then the subgroup $K < \pi_1 X$ corresponding to the cover has finite index and $\gamma \notin K$.

Let $(\tilde{X}, \tilde{x}) \to (X, x)$ be the based universal cover of X and let $\tilde{\gamma}$ be the based lift of γ. Let $Y_{\tilde{\gamma}} \subset \tilde{X}$ be the smallest convex subcomplex of $\tilde{X}$ containing $\tilde{\gamma}$. Note that $Y_{\tilde{\gamma}}$ is obtained by intersecting all halfspaces in $\tilde{X}$ containing $\tilde{\gamma}$, and that a hyperplane of $\tilde{X}$ intersects $Y_{\tilde{\gamma}}$ if and only if it intersects $\tilde{\gamma}$ - there are only finitely many such hyperplanes, so $Y_{\tilde{\gamma}}$ is finite. The restriction of the covering to $Y_{\tilde{\gamma}}$ is a local isometry $Y_{\tilde{\gamma}} \to X$, so we can form the imitator cover $\mathsf{C}(Y_{\tilde{\gamma}}, X; \tilde{x}, x) \to X$. This is a finite cover, and γ has a lift based at $(\tilde{x}, x)$ with distinct endpoints (because the projection of this lift to $Y_{\tilde{\gamma}}$ is $\tilde{\gamma}$). $\square$

Another important consequence of imitator covers is subgroup separability, which appears as [44, Corollary 7.9].

Proposition 5.2.2. *Let X be a finite virtually special cube complex. Then any convex subgroup of $\pi_1 X$ is separable.*

Proof. By Lemma 2.4.27 we may consider a convex subgroup defined by a based local isometry $\phi : (Y, y) \to (X, x)$. Take a finite directly special cover $(\hat{X}, \hat{x}) \to (X, x)$ and a based elevation $(\hat{Y}, \hat{y}) \to (\hat{X}, \hat{x})$ of ϕ. $\pi_1(\hat{Y}, \hat{y})$ is a finite union of cosets of $\pi_1(Y, y)$ (viewing both as subgroups of $\pi_1(X, x)$), and separability is equivalent to being closed in the profinite topology (Definition 2.5.4), so it suffices to prove that $\pi_1(\hat{Y}, \hat{y})$ is separable.

Now consider the canonical completion $\mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x})$. We have the inclusion $\hat{Y} \hookrightarrow \mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x})$ inducing a group inclusion $\pi_1(\hat{Y}, \hat{y}) \hookrightarrow \pi_1(\mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x}), (\hat{y}, \hat{x}))$, and the canonical retraction

$$r : \mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x}) \to (\hat{Y}, \hat{y})$$

inducing a group retraction

$$\pi_1(\mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x}), (\hat{y}, \hat{x})) \to \pi_1(\hat{Y}, \hat{y}).$$

These groups are residually finite by Proposition 5.2.1, so it follows from Lemma 5.2.3 below that $\pi_1(\hat{Y}, \hat{y})$ is separable in $\pi_1(\mathsf{C}(\hat{Y}, \hat{X}; \hat{y}, \hat{x}), (\hat{y}, \hat{x}))$, and hence also in $\pi_1(X, x)$ (finite-index subgroups are always closed in the profinite topology). □

The following lemma is proved for instance in [50, Lemma 3.9], our proof is slightly different and leads to generalisations with more subgroups (Lemmas 5.2.9 and 5.2.11).

Lemma 5.2.3. *Let $\rho : G \twoheadrightarrow K < G$ be a retraction of a residually finite group G. Then K is separable in G.*

Proof. Let $g \in G - K$. The goal is to find a finite quotient of G that separates g from K. We know that $\rho(g) \in K$, so $g \neq \rho(g)$, and by residual finiteness of G there exists a finite quotient $q : G \to \bar{G}$ with

$$q(g) \neq q\rho(g). \tag{5.2.1}$$

Now consider another finite quotient of G, defined by

$$t : G \to \bar{G} \times \bar{G}$$
$$h \mapsto (q(h), q\rho(h)).$$

We claim that $t(g) \notin t(K)$, so suppose for contradiction that $t(g) = t(k)$ with $k \in K$. Then

$$(q(g), q\rho(g)) = (q(k), q\rho(k)). \tag{5.2.2}$$

But $\rho(k) = k$, so (5.2.2) implies that $q(g) = q\rho(g)$, contradicting (5.2.1). □

5.2.2 Entrapment

The following will be a key lemma used throughout the chapter.

Lemma 5.2.4. *(Subcomplex Entrapment)*
Let $\phi : Y \to X$ be a local isometry of directly special cube complexes, and let $Z \subset X$ be a locally convex subcomplex that does not inter-osculate with hyperplanes of X. If we consider an imitator in Y and a walker in X starting at positions $(y, x) \in \phi^{-1}(Z) \times Z$, then the imitator stays inside $\phi^{-1}(Z)$ as long as the walker stays inside Z. In particular, if Y is a locally convex subcomplex of X, then the imitator stays inside $Y \cap Z$.

Proof. It suffices to consider the walker traversing a single edge $e \in \operatorname{link}(x)$ inside Z. If the imitator stays at y then we are done, otherwise they traverse an edge f with $\phi(f)||e$. Suppose for contradiction that f is not in $\phi^{-1}(Z)$. Then $H(\phi(f)) = H(e)$ osculates with Z at $(\phi(y); \phi(f))$. But $H(e)$ and Z also intersect because e is in Z, hence they inter-osculate, contrary to the assumption of the lemma. □

The following special case of subcomplex entrapment allows us to trap imitators inside hyperplane carriers.

Lemma 5.2.5. *(Hyperplane Entrapment)*
Let $\phi : Y \to X$ be a local isometry of directly special cube complexes. Let H be a hyperplane in Y, and suppose that the imitator and walker start at positions $(y, x) \in N(H) \times N(\phi[H])$. If the walker stays inside $N(\phi[H])$ then the imitator stays inside $N(H)$.

Proof. We apply Subcomplex Entrapment with $Z = N(\phi[H])$. This ensures that the imitator stays inside $\phi^{-1}(N(\phi[H]))$ as long as the walker stays inside $N(\phi[H])$. But we know that

$$\phi^{-1}(N(\phi[H])) = \bigcup_{\phi[H']=\phi[H]} N(H'),$$

and this is a disjoint union because any vertex in an intersection of hyperplane carriers would map to a point of self-intersection or self-osculation of $\phi[H]$. Hence the imitator remains in $N(H)$. □

As an application of Hyperplane Entrapment we have the following proposition. Items (1) and (2) of the proposition are contained in [44, Proposition 6.5] but item (3) is new.

Proposition 5.2.6. *Let $\phi : Y \to X$ be a local isometry of directly special cube complexes and let $j : Y \hookrightarrow \mathsf{C}(Y, X)$ be the canonical completion. Then we have the following:*

(1) If H_1 and H_2 are distinct hyperplanes in Y, then $j[H_1]$ and $j[H_2]$ are distinct hyperplanes in $\mathsf{C}(Y, X)$.

(2) If H_1 and H_2 are non-intersecting hyperplanes in Y, then $j[H_1]$ and $j[H_2]$ are non-intersecting hyperplanes in $\mathsf{C}(Y, X)$.

(3) No hyperplane in $\mathsf{C}(Y, X)$ inter-osculates with Y.

Proof. (1) Let H_1 and H_2 be distinct hyperplanes in Y, and suppose for contradiction that $j[H_1] = j[H_2] = \dot{H}$. Let $y_i \in N(H_i)$ be vertices, and let $\dot{\gamma}$ be a path in $N(\dot{H})$ from $j(y_1)$ to $j(y_2)$. By diagram (5.1.1), the hyperplane $\dot{H}$ projects to the hyperplane $H := \phi[H_1] = \phi[H_2]$ in X, and the path $\dot{\gamma}$ descends to a path γ in $N(H)$ from $x_1 := \phi(y_1)$ to $x_2 := \phi(y_2)$. By construction, if imitator and walker start at positions (y_1, x_1), and the walker travels along γ, then the imitator travels along $\delta = \delta(\gamma, y_1)$ to y_2. But the imitator must remain inside $N(H_1)$ by Hyperplane Entrapment, so $y_2 \in N(H_1) \cap N(H_2)$. Hence the hyperplanes H_1 and H_2 either intersect or osculate at y_2; this projects to a self-intersection or self-osculation of H, contradicting the direct specialness of X.

(2) Let H_1 and H_2 be non-intersecting hyperplanes in Y, and suppose for contradiction that the $\dot{H}_i := j[H_i]$ intersect. $\dot{H}_1$ and $\dot{H}_2$ must be distinct by (1), so they must intersect at a vertex $\dot{x} = (y, x) \in \mathsf{C}(Y, X)^0$. For $i = 1, 2$, let $y_i \in N(H_i)$ be a vertex, and let $\dot{\gamma}_i$ be paths in $N(\dot{H}_i)$ from $j(y_i)$ to $\dot{x}$. The hyperplanes $\dot{H}_i$ project to distinct intersecting hyperplanes $\phi[H_i]$ in X, and the paths $\dot{\gamma}_i$ project to paths γ_i in $N(\phi[H_i])$ from $x_i := \phi(y_i)$ to x. By construction, if imitator and walker start at positions (y_i, x_i), and the walker travels along γ_i, then the imitator travels to y. Hyperplane Entrapment implies that $y \in N(H_1) \cap N(H_2)$, hence H_1 and H_2 osculate at y. But this projects to an osculation of $\phi[H_1]$ and $\phi[H_2]$ in X, hence these hyperplanes inter-osculate, a contradiction.

(3) Suppose for contradiction that H is a hyperplane in Y such that $\dot{H} := j[H]$ osculates with Y at $(y; \dot{e}) \in Y^0 \times \mathsf{C}(Y, X)^1$, where $\dot{e}$ is incident at $j(y) = (y, \phi(y)) = (y, x) \in \mathsf{C}(Y, X)^0$ and projects to $e \in X^1$. Let $y' \in N(H)$ be a vertex and $\dot{\gamma}$ a path in $N(\dot{H})$ from $j(y') = (y', x')$ to (y, x). This projects to a path γ in $N(\phi[H])$ from x' to x. By construction, if imitator and walker start at positions (y', x'), and the walker travels along γ, then the imitator travels to y. Hyperplane Entrapment implies that $y \in N(H)$, so there is an edge $f \in \text{link}(y)$ dual to H. Since $\dot{H}$ osculates with Y at $(y; \dot{e})$, we know that $j(f) \neq \dot{e}$, so projecting to X we get distinct edges $\phi(f), e \in \text{link}(x)$ that are both dual to $\phi[H]$. Therefore $\phi[H]$ either self-intersects or self-osculates at x, a contradiction.

□

We get the following corollary. This will be used throughout the chapter to get our hands on subcomplexes that do not inter-osculate with hyperplanes, which then allows us to use Subcomplex Entrapment.

Corollary 5.2.7. *Let $Y_1, ..., Y_n \to X$ be local isometries of finite virtually special cube complexes. Then there is a finite directly special regular cover $\hat{X} \to X$ such that all elevations of the Y_i to $\hat{X}$ are embedded and do not inter-osculate with hyperplanes of $\hat{X}$.*

Proof. Firstly, we may assume that X is already directly special, as otherwise we can pass to a finite directly special cover of X and replace $Y_1, ..., Y_n$ with all their elevations to this cover. Next, apply Proposition 5.2.6 to each Y_i in turn to obtain finite covers $X_i \to X$, each containing an elevation of Y_i which is embedded and does not inter-osculate with hyperplanes. The property of an elevation being embedded and not inter-osculating with hyperplanes passes to further covers. The desired cover $\hat{X} \to X$ is any finite regular cover of X that factors through all the covers $X_i \to X$. The regularity of the cover means that one elevation of Y_i will have the desired property if and only if all elevations of Y_i do, so we are done. □

5.2.3 Double coset separability

Having proven that convex subgroups are separable in Proposition 5.2.2, we now turn our attention to double cosets of convex subgroups. This Proposition was proven by Oregón-Reyes in [67, Theorem A.1]; the proof below is similar but recast into the language of imitator covers.

Proposition 5.2.8. *Let X be a finite virtually special cube complex, and let $K_1, K_2 < G := \pi_1 X$ be convex subgroups. Then for any $g \in G$ the double coset $K_1 g K_2$ is separable in G.*

Proof. As separability is equivalent to being closed in the profinite topology, left and right translates of separable subsets are still separable. So it is enough to prove that $g^{-1}K_1 g K_2$ is separable. But conjugates of convex subgroups are still convex, so we may assume from the start that $g = 1$.

Choose a basepoint $x \in X$ and write $G = \pi_1(X, x)$. By Lemma 2.4.27 there are based local isometries of finite cube complexes $(Y_i, y_i) \to (X, x)$ inducing the subgroups K_i. By Corollary 5.2.7, there is a finite directly special cover $(\hat{X}, \hat{x}) \to (X, x)$ such that the based elevations $(\hat{Y}_i, \hat{y}_i) \hookrightarrow (\hat{X}, \hat{x})$ of $(Y_i, y_i) \to (X, x)$ are embedded and do not inter-osculate with hyperplanes of $\hat{X}$. We will consider the $\hat{Y}_i$ as subcomplexes of $\hat{X}$ with $\hat{y}_i = \hat{x}$. If $\hat{G} < G$ is the subgroup corresponding to $(\hat{X}, \hat{x}) \to (X, x)$ then $\hat{K}_i := K_i \cap \hat{G}$ is the subgroup corresponding to $\hat{Y}_i \subset \hat{X}$. And the double coset $K_1 K_2$ is a finite union of translates of double cosets $g_1 \hat{K}_1 \hat{K}_2 g_2$ for $g_i \in K_i$, so it is enough to prove that $\hat{K}_1 \hat{K}_2$ is separable in $\hat{G}$ (a finite union of closed sets is closed).

Let $(\dot{X}, \dot{x}) \to (\hat{X}, \hat{x})$ be the imitator cover of $\hat{Y}_1 \hookrightarrow \hat{X}$ based at $\dot{x} = (\hat{x}, \hat{x})$, and let $\dot{G} < \hat{G}$ be the corresponding subgroup. Let $\dot{K}_i := \hat{K}_i \cap \dot{G}$. Note that $\dot{K}_1 = \hat{K}_1$ since $\dot{X}$ is the canonical completion of $\hat{Y}_1 \hookrightarrow \hat{X}$. It now suffices to prove that $\hat{K}_1 \dot{K}_2$ is separable in $\dot{G}$. This will follow from Lemma 5.2.9 below if we can construct a group retraction $\rho : \dot{G} \to \hat{K}_1$ that satisfies $\rho(\dot{K}_2) < \dot{K}_2$.

This group retraction will just be the one induced by the canonical retraction $r : (\dot{X}, \dot{x}) \to (\hat{Y}_1, \hat{x})$ composed with the canonical completion $j : (\hat{Y}_1, \hat{x}) \to (\dot{X}, \dot{x})$. A loop in $\dot{X}$ based at $\dot{x}$ is a pair of loops $(\hat{\delta}, \hat{\gamma})$ in $\hat{X}$ based at $\hat{x}$ taken by the imitator and walker respectively, and ρ is defined by $\rho(\hat{\delta}, \hat{\gamma}) := (\hat{\delta}, \hat{\delta})$. Elements of $\dot{Y}_2$ correspond to loops $(\hat{\delta}, \hat{\gamma})$ with $\hat{\gamma}$ in $\hat{Y}_2$, but then Subcomplex Entrapment implies that $\hat{\delta}$ is in $\hat{Y}_1 \cap \hat{Y}_2$, and so $\rho(\dot{K}_2) < \dot{K}_2$ as required. □

The following lemma is needed to complete the proof of Proposition 5.2.8.

Lemma 5.2.9. *Let $\rho : G \twoheadrightarrow K_1 < G$ be a retraction of a group G, and let $K_2 < G$ be a separable subgroup with $\rho(K_2) < K_2$. Then the double coset $K_1 K_2$ is separable in G.*

Proof. Let $g \in G - K_1 K_2$. The goal is to find a finite quotient of G that separates g from $K_1 K_2$. We know that $\rho(g) \in K_1$, so $g \notin \rho(g) K_2$, and by separability of K_2 there exists a finite quotient $q : G \to \bar{G}$ with

$$q(g) \notin q\rho(g) q(K_2). \tag{5.2.3}$$

Now consider another finite quotient of G, defined by

$$\begin{aligned} t : G &\to \bar{G} \times \bar{G} \\ h &\mapsto (q(h), q\rho(h)). \end{aligned}$$

We claim that $t(g) \notin t(K_1 K_2)$, so suppose for contradiction that $t(g) = t(k_1 k_2)$ with $k_i \in K_i$. Then

$$(q(g), q\rho(g)) = (q(k_1 k_2), q\rho(k_1 k_2)). \tag{5.2.4}$$

But $\rho(k_1) = k_1$ and $\rho(k_2) \in K_2$, so $q(k_1 k_2) \in q\rho(k_1 k_2) q(K_2)$, contradicting (5.2.3) and (5.2.4). □

5.2.4 Triple coset separability

We can push the arguments of Proposition 5.2.8 slightly further to also obtain separability of triple cosets.

Proposition 5.2.10. *Let X be a finite virtually special cube complex, and let $K_1, K_2, K_3 < G := \pi_1 X$ be convex subgroups. Then for any $g_1, g_2 \in G$ the triple coset $K_1 g_1 K_2 g_2 K_3$ is separable in G.*

Proof. Conjugates of convex subgroups are convex and left and right translates of separable subsets are separable, so we may assume that $g_1 = g_2 = 1$.

Choose a basepoint $x \in X$ and write $G = \pi_1(X, x)$. By Lemma 2.4.27 there are based local isometries of finite cube complexes $(Y_i, y_i) \to (X, x)$ inducing the subgroups K_i. By Corollary 5.2.7, there is a finite directly special regular cover $(\hat{X}, \hat{x}) \to (X, x)$ such that the based elevations $(\hat{Y}_i, \hat{y}_i) \hookrightarrow (\hat{X}, \hat{x})$ of $(Y_i, y_i) \to (X, x)$ are embedded and do not inter-osculate with hyperplanes of $\hat{X}$. We will consider the $\hat{Y}_i$ as subcomplexes of $\hat{X}$ with $\hat{y}_i = \hat{x}$. If $\hat{G} \triangleleft G$ is the subgroup corresponding to $(\hat{X}, \hat{x}) \to (X, x)$ then $\hat{K}_i := K_i \cap \hat{G}$ is the subgroup corresponding to $\hat{Y}_i \subset \hat{X}$. The triple coset $K_1 K_2 K_3$ is a finite union of triple cosets $g_1 \hat{K}_1 \hat{K}_2 g_2 \hat{K}_3 g_3$ with $g_i \in K_i$, so it is enough to prove that the triple cosets

$$\hat{K}_1 \hat{K}_2 g_2 \hat{K}_3 g_2^{-1}$$

are separable for $g_2 \in K_2$ (again using the fact that left and right translates of separable subsets are separable).

The subgroup $K_3' := g_2 \hat{K}_3 g_2^{-1}$ corresponds to the based elevation $(Y_3', x') \hookrightarrow (\hat{X}, x')$, where x' is the endpoint of the path $\hat{\delta}_2$ in $\hat{Y}_2$ based at $\hat{x}$ which represents the element g_2 (see Remark 2.4.26). Since $\hat{X} \to X$ is regular, we know that Y_3' is embedded in $\hat{X}$ and does not inter-osculate with any hyperplanes.

Let $(\dot{X}, \dot{x}) \to (\hat{X}, \hat{x})$ be the imitator cover of $\hat{Y}_2 \hookrightarrow \hat{X}$ based at $\dot{x} = (\hat{x}, \hat{x})$, and let $\dot{G} < \hat{G}$ be the corresponding subgroup. Let $\dot{K}_1 := \hat{K}_1 \cap \dot{G}$ and $\dot{K}_3' := K_3' \cap \dot{G}$. It now suffices to prove that $\dot{K}_1 \hat{K}_2 \dot{K}_3'$ is separable in $\dot{G}$. This will follow from Lemma 5.2.11 below (and the separability of $\dot{K}_1$-$\dot{K}_3'$ double cosets) if we can construct a group retraction $\rho : \dot{G} \to \hat{K}_2$ that satisfies $\rho(\dot{K}_1) < \dot{K}_1$ and $\rho(\dot{K}_3') < \dot{K}_3'$.

This group retraction will just be the one induced by the canonical retraction $r : (\dot{X}, \dot{x}) \to (\hat{Y}_2, \hat{x})$ composed with the canonical completion $j : (\hat{Y}_2, \hat{x}) \to (\dot{X}, \dot{x})$. A loop in $\dot{X}$ based at $\dot{x}$ is a pair of loops $(\hat{\delta}, \hat{\gamma})$ in $\hat{X}$ based at $\hat{x}$ taken by the imitator and walker respectively, and ρ is defined by $\rho(\hat{\delta}, \hat{\gamma}) := (\hat{\delta}, \hat{\delta})$. Elements of $\dot{K}_1$ correspond to loops $(\hat{\delta}, \hat{\gamma})$ with $\hat{\gamma}$ in $\hat{Y}_1$, but then Subcomplex Entrapment implies that $\hat{\delta}$ is in $\hat{Y}_1 \cap \hat{Y}_2$, and so $\rho(\dot{K}_1) < \dot{K}_1$. Elements of $\dot{K}_3'$ correspond to loops $(\hat{\delta}_2 \cdot \delta' \cdot \hat{\delta}_2^{-1}, \hat{\delta}_2 \cdot \gamma' \cdot \hat{\delta}_2^{-1})$ with γ' a loop in Y_3' based at x', but once again Subcomplex Entrapment implies that δ' is in $\hat{Y}_2 \cap Y_3'$, and so $\rho(\dot{K}_3') < \dot{K}_3'$. □

Lemma 5.2.11. *Let $\rho : G \twoheadrightarrow K_2 < G$ be a retraction of a group G, and let $K_1, K_3 < G$ be subgroups whose double cosets are separable and satisfy $\rho(K_i) < K_i$. Then the triple coset $K_1 K_2 K_3$ is separable in G.*

Proof. Let $g \in G - K_1 K_2 K_3$. The goal is to find a finite quotient of G that separates g from $K_1 K_2 K_3$. We know that $\rho(g) \in K_2$, so $g \notin K_1 \rho(g) K_3$, and by separability of $K_1 \rho(g) K_3$ there exists a finite quotient $q : G \to \bar{G}$ with

$$q(g) \notin q(K_1)\rho(g)q(K_3). \tag{5.2.5}$$

Now consider another finite quotient of G, defined by

$$\begin{aligned} t : G &\to \bar{G} \times \bar{G} \\ h &\mapsto (q(h), q\rho(h)). \end{aligned}$$

We claim that $t(g) \notin t(K_1K_2K_3)$, so suppose for contradiction that $t(g) = t(k_1k_2k_3)$ with $k_i \in K_i$. Then

$$(q(g), q\rho(g)) = (q(k_1k_2k_3), q\rho(k_1k_2k_3)). \tag{5.2.6}$$

But $\rho(k_2) = k_2$ and $\rho(k_i) \in K_i$ for $i = 1, 3$, so

$$\begin{aligned} q(k_1k_2k_3) &= q(k_1)q\rho(k_1)^{-1}q\rho(k_1)q(k_2)q\rho(k_3)q\rho(k_3)^{-1}q(k_3) \\ &\in q(K_1)q\rho(k_1k_2k_3)q(K_3), \end{aligned}$$

contradicting (5.2.5) and (5.2.6). □

5.3 Trivial wall projections

In this section we prove Theorem 1.2.7 about elevations of subcomplexes with trivial wall projections. First we must recall the notion of projection maps in CAT(0) cube complexes.

5.3.1 Projection maps

Definition 5.3.1. (Projection maps)
Let X be a CAT(0) cube complex and let $Y \subset X$ be a convex subcomplex. The *projection to* Y is the combinatorial map $\Pi : X \to Y$ that sends each vertex $x \in X^0$ to the unique closest vertex in Y with respect to the combinatorial metric - this is well-defined by [44, Lemma 13.8].

These projection maps can be used to construct "bridges" between convex subcomplexes as follows. This theorem appears in lecture notes of Hagen as [39, Theorem 1.22]. There are analogues of this theorem in the more general contexts of normal binary spaces [91, Theorems 2.5 and 2.6] and gated sets [28], and a similar phenomenon also occurs for general CAT(0) spaces [15, II.2.12(2)].

Theorem 5.3.2. *(Bridge Theorem)*
Let X be a CAT(0) cube complex and let $Y_1, Y_2 \subset X$ be convex subcomplexes. Let $\Pi_i : X \to Y_i$ be the projection to Y_i. Then:

(1) A hyperplane H crosses $\Pi_1(Y_2)$ if and only if H crosses Y_1 and Y_2. A hyperplane H separates Y_1, Y_2 if and only if H separates $\Pi_1(Y_2)$, $\Pi_2(Y_1)$.

(2) *There is a convex product subcomplex $A \times B \subset X$ and vertices $a_1, a_2 \in A$ such that $a_1 \times B = \Pi_1(Y_2)$ and $a_2 \times B = \Pi_2(Y_1)$.*

Proposition 5.3.3. *Let $G \curvearrowright X$ be a cubulated group with free G-action, and let $K_1, K_2 < G$ be convex subgroups that act cocompactly on convex subcomplexes $Y_1, Y_2 \subset X$ respectively. Let $\Pi_i : X \to Y_i$ be the projection to Y_i. Then:*

(1) *The projection $\Pi_1(Y_2)$ is finite if and only if $K_1 \cap K_2 = \{1\}$.*

(2) *If X/G is directly special and Y_2/K_2 embeds in X/G, then $K_1 \cap K_2 = \{1\}$ implies that $\Pi_1(Y_2)$ embeds in Y_1/K_1.*

Proof. Suppose $\Pi_1(Y_2)$ is finite. $K_1 \cap K_2$ stabilises Y_1 and Y_2, so it must also stabilise $\Pi_1(Y_2)$. Since $K_1 \cap K_2$ acts freely on $\Pi_1(Y_2)$ we deduce that $K_1 \cap K_2$ is finite. But G is torsion free (because it acts freely on a CAT(0) space), so $K_1 \cap K_2 = \{1\}$.

Conversely, suppose $\Pi_1(Y_2)$ is infinite. Let $x \in \Pi_1(Y_2)$, and take an infinite sequence of distinct elements $g_n \in K_1$ with $g_n x \in \Pi_1(Y_2)$. It follows from Theorem 5.3.2 that $\Pi_1(Y_2)$ and $\Pi_2(Y_1)$ lie within bounded neighbourhoods of each other, and so by the cocompactness of the K_2-action on Y_2 we can pick elements $h_n \in K_2$ such that $d(g_n x, h_n x)$ is uniformly bounded. Therefore the points $h_n^{-1} g_n x$ lie in a bounded neighbourhood of x, and by local finiteness of X there exist $n \neq m$ with $h_n^{-1} g_n x = h_m^{-1} g_m x$. Thus $1 \neq g_n g_m^{-1} = h_n h_m^{-1} \in K_1 \cap K_2$. This completes the proof of (1).

We now prove (2), so suppose that X/G is directly special, Y_2/K_2 embeds in X/G and $\Pi_1(Y_2)$ doesn't embed in Y_1/K_1. Hence there exist vertices $x_1, y_1 \in \Pi_1(Y_2)$ and $g \in K_1$ with $g x_1 = y_1$. Now apply Theorem 5.3.2(2). Suppose $x_1 = (a_1, b)$ and $y_1 = (a_1, c)$. Let γ be a combinatorial geodesic in A from a_1 to a_2. Then $\gamma \times \{b\}$ and $\gamma \times \{c\}$ are combinatorial geodesics in X starting at x_1 and y_1 respectively, which cross the same sequence of hyperplanes. Their images in X/G are paths with the same starting point that again cross the same sequence of hyperplanes. Since hyperplanes in X/G do not self-intersect or self-osculate, we deduce that $\gamma \times \{b\}$ and $\gamma \times \{c\}$ project to the same path in X/G. Then $g(\gamma \times \{b\}) = \gamma \times \{c\}$, and in particular $g(a_2, b) = (a_2, c)$. But $(a_2, b), (a_2, c) \in Y_2$ are vertices that map to the same vertex in X/G, hence $g \in K_2$ and $K_1 \cap K_2 \neq \{1\}$. □

Remark 5.3.4. In Proposition 5.3.3 (2) we didn't use the full strength of direct specialness of X/G, we only used the fact that hyperplanes don't self-intersect or self-osculate (sometimes known as weak specialness).

5.3.2 Wall projections

We now recall the notion of wall projection, due to Haglund–Wise [45, Definition 3.14], which will play a key role later in limiting the movement of imitators.

Definition 5.3.5. (Wall projection)
Let X be a cube complex with subcomplexes Y_1 and Y_2. We define $\mathrm{WProj}_X(Y_1 \to Y_2)$, the *wall projection of* Y_1 *onto* Y_2, to equal the union of Y_2^0 together with all cubes of Y_2 whose edges are parallel to edges of Y_1. We say the wall projection is *trivial* if any closed loop of $\mathrm{WProj}_X(Y_1 \to Y_2)$ is homotopically trivial inside X.

Remark 5.3.6. If X is NPC and Y_2 is locally convex then $\mathrm{WProj}_X(Y_1 \to Y_2)$ is also locally convex. So in this setting the wall projection $\mathrm{WProj}_X(Y_1 \to Y_2)$ is trivial if and only if its components are simply connected.

Remark 5.3.7. Consider the walker-imitator construction for the inclusion $Y_2 \hookrightarrow X$. It follows immediately from the definitions that if the walker stays in Y_1 then the imitator stays in $\mathrm{WProj}_X(Y_1 \to Y_2)$.

Remark 5.3.8. If $\hat{Y}_1$ and $\hat{Y}_2$ are elevations of Y_1 and Y_2 to a finite cover $\mu : \hat{X} \to X$, then $\mathrm{WProj}_{\hat{X}}(\hat{Y}_1 \to \hat{Y}_2)$ is contained inside the union of the elevations of $\mathrm{WProj}_X(Y_1 \to Y_2)$ to $\hat{Y}_2$. This is because edges $e_i \in \hat{Y}_i^1$ being parallel implies that $\mu(e_i) \in Y_i^1$ are parallel. In particular, $\mathrm{WProj}_{\hat{X}}(\hat{Y}_1 \to \hat{Y}_2)$ is trivial if $\mathrm{WProj}_X(Y_1 \to Y_2)$ is trivial.

5.3.3 Proof of Theorem 1.2.7

We now come to the main result of the section. This generalises [45, Proposition 5.1], which is the equivalent statement for graphs. It should also be compared to [45, Corollary 5.8]; our result is less general in that it says nothing about wall projections of elevations of Y_1 onto other elevations of Y_1, but more general in that it does not assume hyperbolicity of $\pi_1 X$. Instead of employing hyperbolic arguments, our proof makes use of the separability of triple cosets of convex subgroups.

Theorem 1.2.7. *(Elevating to trivial wall projections)*
Let $Y_1, Y_2 \to X$ be local isometries of finite virtually special cube complexes, and let $K_1, K_2 < G := \pi_1 X$ be the corresponding subgroups (well-defined up to conjugacy). Suppose that K_1 has trivial intersection with every conjugate of K_2. Then there is a finite directly special cover $\hat{X} \to X$ such that all elevations of Y_1 and Y_2 are embedded, and each elevation of Y_1 has trivial wall projection onto each elevation of Y_2.

Proof. We may assume that Y_1 and Y_2 are already embedded in X, as otherwise we can pass to a finite cover $X' \to X$ in which all their elevations are embedded, and then apply the proposition to each pair of elevations Y_1' and Y_2' of Y_1 and Y_2 in X'. We can then take $\hat{X} \to X$ to be a finite cover factoring through all these covers, and by Remark 5.3.8 each elevation of Y_1 to $\hat{X}$ will have trivial wall projection onto each elevation of Y_2. Similarly, we may assume that X is already directly special.

Now consider the universal cover $\tilde{X} \to X$, and let $\tilde{Y}_2$ be an elevation of Y_2. Let $\Pi : \tilde{X} \to \tilde{Y}_2$ be the projection to $\tilde{Y}_2$. Pick a vertex $x \in Y_2$ and take a lift $\tilde{x} \in \tilde{Y}_2$. Let $\Omega(\tilde{x})$ be the collection of elevations $\tilde{Y}_1$ of Y_1 such that $\tilde{x} \in \Pi(\tilde{Y}_1)$. Proposition 5.3.3(2) (with Remark 2.4.26) implies that the projections $\Pi(\tilde{Y}_1)$ are uniformly bounded for $\tilde{Y}_1 \in \Omega(\tilde{x})$, so they are all contained in some finite convex subcomplex $W(\tilde{x}) \subset \tilde{Y}_2$. Let $\mathcal{B}(\tilde{x})$ be the finite collection of hyperplanes dual to edges $e \in \tilde{Y}_2$ that have exactly one endpoint in $W(\tilde{x})$ - so a path in $\tilde{Y}_2$ based at $\tilde{x}$ leaves $W(\tilde{x})$ if and only if it crosses one of the hyperplanes in $\mathcal{B}(\tilde{x})$. For $\tilde{Y}_1 \in \Omega(\tilde{x})$ and $H \in \mathcal{B}(\tilde{x})$, we know that H does not intersect $\Pi(\tilde{Y}_1)$, and so by Theorem 5.3.2(1) H does not intersect $\tilde{Y}_1$ either.

The key step will be to use triple coset separability to prove the following claim.

<u>Claim:</u> There exists a finite cover $\hat{X} \to X$ such that, for all $\tilde{Y}_1 \in \Omega(\tilde{x})$ and $H \in \mathcal{B}(\tilde{x})$, their images in $\hat{X}$ do not intersect.

<u>Proof:</u> The property of the claim passes to further covers, so we can prove the claim by working with one hyperplane H from the finite collection $\mathcal{B}(\tilde{x})$ at a time. However, the collection $\Omega(\tilde{x})$ is infinite; to deal with this we will express $\Omega(\tilde{x})$ as a finite union of sub-collections, and prove the claim one sub-collection at a time. Let H' be a hyperplane that intersects $W(\tilde{x})$ and suppose that it also crosses some $\tilde{Y}_1 \in \Omega(\tilde{x})$. Let $\tilde{y} \in N(H') \cap \tilde{Y}_1$ be a vertex, and say that it projects to $y \in X^0$. Now let $\Omega(\tilde{x}, H', y) \subset \Omega(\tilde{x})$ consist of all elevations $\tilde{Y}_1$ intersecting H' and containing a vertex $\tilde{y} \in N(H') \cap \tilde{Y}_1$ that projects to y. It follows from Theorem 5.3.2(1) that $\Omega(\tilde{x})$ is a finite union of these $\Omega(\tilde{x}, H', y)$.

We now prove the claim for a given $H \in \mathcal{B}(\tilde{x})$ and for a given sub-collection $\Omega(\tilde{x}, H', y)$. Pick one elevation $\tilde{Y}_1 \in \Omega(\tilde{x}, H', y)$ and let $G_{H'}$ be the stabiliser of H'. Note that $\Omega(\tilde{x}, H', y) = \{h' \cdot \tilde{Y}_1 \mid h' \in G_{H'}\}$. The finite cover $\hat{X} \to X$ will be defined by a finite index normal subgroup $\hat{G} \triangleleft G$, so our task is to pick this $\hat{G}$ such that there do not exist $g \in \hat{G}$ and $h' \in G_{H'}$ with

$$H \cap gh'\tilde{Y}_1 \neq \emptyset. \tag{5.3.1}$$

Assume that K_1 is the G-stabiliser of our chosen elevation $\tilde{Y}_1$, and write G_H for the stabiliser of H. Suppose that any R-ball in $\tilde{Y}_1$ (resp. H) has K_1-translates (resp. G_H-translates) that cover $\tilde{Y}_1$ (resp. H). If $g \in G$ and $h' \in G_{H'}$ satisfy (5.3.1), then there exists $k \in K_1$ such that $d(H, gh'k\tilde{y}) \leq R$, where $\tilde{y} \in N(H') \cap \tilde{Y}_1$ is the vertex that projects to y from the definition of $\Omega(\tilde{x}, H', y)$. Then there also exists $h \in G_H$ such that $d(\tilde{x}, hgh'k\tilde{y}) \leq 2R + d(\tilde{x}, H)$. Writing $g' := hgh'k$, we know that g' also satisfies (5.3.1). However, no element of the triple coset $G_H G_{H'} K_1$ satisfies (5.3.1); so $g' \notin G_H G_{H'} K_1$, and by triple coset separability (Proposition 5.2.10) we can choose $\hat{G}$ so that

$$g' \notin \hat{G} G_H G_{H'} K_1, \tag{5.3.2}$$

hence $g \notin \hat{G}$. There are only finitely many $g' \in G$ with $d(\tilde{x}, g'\tilde{y}) \leq 2R + d(\tilde{x}, H)$, so we can satisfy (5.3.2) for all possible g' that arise in this way, thus proving the claim. ■

It remains to deduce the proposition from the claim. We have a covering $\tilde{X} \to \hat{X}$ with deck group $\hat{G}$; let $\hat{Y}_2$ be the image of $\tilde{Y}_2$ and let $\hat{x}$ be the image of $\tilde{x}$. Just as we assumed that Y_1 and Y_2 are already embedded in X, we may assume that the map $W(\tilde{x}) \to X$ is an embedding, hence the map $W(\tilde{x}) \to \hat{X}$ is also an embedding. Considering $W(\tilde{x})$ as a subcomplex of $\hat{X}$, we know that $\hat{x} \in W(\tilde{x}) \subset \hat{Y}_2$. Let $\hat{Y}_1$ be an elevation of Y_1 to $\hat{X}$. We now make a second claim.

Claim: The component of $\mathrm{WProj}_{\hat{X}}(\hat{Y}_1 \to \hat{Y}_2)$ containing $\hat{x}$ lies inside $W(\tilde{x})$.

Proof: Suppose not. Then there exists a path $\hat{\gamma}$ in $\hat{Y}_2$ based at $\hat{x}$ that leaves $W(\tilde{x})$ and is such that every edge of $\hat{\gamma}$ is dual to some edge of $\hat{Y}_1$. Say the first edge of $\hat{\gamma}$ is dual to a hyperplane $\hat{H}'$ and the first edge of $\hat{\gamma}$ leaving $W(\tilde{x})$ is dual to a hyperplane $\hat{H}$ (note these edges might be the same). So $\hat{H}'$ and $\hat{H}$ both intersect $\hat{Y}_1$ and $\hat{Y}_2$. Lifting $\hat{\gamma}$ to a path γ in $\tilde{X}$ based at $\tilde{x}$, we see that $\hat{H}'$ and $\hat{H}$ lift to hyperplanes H' and H crossing $\tilde{Y}_2$, with $H \in \mathcal{B}(\tilde{x})$. We can also choose an elevation $\tilde{Y}_1$ of $\hat{Y}_1$ that intersects H'. Since $\tilde{Y}_1$ is also an elevation of Y_1, and since $\tilde{x} \in N(H')$, we see that $\tilde{Y}_1 \in \Omega(\tilde{x})$. But then we contradict the first claim because the images of $\tilde{Y}_1$ and H in $\hat{X}$ do intersect. ■

$W(\tilde{x})$ is a convex subcomplex of $\tilde{X}$, hence simply connected, so any loop in $\mathrm{WProj}_{\hat{X}}(\hat{Y}_1 \to \hat{Y}_2)$ based at $\hat{x}$ is homotopically trivial inside $\hat{X}$. As $\hat{X} \to X$ is a regular cover, the same is true for any elevations $\hat{Y}_1$ and $\hat{Y}_2$ of Y_1 and Y_2, and any lift $\hat{x}$ of x contained inside $\hat{Y}_2$. Running the entire argument for each vertex $x \in Y_2$ gives us a finite collection of finite covers of X, and any finite cover $\hat{X} \to X$ factoring through all of them will satisfy the conclusions of the proposition. □

5.4 Commanding elements and subgroups

In this section we recall the notion of a group commanding a set of elements, and generalise it to commanding subgroups. These definitions are motivated by existing theorems in the literature, most of which are designed as tools for building finite covers of graphs of groups. In particular, the notion of commanding elements is inspired by that of omnipotence (as discussed in the introduction), and we will see that Wise's Malnormal Special Quotient Theorem gives rise to many examples of commanding subgroups. We also give a general proposition about how commanding subgroups can be used to deduce profinite and separability properties for graphs of groups.

5.4.1 Definitions and examples

Definition 5.4.1. (Commanding group elements)
A group G *commands* a set of elements $\{g_1, ..., g_n\} \subset G$ if there exists an integer $N > 0$ such that for any integers $r_1, ..., r_n > 0$ there exists a homomorphism to a finite group $G \to \bar{G}, g \mapsto \bar{g}$ such that the order of $\bar{g}_i$ is Nr_i. If this can always be done with $\langle \bar{g}_i \rangle \cap \langle \bar{g}_j \rangle = \{1\}$ for all $i \neq j$ then we say that G *strongly commands* $\{g_1, ..., g_n\}$.

There is a natural generalisation to commanding subgroups, as follows.

Definition 5.4.2. (Commanding subgroups)
A group G *commands* a collection of subgroups $(P_1, ..., P_n)$ if there exist finite index subgroups $\dot{P}_i < P_i$ such that, for any choice of finite index subgroups $P_i' < \dot{P}_i$ with $P_i' \lhd P_i$, there exists a finite-index normal subgroup $G' \lhd G$ such that $P_i \cap G' = P_i'$ for $1 \leq i \leq n$. If this can always be done with $P_i G' \cap P_j G' = G'$ for any $i \neq j$ then we say that G *strongly commands* $(P_1, ..., P_n)$. Note that the collection of subgroups $(P_1, ..., P_n)$ can contain duplicates; but if G strongly commands them then any duplicates must be trivial subgroups; or if G is residually finite and commands $(P_1, ..., P_n)$ then any duplicates must be finite subgroups.

Remark 5.4.3. G (strongly) commands a set of infinite order elements $\{g_1, ..., g_n\}$ if and only if it (strongly) commands the subgroups they generate $(\langle g_1 \rangle, ..., \langle g_n \rangle)$.

Remark 5.4.4. G commands $(P_1, ..., P_n)$ implies that G commands $(P_1, ..., gP_ig^{-1}, ..., P_n)$ for any $g \in G$ and $1 \leq i \leq n$.

The following is an easy example of commanding subgroups.

Proposition 5.4.5. *Let A be a finitely generated free abelian group. Then A commands a collection of subgroups $(A_1, ..., A_n)$ if and only if they are linearly independent (i.e. $\sum_i a_i = 0$ for $a_i \in A_i$ implies $a_i = 0$ for all i). Moreover, in this case A will strongly command them.*

Proof. Suppose $(A_1, ..., A_n)$ are linearly independent. We will show that A strongly commands (A_i) with $\dot{A}_i = A_i$. Let $A_i' < A_i$ be finite index subgroups. Put $B := A_1' + ... + A_n'$ and let $A/B = \mathbb{Z}^k \oplus F$ with F finite. Let $C < A$ be a subgroup that intersects B trivially and surjects onto the $\mathbb{Z}^k$ factor of A/B. Then $A' := B + C$ has finite index in A, and $A_i \cap A' = A_i'$ by the linear independence of (A_i). Moreover, for $j \neq k$ we have

$$\begin{aligned}(A' + A_j) \cap (A' + A_k) &= (B + A_j) \cap (B + A_k) + C \\ &= B + C \\ &= A',\end{aligned}$$

by linear independence of (A_i).

Conversely, suppose (A_i) are not linearly independent. Then there exist $a_i \in A_i$ not all zero with $\sum_i a_i = 0$. Let $\dot{A}_i < A_i$ be finite-index subgroups. We may assume that $a_i \in \dot{A}_i$ by multiplying the a_i by a large integer (remember that A is torsion-free). Let $a_j \neq 0$, and set $A'_i = \dot{A}_i$ for $i \neq j$ and $A'_j < \dot{A}_j$ some finite-index subgroup with $a_j \notin A'_j$. Then there is no finite-index subgroup $A' < A$ with $A_i \cap A' = A'_i$, because it would contain the a_i for $i \neq j$, and so would also contain $a_j = -\sum_{i\neq j} a_i$ which is in $A_j - A'_j$. □

The following deep theorem is a consequence of Wise's Malnormal Special Quotient Theorem (almost malnormality is defined in Definition 2.8.5).

Theorem 1.2.5. *Every virtually special hyperbolic group commands every almost malnormal collection of quasiconvex subgroups.*

Proof. Let G be a virtually special hyperbolic group, and let $\{P_1, ..., P_n\}$ be an almost malnormal collection of quasiconvex subgroups. Then by Wise's Malnormal Special Quotient Theorem [96, Theorem 12.3], there are finite index subgroups $\dot{P}_i < P_i$ such that, for any choice of finite index subgroups $P'_i < \dot{P}_i$ with $P'_i \triangleleft P_i$, the quotient

$$\bar{G} := G/\langle\langle P'_1, ..., P'_n\rangle\rangle$$

is virtually special and hyperbolic. By Remark 2.4.24, $\bar{G}$ contains a finite-index torsion-free subgroup $\bar{G}' \triangleleft \bar{G}$. Let $G' \triangleleft G$ be the preimage of $\bar{G}'$ in G. Each subgroup P_i will have finite image in $\bar{G}$, so this image will intersect $\bar{G}'$ trivially, thus $P_i \cap G' = \ker(P_i \to \bar{G})$. G is hyperbolic relative to $\{P_1, ..., P_n\}$ by Theorem 2.8.6, so we may apply [69, Theorem 1.1(1)]. This gives us a finite set of non-trivial elements $\mathcal{F} \subset G$ (which only depends on the P_i), such that $\ker(P_i \to \bar{G}) = P'_i$ for each i if $P'_i \cap \mathcal{F} = \emptyset$ for each i. So we are done if we replace the $\dot{P}_i$ by smaller subgroups that miss the finite set $\mathcal{F}$, which is possible because G is residually finite. □

We obtain an even more general theorem if we instead use Einstein's relatively hyperbolic version of the Malnormal Special Quotient Theorem [31, Theorem 2] in the above proof.

Theorem 5.4.6. *Let G be a virtually special group that is hyperbolic relative to subgroups $\{P_1, ..., P_n\}$. Then G commands $(P_1, ..., P_n)$.*

Remark 5.4.7. In Theorem 5.4.6, if G admits a cubulation $G \curvearrowright X$, then the subgroups P_i are necessarily convex [37, Remark 1.3]. Also, the assumption that $G \curvearrowright X$ is virtually special can be reduced to certain assumptions about how the subgroups P_i interact with X [37, 67].

The following conjecture is a natural extension of Theorem 5.4.6. We prove some weaker versions of this in Section 5.6 (Theorems 5.6.1 and 5.6.4).

Conjecture 1.2.6. *Every virtually special cubulated group $G \curvearrowright X$ commands every almost malnormal collection of convex subgroups.*

5.4.2 Application to graphs of groups

Commanding subgroups is a very useful property when trying to build finite covers of graphs of groups (or graphs of spaces), which in turn can be used to prove separability or commensurability properties for groups. This idea has occurred several times before in the literature, even though the terminology of commanding is new to this thesis. We already discussed in the introduction how commanding elements is related to omnipotence, and Wise actually defined omnipotence in order to characterise when graphs of free groups with cyclic edge groups are subgroup separable [97]. In a far more general setting, Wise used his Malnormal Special Quotient Theorem in order to prove his Quasi-convex Hierarchy Theorem about graphs of hyperbolic virtually special groups [96], albeit this didn't go via any analogue of Theorem 1.2.5. In addition, Wise effectively proved that finite-volume hyperbolic 3-manifold groups command their cusp subgroups [96, Corollary 16.15] (note that Wise forgot to specify the dimension), and this was subsequently used by Behrstock and Neumann to prove a quasi-isometric rigidity result for certain non-geometric 3-manifold groups [10] - in this case the graph of spaces comes from the geometric decomposition of the 3-manifold. Also, the author used Theorem 1.2.5 in his proof of Agol's theorem on hyperbolic cubulated groups [83], and he used Theorem 5.4.6 in joint work with Woodhouse to prove a quasi-isometric rigidity result for graphs of virtually free groups [85].

The above papers always used commanding (or some version of it) together with properties about the specific groups involved. In Proposition 5.4.8 we give a very general statement about how commanding subgroups can be used to deduce profinite and separability properties for graphs of groups. See Section 2.6.1 for the relevant definitions.

Proposition 5.4.8. *Let (G, Γ) be a graph of groups decomposition, and suppose that each vertex group commands its collection of incident edge groups. Then we have the following:*

(1) For any vertex group $G_v < G$, the profinite topology on G_v coincides with the subspace topology on G_v induced by the profinite topology on G.

(2) If every incident edge group is separable in its corresponding vertex group, then the vertex groups are separable in G.

(3) If every vertex group is residually finite and every incident edge group is separable in its corresponding vertex group, then G is residually finite.

Remark 5.4.9. Before giving the proof, we remark that the properties appearing in Proposition 5.4.8 have been well studied in contexts other than commanding. Tavgen' and Zalesskiĭ considered properties (1) and (2) when proving conjugacy separability for certain amalgams of polycyclic groups [103]. Hempel used similar ideas to prove that fundamental groups of Haken 3-manifolds are residually finite [48]. In [75], a graph of groups is called *efficient* precisely if the conclusions hold for properties (1)–(3), and this is needed if one wants to build the profinite completion of a graph of groups by piecing together the profinite completions of its vertex groups.

Proof. The vertex groups command their incident edge groups, so by definition there are finite-index subgroups $\dot{G}_e \lhd G_e$ for $e \in E\Gamma$ such that, for any choice of finite-index subgroups $G'_e < \dot{G}_e$ with $G'_e \lhd G_e$, there exist finite-index normal subgroups $G'_v \lhd G_v$ for $v \in V\Gamma$ with

$$\zeta_e(G_e) \cap G'_{\tau(e)} = \zeta_e(G'_e) \tag{5.4.1}$$

for $e \in E\Gamma$. Moreover, given finite-index normal subgroups $\hat{G}_v \lhd G_v$, if we choose the G'_e such that $\zeta_e(G'_e) < \hat{G}_{\tau(e)}$, then we may assume that $G'_v < \hat{G}_v$ for all $v \in V\Gamma$ (by replacing the G'_v with $G'_v \cap \hat{G}_v$ if necessary). Furthermore, if $G'_e = G'_{\bar{e}}$ for all $e \in E\Gamma$, then we can construct a finite cover of the graph of groups, with copies of the G'_v and G'_e as vertex and edge groups, and with edge morphisms obtained by composing the existing edge morphisms ζ_e with inner automorphisms of the original vertex groups G_v. This gives a valid cover of graphs of groups precisely because of equation (5.4.1). We are now ready to prove (1) and (2).

(1) Fix a vertex group G_v. Let $\hat{G}_v \lhd G_v$ be a finite-index normal subgroup. Take a finite-index subgroup $G' < G$ corresponding to a finite cover of graphs of groups as above, such that $G'_v < \hat{G}_v$. Observe that $G' \cap G_v = G'_v$ is closed in the subspace topology on G_v induced by the profinite topology on G. But G'_v has finite index in G_v, so it is also open in the subspace topology. Hence $\hat{G}_v$ is open in the subspace topology too. By definition, the cosets of finite-index subgroups of G_v form a basis for the profinite topology, thus all open subsets of G_v in the profinite topology are also open in the subspace topology. On the other hand, the intersection of any coset of a finite-index subgroup of G with G_v is a coset of a finite index subgroup of G_v, so all open subsets of G_v in the subspace topology are also open in the profinite topology on G_v.

(2) Now suppose the vertex groups are residually finite. Fix a vertex group G_v and an element $g \in G - G_v$. We must find a finite-index subgroup $G' < G$ with $g \notin G'G_v$ (as then $G'g$ is an open neighbourhood of g distinct from G_v). Now let (X, Γ) be a graph of spaces corresponding to (G, Γ), and let γ be a loop in X corresponding to g. We may assume that the basepoint is in the vertex space X_v. Suppose that γ projects to an edge loop $\delta = (e_1, ..., e_n)$ in Γ with vertex sequence $v = v_0, v_1, ..., v_n = v$. After homotoping, we may assume that γ is a concatenation of standard paths that go across the edge spaces X_{e_i}, together with based loops in the vertex spaces corresponding to elements $g_i \in G_{z_i}$. We may also assume that γ is *reduced* in the sense that $g_i \notin \zeta_{e_i}(G_{e_i})$ whenever $e_i = \bar{e}_{i+1}$.

By separability of the incident edge groups, there exist finite-index normal subgroups $\hat{G}_v \lhd G_v$ such that

$$g_i \notin \zeta_{e_i}(G_{e_i})\hat{G}_{z_i}, \tag{5.4.2}$$

whenever $e_i = \bar{e}_{i+1}$. Now construct a finite cover (G', Γ') of the graph of groups decomposition (G, Γ) as described above, with $G'_v < \hat{G}_v$ for all $v \in V\Gamma$, and let $X' \to X$ be the corresponding finite cover of graphs of spaces. Let γ lift to a based path γ' in X', and suppose this projects to a path δ' in Γ'. It follows from (5.4.2) that δ' never backtracks, and if we construct (G', Γ') appropriately we can ensure that δ' is embedded. The subgroup $G'G_v$ corresponds to based paths in X' that start and finish at the same vertex space, so we conclude that $g \notin G'G_v$ as required (note that $n \geq 1$ since $g \notin G_v$).

(3) For any vertex group G_v, (2) implies that G_v is closed in the profinite topology on G. Then it follows from (1) and the residual finiteness of G_v that the trivial subgroup is also closed in G, hence G is residually finite.

□

5.4.3 Passing to finite-index subgroups

When we prove Theorem 1.2.3 in Section 5.6 we will first show that a certain finite-index subgroup $\hat{G} \lhd G$ (strongly) commands a certain collection of elements, so we need the following lemma to transfer back to the whole of G.

Lemma 5.4.10. *Let G be a finitely generated group with independent elements $\{g_1, ..., g_n\}$. Assume G is balanced - meaning that $hg^kh^{-1} = g^l$ for g infinite order implies that $k = \pm l$. Suppose that $\hat{G} \lhd G$ is a finite-index normal subgroup and suppose that M_i is the order of g_i in the quotient $G/\hat{G}$. We define an equivalence relation on the G-conjugacy class of $h_i := g_i^{M_i}$, where $g \sim g'$ if $\langle g\rangle \cap \langle hg'h^{-1}\rangle \neq \{1\}$ for some $h \in \hat{G}$, and let $\{h_{ij}\}_{j\in J_i}$ be a set of representatives of the $\sim$-classes (with one h_{ij} equal to h_i). Then $\{h_{ij}\}$ is a finite independent set of elements, and G (strongly) commands $\{g_i\}$ if $\hat{G}$ (strongly) commands $\{h_{ij}\}$.*

Proof. The G-conjugacy class of h_i divides into finitely many $\hat{G}$-conjugacy classes, and each $\sim$-class will be a union of $\hat{G}$-conjugacy classes - in particular $\{h_{ij}\}$ is finite. The independence of $\{h_{ij}\}$ follows from the independence of $\{g_i\}$ and the definition of $\sim$.

For any $g \in G$ and any i, $g^{-1}h_ig \sim h_{ij}$ for some $j \in J_i$, so there exists $h \in \hat{G}$ with

$$\langle h_{ij}\rangle \cap \langle h^{-1}g^{-1}h_igh\rangle \neq \{1\}.$$

Now h_{ij} is a G-conjugate of h_i and G is balanced, so there exists an integer $R > 0$ such that

$$h_{ij}^{\pm R} = h^{-1}g^{-1}h_i^Rgh. \tag{5.4.3}$$

But there are finitely many $\hat{G}$-conjugacy classes of each h_i, so this R can be chosen uniformly. More precisely, there is an integer $R > 0$ such that for any $g \in G$ and any i, there exists $h \in \hat{G}$ and $j \in J_i$ satisfying (5.4.3).

Now suppose that $\hat{G}$ commands $\{h_{ij}\}$. Then there exists an integer $N > 0$ such that for any integers $r_{ij} > 0$ there exists a finite-index normal subgroup $\hat{G}' \triangleleft \hat{G}$ such that h_{ij} has order Nr_{ij} in $\hat{G}/\hat{G}'$. Given positive integers $r_i > 0$, set

$$r_{ij} := r_i R \prod_{k \neq i} M_k,$$

and take $\hat{G}' \triangleleft \hat{G}$ as above. Then define

$$G' := \bigcap_{g \in G} g\hat{G}'g^{-1},$$

which is a finite-index normal subgroup of G (as G is finitely generated). Denote the quotient map $G \to G/G'$ by $g \mapsto \bar{g}$. The following claim establishes that G commands $\{g_i\}$.

Claim: $\bar{g}_i$ has order $(NR\prod_k M_k)r_i$.

Proof: For a power g_i^l to be in G' we would certainly need $g_i^l \in \hat{G}$, so we'd need M_i to divide l. Thus it remains to show that $\bar{h}_i = \bar{g}_i^{M_i}$ has order $T_i := (NR\prod_{k\neq i} M_k)r_i$.

As h_i is equal to one of the h_{ij}, we know that $T_i = Nr_{ij}$ is the order of h_i in $\hat{G}/\hat{G}'$, so it remains to show that $h_i^{T_i}$ lands in every conjugate of $\hat{G}'$. Indeed, let $g \in G$, and take $h \in \hat{G}$ and $j \in J_i$ satisfying (5.4.3). Since R divides T_i, we have

$$\begin{aligned} h_i^{T_i} &= ghh_{ij}^{\pm T_i}h^{-1}g^{-1} \\ &\in gh\hat{G}'h^{-1}g^{-1} \\ &= g\hat{G}'g^{-1}. \end{aligned}$$

■

If $\hat{G}$ strongly commands $\{h_{ij}\}$ then we can assume that the subgroups $\langle h_{ij}\rangle$ have pairwise trivial intersection in $\hat{G}/\hat{G}'$, so in particular the subgroups $\langle h_i\rangle$ have pairwise trivial intersection in $\hat{G}/\hat{G}'$. To show that G strongly commands $\{g_i\}$ we must prove the following claim.

Claim: The subgroups $\langle \bar{g}_i\rangle$ have pairwise trivial intersection.

Proof: Suppose for contradiction there is $i \neq j$ with $\langle \bar{g}_i\rangle \cap \langle \bar{g}_j\rangle$ non-trivial, say it contains an element of order p, with p prime. Then p divides $|\langle \bar{g}_i\rangle| = M_iT_i$ and $|\langle \bar{g}_j\rangle| = M_jT_j$.

Since M_i divides T_j, we deduce that p divides at least one of T_i and T_j, say it divides T_i. Then

$$1 \neq \bar{g}_i^{M_i T_i/p} = \bar{h}_i^{T_i/p} \in \langle \bar{g}_i \rangle \cap \langle \bar{g}_j \rangle \cap \hat{G}/G'.$$

But $\langle \bar{g}_j \rangle \cap \hat{G}/G' = \langle \bar{h}_j \rangle$, hence $\langle \bar{h}_i \rangle \cap \langle \bar{h}_j \rangle \neq \{1\}$, so $\langle h_i \rangle$ and $\langle h_j \rangle$ also have non-trivial intersection in $\hat{G}/\hat{G}'$, contrary to our assumption. ■

□

5.5 Multiple imitators

In this section we construct imitator covers with multiple imitators, which will be needed in the next section. First we must define cube complex pairs.

5.5.1 Cube complex pairs

When working with fundamental groups of cube complexes, the geometric model for a collection of convex subgroups will be given by a cube complex pair. The definition given below follows [3, Section 5.1]. Also, if working with graphs of cube complexes, the edge spaces incident to a vertex space will correspond exactly to the peripheral subcomplexes defined below. Having vertex groups that command incident edge groups is a powerful tool for constructing finite covers of graphs of cube complexes, as in Proposition 5.4.8.

Definition 5.5.1. (Cube complex pairs)
For $i = 1, ..., n$, let $\phi_i : Z_i \to X$ be local isometries of NPC cube complexes. Writing $\mathcal{Z} = \sqcup_i Z_i$, we call $(X, \mathcal{Z})$ a *cube complex pair*. We say that $(X, \mathcal{Z})$ is *finite* (resp. *directly special*) if X and $\mathcal{Z}$ are finite (resp. directly special). The mapping cylinder of $\Phi = \sqcup_i \phi_i$,

$$C_\Phi = X \sqcup (\mathcal{Z} \times [0,1])/\{(z,1) \sim \Phi(z)\}, \tag{5.5.1}$$

naturally has the structure of an NPC cube complex. We call C_Φ the *augmented cube complex* based on the pair $(X, \mathcal{Z})$.

There are canonical inclusions $X \hookrightarrow C_\Phi$ and $\mathcal{Z} \hookrightarrow C_\Phi$ (the second via $z \mapsto (z, 0)$). The subset $X \subset C_\Phi$ is a deformation retract of C_Φ. The components Z_i of $\mathcal{Z} \subset C_\Phi$ are referred to as *peripheral subcomplexes.*

Any cover $\hat{X} \to X$ gives rise to a cover $C_{\hat{\Phi}} \to C_\Phi$, where $\hat{\Phi} : \hat{\mathcal{Z}} \to \hat{X}$ is made up of all the elevations of the maps ϕ_i. We say that $(\hat{X}, \hat{\mathcal{Z}}) \to (X, \mathcal{Z})$ is a *cover of cube complex pairs*, and we call it *regular* if $\hat{X} \to X$ is regular.

Remark 5.5.2. If $(\hat{X}, \hat{\mathcal{Z}}) \to (X, \mathcal{Z})$ is a regular cover of cube complex pairs then the cover of augmented cube complexes $C_{\hat{\Phi}} \to C_\Phi$ is also regular. We then get an action of $\pi_1 X / \pi_1 \hat{X} = \pi_1 C_\Phi / \pi_1 C_{\hat{\Phi}}$ on $C_{\hat{\Phi}}$ by deck transformations, preserving the subsets $\hat{X}, \hat{\mathcal{Z}} \subset C_{\hat{\Phi}}$.

5.5.2 Imitator covers with multiple imitators

For the rest of this section fix a finite directly special cube complex pair $(X, \mathcal{Z})$.

Construction 5.5.3. (Imitator covers with multiple imitators)
We can consider multiple imitators wandering around the peripheral subcomplexes while a single walker wanders around in X. All the imitators try to copy the walker in the same way as in Construction 5.1.1, with respect to the map $\Phi : \mathcal{Z} \to X$. Let $\mathcal{I}$ denote the (finite) set of imitators. As in Construction 5.1.4, the movements of imitators and walker define a finite cover $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ of X, whose vertices represent all possible positions of imitators and walker on the vertices of $\mathcal{Z}$ and X respectively. The only restriction we impose is that two imitators cannot occupy the same vertex - note that no movement of the walker can force two imitators to move onto the same vertex because their movements are reversible (see Construction 5.1.1). We write a vertex of $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ as a pair (θ, x), where $\theta : \mathcal{I} \to \mathcal{Z}^0$ is an injection and $x \in X^0$. For each $I \in \mathcal{I}$ we will also have a cellular map $r_I : \mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X) \to \mathcal{Z}$ by taking the position of the imitator I.

As in Construction 5.1.1, if one of the imitators and the walker are at positions (z, x), then the imitator travels along the path $\delta(\gamma, z)$ as the walker travels along the path γ. If γ is a loop in X based at x that lifts to a path $\dot{\gamma}$ in $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ based at (θ, x), then $\dot{\gamma}$ closes up as a loop if and only if the paths $\delta(\gamma, \theta(I))$ close up as loops for all $I \in \mathcal{I}$.

We let $\mathsf{C}(\mathcal{Z}, X; \theta, x)$ denote the component of $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ containing (θ, x), and we call it the *imitator cover of* $(X, \mathcal{Z})$ *based at* (θ, x) (note that the set $\mathcal{I}$ is implicit in the map θ). We will usually work with imitator covers $\mathsf{C}(\mathcal{Z}, X; \theta, x)$ for which $|\mathcal{I}| = |\mathcal{Z}^0|$, so that every vertex in the peripheral subcomplexes is occupied by an imitator, and we call such imitator covers *full*.

Remark 5.5.4. The imitator cover of $(X, \mathcal{Z})$ based at (θ, x) can alternatively be defined as the smallest based cover of X that factors through the based covers $\mathsf{C}(Z_i, X; \theta(I), x)$ (in the notation of Construction 5.1.4) for each peripheral subcomplex $Z_i \subset \mathcal{Z}$ and each imitator $I \in \mathcal{I}$ with $\theta(I) \in Z_i$.

Lemma 5.5.5. *The permutation group of the imitator set* $\mathrm{Sym}(\mathcal{I})$ *acts on* $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$*, by* $\alpha \cdot (\theta, x) := (\theta\alpha^{-1}, x)$ *for* $\alpha \in \mathrm{Sym}(\mathcal{I})$*, and this action preserves each fibre of the covering* $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X) \to X$*. Moreover, the stabiliser of a full imitator cover* $\dot{X} = \mathsf{C}(\mathcal{Z}, X; \theta, x) \subset \mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ *acts as the full deck group of* $\dot{X} \to X$ *- in particular* $\dot{X} \to X$ *is regular.*

Proof. The action of $\mathrm{Sym}(\mathcal{I})$ on $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$ is well-defined as relabelling imitators doesn't affect the construction of imitator covers, and the action clearly respects the covering $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X) \to X$. For the second part, we observe that any vertex of $\dot{X}$ that maps to $x \in X$ takes the form (θ', x), and $\alpha \cdot (\theta, x) := (\theta', x)$ for $\alpha = (\theta')^{-1}\theta \in \mathrm{Sym}(\mathcal{I})$ (note that θ, θ' are bijections because $\dot{X}$ is a full imitator cover). Since $\dot{X}$ is a component of $\mathsf{C}_{\mathcal{I}}(\mathcal{Z}, X)$, such an α must stabilise $\dot{X}$ and define a deck transformation of $\dot{X} \to X$. $\square$

Remark 5.5.6. It follows from Lemma 5.5.5 that the action of $\pi_1(X, x)$ on a full imitator cover $\dot{X} \to X$ by deck transformations factors through the action of $\text{Sym}(\mathcal{I})$, hence we get a homomorphism $\pi_1(X, x) \to \text{Sym}(\mathcal{I})$. In fact we can describe this homomorphism in terms of imitator movements: a loop γ in X maps to the permutation $\alpha_\gamma \in \text{Sym}(\mathcal{I})$, where for each imitator $I \in \mathcal{I}$, the walker traversing γ causes imitator $\alpha_\gamma(I)$ to replace imitator I at the vertex $\theta(I)$.

Remark 5.5.7. Any two full imitator covers of $(X, \mathcal{Z})$ are isomorphic as covers of X, so we can refer to it as *the full imitator cover of* $(X, \mathcal{Z})$. Indeed, if the first imitator cover is based at (θ_1, x) then we can choose the basepoint of the second to have the same X-coordinate, say (θ_2, x); and then the action of $\alpha = \theta_2^{-1}\theta_1 \in \text{Sym}(\mathcal{I})$ from Lemma 5.5.5 induces an isomorphism $\mathsf{C}(\mathcal{Z}, X; \theta_1, x) \to \mathsf{C}(\mathcal{Z}, X; \theta_2, x)$

Remark 5.5.8. The full imitator cover of $(X, \mathcal{Z})$ can alternatively be characterised as the smallest regular cover of X that factors through every component of every cover $\mathsf{C}(Z_i, X)$ for $Z_i \subset \mathcal{Z}$ a peripheral subcomplex. In other words, it corresponds to the normal subgroup of $\pi_1(X, x)$ obtained by intersecting all subgroups induced by based single imitator covers supported on the peripheral subcomplexes.

5.6 Achieving command

In this section we prove Theorem 1.2.3, which states that a virtually special cubulated group strongly commands any independent set of convex elements. We also prove Theorems 5.6.1 and 5.6.4, which are related to, but weaker than, the notion of commanding subgroups. Finally, we investigate when Theorem 1.2.3 can be applied to right-angled Artin groups.

5.6.1 Non-normal commanding

The following theorem is a non-normal version of commanding subgroups since it does not require the finite-index subgroup G' to be normal in G. Interpreted in terms of covers, this provides a way of controlling based elevations.

Theorem 5.6.1. *(Non-normal commanding)*
Let X be a finite virtually special cube complex, and let $P_1, ..., P_n < G := \pi_1 X$ be convex subgroups with pairwise trivial intersections. Then there exist finite-index subgroups $\dot{P}_i < P_i$ such that for any further finite-index subgroups $P_i' < \dot{P}_i$ there exists a finite-index $G' < G$ with $P_i \cap G' = P_i'$.

Proof. Pick a basepoint $x \in X$ and write $G = \pi_1(X, x)$. By Lemma 2.4.27 there exist local isometries of finite cube complexes $\phi_i : (Z_i, z_i) \to (X, x)$ corresponding to the subgroups P_i. By Corollary 5.2.7, there is a finite directly special cover $(\hat{X}, \hat{x}) \to (X, x)$ such that the based elevations $\hat{\phi}_i : (\hat{Z}_i, \hat{z}_i) \to (\hat{X}, \hat{x})$ of the ϕ_i are embedded and do not inter-osculate with hyperplanes of $\hat{X}$. If $\hat{G} < G$ is the subgroup corresponding to $(\hat{X}, \hat{x}) \to (X, x)$ then $\hat{P}_i := P_i \cap \hat{G}$ is the subgroup corresponding to $(\hat{Z}_i, \hat{z}_i) \to (X, x)$.

Now consider the imitator cover $(\dot{X}, \dot{x}) \to (\hat{X}, \hat{x})$ consisting of peripheral subcomplexes $\hat{Z}_i$, imitators $I_1, ..., I_n$ starting at $\hat{z}_1, ..., \hat{z}_n$ respectively and the walker starting at $\hat{x}$. So the basepoint is $\dot{x} = (\theta, \hat{x})$, where $\theta(I_i) = \hat{x}$ for all i. Let $\dot{G} < \hat{G}$ be the subgroup corresponding to $\dot{X}$ and let $\dot{P}_i := \hat{P}_i \cap \dot{G}$. We have maps $r_{I_i} : (\dot{X}, \dot{x}) \to (\hat{Z}_i, \hat{z}_i)$ giving the positions of the different imitators, and these induce homomorphisms $\rho_i : \dot{G} \to \hat{P}_i$. A loop in $\dot{X}$ based at $\dot{x}$ consists of the walker traversing a loop $\hat{\gamma}$ in $\hat{X}$ based at $\hat{x}$ while imitator I_i traverses a loop $\hat{\delta}_i$ in $\hat{Z}_i$ based at $\hat{z}_i$ for each i. The inclusion $\dot{G} \hookrightarrow \hat{G}$ is given by projecting to the walker (i.e. taking the loop $\hat{\gamma}$) while the inclusions $\hat{P}_i \hookrightarrow \hat{G}$ are induced by the maps $\hat{\phi}_i$. Elements of $\dot{P}_i$ correspond to loops in $\dot{X}$ such that $\hat{\gamma} = \hat{\phi}_i \hat{\delta}_i$, so we have

$$\rho_i|_{\dot{P}_i} = \mathrm{id}_{\dot{P}_i}. \tag{5.6.1}$$

If $\hat{\gamma} = \hat{\phi}_i \hat{\delta}_i$, then by Subcomplex Entrapment we have that $\hat{\delta}_j$ is contained in $\hat{\phi}_j^{-1} \hat{\phi}_i(\hat{Z}_i)$ for all $i \neq j$, and so

$$\rho_j(\dot{P}_i) \subset \hat{P}_i \cap \hat{P}_j = \{1\}. \tag{5.6.2}$$

Given finite-index subgroups $P'_i < \dot{P}_i$, we can define a finite-index subgroup $G' < \dot{G}$ by

$$G' := \bigcap_{i=1}^{n} \rho_i^{-1}(P'_i),$$

and by (5.6.1) and (5.6.2) we have

$$P_i \cap G' = \dot{P}_i \cap G' = P'_i,$$

as required. $\square$

5.6.2 Imitator homomorphisms

Next we need the following lemma about the regularity of imitator covers.

Lemma 5.6.2. *Let $(\hat{X}, \hat{\mathcal{Z}}) \to (X, \mathcal{Z})$ be a regular cover of cube complex pairs with $\hat{X}$ directly special. Then the full imitator cover $\dot{X}$ of $(\hat{X}, \hat{\mathcal{Z}})$ is a regular cover of X. Moreover, the deck group action of $G = \pi_1 X$ on $\dot{X}$ takes the form*

$$g(\theta, \hat{x}) = (g\theta\alpha_g^{-1}, g\hat{x}), \tag{5.6.3}$$

for $g \in G$, where $g\theta$ and $g\hat{x}$ refer to the deck group action of G on $(X, \mathcal{Z})$ from Remark 5.5.2, and $G \to \mathrm{Sym}(\mathcal{I})$; $g \mapsto \alpha_g$ is an action of G on $\mathcal{I}$ that depends on the cover $\dot{X} \to X$.

Proof. Firstly, we consider the imitator cover $\mathsf{C}_{\mathcal{I}}(\hat{\mathcal{Z}}, \hat{X}) \to \hat{X}$, of which $\dot{X}$ is one component. We have an action of $G/\hat{G} \times \mathrm{Sym}(\mathcal{I})$ on $\mathsf{C}_{\mathcal{I}}(\hat{\mathcal{Z}}, \hat{X})$ given by

$$(g\hat{G}, \alpha) \cdot (\theta, \hat{x}) = (g\theta\alpha^{-1}, g\hat{x}), \tag{5.6.4}$$

where $g\theta$ and $g\hat{x}$ once again refer to the deck group action of G on $(\hat{X}, \hat{\mathcal{Z}})$ from Remark 5.5.2, and $\hat{G} \triangleleft G$ is the subgroup corresponding to $\hat{X}$. Note that the action (5.6.4) is free and respects the covering $\dot{X} \to X$.

We will now show that a subgroup of $G/\hat{G} \times \mathrm{Sym}(\mathcal{I})$ stabilises the component $\dot{X} \subset \mathsf{C}_{\mathcal{I}}(\hat{\mathcal{Z}}, \hat{X})$ and defines a full group of deck transformations for the cover $\dot{X} \to X$, demonstrating that this cover is regular. Indeed let $(\theta_1, \hat{x}_1), (\theta_2, \hat{x}_2) \in \dot{X}$ be vertices that both map to a vertex $x \in X$. We must find $(g\hat{G}, \alpha) \in G/\hat{G} \times \mathrm{Sym}(\mathcal{I})$ with $(g\hat{G}, \alpha) \cdot (\theta_1, \hat{x}_1) = (\theta_2, \hat{x}_2)$. By (5.6.4) this equivalent to $g\theta_1\alpha^{-1} = \theta_2$ and $g\hat{x}_1 = \hat{x}_2$; but we can choose $g \in G$ that satisfies the second equation since $\hat{x}_1, \hat{x}_2 \in \hat{X}$ both map to $x \in X$, and then we can satisfy the first equation by setting $\alpha = \theta_2^{-1} g\theta_1$.

We deduce that there is a deck group action of G on $\dot{X}$ which factors through a homomorphism $\alpha : G \to G/\hat{G} \times \mathrm{Sym}(\mathcal{I})$. The cover $\dot{X} \to \hat{X}$; $(\theta, \hat{x}) \mapsto \hat{x}$ must be equivariant with respect to the deck group actions of G, so the projection of α to the $G/\hat{G}$ factor must recover the quotient map $G \to G/\hat{G}$. Denoting the projection of α to the $\mathrm{Sym}(\mathcal{I})$ factor by $g \mapsto \alpha_g$ then recovers equation (5.6.3), as required. □

The following technical lemma is a crucial step towards proving Theorem 1.2.3. It involves building a full imitator cover and defining a homomorphism for each imitator corresponding to their movements. Various properties of these homomorphisms are then established, in particular we make use of the trivial wall projections from Theorem 1.2.7.

Lemma 5.6.3. *(Imitator homomorphisms)*
Let X be a finite virtually special cube complex, and let $P_1, ..., P_n < G := \pi_1 X$ be convex subgroups such that all conjugates of P_i and P_j have trivial intersection if $i \neq j$. Then there is a finite-index normal subgroup $\dot{G} \triangleleft G$, and for each i there is a finite collection Λ_i of homomorphisms $\lambda : \dot{G} \to P_i$ such that:

(1) Λ_i contains a homomorphism which is the identity on $\dot{P}_i := P_i \cap \dot{G}$.

(2) $\lambda(\dot{P}_j) = \{1\}$ for $\lambda \in \Lambda_i$ and $i \neq j$.

(3) For any $\lambda \in \Lambda_i$ and $g \in G$ there is some $\lambda' \in \Lambda_i$ and $p \in P_i$ such that $p\lambda(\dot{g})p^{-1} = \lambda'(g\dot{g}g^{-1})$ for all $\dot{g} \in \dot{G}$.

Proof. Pick a basepoint $x \in X$ and write $G = \pi_1(X, x)$. By Lemma 2.4.27 there exist local isometries of finite cube complexes $\phi_i : (Z_i, z_i) \to (X, x)$ corresponding to the subgroups P_i. This makes X into a cube complex pair $(X, \mathcal{Z})$. By Corollary 5.2.7, there is a finite regular cover $(\hat{X}, \hat{\mathcal{Z}}) \to (X, \mathcal{Z})$, with $\hat{X}$ directly special, such that $\hat{\phi} : \hat{Z} \to \hat{X}$

is an embedding that does not inter-osculate with hyperplanes of $\hat{X}$ for each peripheral subcomplex $\hat{Z} \subset \hat{\mathcal{Z}}$ (where $\hat{\phi}$ denotes the restriction of $\hat{\Phi} : \hat{\mathcal{Z}} \to \hat{X}$ to $\hat{Z}$). By applying Theorem 1.2.7, we can also assume that any elevation of Z_i has trivial wall projection onto any elevation of Z_j for $i \neq j$.

Pick $\hat{x} \in \hat{X}$ a lift of x and let $\dot{X} = \mathsf{C}(\hat{\mathcal{Z}}, \hat{X}; \theta, \hat{x}) \to \hat{X}$ be the full imitator cover of $(\hat{X}, \hat{\mathcal{Z}})$. Let $\dot{G} \triangleleft G$ be the subgroup corresponding to $\dot{X}$, which is normal by Lemma 5.6.2, and let $\dot{x} := (\theta, \hat{x})$. We will make the identification $\dot{G} = \pi_1(\dot{X}, \dot{x})$ throughout. For an imitator $I \in \mathcal{I}$ contained in a peripheral subcomplex $\hat{Z}_I \subset \hat{\mathcal{Z}}$, we have the projection map

$$r_I : (\dot{X}, \dot{x}) \to (\hat{Z}_I, \theta(I)),$$

which induces a homomorphism $\rho_I : \dot{G} \to \pi_1(\hat{Z}_I, \theta(I))$. If $C_{\hat{\Phi}} \to C_\Phi$ restricts to $(\hat{Z}_I, \theta(I)) \to (Z_i, z_i')$, then for any path δ in Z_i from z_i to z_i' there is an embedding of fundamental groups

$$\pi_1(\hat{Z}_I, \theta(I)) \xhookrightarrow{\nu_I} \pi_1(Z_i, z_i') \xrightarrow[\sim]{\tau_\delta} \pi_1(Z_i, z_i) \cong P_i, \tag{5.6.5}$$

where τ_δ is the isomorphism defined by the concatenation map $\gamma \mapsto \delta * \gamma * \delta^{-1}$. Hence ρ_I induces a homomorphism $\lambda_I : \dot{G} \to P_i$. We let Λ_i denote the collection of all λ_I such that $\hat{Z}_I$ covers Z_i. It remains to verify properties (1)–(3) from the lemma.

(1) Let $\hat{Z} \subset \hat{\mathcal{Z}}$ be the peripheral subcomplex with vertex $\hat{z}$ such that $\hat{\phi} : (\hat{Z}, \hat{z}) \hookrightarrow (\hat{X}, \hat{x})$ is the based elevation of $\phi_i : (Z_i, z_i) \to (X, x)$, and consider the imitator $I \in \mathcal{I}$ with $\theta(I) = \hat{z} \in \hat{Z}$. We claim that λ_I restricts to the identity on $\dot{P}_i$. Considering $\dot{P}_i$ as a subgroup of P_i, it consists of those loops $\gamma \in \pi_1(Z_i, z_i)$ such that $\phi_i \gamma$ lifts to a loop $\dot{\gamma} \in \pi_1(\dot{X}, \dot{x})$. If $\hat{\gamma} \in \pi_1(\hat{Z}, \hat{z})$ is the lift of γ to $\hat{Z}$, then the loop $\dot{\gamma}$ consists of the imitator I traversing the loop $\hat{\gamma}$ and the walker traversing the loop $\hat{\phi}\hat{\gamma}$. So $\rho_I(\gamma) = \hat{\gamma}$ and $\lambda_I(\gamma) = \gamma$ (assuming we choose δ from (5.6.5) to be the trivial path at z_i).

(2) As above, an element of $\dot{P}_j$ is a loop $\gamma \in \pi_1(Z_j, z_j)$ that lifts to a loop $\dot{\gamma} \in \pi_1(\dot{X}, \dot{x})$, and $\dot{\gamma}$ consists of the walker traversing a loop in the image of some peripheral subcomplex $\hat{\phi} : \hat{Z} \to \hat{X}$ an elevation of ϕ_j. $\lambda_I \in \Lambda_i$ corresponds to an imitator I in a peripheral subcomplex $\hat{\phi}_I : \hat{Z}_I \to \hat{X}$ an elevation of ϕ_i. By construction of $\hat{X}$, we know that $\mathrm{WProj}_{\hat{X}}(\hat{\phi}(\hat{Z}) \to \hat{\phi}_I(\hat{Z}_I))$ is trivial, so the loop $r_I \dot{\gamma}$ traversed by I will be null-homotopic, and $\lambda_I(\gamma) = 1$.

(3) Let $\lambda_I \in \Lambda_i$ and $g \in G$. Let γ_g be a loop in X based at x in the homotopy class g, and let $\dot{\gamma}_g$ be a lift to $\dot{X}$ based at $\dot{x}$. By Lemma 5.6.2 we have an action of G on $\dot{X}$ by deck transformations, with $g\dot{x} = (g\theta\alpha_g^{-1}, g\hat{x})$, where $g\theta$ and $g\hat{x}$ refer to the deck

group action of G on $\hat{X}, \hat{\mathcal{Z}} \subset C_{\hat{\Phi}}$ from Remark 5.5.2, and $\alpha_g \in \mathrm{Sym}(\mathcal{I})$ comes from an action of G on the set of imitators $\mathcal{I}$. Writing $J = \alpha_g(I)$, we know that

$$\begin{aligned} r_J g(\dot{x}) &= r_J(g\theta\alpha_g^{-1}, g\hat{x}) \\ &= g\theta\alpha_g^{-1}(J) \\ &= g\theta(I) \\ &= g r_I(\theta, \hat{x}) \\ &= g r_I(\dot{x}). \end{aligned}$$

This actually holds for all vertices in $\dot{X}$, so we also know that

$$r_J g\dot{\gamma} = g r_I \dot{\gamma}, \tag{5.6.6}$$

for any path $\dot{\gamma}$ in $\dot{X}$. In particular, the peripheral subcomplexes containing imitators I and J satisfy $\hat{Z}_J = g\hat{Z}_I$, so both are elevations of Z_i. We get the following commutative diagram of homomorphisms, where the horizontal maps are defined on the right, with $\theta(I)$ and $\hat{\eta} := r_J\dot{\gamma}_g$ descending to a vertex z_i' and a loop η in Z_i.

$$\begin{array}{ccccc} \pi_1(\dot{X}, \dot{x}) & \longrightarrow & \pi_1(\dot{X}, \dot{x}) & \dot{\gamma} \longmapsto & \dot{\gamma}_g * (g\cdot\dot{\gamma}) * \dot{\gamma}_g^{-1} \\ \downarrow{\scriptstyle \rho_I} & & \downarrow{\scriptstyle \rho_J} & & \\ \pi_1(\hat{Z}_I, \theta(I)) & \longrightarrow & \pi_1(\hat{Z}_J, g\theta(I)) & \hat{\gamma} \longmapsto & \hat{\eta} * (g\hat{\gamma}) * \hat{\eta}^{-1} \\ \downarrow{\scriptstyle \nu_I} & & \downarrow{\scriptstyle \nu_J} & & \\ \pi_1(Z_i, z_i') & \xrightarrow{\tau_\eta} & \pi_1(Z_i, z_i') & \gamma \longmapsto & \eta * \gamma * \eta^{-1} \end{array} \tag{5.6.7}$$

The upper square commutes by (5.6.6) and the lower square commutes because g acts as a deck transformation of $C_{\hat{\Phi}} \to C_\Phi$. Identifying $\dot{G} = \pi_1(\dot{X}, \dot{x})$, the top map is the automorphism $\dot{G} \to \dot{G}$ given by g-conjugation. And up to choices of maps τ_δ from (5.6.5), the compositions $\nu_I\rho_I$ and $\nu_J\rho_J$ correspond to the homomorphisms $\lambda_I, \lambda_J : \dot{G} \to P_i$. Changing the choice of map τ_δ corresponds to conjugating by an element of P_i, so property (3) follows from diagram (5.6.7).

□

5.6.3 Towards commanding subgroups

Lemma 5.6.3 can be used to deduce the following theorem, which is a weak form of a virtually special cubulated group $G \curvearrowright X$ commanding a malnormal collection of convex subgroups (P_i). Like the commanding property, it allows us to construct a finite-index normal subgroup $G' \triangleleft G$ and control the intersections $P_i \cap G'$ independently; but, unlike commanding, $P_i \cap G'$ must belong to a fixed sequence (P_{il}) of subgroups of P_i. The author believes upgrading this theorem to the full strength of commanding is a difficult task, but there are intermediate upgrades that would already be very useful; for instance if each sequence (P_{il}) only depended on P_i as an abstract group, then the theorem would provide a powerful tool for constructing finite covers of graphs of groups for which each vertex group satisfies the assumptions of the theorem with respect to its incident edge groups (see Proposition 5.4.8).

Theorem 5.6.4. *Let X be a finite virtually special cube complex, and let $P_1, ..., P_n < G := \pi_1 X$ be convex subgroups such that all conjugates of P_i and P_j have trivial intersection if $i \neq j$. Then for each i there is a descending sequence of finite-index subgroups $P_i > P_{i1} > P_{i2} > \cdots$ with trivial intersection, such that for any map $\sigma : \{1, ..., n\} \to \mathbb{N}$ there exists a finite-index normal subgroup $G' \triangleleft G$ with $P_i \cap G' = P_{i\sigma(i)}$.*

Proof. Apply Lemma 5.6.3. For each i, choose a descending sequence of finite-index subgroups $\dot{P}_i > \dot{P}_{i1} > \dot{P}_{i2} > ...$ with trivial intersection, that are all normal in P_i (such a sequence exists by residual finiteness of P_i). Set

$$P_{il} := \dot{P}_i \cap \bigcap_{\lambda \in \Lambda_i} \lambda^{-1}(\dot{P}_{il}).$$

Λ_i is finite, so P_{il} has finite index in P_i, and the sequence (P_{il}) is clearly descending. Lemma 5.6.3(1) implies that $P_{il} < \dot{P}_{il}$, so $\cap_l P_{il} = \{1\}$. Given $\sigma : \{1, ..., n\} \to \mathbb{N}$, set

$$G' := \bigcap_{i=1}^{n} \bigcap_{\lambda \in \Lambda_i} \lambda^{-1}(\dot{P}_{i\sigma(i)}) < \dot{G}.$$

The collections Λ_i are finite, so G' has finite index in $\dot{G}$, hence also in G. Given $g \in G$ and $\lambda \in \Lambda_i$, Lemma 5.6.3(3) implies that there exists $\lambda' \in \Lambda_i$ with

$$g\lambda^{-1}(\dot{P}_{i\sigma(i)})g^{-1} = (\lambda')^{-1}(\dot{P}_{i\sigma(i)}),$$

so we deduce that G' is normal in G. Finally, by Lemma 5.6.3(2) we have

$$\begin{aligned} P_i \cap G' &= \dot{P}_i \cap G' \\ &= \dot{P}_i \cap \bigcap_{\lambda \in \Lambda_i} \lambda^{-1}(\dot{P}_{i\sigma(i)}) \\ &= P_{i\sigma(i)}. \end{aligned}$$

□

5.6.4 Proof of Theorem 1.2.3

Recall that a collection of group elements $\{g_1, ..., g_n\}$ is *independent* if the elements g_i have infinite order and no non-zero power of g_i is conjugate to a non-zero power of g_j for $i \neq j$.

Theorem 1.2.3. *Every virtually special cubulated group $G \curvearrowright X$ strongly commands every independent set of convex elements.*

Proof. Let $\{g_1, ..., g_n\}$ be an independent set of convex elements. Given a finite-index torsion-free subgroup $\hat{G} \triangleleft G$, we can choose an independent set of elements $\{h_{ij}\}$ in $\hat{G}$ as in Lemma 5.4.10, and reduce to the problem of whether $\hat{G}$ strongly commands $\{h_{ij}\}$. Lemma 5.4.10 requires G to be balanced, but this follows from Corollary 2.3.7. Each h_{ij} is conjugate to a power of g_i, so is convex. To simplify notation, we will henceforth assume that G is already torsion-free and prove that it strongly commands $\{g_i\}$.

Since G acts freely on X, we can now apply Lemma 5.6.3 to X/G and the subgroups $P_i := \langle g_i \rangle$. This gives us a finite-index normal subgroup $\dot{G} \triangleleft G$, and for each i a finite collection Λ_i of homomorphisms $\lambda : \dot{G} \to P_i$. Suppose $P_i \cap \dot{G} = \langle g_i^{M_i} \rangle$, and write $\lambda(g_i^{M_i}) = g_i^{M_\lambda}$ for $\lambda \in \Lambda_i$. The three properties of Lemma 5.6.3 become the following:

(1) For each i, there is $\lambda \in \Lambda_i$ with $M_\lambda = M_i$.

(2) $\lambda(g_j^{M_j}) = 1$ for $\lambda \in \Lambda_i$ and $i \neq j$.

(3) For any $\lambda \in \Lambda_i$ and $g \in G$ there is some $\lambda' \in \Lambda_i$ such that $\lambda(\dot{g}) = \lambda'(g\dot{g}g^{-1})$ for all $\dot{g} \in \dot{G}$.

Let $N = \prod_i M_i$, and let $d_i > 0$ be the greatest common divisor of $\{M_\lambda \mid \lambda \in \Lambda_i\}$ (well-defined because (1) guarantees that at least one M_λ is non-zero). Given integers $r_i > 0$ we will now define a finite quotient G/G' with quotient map $g \mapsto \bar{g}$, such that $\bar{g}_i$ has order Nr_i, and such that the subgroups $\langle \bar{g}_i \rangle$ have pairwise trivial intersections. Indeed we define $G' \triangleleft G$ by

$$G' := \bigcap_{i=1}^{n} \bigcap_{\lambda \in \Lambda_i} \lambda^{-1}(\langle g_i^{Nd_i r_i / M_i} \rangle).$$

The collections Λ_i are finite, so G' has finite index in $\dot{G}$, hence also in G. And property (3) ensures that G' is normal in G. It remains to prove the following two claims.

<u>Claim:</u> $\bar{g}_i$ has order Nr_i.

<u>Proof:</u> The only powers of g_i that land in $\dot{G}$ are powers of $g_i^{M_i}$; since $G' < \dot{G}$, we must therefore find the smallest integer $l > 0$ with $\bar{g}_i^{M_i l} = 1$. Property (2) tells us that $\lambda(g_i^{M_i l}) = 1$ for $\lambda \in \Lambda_j$ and $i \neq j$, so $\bar{g}_i^{M_i l} = 1$ is equivalent to $\lambda(g_i^{M_i l}) \in \langle g_i^{Nd_i r_i / M_i} \rangle$ for all $\lambda \in \Lambda_i$. As $\lambda(g_i^{M_i l}) = g_i^{M_\lambda l}$, this is equivalent to $Nd_i r_i / M_i$ dividing $M_\lambda l$ for all $\lambda \in \Lambda_i$, which in turn is equivalent to Nr_i / M_i dividing l. Thus $\bar{g}_i$ has order Nr_i as required. ■

<u>Claim:</u> The subgroups $\langle \bar{g}_i \rangle$ have pairwise trivial intersections.

<u>Proof:</u> Suppose for contradiction there is $i \neq j$ with $\langle \bar{g}_i \rangle \cap \langle \bar{g}_j \rangle$ non-trivial, say it contains an element of order p, with p prime. Then p divides $|\langle \bar{g}_i \rangle| = Nr_i$ and $|\langle \bar{g}_j \rangle| = Nr_j$.

Since M_iM_j divides N, we deduce that p divides at least one of Nr_i/M_i and Nr_j/M_j, say $pl_i = Nr_i/M_i$ for some integer l_i. Then

$$1 \neq \bar{g}_i^{M_i l_i} \in \langle \bar{g}_i \rangle \cap \langle \bar{g}_j \rangle \cap \dot{G}/G'.$$

But $\langle \bar{g}_j \rangle \cap \dot{G}/G' = \langle \bar{g}_j^{M_j} \rangle$, so $\bar{g}_i^{M_i l_i} = \bar{g}_j^{M_j l_j}$ for some integer l_j. Finally, we show that

$$g_i^{M_i l_i} G' = g_j^{M_j l_j} G' \subset \bigcap_{\lambda \in \Lambda_k} \lambda^{-1}(\langle g_k^{N d_k r_k / M_k} \rangle) \tag{5.6.8}$$

for each k, hence contradicting $\bar{g}_i^{M_i l_i} \neq 1$. Indeed, property (2) tells us that $\lambda(g_i^{M_i l_i}) = 1$ for $\lambda \in \Lambda_k$ and $k \neq i$, so (5.6.8) holds for $k \neq i$, and then the same argument applied to $g_j^{M_j l_j}$ establishes (5.6.8) for $k \neq j$. ■

□

5.6.5 Right-angled Artin groups

Right-angled Artin groups (or RAAGs) are a good source of virtually special cubulated groups that are well understood but not hyperbolic, so they are the perfect setting to apply Theorem 1.2.3. In this subsection we investigate when elements of RAAGs are independent or convex, and in Theorem 5.6.9 we deduce that RAAGs command random collections of elements.

Given a finite simplicial graph Γ, the *right-angled Artin group* (or *RAAG*) G_Γ has generators corresponding to the vertex set V of Γ and relations $[v_1, v_2]$ whenever $v_1, v_2 \in V$ are adjacent in Γ. There are very explicit solutions to the word and conjugacy problems in G_Γ, which we now describe. These solutions are implicit in the work of Servatius [82] and Hermiller-Meier [49]. Write $V^\pm$ for the set of generators and their inverses. An element $g \in G_\Gamma$ can be represented by a word w on the letters $V^\pm$, and we say that w is *reduced* if its length $|w|$ is minimal among all words representing g. A solution to the word problem is given by the following proposition.

Proposition 5.6.5. *If w, w' are words representing $g \in G_\Gamma$, with w' reduced, then we can transform w into w' via a sequence of the following two moves:*

(M1) Remove a subword $v^{-1}v$ with $v \in V^\pm$.

(M2) Replace a subword v_1v_2 with v_2v_1 where $v_1, v_2 \in V^\pm$ correspond to adjacent vertices.

In particular, it follows that $v_1, v_2 \in V^\pm$ commute in G_Γ if and only if their corresponding vertices are equal or adjacent. It also follows that any two reduced words representing g differ by (M2) moves.

We now turn to the conjugacy problem. We say that a word w representing $g \in G_\Gamma$ is *cyclically reduced* if it has minimal length among words representing conjugates of g.

Proposition 5.6.6. *If w, w' are words representing conjugate elements in G_Γ, with w' cyclically reduced, then we can transform w into w' using the moves (M1), (M2) and the following cyclic permutation move:*

(M3) Replace the word vw with wv, where $v \in V^\pm$.

Moreover, w is cyclically reduced if and only if it cannot be transformed via (M2) moves to a word of the form vuv^{-1} with $v \in V^\pm$.

In particular, any two cyclically reduced words representing conjugate elements differ by (M2) and (M3) moves. And it follows from the moreover part of Proposition 5.6.6 that any power of a cyclically reduced word is cyclically reduced.

We are now ready to describe an efficient algorithm to check independence of elements.

Proposition 5.6.7. *There is an explicit algorithm that determines whether a collection of elements in G_Γ is independent.*

Proof. It suffices to determine whether a pair of non-trivial elements $\{g_1, g_2\}$ is independent. Let w_1, w_2 be cyclically reduced words representing conjugates of g_1, g_2. Suppose that $\{g_1, g_2\}$ is not independent. Then there exist non-zero powers w_1^m, w_2^n that represent conjugate elements. We know that w_1^m, w_2^n are also cyclically reduced, hence they differ by (M2) and (M3) moves - in particular $m|w_1| = |w_1^m| = |w_2^n| = n|w_2|$. Furthermore, if m, n have a common factor, say $m = Md, n = Nd$, then we deduce that w_1^M, w_2^N also represent conjugate elements (we always have $g^d = h^d$ implies $g = h$ for elements of G_Γ by Proposition 5.6.5). Thus we obtain the following algorithm:

(1) Find cyclically reduced words w_1, w_2 representing conjugates of g_1, g_2 using using (M1)–(M3) moves.

(2) Take non-zero integers m, n with $m|w_1| = n|w_2|$.

(3) Check whether w_1^m can be transformed into either w_2^n or w_2^{-n} via (M2)–(M3) moves. $\{g_1, g_2\}$ is independent if and only if the answer is no. □

Next we turn to the cubical geometry of RAAGs. G_Γ is the fundamental group of a finite special cube complex S_Γ, called the Salvetti complex - see [44]. Let X_Γ be the universal cover of S_Γ, so X_Γ is a CAT(0) cube complex with a free cocompact G_Γ action. Moreover, the one-skeleton of X_Γ is precisely the Cayley graph of G_Γ with respect to the generators V. So there is an identification $G_\Gamma = X_\Gamma^0$, and each oriented edge in X_Γ is labelled by an element of $V^\pm$. Furthermore, the squares in X_Γ correspond to commuting relations in G_Γ.

We also have geometric interpretations for reduced and cyclically reduced words. For a reduced word w representing an element $g \in G_\Gamma$, there is a combinatorial geodesic γ_w

in X_Γ from 1 to g whose edge labels spell the word w. If in addition w is cyclically reduced, then all powers w^n are reduced. Then any segment of the union $\gamma = \cup_{i\in\mathbb{Z}} g^i\gamma_w$ is contained in a translate of some γ_{w^n}, so γ is a bi-infinite combinatorial geodesic admitting a $\langle g\rangle$-action by translations.

For w a word on $V^{\pm}$, we define the *support* of w to be

$$\operatorname{supp}(w) := \{v \in V \mid v \text{ or } v^{-1} \text{ appears in } w\}.$$

We are now ready to prove a criterion for elements of G_Γ to be convex with respect to X_Γ. Note that convexity is invariant under conjugation, so it is enough to consider cyclically reduced words.

Proposition 5.6.8. *Let w be a cyclically reduced word representing $g \in G_\Gamma$. Then we have the following dichotomy:*

(1) g is convex with respect to X_Γ.

(2) There is a non-trivial partition $\operatorname{supp}(w) = A \sqcup B$*, such that elements of A commute with elements of B.*

Proof. By Behrstock–Charney [7, Lemma 5.1] we have that the centraliser $C(g)$ is non-cyclic if and only if g is contained in a join subgroup of G_Γ. A join subgroup is a subgroup of the form $\langle A \sqcup B\rangle$ for $\emptyset \neq A, B \subset V$ disjoint subsets such that elements of A commute with elements of B. In particular, for any $n > 0$ we have that $C(g)$ is cyclic if and only if $C(g^n)$ is cyclic (remember that powers of cyclically reduced words are cyclically reduced).

We now prove that w must satisfy at least one of the cases (1) and (2). We consider two cases. Firstly suppose that $C(g)$ is cyclic. Applying the Cubical Flat Torus Theorem of Wise–Woodhouse [99, Theorem 3.6] to the *highest* virtually abelian subgroup $\langle g\rangle$ (meaning that no power of g is contained in a $\mathbb{Z}^2$ subgroup), we deduce that g is convex with respect to X_Γ. Secondly suppose that $C(g)$ is non-cyclic. Then as above we know that g lies in a join subgroup $\langle A \sqcup B\rangle$, and by Proposition 5.6.6 we deduce that $\operatorname{supp}(w) \subset A \sqcup B$. If $\operatorname{supp}(w)$ intersects both A and B then w satisfies (2) (shrinking A or B if necessary). Otherwise, if $\operatorname{supp}(w) \subset A$, then $g \in \langle A\rangle \cong G_{\Gamma[A]}$ (where $\Gamma[A]$ is the induced subgraph of A); moreover there is a canonical embedding of $X_{\Gamma[A]}$ in X_Γ as a convex subcomplex which is equivariant with respect to the isomorphism $\langle A\rangle \cong G_{\Gamma[A]}$. Since $\Gamma[A]$ has fewer vertices than Γ, we may conclude by induction that at least one of the cases (1) and (2) is satisfied.

It remains to show that cases (1) and (2) are mutually exclusive. Indeed, suppose we are in case (2), then we may reorder the word w using (M2) moves so that it takes the form $w = w_1w_2$, where $\operatorname{supp}(w_1) = A$ and $\operatorname{supp}(w_2) = B$. The words w_1, w_2 give rise to bi-infinite combinatorial geodesics $\gamma_1, \gamma_2 \subset X_\Gamma$ through 1. The labels of γ_1 commute with the labels of γ_2, so we have a product subcomplex $Y = \gamma_1 \times \gamma_2 \subset X_\Gamma$, and the geodesic γ zig-zags across the diagonal of Y. Importantly, the hyperplanes that cross Y are precisely the hyperplanes that cross γ, so any convex subcomplex containing γ must also contain Y. Now any $\langle g\rangle$-invariant convex subcomplex $Z \subset X_\Gamma$ must contain γ in a bounded neighbourhood, so Z has a cubical thickening that contains γ, and hence also contains Y. As a result, Z cannot be a quasiline, and g cannot be convex. □

Armed with the previous two propositions, we can now show that RAAGs command collections of elements produced by randomly picking long words on $V^{\pm}$.

Theorem 5.6.9. *(RAAGs command random elements)*
Fix n, and let $w_1, ..., w_n$ be words on $V^{\pm}$ chosen uniformly at random from among all freely reduced words of length at most L, and suppose they represent elements $g_1, ..., g_n \in G_\Gamma$. Then G_Γ commands $\{g_1, ..., g_n\}$ with probability converging to 1 as $L \to \infty$ - unless Γ is a complete graph with fewer than n vertices.

Proof. If Γ is a complete graph with at least n vertices then G_Γ is a free abelian group of rank at least n, and the theorem follows easily from Proposition 5.4.5.

Now suppose that Γ is not complete. For $v \in V$ let $\pi_v : G_\Gamma \to \mathbb{Z}$ be the homomorphism that maps $v \mapsto 1$ and $v' \mapsto 0$ for $v \neq v' \in V$. Note that $\pi_v(g_i)$ is the number of occurrences of v in the word w_i subtract the number of occurrences of v^{-1}. We will show that G_Γ commands $\{g_1, ..., g_n\}$ provided the following two conditions are satisfied. These conditions are clearly satisfied for long random words $w_1, ..., w_n$, i.e. with probability converging to 1 as $L \to \infty$.

(1) $\pi_v(g_i) \neq 0$ for all $v \in V$ and $1 \leq i \leq n$.

(2) For each pair $\{v, v'\} \subset V$, the ratios $(\pi_v(g_i)/\pi_{v'}(g_i) \mid 1 \leq i \leq n)$ are distinct.

Take a maximal partition $V = \sqcup_{k=1}^{q} V_k$ such that any two vertices in distinct V_k are joined by an edge. If the induced subgraph of V_k is Γ_k, then we get a product decomposition

$$G_\Gamma = \prod_{k=1}^{q} G_{\Gamma_k}.$$

Since Γ is not complete, $|V_k| > 1$ for some k, so suppose $|V_1| > 1$. For $v \in V_1$ we can define a homomorphism $\pi_v : G_{\Gamma_1} \to \mathbb{Z}$ as we did for G_Γ, and for $\mu : G_\Gamma \to G_{\Gamma_1}$ the projection map we have $\pi_v = \pi_v \mu : G_\Gamma \to \mathbb{Z}$. Hence conditions (1) and (2) apply to $\mu(g_1), ..., \mu(g_n)$ and vertices in V_1.

The homomorphisms π_v are conjugation invariant, so condition (1) implies that any cyclically reduced word representing a G_{Γ_1}-conjugate of $\mu(g_i)$ will have support equal to the whole of V_1. By Proposition 5.6.8 and the maximality of the partition $\sqcup_k V_k$, we deduce that the elements $\mu(g_i)$ are convex with respect to X_{Γ_1}. The ratios $\pi_v(g_i)/\pi_{v'}(g_i)$ are preserved by taking conjugates and powers, so the collection $\{\mu(g_1), ..., \mu(g_n)\}$ is independent by condition (2). Thus G_{Γ_1} commands $\{\mu(g_1), ..., \mu(g_n)\}$ by Theorem 1.2.3. By composing μ with suitable homomorphisms from G_{Γ_1} to finite groups, it follows that G_Γ commands $\{g_1, ..., g_n\}$. □

5.7 Virtually connected intersections

The aim of this section is to prove Theorem 1.2.8, which we restate below for the reader's convenience. This generalises a theorem of Haglund and Wise [45, Theorem 4.25] to the non-hyperbolic setting.

Theorem 1.2.8. *(Virtually connected intersections)*
For $i = 1, ..., n$ let $(Z_i, z_i) \to (X, x)$ be based local isometries of finite virtually special cube complexes. Then there is a finite cover $(\dot{X}, \dot{x}) \to (X, x)$ such that the based elevations $\dot{Z}_i$ of Z_i are embedded in $(\dot{X}, \dot{x})$ and the intersections $\cap_{i\in E}\dot{Z}_i$ are connected for every non-empty $E \subset \{1, ..., n\}$. Moreover, if the Z_i are embedded in X and do not inter-osculate with hyperplanes of X, then we may assume that the full intersection $\cap_{i=1}^n Z_i$ is isomorphic to its based elevation.

5.7.1 Hierarchies of imitators

In order to prove Theorem 1.2.8 we generalise the imitator covers from Construction 5.5.3 to ones obtained from hierarchies of imitators.

Construction 5.7.1. (Hierarchies of imitators)
Let $(X, \mathcal{Z})$ be a directly special cube complex pair. A *hierarchy of imitators* $(\mathcal{I}, \xi)$ is a set of imitators $\mathcal{I}$ who live in $\mathcal{Z}$ as before, but now an imitator $I \in \mathcal{I}$ might copy another imitator instead of the walker - we say imitator I copies $\xi(I) \in \mathcal{I} \sqcup \{W\}$, where W is the walker. The way I copies $\xi(I)$ is essentially the same as before; if $\xi(I) = W$ then it is exactly as before, otherwise we have the following rule: if I and $\xi(I)$ are at vertices $z_1, z_2 \in \mathcal{Z}$ and $\xi(I)$ traverses the edge $f_2 \in \text{link}(z_2)$, then I traverses the edge $f_1 \in \text{link}(z_1)$ whose image in X is parallel to that of f_2, if such an edge exists, otherwise I remains at z_1. We assume that each imitator $I \in \mathcal{I}$ defines a finite *imitator chain* $I = I_0, I_1, ..., I_m \in \mathcal{I}$, where $\xi(I_s) = I_{s+1}$ for $0 \leq s < m$ and $\xi(I_m) = W$; so each movement of the walker causes all of the imitators to move (or not move) by iteratively following the aforementioned movement rule.

Lemma 5.1.2 implies that $\xi(I)$ going round the boundary of a square causes I to return to their starting point, hence the movements of imitators and walker define a cover $\mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X) \to X$ as in Construction 5.5.3. Unlike before, we allow $\mathcal{I}$ to be infinite and we allow multiple imitators to occupy the same vertex, so in general $\mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X)$ will be infinite. As before, for $I \in \mathcal{I}$ we have a cellular map $r_I : \mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X) \to \mathcal{Z}$ that projects to the position of imitator I. In keeping with the notation of Constructions 5.1.1 and 5.5.3, if imitators and walker start at positions (θ, x) and the walker traverses a path γ, then I traverses a path $\delta_I(\gamma, \theta)$.

Given a vertex $(\theta, x) \in \mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X)$, we let $\mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X; \theta, x)$ denote the component of $\mathsf{C}_{\mathcal{I},\xi}(\mathcal{Z}, X)$ containing (θ, x), and call this the *ξ-hierarchical imitator cover of $(X, \mathcal{Z})$ based at (θ, x)*.

5.7.2 Proof of Theorem 1.2.8

Using Corollary 5.2.7, we reduce to the case where the Z_i are embedded in X and do not inter-osculate with hyperplanes. Let $(X, \mathcal{Z})$ denote the corresponding cube complex pair. We will consider the Z_i as subcomplexes of X rather than giving names to the inclusions $Z_i \hookrightarrow X$. Fix a ξ-hierarchical imitator cover $(\dot{X}, \dot{x}) \to (X, x)$ of $(X, \mathcal{Z})$ based at $\dot{x} = (\theta, x)$, with $\theta(I) = x$ for all $I \in \mathcal{I}$ - we can think of this as all imitators and the walker starting at the vertex x. Say that each imitator $I \in \mathcal{I}$ is contained in the subcomplex $Z_{\kappa(I)}$, and write

$$\boldsymbol{\kappa}(I) := (\kappa(I_0), \kappa(I_1), ..., \kappa(I_m)), \tag{5.7.1}$$

where $I = I_0, I_1, ..., I_m \in \mathcal{I}$ is the imitator chain of I. If imitator I starts at a vertex z and $\xi(I)$ traverses a path γ, then the path taken by I only depends on γ, z and $i := \kappa(I)$, call this path $\delta_i(\gamma, z)$ (this is really the same as $\delta(\gamma, z)$ from Construction 5.1.1, we just need the parameter i to remember which subcomplex the imitator is in). Note that

$$\delta_i(\gamma, z) = \delta_I(\gamma, \theta') \tag{5.7.2}$$

if $\kappa(I) = (i)$ and $\theta'(I) = z$. We also have a recursive identity

$$\delta_I(\gamma, \theta') = \delta_{\kappa(I)}(\delta_{\xi(I)}(\gamma, \theta'), \theta'(I)), \tag{5.7.3}$$

for all imitators $I \in \mathcal{I}$, points $(\theta', x') \in \dot{X}$ and paths γ in X based at x'. We now prove some lemmas that will help us to understand the movements of imitators in the hierarchy.

Lemma 5.7.2. *For $I \in \mathcal{I}$, the map $r_I : \dot{X} \to Z_{\kappa(I)}$ only depends on the multi-index $\boldsymbol{\kappa}(I)$.*

Proof. Since $\dot{x} \in \dot{X}$, the lemma in equivalent to: if imitators and the walker all start at x, then the movements of I only depend on $\boldsymbol{\kappa}(I)$ and the movements of the walker. This follows from (5.7.2) and (5.7.3) applied to $(\theta', x') = (\theta, x)$. □

Lemma 5.7.3. *If $\boldsymbol{\kappa}(I) = (i_0, ..., i_m)$, then $r_I(\dot{X}) \subset \cap_s Z_{i_s}$.*

Proof. We prove this by induction on m. It is immediate for $m = 0$, so suppose $m > 0$. Writing $I_1 := \xi(I)$, we know that $\boldsymbol{\kappa}(I_1) = (i_1, ..., i_m)$, so imitator I_1 must stay in the subcomplex $\cap_{s\geq 1} Z_{i_s}$ by induction. It is enough to consider paths in $\dot{X}$ that start at $\dot{x}$, and these will correspond to movements of imitators where I and I_1 both start at $x \in \cap_s Z_{i_s}$. Since the subcomplexes Z_i do not inter-osculate with hyperplanes, and since I is copying the movements of I_1, it follows from Subcomplex Entrapment (Lemma 5.2.4) that I stays inside $\cap_s Z_{i_s}$. □

Lemma 5.7.4. *If $\boldsymbol{\kappa}(I) = (i_0, ..., i_m)$, and $i_0 = i_s =: i$ for $s > 0$, then $r_I = r_{I_1} : \dot{X} \to Z_i$, where $I_1 = \xi(I)$.*

Proof. Once again, it is enough to consider paths in $\dot{X}$ that start at $\dot{x}$, and these will correspond to movements of imitators where I and I_1 both start at $x \in Z_i$. I copies the movements of I_1 subject to staying in Z_i, but we know that I_1 also stays in Z_i by Lemma 5.7.3, hence the movements of I and I_1 are identical. □

Lemma 5.7.5. *$\dot{X} \to X$ is a finite cover.*

Proof. Lemma 5.7.4 implies that there are only finitely many distinct maps r_I. The result follows because a vertex of $\dot{X}$ is determined by its images under the projections r_I and the projection to X. □

We now make the additional assumption that every finite sequence of integers in $\{1, ..., n\}$ is realised by some $\boldsymbol{\kappa}(I)$. It is straightforward to construct such a hierarchy of imitators - in fact we could construct a canonical one in which $\mathcal{I}$ is in bijection with the finite sequences of integers in $\{1, ..., n\}$. We now prove the main part of Theorem 1.2.8.

Lemma 5.7.6. *If $\dot{Z}_i$ is the based elevation of Z_i to $(\dot{X}, \dot{x})$, then $\cap_{i\in E} \dot{Z}_i$ is connected for any $\emptyset \neq E \subset \{1, ..., n\}$.*

Proof. We induct on $|E|$. The case $|E| = 1$ is immediate, so suppose $|E| > 1$ and write $E = F \sqcup \{j, k\}$. We know that $\dot{x} \in \cap_{i\in E} \dot{Z}_i$, so consider another vertex $\dot{x}' \in \cap_{i\in E} \dot{Z}_i$. Our task is to find a path in $\cap_{i\in E} \dot{Z}_i$ from $\dot{x}$ to $\dot{x}'$.

By induction we know that $\cap_{i\in F\sqcup\{j\}} \dot{Z}_i$ and $\cap_{i\in F\sqcup\{k\}} \dot{Z}_i$ are connected, so let $\dot{\gamma}_j$ and $\dot{\gamma}_k$ be paths from $\dot{x}$ to $\dot{x}'$ in these respective intersections. Suppose these project to paths γ_j and γ_k in X. Let $J \in \mathcal{I}$ be an imitator with $\boldsymbol{\kappa}(J) = (j)$, and put $\gamma := r_J \dot{\gamma}_k$, which is the path taken by J as the walker traverses γ_k. The proof concludes with the following claim.

<u>Claim:</u> γ lifts to a path $\dot{\gamma}$ from $\dot{x}$ to $\dot{x}'$ in $\cap_{i\in E} \dot{Z}_i$.

Proof: The path γ_k stays in the intersection $\cap_{i\in F\sqcup\{k\}}Z_i$, so γ stays in $\cap_{i\in E}Z_i$ by Subcomplex Entrapment. If $\dot\gamma$ is the lift of γ based at $\dot x$, it follows that $\dot\gamma$ lies in $\cap_{i\in E}\dot Z_i$. So it remains to show that $\dot\gamma$ ends at $\dot x'$.

Any point in $\dot X$ is determined by its images under the projections r_I for $I \in \mathcal{I}$ and the projection to X (the walker coordinate). The path $\dot\gamma_j$ ends at $\dot x'$, so it suffices to show that $r_I\dot\gamma$ and $r_I\dot\gamma_j$ have the same endpoint for all $I \in \mathcal{I}$ (the walker coordinate projections γ and γ_j both lie in Z_j so are equal to the paths $r_J\dot\gamma$ and $r_J\dot\gamma_j$ respectively).

Let $I \in \mathcal{I}$ with imitator chain $I = I_0, I_1, ..., I_m$. Let $I' \in \mathcal{I}$ be another imitator with $\boldsymbol{\kappa}(I') = (\boldsymbol{\kappa}(I), j)$ and imitator chain $I' = I'_0, I'_1, ..., I'_{m+1}$. If the walker traverses γ_j, then imitator I'_{m+1} also traverses γ_j (as γ_j is in Z_j). Since $\kappa(I_s) = \kappa(I'_s)$ for $0 \le s \le m$, (5.7.3) implies that I_s and I'_s follow the same path as the walker traverses γ_j - in particular

$$r_I\dot\gamma_j = r_{I'}\dot\gamma_j. \tag{5.7.4}$$

Observe that $\boldsymbol{\kappa}(I'_{m+1}) = (j) = \boldsymbol{\kappa}(J)$, so I_{m+1} and J both follow the path γ if the walker traverses γ_k. And iterating up the imitator chains again, we see that $r_{I_s}\dot\gamma = r_{I'_s}\dot\gamma_k$ for $0 \le s \le m$ - in particular

$$r_I\dot\gamma = r_{I'}\dot\gamma_k. \tag{5.7.5}$$

But $r_{I'}\dot\gamma_j$ and $r_{I'}\dot\gamma_k$ both end at $r_{I'}(\dot x')$, so (5.7.4) and (5.7.5) imply that $r_I\dot\gamma$ and $r_I\dot\gamma_j$ also end at $r_{I'}(\dot x')$, as required. ■

□

Finally, the assertion of the last sentence of Theorem 1.2.8 follows easily from way we constructed $\dot X$ with imitators: any loop γ in $\cap_{i=1}^n Z_i$ based at x lifts to a loop in $\dot X$ based at $\dot x$ in which all the imitators and walker traverse γ together. □

Chapter 6

Agol's theorem on hyperbolic cubulations

In [1] Agol proved that hyperbolic cubulated groups are virtually special (Theorem 1.2.1). The aim of this chapter is to make the proof accessible to a wider audience; we retain the underlying ideas and constructions of Agol, but substantially change or add to many parts of the argument to give a more transparent and detailed account. Throughout this chapter we will always write "Agol's paper" to refer to [1].

6.1 Virtually special hyperbolic groups

We will now state some powerful theorems showing that hyperbolic groups enjoy some strong properties if they are fundamental groups of virtually special cube complexes. The first of these theorems, due to Wise, characterises the hyperbolic fundamental groups of compact NPC virtually special cube complexes using the following quasiconvex hierarchy of hyperbolic groups.

Definition 6.1.1. Let $\mathcal{QVH}$ denote the smallest class of hyperbolic groups that is closed under the following operations:

(1) $1 \in \mathcal{QVH}$.

(2) If $G = A *_B C$ and $A, C \in \mathcal{QVH}$ and B is finitely generated and quasiconvex in G, then $G \in \mathcal{QVH}$.

(3) If $G = A*_B$ and $A \in \mathcal{QVH}$ and B is finitely generated and quasiconvex in G, then $G \in \mathcal{QVH}$.

(4) Let $\hat{G} < G$ with $|G : \hat{G}| < \infty$. If $\hat{G} \in \mathcal{QVH}$ then $G \in \mathcal{QVH}$.

Theorem 6.1.2. *(Wise, 2011)[96, Theorem 13.3]*
A torsion-free hyperbolic group is in $\mathcal{QVH}$ if and only if it is the fundamental group of a compact virtually special cube complex.

In particular, this implies that a torsion-free hyperbolic group is virtually special if and only if it is the fundamental group of a compact virtually special cube complex. Theorem 6.1.2 is even more powerful when used in combination with the following theorem, which appeared in Haglund and Wise's original paper on special cube complexes (subgroup separability is defined in Definition 2.5.1).

Theorem 6.1.3. *(Haglund–Wise, 2008)[44, Theorems 1.3 and 1.4]*
Let X be a compact NPC cube complex with $\pi_1(X)$ hyperbolic. Then X is virtually special if and only if every quasiconvex subgroup of $\pi_1(X)$ is separable.

Corollary 6.1.4. *Let X and Y be compact NPC cube complexes with $\pi_1(X) \cong \pi_1(Y)$ hyperbolic. Then X is virtually special if and only if Y is virtually special.*

These theorems will be used in Section 6.7 in the following way. They allow us to glue together two compact NPC virtually special cube complexes X_1 and X_2 along locally convex subcomplexes $Z_1 \cong Z_2$, to produce a larger virtually special cube complex X - provided that X_1, X_2 and X have hyperbolic fundamental groups. We first argue that the composite X is NPC, and then by considering universal covers we argue that $\pi_1(Z_1)$ is quasiconvex in $\pi_1(X)$ (these things will be explained more carefully later in the chapter). We may then apply Theorem 6.1.2 to deduce that $\pi_1(X) = \pi_1(Y)$ for some compact NPC virtually special cube complex Y, and finally Corollary 6.1.4 implies that X itself is virtually special.

One problem we will face in Section 6.7 is that we might need to pass to finite covers of X_1 and X_2 in order to get isomorphic subcomplexes that we can glue along. For this we will use the fact that virtually special hyperbolic groups command almost malnormal collections of quasiconvex subgroups (Theorem 1.2.5). We restate what this means below.

Theorem 6.1.5. *Let G be a virtually special hyperbolic group. Let $\{K_1, \dots, K_m\}$ be an almost malnormal collection of quasiconvex subgroups. Then there exist finite index subgroups $\dot{K}_i \triangleleft K_i$, such that for any further finite index subgroups $\hat{K}_i < \dot{K}_i$, with $\hat{K}_i \triangleleft K_i$, there exists a finite index subgroup $\hat{G} \triangleleft G$ with $\hat{G} \cap K_i = \hat{K}_i$.*

The last big theorem we state in this section, due to Agol, Groves and Manning, appears in the appendix of Agol's paper. We will use this theorem in the next section to take a quotient of the CAT(0) cube complex X from Theorem 1.2.1 that makes the hyperplanes finite.

Theorem 6.1.6. *(Agol–Groves–Manning, 2013)[1, Theorem A.1]*
Let G be a hyperbolic group and $K < G$ a quasiconvex virtually special subgroup. Then for any $g \in G - K$, there is a hyperbolic group $\mathcal{G}$ and a homomorphism $\phi : G \to \mathcal{G}$ such that $\phi(g) \notin \phi(K)$ and $\phi(K)$ is finite.

6.2 Making hyperplanes finite

For the rest of this chapter let G be a hyperbolic group acting properly and cocompactly on a CAT(0) cube complex X as in Theorem 1.2.1. The object of this section is to construct a quotient map $X \to \mathcal{X}$ such that hyperplanes in $\mathcal{X}$ are finite, and so that distinct hyperplanes in X which are close together map to distinct hyperplanes in $\mathcal{X}$. The quotient complex $\mathcal{X}$ will be important for defining the local colouring data used in later sections. This section is based on §4 of Agol's paper, but with considerably more detail added.

We may assume X is unbounded since the theorem is trivial otherwise. For $x, y \in X$ we will use $[x, y]$ to denote the unique geodesic segment between them (with respect to the CAT(0) metric, as always).

Remark 6.2.1. By passing to the cubical subdivision of X, we can assume that for every hyperplane H in X, G_H does not exchange the sides of H, so $gH^{\pm} = H^{\pm}$ for all $g \in G_H$. By appropriate choice of labelling we can also assume that $gH^{\pm} = (gH)^{\pm}$ for all $g \in G$.

Proposition 6.2.2. *X is finite dimensional, locally finite and δ-hyperbolic (for some δ).*

Proof. X is finite dimensional because G acts cocompactly on it. Now suppose X is not locally finite and that $x \in X$ is a vertex contained in infinitely many cubes. By cocompactness there is a cube C and $A \subset G$ such that $A \cdot C$ is an infinite family of cubes each containing x; but then C must have a vertex x' such that $gx' = x$ for infinitely many $g \in A$, contradicting properness at x. Lastly, by Švarc–Milnor, for any $x \in X$ the map $G \to X$, $g \mapsto gx$ is a quasi-isometry; hyperbolicity is a quasi-isometry invariant for geodesic spaces, and so X is δ-hyperbolic for some δ. □

Lemma 6.2.3. *If $C \subset X$ is a closed convex subspace, $p : X \to C$ the closest point projection map (Proposition 2.3.8), and $A \subset X$ another convex subspace with $p(A)$ unbounded, then $d(A, C) < 2\delta$. Moreover, the 2δ-neighbourhood $N_{2\delta}(A) \cap C$ will be unbounded.*

Proof. Let $x, y \in A$ and suppose $z \in [p(x), p(y)] \subset C$ with $d(z, p(x)), d(z, p(y)) > 4\delta$. We must have $d(z, [x, p(x)]) \geq 2\delta$ as otherwise z would be closer to x than $p(x)$ is. Similarly we must have $d(z, [y, p(y)]) \geq 2\delta$. By applying Remark 2.8.3 to the geodesic quadrilateral in Figure 6.1, we deduce that $d(z, [x, y]) < 2\delta$. Thus $N_{2\delta}(A) \cap [p(x), p(y)]$ contains all of $[p(x), p(y)]$ except possibly the end segments of length 4δ. As we can have $d(p(x), p(y))$ arbitrarily large, we see that $N_{2\delta}(A) \cap C$ must be unbounded.

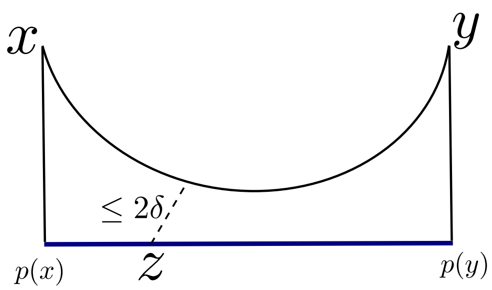

Figure 6.1: Geodesic quadrilateral $x, p(x), p(y), y$.

□

The remainder of this chapter will be a proof of Theorem 1.2.1 by induction on $\dim X$ (the case $\dim X = 0$ is trivial), so from now on assume that the theorem holds for lower dimensional cases - we will need this in the next two lemmas.

Lemma 6.2.4. *Either we can choose orbit representatives $H_1, ..., H_m$ for the hyperplanes of X such that $d(H_i, H_j) > 3\delta$ for all $1 \leq i < j \leq m$, or Theorem 1.2.1 holds.*

Proof. G is hyperbolic, so contains an infinite order element b [15, Γ.2.22]. All elements of G give semi-simple isometries of X (Proposition 2.3.5), so b acts hyperbolically on X, and acts by translations on an axis γ in X (a geodesic line in X) by Theorem 2.3.6. Let $p : X \to \gamma$ be the closest point projection map to γ.

First suppose there is a hyperplane H, such that $p(gH)$ is unbounded for all $g \in G$. By Lemma 6.2.3, we know that $\gamma \cap N_{2\delta}(gH)$ is unbounded for all $g \in G$. Because we are in a CAT(0) space, and γ and gH are convex, we see that $\gamma \cap N_{2\delta}(gH)$ is convex, so contains an infinite subinterval of γ. We deduce that there are only finitely many distinct translates gH, else infinitely many of them would be within 2δ of some point on γ, contradicting local finiteness of X. This means that the stabiliser G_H is finite index in G. But G_H acts properly cocompactly on the CAT(0) cube complex H, so by the lower dimensional case of theorem 1.2.1 there is a finite index subgroup $G' < G_H$ acting freely on H such that H/G' is special. Then G' also acts freely on X, for if $g \in G'$ fixed $x \in X$ then g would also fix $p(x)$. Then X/G' is virtually special by Corollary 6.1.4, and Theorem 1.2.1 holds by replacing G' with a further finite index subgroup.

Conversely, suppose that for every hyperplane H there exists $g \in G$ with $p(gH)$ bounded. Let $H_1, ..., H_m$ be orbit representatives for the hyperplanes of X such that $p(H_i)$ is bounded for each i. Each $p(H_i)$ is contained in a finite subinterval of γ, and b acts as a translation along γ, so we may choose $n_1, ..., n_m \in \mathbb{Z}$ such that $d(p(b^{n_i}H_i), p(b^{n_j}H_j)) = d(b^{n_i}p(H_i), b^{n_j}p(H_j)) > 3\delta$ for all $1 \leq i < j \leq m$. But p is distance non-increasing by Proposition 2.3.8, thus $d(b^{n_i}H_i, b^{n_j}H_j) > 3\delta$ for all $1 \leq i < j \leq m$, as required. □

Henceforth we will assume that we are in the first scenario of Lemma 6.2.4, so for the remainder of this section let $H_1, ..., H_m$ be orbit representatives for the hyperplanes of X such that $d(H_i, H_j) > 3\delta$ for all $1 \leq i, j \leq m$. We are now ready for the main technical lemma of this section, which will be used in Lemma 6.2.8 to produce a quotient of X with finite hyperplanes. Note that the homomorphism $\phi : G \to \mathcal{G}$ will be a product of homomorphisms obtained from Theorem 6.1.6, so in particular $\mathcal{G}$ might not be hyperbolic.

Lemma 6.2.5. *For any $R > 1$ large enough so that $G \cdot B = X$ for any R-ball B in X, there exists a surjective homomorphism $\phi : G \to \mathcal{G}$ with kernel N and $K_i \triangleleft G_{H_i}$ finite index such that*

(1) $\phi(K_i)$ are all finite,

(2) if $g \in G - K_i$ with $d(gH_i, H_i) \leq 2R$ then $\phi(g) \notin \phi(K_i)$,

(3) N is torsion-free (so acts freely on X).

The proof will use the following variant of the ping-pong lemma. Some of the arguments will be closely related to the results in [35].

Lemma 6.2.6. *(Ping-pong Lemma)*
Let K be a group that acts on a set Y. If $Y_1, ..., Y_n \subset Y$ and $K_1, ..., K_n < K$ and $y_0 \in Y - \cup_i Y_i$ are such that $kY_j \subset Y_i$ and $ky_0 \in Y_i$ whenever $1 \neq k \in K_i$ and $j \neq i$, then K splits as a free product $K = K_1 * ... * K_n$.

Proof. A product $k = k_1 \cdots k_l$ with $1 \neq k_i \in K_{m_i}$ and $m_i \neq m_{i+1}$ clearly maps y_0 into Y_{m_1}, and so is not the identity. □

Proof of Lemma 6.2.5.

Let $p_i : X \to H_i$ be the closest point projection map to the hyperplane H_i. As the H_i are at least 3δ apart from each other, Lemma 6.2.3 tells us that the images $p_j(H_i)$ for $i \neq j$ are all bounded. By Remark 2.4.4 and induction on the lower dimensional cases of Theorem 1.2.1, for each i there exists $K_i < G_{H_i}$ finite index acting freely on H_i with H_i/K_i directly special. Define bounded subspaces

$$A_i := N_{14\delta+2R+1}(\cup_{j\neq i} p_i(H_j)).$$

Theorem 6.1.3 tells us in particular that K_i is residually finite, so by replacing K_i with a further finite index subgroup we can assume that $d(kA_i, A_i) > 1$ for all $1 \neq k \in K_i$. We can also assume that $K_i \triangleleft G_{H_i}$ by intersecting it with its finitely many conjugates, and H_i/K_i will still be directly special by Remark 2.4.15. Note that some H_i might be finite and have $A_i = H_i$ - in these cases K_i will be trivial.

Let $X_i := p_i^{-1}(H_i - A_i)$. Pick $x_0 \in p_1(H_2) \subset A_1$ and note that x_0 is not in any of the sets X_i. The next part of the proof does ping-pong with x_0, K_i and X_i to prove that we get a free splitting $K := \langle K_1, ..., K_m \rangle \cong K_1 * ... * K_m$. By ignoring the i for which K_i is trivial we can assume that the sets $H_i - A_i$ and X_i are non-empty.

Claim: $p_i(X_j) \subset A_i$ for $j \neq i$.

Proof: Let $x \in X_j$ and suppose for contradiction that $p_i(x) \notin A_i$.

We have the geodesic pentagon shown, where y is any point in $p_j(H_i)$ and z is a point on $[p_i(x), p_i(y)] \cap A_i$. We have defined A_i to include a 14δ buffer zone around $\cup_{j \neq i} p_i(H_j)$, and $p_i(x) \notin A_i$, therefore we can choose z to satisfy $d(z, p_i(x)), d(z, p_i(H_j)) > 7\delta$. By Remark 2.8.3, z is within 3δ of one of the other sides of the pentagon, we now check each of these four sides in turn:

(a) If $z' \in [x, p_j(x)]$ then $p_j(z) = p_j(x) \notin A_j$, and so $d(p_j(z), p_j(z')) \geq d(p_j(H_i), p_j(x))$, which is greater than 3δ because of the buffer zone in A_j. But p is distance non-increasing, so $d(z, z') > 3\delta$.

(b) $d(z, H_j) \geq d(H_i, H_j) > 3\delta$ by choice of the hyperplanes H_l.

(c) If $z' \in [y, p_i(y)]$ then

$$\begin{aligned} d(y,z) &\leq d(y,z') + d(z',z) \\ &= d(y, p_i(y)) - d(p_i(y), z') + d(z', z) \\ &\leq d(y, p_i(y)) - d(z, p_i(y)) + 2d(z, z') \\ &\leq d(y, p_i(y)) - 7\delta + 2d(z, z') \\ &\leq d(y, z) - 7\delta + 2d(z, z'), \end{aligned}$$

(the last inequality by definition of p_i). This implies $d(z, z') > 3\delta$.

(d) The same argument as 3. shows that z cannot be within 3δ of $[x, p_i(x)]$.

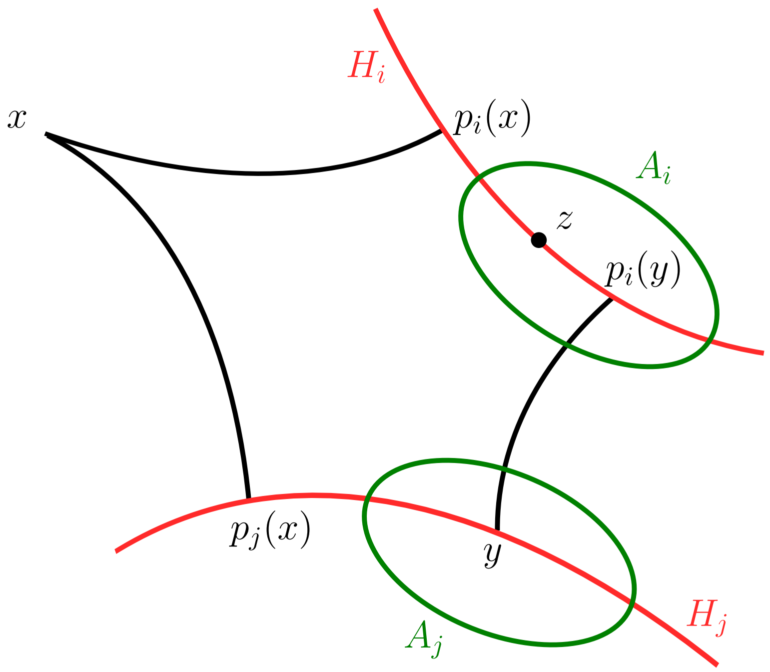

Figure 6.2: Geodesic pentagon $y, p_i(y), p_i(x), x, p_j(x)$.

We conclude that z is not within 3δ of one of the other sides of the pentagon, contradicting Remark 2.8.3. The claim follows. ∎

Claim: For $j \neq i$ and $1 \neq k \in K_i$ we have $kX_j \subset X_i$, $kH_j \subset X_i$ and $kx_0 \in X_i$.

Proof: Let $x \in X_j \cup H_j$. By the previous claim we have $p_i(x) \in A_i$, so $p_i(kx) = kp_i(x) \in kA_i$, hence $p_i(kx) \notin A_i$ and so $kx \in X_i$. Additionally, $p_i(kx_0) = kp_i(x_0) \in kA_i$ and so $kx_0 \in X_i$. ∎

This last claim allows us to do ping-pong, as in Lemma 6.2.6, to obtain the desired splitting $K \cong K_1 * ... * K_m$.
Claim: $K < G$ is quasiconvex.

Proof: $G \to X, g \mapsto gx_0$ is a quasi-isometry, so it suffices to show that $K \cdot x_0$ is quasiconvex in X.

Let $k = k_1 k_2 ... k_l$ with $1 \neq k_i \in K_{n_i}$ and $n_i \neq n_{i+1}$. Put $g_i = k_1 ... k_i$ and $g_0 = 1$. Our strategy will be to show that all points on the geodesic $[x_0, kx_0]$ are close to one of the hyperplanes $g_i H_{n_i}$ for $1 \leq i \leq l$.

First we consider projections to such a hyperplane. Let $q_i : X \to g_i H_{n_i}$ be the closest point projection map to $g_i H_{n_i}$. For $x \in X$, $q_i(x)$ is the closest point on $g_i H_{n_i}$ to x, left multiplying by g_i^{-1} then tells us that $g_i^{-1} q_i(x)$ is the closest point on H_{n_i} to $g_i^{-1} x$, so $g_i^{-1} q_i(x) = p_{n_i}(g_i^{-1} x)$. Therefore

$$q_i = g_i p_{n_i} g_i^{-1}. \tag{6.2.1}$$

We can then compute for $1 \le i < l$,

$$\begin{aligned} q_i(kx_0) &= g_i p_{n_i}(g_i^{-1} k x_0) \\ &= g_i p_{n_i}(k_{i+1} ... k_l x_0) \\ &\in g_i p_{n_i}(X_{n_{i+1}}) && \text{by the second claim,} \\ &\subset g_i A_{n_i} && \text{by the first claim.} \end{aligned} \tag{6.2.2}$$

Similarly,

$$q_l(kx_0) = kp_{n_l}(x_0) \in kA_{n_l}. \tag{6.2.3}$$

Next observe that $g_i H_{n_i} = g_{i-1} H_{n_i}$, and so analogously to (6.2.1) we have $q_i = g_{i-1} p_{n_i} g_{i-1}^{-1}$ $(1 \le i \le l)$. We then compute for $1 < i \le l$,

$$\begin{aligned} q_i(x_0) &= g_{i-1} p_{n_i}(g_{i-1}^{-1} x_0) \\ &= g_{i-1} p_{n_i}(k_{i-1}^{-1} ... k_1^{-1} x_0) \\ &\in g_{i-1} p_{n_i}(X_{n_{i-1}}) && \text{by the second claim,} \\ &\subset g_{i-1} A_{n_i} && \text{by the first claim.} \end{aligned} \tag{6.2.4}$$

And similarly

$$q_1(x_0) = p_{n_1}(x_0) \in A_{n_1}. \tag{6.2.5}$$

We now consider the concatenation of geodesics joining the following points pairwise in order.

$$x_0, q_1(x_0), q_1(kx_0), q_2(x_0), q_2(kx_0), ..., q_l(x_0), q_l(kx_0), kx_0$$

Call this path γ, and refer to the above points as the vertices of γ. Recalling that $x_0 \in A_1$, we can bound every other gap between consecutive vertices as follows.

$$D := \operatorname{diam}(\cup A_j) \ge \begin{cases} d(x_0, q_1(x_0)), & \text{by (6.2.5)} \\ d(q_i(kx_0), q_{i+1}(x_0)), & \text{for } 1 \le i < l, \text{ by (6.2.2) and (6.2.4)} \\ d(q_l(kx_0), kx_0), & \text{by (6.2.3)} \end{cases}$$

The other gaps between consecutive vertices are spanned by segments $\gamma_i := [q_i(x_0), q_i(kx_0)] \subset g_i H_{n_i}$. Since K_{n_i} acts cocompactly on H_{n_i} and $g_i \in K$, we deduce that each γ_i is contained within $N_M(K \cdot x_0)$ for some constant M that is independent of k. Hence $\gamma \subset N_{M+D}(K \cdot x_0)$.

To complete the proof of the claim it remains to show that $\sigma := [x_0, kx_0] \subset N_L(\gamma)$ for some constant L that is independent of k. Consider $z \in \gamma_i$ at least 5δ away from the endpoints of γ_i. By Remark 2.8.3, z is within 2δ of one of the other sides of the geodesic quadrilateral shown in Figure 6.3, so it must be within 2δ of σ - otherwise it contradicts the definition of closest point projection. Therefore $\gamma_i \subset N_{7\delta}(\sigma)$ and $\gamma \subset N_{D+7\delta}(\sigma)$.

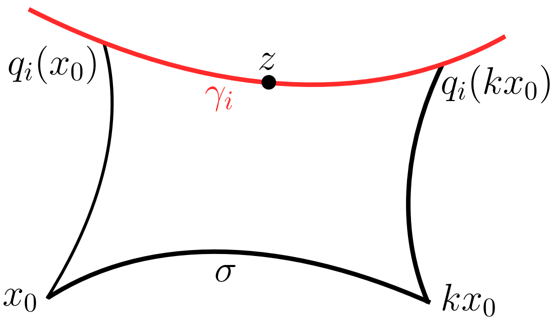

Figure 6.3: Geodesic quadrilateral $x_0, kx_0, q_i(kx_0), q_i(x_0)$.

Finally note that projection from γ to σ is continuous; and, as the paths share end-points x_0 and kx_0, the image is the whole of σ. So in fact $\sigma \subset N_{D+7\delta}(\gamma)$. ■

By Proposition 2.4.12, $K \cong K_1 * ... * K_m$ is the fundamental group of a directly special cube complex. By taking direct products of the homomorphisms in 6.1.6, we deduce that for any finite $A \subset G - K$ there is a quotient homomorphism $\phi : G \to \mathcal{G}$ such that $\phi(A) \cap \phi(K) = \emptyset$ and $\phi(K)$ is finite. We now show that conclusions (1)–(3) of the lemma can be satisfied by a certain choice of A.

(1) $\phi(K_i) < \phi(K)$ so must be finite.

(2) For each i the collection of double cosets

$$\mathcal{A}_i := \{K_i g K_i \,|\, g \in G,\, d(gH_i, H_i) \leq 2R\} - \{K_i\}$$

is finite. To see this, fix $y \in H_i$ and consider $g \in G$ with $d(gH_i, H_i) \leq 2R$ - say $x, x' \in H_i$ satisfy $d(gx', x) \leq 2R$. Suppose that $Q > 0$ with $H_i \subset K_i \cdot B_Q(y)$. Now pick $k, k' \in K_i$ so that $d(kx, y), d(k'y, x') < Q$; then $d(kgk'y, y) \leq d(kgk'y, kgx') + d(kgx', kx) + d(kx, y) < 2Q + 2R$. The finiteness of $\mathcal{A}_i$ then follows because X is locally finite and the action of G is proper.

We now claim that $g \in G - K_i$ with $d(gH_i, H_i) \leq 2R$ cannot have $g \in K$. Indeed, if $g = k_1 k_2 ... k_l$ with $k_j \in K_{n_j}$ and $n_j \neq n_{j+1}$, then by the second claim we see that $gH_i \subset X_{n_1}$ and $p_{n_1}(gH_i) \subset H_{n_1} - A_{n_1}$. Removing k_1 if necessary we can assume that $n_1 \neq i$. But there exists $x \in gH_i$ with $d(x, H_i) \leq 2R$, and as p_{n-1} is distance non-increasing we deduce that $d(p_{n_1}(x), p_{n_1}(H_i)) \leq 2R$. So $p_{n_1}(x) \in A_{n_1}$, a contradiction.

Given these two facts, we can ensure that A contains representatives for all of the double cosets in the $\mathcal{A}_i$, and this ensures that $\phi(g) \notin \phi(K_i)$ for any $g \in G - K_i$ with $d(gH_i, H_i) \leq 2R$.

(3) If $g \in G$ is a torsion element then by Theorem 2.3.6(1) it has a fixed point $x \in X$. By assumption of the lemma, there exists $h \in G$ with $hx \in B_R(x_0)$. Then $d(hgh^{-1}x_0, x_0) < 2R$. Therefore there is a finite set $\mathcal{T}$ of representatives for conjugacy classes of torsion elements in G. Each K_i is torsion-free because it acts freely on H_i, so K is also torsion-free and $K \cap \mathcal{T} = \emptyset$. Adding $\mathcal{T}$ to A will ensure that N is torsion-free.

□

The point of Lemma 6.2.5 is that it allows us to define the following quotient complex.

Definition 6.2.7. (Quotient Complex $\mathcal{X}$)
As a result of Lemma 6.2.5, we can define the NPC cube complex $\mathcal{X} := X/N$. The value of R we use will be some constant large enough to satisfy the cocompactness condition of the lemma, and we also require $R \geq \delta + 2\sqrt{\dim X}$ (this inequality will be demystified in Section 6.6). The metric on $\mathcal{X}$ will be denoted d, the same as for X.

As was the aim of this section, this quotient complex satisfies the following properties.

Lemma 6.2.8. *(Properties of $\mathcal{X}$)*

(1) There are natural cocompact actions of G and $\mathcal{G}$ on $\mathcal{X}$.

(2) All hyperplanes of $\mathcal{X}$ are finite.

(3) For any hyperplane H in X, the R-neighbourhood $N_R(H)$ quotiented by $N \cap G_H$ embeds in $\mathcal{X}$. In particular this implies that all hyperplanes of $\mathcal{X}$ are embedded, and that distinct hyperplanes in X which are less than R apart map to distinct hyperplanes in $\mathcal{X}$.

Proof. (1) holds because N is normal in G and G acts cocompactly on X. By Lemma 6.2.5(1) we know that, for any $g \in G$, $N \cap gG_{H_i}g^{-1}$ has finite index in $gG_{H_i}g^{-1} = G_{gH_i}$ and so acts cocompactly on gH_i - property (2) follows. Lemma 6.2.5(2) tells us that the R-neighbourhood $N_R(H_i)$ quotiented by $N \cap K_i = N \cap G_{H_i}$ embeds in $\mathcal{X}$ - property (3) follows by considering translates of the H_i and conjugates of the K_i. □

To finish the section we introduce some notation.

Notation 6.2.9. The quotient map $m : X \to \mathcal{X}$ will send a vertex x (resp. an edge e and hyperplane H) to a vertex $\bar{x}$ (resp. an edge $\bar{e}$ and hyperplane $\overline{H}$). And $\overline{H}(\bar{e})$ will denote the hyperplane dual to $\bar{e}$. By Remark 6.2.1 we know that hyperplanes in $\mathcal{X}$ are two-sided and that no element of $\mathcal{G}$ can exchange the sides of any hyperplane. We can define $\overline{H}^{\pm} := \overline{(H^{\pm})}$.

6.3 Invariant colouring measures

This section is preparatory. More specifically, before we can start colouring hyperplanes in the next section, we need to establish some theory about colouring graphs. This section is essentially the same as §5 of Agol's paper, but with a little more detail regarding weak* convergence.

Definition 6.3.1. (Colourings)
An *n-colouring* of a graph Γ is a map $c : V\Gamma \to [n] := \{1, ..., n\}$ such that $c(v_1) \neq c(v_2)$ whenever $\{v_1, v_2\} \in E\Gamma$. Let $C_n(\Gamma)$ denote the set of n-colourings. If vertex degrees are bounded by L then it is clear that $C_{L+1}(\Gamma) \neq \emptyset$.

Suppose a group K acts on Γ. Then we have an action of K on $C_n(\Gamma)$ by $k : c \mapsto c \circ k^{-1}$.

Definition 6.3.2. (Colourings as a measurable space)
Consider $C_n(\Gamma)$ as a closed subspace of $[n]^{V\Gamma}$ with the product topology. We will consider $[n]^{V\Gamma}$ as a measurable space with σ-algebra generated by the sets $A_{v,j} := \{f \in [n]^{V\Gamma} | \, f(v) = j\}$ for $v \in V\Gamma$ and $j \in [n]$ - note that if $V\Gamma$ is countable then this is also the σ-algebra generated by the open subsets of $[n]^{V\Gamma}$. The action of K on $C_n(\Gamma)$ extends naturally to $[n]^{V\Gamma}$ by $k : f \mapsto f \circ k^{-1}$. These are measurable functions so the family of measurable subsets of $[n]^{V\Gamma}$ is K-invariant.

Theorem 6.3.3. *Suppose a group K acts cocompactly on a countable graph Γ with vertex degrees bounded by L. Then there exists an K-invariant probability measure μ on $C_{L+1}(\Gamma)$.*

Proof. Let $M_K(n)$ denote the set of K-invariant probability measures on $[n]^{V\Gamma}$. Note that $M_K(n)$ is non-empty because it contains the measure μ_n which is the product of uniform measures on $[n]$. We have that $C_{L+1}(\Gamma) \subset [L+1]^{V\Gamma}$, so our task is to find $\mu \in M_K(L+1)$ with $\mu(C_{L+1}(\Gamma)) = 1$ (thus making μ an K-invariant probability measure on $C_{L+1}(\Gamma)$).

For an edge $e = \{v_1, v_2\} \in E\Gamma$ let $B_e(n) := \{f \in [n]^{V\Gamma} | \, f(v_1) = f(v_2)\}$; it is clear that $[n]^{V\Gamma}$ splits as a disjoint union $[n]^{V\Gamma} = C_n(\Gamma) \sqcup \bigcup_{e \in E\Gamma} B_e(n)$, so our task reduces to finding $\mu \in M_K(L+1)$ with $\mu(B_e(L+1)) = 0$ for all $e \in E\Gamma$. Let $\{e_1, ..., e_m\} \subset E\Gamma$ be a complete set of orbit representatives for the action of K on $E\Gamma$; since $\nu(B_{ke}(L+1)) = \nu(kB_e(L+1)) = \nu(B_e(L+1))$ for all $\nu \in M_K(L+1)$, it suffices to find $\mu \in M_K(L+1)$ with $\mu(B_{e_i}(L+1)) = 0$ for $i = 1, ..., m$.

For $\nu \in M_K(n)$ define $\text{weight}(\nu) := \sum_i \nu(B_{e_i}(n))$. We will use a limiting argument to construct $\mu \in M_K(L+1)$ with zero weight. As mentioned above, there is a K-invariant probability measure μ_n on $[n]^{V\Gamma}$ which is the product of uniform measures on $[n]$. Note that μ_n is the unique measure with $\mu_n(A_{v,j}) = 1/n$ for every $A_{v,j}$. Also, $\mu_n(B_e(n)) = \mu_n(\cup_j (A_{v,j} \cap A_{w,j})) = 1/n$ for $e = \{v, w\} \in E\Gamma$, so $\text{weight}(\mu_n) = m/n$.

For $n > L+1$ define a map $p_n : [n]^{V\Gamma} \to [n-1]^{V\Gamma}$ by

$$p_n(c)(v) := \begin{cases} c(v), & c(v) < n \\ \min([n-1] - \{c(u) | \, \{u, v\} \in E\Gamma\}), & c(v) = n \end{cases}$$

for $c \in [n]^{V\Gamma}$ and $v \in V\Gamma$. In other words p_n changes the colour of each vertex coloured n to the smallest colour not used by its neighbours, and leaves other vertices with the same colour. This is well-defined because vertex degrees are at most L. It is clear that p_n is K-equivariant and continuous, so it induces a well-defined push-forward $p_{n*} : M_K(n) \to M_K(n-1)$ given by $p_{n*}(\nu)(A) = \nu(p_n^{-1}(A))$ for $\nu \in M_K(n)$ and $A \subset [n-1]^{V\Gamma}$ measurable. Furthermore, for $\{v_1, v_2\} \in E\Gamma$ and $c \in [n]^{V\Gamma}$, if $p_n(c)(v_1) = p_n(c)(v_2)$ then $c(v_1) = c(v_2)$. Therefore $p_n^{-1}(B_e(n-1)) \subset B_e(n)$ for any edge e; consequently $\text{weight}(p_{n*}(\nu)) \leq \text{weight}(\nu)$ for any $\nu \in M_K(n)$.

Now define $P_{n*} = p_{L+2*} \circ p_{L+3*} \circ \cdots \circ p_{n*} : M_K(n) \to M_K(L+1)$. We will then have that $\text{weight}(P_{n*}(\mu_n)) \leq \text{weight}(\mu_n) = m/n \to 0$ as $n \to \infty$.

By Prokhorov's Theorem [74] the set of all probability measures on $[L+1]^{V\Gamma}$ is compact metrizable in the weak* topology (measures (ν_n) converge to ν in the weak* topology if and only if $\int \alpha \, d\nu_n \to \int \alpha \, d\nu$ for every continuous map $\alpha : [L+1]^{V\Gamma} \to \mathbb{R}$). $M_K(L+1)$ is a closed subspace with respect to this topology; to see this suppose a sequence (ν_n) in $M_K(L+1)$ converges to a measure ν. Let Σ be the algebra generated by the $A_{v,j}$ (the smallest family containing the $A_{v,j}$ that is closed under finite union and complementation), all sets in Σ are clopen in $[L+1]^{V\Gamma}$ so by considering characteristic functions we have $\nu_n(A) \to \nu(A)$ for every $A \in \Sigma$. For $k \in K$ define ν_k by $\nu_k(A) = \nu(kA)$ for measurable sets A; for $A \in \Sigma$ we have $\nu_k(A) = \lim_{n\to\infty} \nu_n(kA) = \lim_{n\to\infty} \nu_n(A) = \nu(A)$, so by Caratheodory's Extension theorem we get $\nu = \nu_k$, thus $\nu \in M_K(L+1)$.

We conclude that $P_{n*}(\mu_n)$ has a convergent subsequence converging to some $\mu \in M_K(L+1)$. The weight is continuous with respect to the weak* topology because $B_{e_i}(L+1) \in \Sigma$, thus $\text{weight}(\mu) = 0$ as required. □

Remark 6.3.4. When we apply this theorem later in the thesis it would be enough for μ to only be defined on the algebra $\{C_{L+1}(\Gamma) \cap A | \, A \in \Sigma\}$, where Σ is the algebra generated by the $A_{v,j}$, rather than having μ defined on all measurable subsets of $C_{L+1}(\Gamma)$. If we had modified the theorem to only require this, then the last part of the proof that argues about convergence would be easier because we wouldn't need weak* convergence or Caratheodory's Extension theorem (the reason is basically that Σ is countable). However keeping the theorem as it is makes for a cleaner statement.

6.4 Colouring hyperplanes

In this section we define the local colouring data that will play a key role in the following sections. This local colouring data will be defined as equivalence classes on the space of all colourings of hyperplanes in $\mathcal{X}$; first we define this space of colourings by building a graph out of the hyperplanes of $\mathcal{X}$. This section follows Definitions 6.6–6.8

of Agol's paper, but with a simplification implied by Agol's ICM notes [2, 7.4] and with some differences in notation.

Definition 6.4.1. (The graph Γ)
Let $\Gamma = \Gamma(\mathcal{X})$ be a graph whose vertices are the hyperplanes of $\mathcal{X}$, and with hyperplanes $\overline{H}_1, \overline{H}_2$ joined by an edge in Γ if $d(\overline{H}_1, \overline{H}_2) \leq R$. We have a natural action of $\mathcal{G}$ on Γ. As $\mathcal{X}$ is locally finite, cocompact and with finite hyperplanes (see Lemma 6.2.8) it follows that the degree of vertices in Γ is bounded by some $L \in \mathbb{N}$.

As in Definition 6.3.1, we have an action of $\mathcal{G}$ on $C_{L+1}(\Gamma)$ by $c \mapsto gc := c \circ g^{-1}$ - this also induces an action of G on $C_{L+1}(\Gamma)$ by $c \mapsto gc := c \circ \phi(g^{-1})$.

Definition 6.4.2. (Equivalent colourings)
For H a hyperplane in X, we will define equivalence classes $[-]_H$ in $C_{L+1}(\Gamma)$ that depend only on the colour of vertices "near" to $\overline{H}$ in Γ. What we mean by vertices "near" to $\overline{H}$ will depend on the colouring $c \in C_{L+1}(\Gamma)$ in question. Specifically, define the equivalence class $[c]_H$ by

$$[c]_H := \{c' \in C_{L+1}(\Gamma) \,|\, c' = c \text{ on the ball of radius } c(\overline{H}) \text{ in } \Gamma \text{ centred at } \overline{H}\}.$$

For $e \in X^1$ that crosses a hyperplane H, we will use $[-]_e$ as an alternative notation for the equivalence class $[-]_H$. We will also use the notation $c(e) := c(\overline{H})$. Note that $c' \in [c]_e$ implies $c'(e) = c(e)$. For $x \in X^0$ we will require the finer equivalence classes $[-]_x$ in $C_{L+1}(\Gamma)$ defined by

$$[c]_x := \cap\{[c]_e \,|\, e \in X^1 \text{ incident at } x\}.$$

We want to combine these families of equivalence relations into two G-invariant equivalence relations, one for edges and one for vertices. Do this as follows. Define equivalence classes on $X^1 \times C_{L+1}(\Gamma)$ by $[e, c] := \{e\} \times [c]_e$. Similarly, define equivalence classes on $X^0 \times C_{L+1}(\Gamma)$ by $[x, c] := \{x\} \times [c]_x$. The action of G on X^1, X^0 and $C_{L+1}(\Gamma)$ induces actions on the product spaces and on the two sets of equivalence classes by $g[e, c] = [ge, gc]$ and $g[x, c] = [gx, gc]$. It is straightforward to check that these are well-defined (first check that $g[c]_H = [gc]_{gH}$).

Remark 6.4.3. For each $e \in X^1$, the classes $[-]_e$ only depend on the colour of vertices in some $(L+1)$-ball of Γ, and so there are only finitely many of these equivalence classes. Similarly, for each $v \in X^0$, there are only finitely many classes $[-]_v$. As there are finitely many G-orbits in X^1 and X^0, there must only be finitely many G-orbits of equivalence classes on $X^1 \times C_{L+1}(\Gamma)$ and $X^0 \times C_{L+1}(\Gamma)$.

Remark 6.4.4. If edges $e, f \in X^1$ are both incident at a vertex x, then $d(\overline{H(e)}, \overline{H(f)}) \leq 1 < R$, so $c(e) = c(\overline{H(e)}) \neq c(\overline{H(f)}) = c(f)$ for any $c \in C_{L+1}(\Gamma)$.

6.5 Starting the gluing construction

In this section we introduce the main construction used to prove Theorem 1.2.1. We will implement the first step of the construction and show how the theorem follows from the final step. The process of going between steps will be left to Sections 6.6 and 6.7. The construction is similar in spirit to §8 of Agol's paper, but we work with subspaces of X rather than orbi-complexes; in other words we work on the space where G acts rather than in the quotient. This will allow us to make use of the CAT(0) geometry of X, as we do in Sections 6.6 and 6.7, and it will save us from having to recall technical background about orbi-complexes. Lemma 6.5.1 is based on the ideas of §7 of Agol's paper.

To prove Theorem 1.2.1 we will inductively construct $\mathcal{V}_{L+1}, \mathcal{V}_L, ..., \mathcal{V}_0$ (with the L from Definition 6.4.1). Each $\mathcal{V}_j$ will be a non-empty collection of triples $(Z, K, (c_x))$ where $Z \subset X$ is a non-empty intersection of half-spaces (so is closed and convex), and for each $x \in Z$ a vertex we have $c_x \in C_{L+1}(\Gamma)$ a colouring. K will be a subgroup of G that acts on Z freely and cocompactly, and $c_{kx} = kc_x$ for $k \in K$. We will permit $\mathcal{V}_j$ to contain duplicates of some triples. Where there is no danger of ambiguity, we will write $Z \in \mathcal{V}_j$ as shorthand for $(Z, K, (c_x)) \in \mathcal{V}_j$.

$\mathcal{V}_j$ will satisfy four properties; before stating these formally we give some loose motivation for them. We will often work with the finite complex Z/K which has universal cover Z (as Z is CAT(0)). Technically Z and Z/K are not quite cube complexes, rather they are something like "cube complexes with boundary hyperplanes" - we'll just have to live with this, but we will get a genuine cube complex once we've finished all the gluing. Think of the colourings c_x as giving information about some of the hyperplanes nearby x, such as which hyperplanes mark the boundary of Z. The rough idea of the construction is to glue the complexes Z/K along hyperplanes coloured j to form $\mathcal{V}_{j-1}$. Neighbouring vertices in Z must agree about colours of nearby hyperplanes, and so vertices next to boundary hyperplanes that will later be glued up must have a potential matching in which colourings are compatible.

Here are the properties that $\mathcal{V}_j$ must satisfy (notation from Definitions 6.4.2 and 6.1.1):

(1) If $e \in X^1$ joins vertices $x, y \in Z \in \mathcal{V}_j$, then $[e, c_x] = [e, c_y]$. So adjacent vertices are equipped with similar colourings.

(2) If $e \in X^1$ joins vertices x, y with $x \in Z \in \mathcal{V}_j$, then

$$y \in Z \Leftrightarrow c_x(e) > j.$$

So hyperplanes in the interior of Z are coloured $> j$ by their neighbouring vertices, whereas hyperplanes on the boundary of Z are coloured $\leq j$.

(3) (Gluing Equations)

Let $f \in X^1$ and $c \in C_{L+1}(\Gamma)$. Define sets

$$\mathcal{V}_j^{\pm}(f,c) := \{(K \cdot e, Z) \,|\, (Z, K, (c_x)) \in \mathcal{V}_j,\, \exists g \in G:\, gZ \cap H(f)^{\pm} \neq \emptyset,\, g[e, c_e] = [f, c]\},$$

where any duplicates of triples $(Z, K, (c_x)) \in \mathcal{V}_j$ are counted separately. In other words, $\mathcal{V}_j^{\pm}(f,c)$ is the set of edges in the complexes Z/K that, modulo the action of G, correspond to f and have colouring in the class $[c]_f$; and the $\pm$ means that Z/K continues on the $\pm$ side of the hyperplane dual to the edge.

The Gluing Equations are given by

$$|\mathcal{V}_j^{+}(f,c)| = |\mathcal{V}_j^{-}(f,c)|,$$

where f ranges over X^1 and c ranges over $C_{L+1}(\Gamma)$. Roughly speaking, these equations will ensure that hyperplanes on the boundary of complexes Z/K that look like $H(f)$ on the $H(f)^+$ side can be matched up with those that look like $H(f)$ on the $H(f)^-$ side, with compatible colourings matched together - but there is more work to be done later to arrange this precisely.

(4) $K \in \mathcal{QVH}$ for any triple $(Z, K, (c_x)) \in \mathcal{V}_j$. This will allow us to make use of the theorems in Section 6.1.

It will be convenient to also define colourings $c_e \in C_{L+1}(\Gamma)$ for edges $e \in X^1$ that intersect Z, such that if e is incident at a vertex $x \in Z$ then $[e, c_e] = [e, c_x]$. This is possible by property (1). We will only ever care about the class $[e, c_e]$, so the colourings (c_e) are not extra data that needs to be added to the triple $(Z, K, (c_x))$.

To start the inductive construction of the $\mathcal{V}_j$, we first define $\mathcal{V}_{L+1}$.

Lemma 6.5.1. *There exists $\mathcal{V}_{L+1}$ satisfying all of the above conditions.*

Proof. Let $\{x_1, ..., x_t\}$ be a complete set of G-orbit representatives in X^0. For each i pick colourings $c_{il} \in C_{L+1}(\Gamma)$ so that we have a partition

$$C_{L+1}(\Gamma) = \bigsqcup_{1 \leq l \leq n_i} [c_{il}]_{x_i}. \tag{6.5.1}$$

These partitions are finite by Remark 6.4.3. For each (i, l) let Z_{il} be the intersection of all half-spaces containing x_i. Note that Z_{il} will be compact, in fact it will be the union of cubes in $\dot{X}$ (the cubical subdivision of X) that contain x_i, so $Z_{il} \cap X^0 = \{x_i\}$. We will then define $\mathcal{V}_{L+1}$ to be the collection of triples $(Z_{il}, \{1\}, c_{il})$.

We must check that this definition of $\mathcal{V}_{L+1}$ satisfies properties (1)–(4) above. Each Z_{il} only contains one vertex of X, so properties (1) and (2) hold vacuously, and (4) is also immediate since the trivial group is in $\mathcal{QVH}$. However (3) might not hold. To rectify this we will make $\mathcal{V}_{L+1}$ contain α_{il} copies of Z_{il} for appropriate integers α_{il}, which we will spend the rest of the proof constructing.

Take $f \in X^1$ with endpoints $x_+ \in H(f)^+$ and $x_- \in H(f)^-$, and take $c \in C_{L+1}(\Gamma)$. Say x_+ is in the orbit of x_i. How can we count $\mathcal{V}^+_{L+1}(f,c)$? Well the contributions will come from precisely the pairs (c_{il}, e), with e incident at x_i, such that there exists $g \in G$ with $gx_i = x_+$ and $g[e, c_{il}] = [f, c]$ - and each pair will contribute α_{il}. Note that $c_{il}(e) = gc_{il}(f) = c(f)$, so by Remark 6.4.4 there is at most one valid choice of edge e for each choice of colouring c_{il} - call this edge e_l. Another way we could try to count $\mathcal{V}^+_{L+1}(f,c)$ is by counting pairs (g,l) with $g \in G$ and $1 \leq l \leq n_i$ such that $gx_i = x_+$ and $g[e_l, c_{il}] = [f,c]$, and let each such (g,l) contribute α_{il}. This method will over-count, but we can measure the extent of this over-counting by the following claim, where $M^+(f,c)$ denotes the total obtained by this over-counting method.

<u>Claim:</u> $M^+(f,c) = |\mathrm{Stab}_G([f,c])||\mathcal{V}^+_{L+1}(f,c)|$

<u>Proof:</u> Fix l such that (c_{il}, e_l) contributes to $\mathcal{V}^+_{L+1}(f,c)$. We claim that (g,l) contributes to $M^+(f,c)$ if and only if $g[e_l, c_{il}] = [f,c]$. The "only if" direction is immediate; for the "if" direction we just need to check that $g[e_l, c_{il}] = [f,c]$ implies $gx_i = x_+$. Indeed we assumed that (c_{il}, e_l) contributes to $\mathcal{V}^+_{L+1}(f,c)$, so there is some $g \in G$ with $ge_l = f$ and $gx_i = x_+$, but then Remark 6.2.1 implies that any $g \in G$ with $ge_l = f$ satisfies $gx_i = x_+$.

Since G acts on the equivalence classes of $X^1 \times C_{L+1}(\Gamma)$, we deduce that there are $|\mathrm{Stab}_G([f,c])|$ elements $g \in G$ such that (g,l) contributes to $M^+(f,c)$. Each (g,l) will contribute α_{il} to $M^+(f,c)$, and (c_{il}, e_l) also contributes α_{il} to $|\mathcal{V}^+_{L+1}(f,c)|$, so the claim follows. ∎

Why on earth is it useful to over-count $\mathcal{V}^+_{L+1}(f,c)$? Well the factor of over-counting we obtained only depends on f and c, so if we over-count $\mathcal{V}^-_{L+1}(f,c)$ in the same way to produce a total $M^-(f,c)$, then the factor of over-counting is the same. Hence the Gluing Equation $|\mathcal{V}^+_{L+1}(f,c)| = |\mathcal{V}^-_{L+1}(f,c)|$ is equivalent to $M^+(f,c) = M^-(f,c)$. The trick now is to solve these transformed Gluing Equations by using the measure from Theorem 6.3.3.

The $M^\pm(f,c)$ are just integer sums of the α_{il}. We want the α_{il} to be non-negative integers (that are not all zero), but as a start we will exhibit positive real numbers α_{il} that solve the Gluing Equations. We do this by taking the measure μ from Theorem 6.3.3 applied to the graph Γ with the action of G, and putting

$$\alpha_{il} = \frac{\mu([c_{il}]_{x_i})}{|\mathrm{Stab}_G(x_i)|}. \tag{6.5.2}$$

Next, observe that $[c]_f$ can be partitioned into $[-]_{x_+}$ equivalence classes, which can be written

$$[c]_f = \bigsqcup_{b \in \beta_+(f,c)} [c_b]_{x_+}, \tag{6.5.3}$$

for some $c_b \in C_{L+1}(\Gamma)$ (and this partition is finite by Remark 6.4.3). A pair (g, l) contributes to $M^+(f, c)$ if and only if $g[x_i, c_{il}] = [x_+, c_b]$ for some $b \in \beta_+(f, c)$, so we can count $M^+(f, c)$ by adding up the contributions from each c_b. There will be $|\mathrm{Stab}_G(x_i)|$ choices of g with $gx_i = x_+$, and for each pair (g, b) there will be a unique l with $g[x_i, c_{il}] = [x_+, c_b]$. As μ is G-invariant, we see from (6.5.2) that each α_{il} only depends on the G-orbit of $[x_i, c_{il}]$, and so the contribution to $M^+(f, c)$ from a given c_b will equal $\alpha_{il}|\mathrm{Stab}_G(x_i)|$ for any l with $[x_+, c_b] \in G \cdot [x_i, c_{il}]$. For any such l, the G-invariance of μ implies that $\mu([c_b]_{x_+}) = \mu([c_{il}]_{x_i})$, hence

$$\begin{aligned} M^+(f, c) &= \sum_{b \in \beta_+(f,c)} \frac{\mu([c_b]_{x_+})}{|\mathrm{Stab}_G(x_i)|} |\mathrm{Stab}_G(x_i)| \\ &= \sum_{b \in \beta_+(f,c)} \mu([c_b]_{x_+}) \\ &= \mu([c]_f). \end{aligned}$$

Again this only depends on f and c, so our clever choices of α_{il} will also give $M^-(f, c) = \mu([c]_f)$. This will hold for all f and c, thus solving the Gluing Equations.

All that remains is to convert this into a non-negative integer solution of the Gluing Equations. Note that, as functions of the α_{il}, $M^\pm(f, c)$ only depend on the G-orbit of $[f, c]$; there are finitely many such orbits by Remark 6.4.3, so there are actually just finitely many Gluing Equations (as equations in the α_{il}). Note that the values of α_{il} from (6.5.2) are not all zero, because, for fixed i, $C_{L+1}(\Gamma)$ can be expressed as a finite partition of $[-]_{x_i}$ equivalence classes as in (6.5.1), and $\mu(C_{L+1}(\Gamma)) = 1$. We can then promote our non-negative real number solution of the Gluing Equations to a non-negative integer solution using the following claim. Moreover, since our real number solution isn't identically zero, we can arrange that the integer solution isn't identically zero either. This is an example of linear programming, a technique which has been widely used in topology, an early instance being Haken's work on normal surface theory [46].

Claim: Let A be an integer matrix defining a linear map $A : \mathbb{R}^n \to \mathbb{R}^m$. If $\exists v \in \ker A - \{0\}$ with non-negative entries, then $\exists w \in \ker A - \{0\}$ with non-negative integer entries.

Proof: Let $v \in \ker A - \{0\}$ have non-negative entries. In fact we may assume that all entries of v are strictly positive (else delete columns in A corresponding to the zero entries of v and solve the claim for this matrix, and reintroduce the zero entries to w afterwards). It suffices to find w with non-negative rational entries since we can multiply out denominators to make the entries integers. Now $A\mathbb{R}^n$ is the closure of $A\mathbb{Q}^n$ so both have the same dimension as vector spaces over $\mathbb{R}$ and $\mathbb{Q}$ respectively; thus $\ker(A)$ and $\ker(A) \cap \mathbb{Q}^n$ also have the same dimensions, and so the former must be the closure of the latter. Therefore we can choose $w \in \ker(A) \cap \mathbb{Q}^n$ to be a rational approximation of v, close enough so that it has positive entries. ■

□

The inductive construction of $\mathcal{V}_L, ..., \mathcal{V}_0$ will be left to Sections 6.6 and 6.7. To close this section we show that Theorem 1.2.1 follows from the existence of $\mathcal{V}_0$.

Proof of theorem 1.2.1, given $\mathcal{V}_0$.

Take some triple $(Z, K, (c_x)) \in \mathcal{V}_0$. Property (2) and the connectedness of X imply that $Z = X$. K acts cocompactly on X so must be finite index in G. K acts freely on X by definition of $\mathcal{V}_0$. Property (4) in conjunction with Theorem 6.1.2 and Corollary 6.1.4 tells us that X/K is virtually special. We can then take $G' < K$ finite index such that X/G' is special. □

6.6 Controlling boundary hyperplanes

To go from $\mathcal{V}_j$ to $\mathcal{V}_{j-1}$ we will glue together the various complexes Z/K along the quotients of certain boundary hyperplanes. In this section we will establish what boundary hyperplanes are and which ones we are gluing along, and we will prove some technical lemmas (to be used later) that control the behaviour of these hyperplanes. Lemma 6.6.3 comes from page 1062 of Agol's paper, and Lemma 6.6.7 comes from page 1063, but both are recast to fit with our definitions. A novel feature of our argument is the notion of portals, which we introduce in Definition 6.6.4; these are the parts of boundary hyperplanes that we want to glue along - they will be used extensively in Section 6.7

For this section fix $(Z, H, (c_x)) \in \mathcal{V}_j$.

6.6.1 The Zipping Lemma

Definition 6.6.1. (Boundary hyperplanes)
For e an edge crossing out of Z we call $H(e)$ a *boundary hyperplane* of Z. Equivalently, boundary hyperplanes are hyperplanes $H(e)$ for e an edge intersecting Z and $c_e(e) \leq j$. Z is an intersection of half-spaces, so if H is a boundary hyperplane then Z is contained in one half-space of H. Let $\partial Z \subset Z$ be the union of all boundary hyperplanes intersected with Z.

Remark 6.6.2. For vertices $x, y \in X^0$ with $x \in Z$, let γ be a shortest edge path from x to y. Then the following are equivalent:

(1) $y \notin Z$,

(2) γ crosses a boundary hyperplane,

(3) y and z are separated by a boundary hyperplane.

Indeed if $y \notin Z$ then the first time γ leaves Z it must cross a boundary hyperplane, so (1) implies (2). And it follows from Proposition 2.4.30 that (2) implies (3). Finally, (3) implies (1) because Z is contained in one half-space of the boundary hyperplane. In particular this shows that any two vertices in Z are connected by an edge path that stays in Z.

If there are two edges crossing different boundary hyperplanes (possibly in different triples of $\mathcal{V}_j$), and these edges give the same colouring equivalence class, then we want to be able to "zip" together these boundary hyperplanes in a colour-compatible way. The following lemma will help us to achieve this, although we won't actually do the zipping until Section 6.7.

Lemma 6.6.3. *(Zipping Lemma)*
Let H be a boundary hyperplane of Z. Then the edges e crossing out of Z with $H = H(e)$ all induce the same class $[c_e]_H$ and hence give the same colour $c_e(\overline{H}) = c_e(e)$ (this can be thought of as the colour of H, and we'll refer to it as such).

Proof. Let S be a square in X with an edge e joining vertices $x_1, x_2 \in Z$, let e_1, e_2 be the other edges incident at x_1, x_2 respectively and suppose they cross out of Z with $H(e_1) = H(e_2) = H$. By property (2) of $\mathcal{V}_j$, $c_{x_1}(e) > j \geq c_{x_1}(e_1)$. $\overline{H(e)}$ and $\overline{H}$ intersect, so are adjacent vertices in Γ. Now $[c_{x_1}]_{H(e)} = [c_e]_{H(e)}$, so c_{x_1} agrees with c_e on the ball of radius $c_{x_1}(e)$ about $\overline{H(e)}$ in Γ. But this ball contains the ball of radius $c_{x_1}(e_1)$ about $\overline{H}$, hence $[c_{x_1}]_H = [c_e]_H$. Similarly $[c_{x_2}]_H = [c_e]_H$. So $[c_{e_1}]_H = [c_{x_1}]_H = [c_{x_2}]_H = [c_{e_2}]_H$.

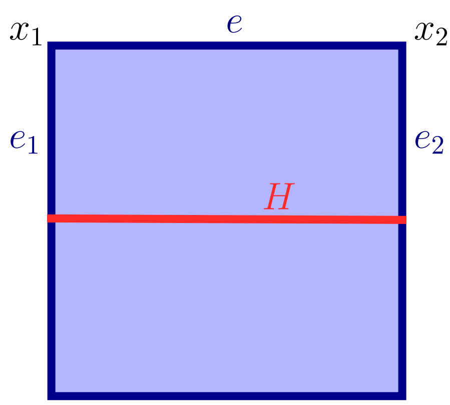

Figure 6.4: The square S.

Z is an intersection of half-spaces, so $H \cap Z$ is an intersection of half-spaces in the induced cube structure on H; so by Remark 6.6.2 (applied to $H \cap Z \subset H$ instead of

$Z \subset X$) any two vertices in H that lie in Z are joined by an edge path in H that stays in Z. Vertices in H that lie in Z correspond to edges dual to H that cross out of Z, and an edge in H that lies in Z corresponds to a square, as in Figure 6.4, joining edges e_1, e_2 dual to H that cross out of Z. Thus the lemma follows from the fact that $[c_{e_1}]_H = [c_{e_2}]_H$. □

6.6.2 j-boundary hyperplanes and portals

Definition 6.6.4. (j-boundary hyperplanes and portals)
If a boundary hyperplane of $Z \in \mathcal{V}_j$ has colour j (in the sense of the Zipping Lemma), call it a *j-boundary hyperplane*. For H a j-boundary hyperplane, let $P(H) := Z \cap H$ be the *portal of H leading to Z*. Whenever we talk about a portal P it will implicitly be equipped with a choice of $Z \in \mathcal{V}_j$ that it leads to. If an edge e dual to H crosses out of Z, say that e is *dual* to $P(H)$. These are shown in Figure 6.5, with j-boundary hyperplanes in red and other boundary hyperplanes in black. Note that portals need not be bounded. Let $\partial_j Z \subset \partial Z$ denote the union of all portals leading to Z.

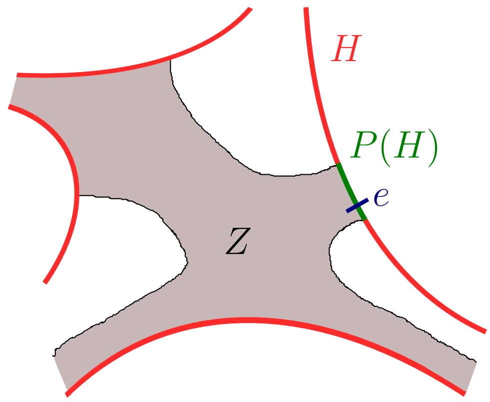

Figure 6.5: The portal of H leading to Z.

As alluded to at the beginning of this section, we will be gluing together the various Z/K along the K-quotients of their portals. To facilitate this we now establish some lemmas that control the behaviour of j-boundary hyperplanes and portals.

Lemma 6.6.5. *Let $H_1, ..., H_n$ be pairwise intersecting hyperplanes, with $Z \cap H_i \neq \emptyset$ for each i, then Z contains a vertex x incident at edges $e_1, ..., e_n$ that are dual to $H_1, ..., H_n$ respectively (and so $e_1, ..., e_n$ form the corner of an n-cube in X).*

Proof. It suffices to prove the lemma for $H_1, ..., H_n$ a maximal family of pairwise intersecting hyperplanes with $Z \cap H_i \neq \emptyset$ for each i. Let $y \in Z$ be a vertex. By Helly's Theorem for cube complexes (Proposition 2.4.31) there is an n-cube C intersected by the hyperplanes $H_1, ..., H_n$, and C has a vertex x such that no H_i separates x and y. Consider a shortest edge path from y to x; if it crosses no boundary hyperplanes then $x \in Z$ and we are done. Suppose it does cross a boundary hyperplane, H say. By Proposition 2.4.30,

H divides X into two half-spaces, one containing x and the other containing y and Z. For each i, part of H_i is in the half-space containing x, but $Z \cap H_i \neq \emptyset$, so we must also have $H \cap H_i \neq \emptyset$, contradicting the maximality of $H_1, ..., H_n$. □

Lemma 6.6.6. *A vertex in Z cannot be incident at distinct edges dual to j-boundary hyperplanes. Moreover, any two j-boundary hyperplanes are disjoint.*

Proof. Suppose there is a vertex $x \in Z$ incident at distinct edges dual to j-boundary hyperplanes H_1 and H_2. By Proposition 2.4.28, $H_1 \neq H_2$. Furthermore, we know from Lemma 6.2.8(3) that H_1 and H_2 map to distinct hyperplanes $\overline{H}_1$ and $\overline{H}_2$ in $\mathcal{X}$. Since H_1 and H_2 are j-boundary hyperplanes, we have that $c_x(\overline{H}_1) = c_x(\overline{H}_2) = j$. But $d(\overline{H}_1, \overline{H}_2) \leq d(H_1, H_2) \leq 1 < R$, contradicting c_x being a colouring in $C_{L+1}(\Gamma)$. For the second part of the lemma, if we have two intersecting j-boundary hyperplanes then we can apply Lemma 6.6.5 to reduce to the first part of the lemma. □

Gluing together the Z/K along K-quotients of portals will form a graph of spaces, and in the corresponding graph of groups the edge groups incident to each vertex group will form a malnormal family by the following lemma.

Lemma 6.6.7. *Stabilisers of distinct portals intersect trivially.*

Proof. Suppose H_0 and H_1 are distinct hyperplanes giving rise to distinct portals $P_0 := Z \cap H_0$ and $P_1 := Z \cap H_1$, and suppose $1 \neq k \in K$ stabilises both of them. Since k has no fixed points in Z, and since $P_0, P_1 \subset Z$ are convex, we deduce that k restricts to hyperbolic isometries of P_0 and P_1 with translation axes a_1, a_2 respectively (Proposition 2.3.5 and Theorem 2.3.6(2)). Any two translation axes of k lie within bounded neighbourhoods of each other, hence we can apply Theorem 2.3.3 to see that a_1 and a_2 bound a flat strip, and by δ-hyperbolicity this strip can have width at most δ. Any point p on P_0 is contained in a cube C of X; and one of the edges of C closest to p will be dual to P_0 and have endpoint x in Z, with $d(p, x) \leq \frac{1}{2}\sqrt{\dim C} \leq \frac{1}{2}\sqrt{\dim X}$. The same is true for P_1, and so there is a path β in Z between vertices $x_0, x_1 \in Z$ of length at most $\delta + \sqrt{\dim X}$ with x_0, x_1 being incident at edges dual to P_0, P_1 respectively. By considering the sequence of cubes that β travels through, there is an edge path γ in Z from x_0 to x_1 with $\gamma \subset N_{\sqrt{\dim X}}(\beta)$. Let γ have edges $e_1, ..., e_n$ and vertices $x_0 = y_0, y_1, ..., y_n = x_1$. Since $R \geq \delta + 2\sqrt{\dim X}$ (and thus the mystery of R is revealed!), we have that $d(\overline{H(e_i)}, \overline{H_0}) \leq d(H(e_i), P_0) \leq R$ for $1 \leq i \leq n$, and so $\overline{H(e_i)}$ and $\overline{H_0}$ are adjacent vertices in Γ.

For $1 \leq i \leq n$, we know from property (1) of $\mathcal{V}_j$ that $[c_{y_{i-1}}]_{e_i} = [c_{y_i}]_{e_i}$, so $c_{y_i}(\overline{H_0}) = c_{y_{i-1}}(\overline{H_0})$. We deduce that $c_{x_1}(\overline{H_0}) = c_{x_0}(\overline{H_0})$. And $c_{x_0}(\overline{H_0}) = j = c_{x_1}(\overline{H_1})$ since H_0 and H_1 are j-boundary hyperplanes. But $d(\overline{H_1}, \overline{H_0}) \leq d(H_1, H_0) \leq \delta \leq R$, and by 6.2.8(3) we know that $\overline{H_0} \neq \overline{H_1}$, hence $\overline{H_0}$ and $\overline{H_1}$ are adjacent vertices in Γ which are given the same colour by c_{x_1}, a contradiction. □

Definition 6.6.8. (Splitting along colourings)

For $\overline{H}$ a hyperplane in $\mathcal{X}$ and $c \in C_{L+1}(\Gamma)$, let $B(\overline{H}, c) := \overline{H} \cap \bigcup c^{-1}([1, j])$ be the intersection of $\overline{H}$ with other hyperplanes in $\mathcal{X}$ that are coloured $\leq j$ by c. Define $\overline{H}$ *split along* c by $\overline{H} - c := \overline{H} - B(\overline{H}, c)$ (this will of course depend on j, but j is fixed for the rest of the chapter so we don't include it in the notation). Working in the cubical subdivision of $\overline{H}$, $\overline{H} - c$ will be a cube complex with some missing faces corresponding to where we have removed $B(\overline{H}, c)$. In general $\overline{H} - c$ will be disconnected, so for a vertex $\bar{x}$ in $\overline{H}$, let $(\overline{H} - c)(\bar{x})$ denote the component of $\overline{H} - c$ containing $\bar{x}$.

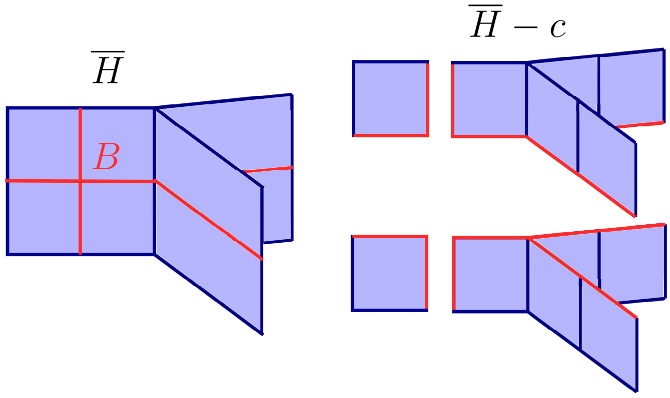

Figure 6.6: $\overline{H}$ split along c.

Lemma 6.6.9. *(Portal covers)*

Let H be a j-boundary hyperplane with portal $P = Z \cap H$ and let e be an edge dual to P. Let x_0 denote the midpoint of e - so x_0 is a vertex of H. Then the quotient map $m : X \to \mathcal{X}$ restricts to a universal covering map

$$m|_{\mathring{P}} : \mathring{P} \to (\overline{H} - c_e)(\bar{x}_0), \tag{6.6.1}$$

where $\mathring{P}$ is the interior of P with respect to the metric topology of H (equivalently $\mathring{P}$ is P minus all boundary hyperplanes $H' \neq H$ that intersect H). Moreover, $(\overline{H} - c_e)(\bar{x}_0) = (\overline{H} - c)(\bar{x})$ for any other $x \in P$ a vertex of H and any $c \in [c_e]_H$ - in particular, by the Zipping Lemma, $(\overline{H} - c_e)(\bar{x}_0)$ is independent of the choice of e dual to P. Furthermore, the group of deck transformations of $m|_{\mathring{P}}$ is $N_P := \{g \in N \,|\, gx_0 \in P\}$ (where N is from Lemma 6.2.5).

Proof. Consider a hyperplane H_1 with $H_1 \cap H \cap Z \neq \emptyset$. By Lemma 6.6.5, there is a vertex $x \in Z$ with edges e_1, e_2 dual to H_1, H respectively. Then

$$c_{e_1}(\overline{H_1}) = c_x(\overline{H_1}) = c_{e_2}(\overline{H_1}) = c_e(\overline{H_1}), \tag{6.6.2}$$

with the third equality due to the Zipping Lemma. By property (2) of $\mathcal{V}_j$, H_1 is a boundary hyperplane if and only if $c_{e_1}(\overline{H_1}) \leq j$. So (6.6.2) implies that H_1 is a boundary hyperplane if and only if $c_e(\overline{H_1}) \leq j$.

The restriction of the quotient map $m : X \to \mathcal{X} = X/N$ certainly defines a map $m|_{\mathring{P}} : \mathring{P} \to \overline{H} \subset \mathcal{X}$ with $m(x_0) = \bar{x}_0$. Now a path γ in P based at x_0 can go anywhere in H except cross over a boundary hyperplane H_1, which by the above arguments is equivalent to γ not crossing a hyperplane H_1 with $c_e(\overline{H_1}) \leq j$, which in turn is equivalent to $m \circ \gamma$ not crossing $B(\overline{H}, c_e)$ (see Definition 6.6.8). This establishes that $m|_{\mathring{P}}$ is a covering. Note that $\mathring{P}$ is equal to H intersected with open half-spaces corresponding to other boundary hyperplanes, so it is convex in X and hence simply connected, making $m|_{\mathring{P}}$ a universal covering.

If $c \in [c_e]_H$, then it follows from the definition of $[-]_H$ that, among the hyperplanes intersecting $\overline{H}$, the hyperplanes coloured $\leq j$ by c are exactly the same as those coloured $\leq j$ by c_e - hence $\overline{H} - c = \overline{H} - c_e$. If we also replace x_0 by a different $x \in P$ then clearly $\bar{x} = m_P(x) \in \overline{H} - c$, and so $(\overline{H} - c_e)(x) = (\overline{H} - c)(x) = (\overline{H} - c)(x_0)$.

Finally, we know that $m : X \to \mathcal{X}$ has N as the group of deck transformations, that $\mathring{P}$ is a component of $m^{-1}((\overline{H} - c_e)(\bar{x}_0))$, and that N acts on these components. So N_P is the stabiliser in N of $\mathring{P}$ (and of P), it preserves the covering $m|_{\mathring{P}}$ and it acts transitively on $m|_{\mathring{P}}^{-1}(\bar{x}_0) = P \cap m^{-1}(\bar{x}_0)$ - thus N_P is exactly the group of deck transformations of $m|_{\mathring{P}}$. □

6.7 Gluing up hyperplanes

We are now ready to start constructing $\mathcal{V}_{j-1}$ from $\mathcal{V}_j$. Our strategy will be to glue together different complexes Z/K along the K-quotients of portals with compatible colourings. To glue together two quotient portals, we need them to be isomorphic as complexes and for this isomorphism to be colour compatible; initially it may not be possible to glue up all the portals, but the idea is to make it possible by passing to finite covers of the Z/K. The arguments in this section are based on Theorem 3.1 from Agol's paper, for example we quote the same theorems of Haglund–Wise (Theorem 6.1.3) and Bestvina–Feighn (Theorem 2.8.7) - but our arguments contain considerably more detail and will appear quite different as they are recast to work for our set-up.

6.7.1 Teleports

Definition 6.7.1. (Compatible portals)
We say that two portals P and P' leading to $(Z, K, (c_x)), (Z', K', (c'_x)) \in \mathcal{V}_j$ respectively are *compatible* if there are edges e and f dual to P and P' respectively such that $[e, c_e] \in G \cdot [f, c'_f]$.

Let P and P' be compatible portals as above, and say they lie in hyperplanes H and H'. Take $g \in G$ and edges e and f dual to P and P' such that $[e, c_e] = g[f, c'_f]$, and let x_0

and y_0 be the midpoints of e and f. So $e = gf$, $H = gH'$ and $x_0 = gy_0$. As N is normal in G, we have the following commuting diagram.

$$\begin{array}{ccc} X & \xrightarrow{g} & X \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle m} \\ \mathcal{X} & \xrightarrow{g} & \mathcal{X} \end{array} \tag{6.7.1}$$

Then g acts on the hyperplane $\overline{H'}$ to produce

$$\begin{aligned} g(\overline{H'} - c'_f)(\bar{y}_0) &= (\overline{H} - gc'_f)(\bar{x}_0) \\ &= (\overline{H} - c_e)(\bar{x}_0) \qquad \text{by Lemma 6.6.9 and the fact that } [c_e]_H = [gc'_f]_H. \end{aligned}$$

As the maps (6.6.1) are coverings for P and P', we deduce that g restricts to a cube isomorphism $\mathring{P}' \to \mathring{P}$ and also $P' \to P$. In fact $P' \to P$ is equivariant with respect to the group isomorphism $N_{P'} \to N_P$; $n \mapsto gng^{-1}$. This can all be put into the following commutative diagram.

$$\begin{array}{ccc} N_{P'} & \xrightarrow[\sim]{g(-)g^{-1}} & N_P \\ \circlearrowright & & \circlearrowright \\ P' & \xrightarrow{g} & P \\ \uparrow & & \uparrow \\ \mathring{P}' & \xrightarrow{g} & \mathring{P} \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle m} \\ (\overline{H'} - c'_f)(\bar{y}_0) & \xrightarrow{g} & (\overline{H} - c_e)(\bar{x}_0) \end{array} \tag{6.7.2}$$

We also have the following lemma and corollary, which are basically consequences of the Zipping Lemma. These will allow us to group together compatible portals into compatibility classes.

Lemma 6.7.2. *(Teleports)*
Portals P and P' leading to $(Z, K, (c_x)), (Z', K', (c'_x)) \in \mathcal{V}_j$ are compatible if and only if there exists $g \in G$ such that

$$\{[e, c_e] \mid e \text{ is dual to } P\} = g\{[f, c'_f] \mid f \text{ is dual to } P'\}. \tag{6.7.3}$$

*In this case we say that P is a g-*teleport *of P' (note that P could be a g-teleport of P' for several different g, and that P' could have several different g-teleports corresponding to portals that lead to different $Z \in \mathcal{V}_j$).*

Proof. Suppose P and P' are compatible. Then there exist edges e and f dual to P and P', and $g \in G$, such that $[e, c_e] = g \cdot [f, c'_f]$. If f_1 is another edge dual to P', then $e_1 := gf_1$

is dual to P because, as we showed above, $g : P' \to P$ is an isomorphism. Suppose that P and P' lie in hyperplanes H and H'. We then have

$$\begin{aligned} g[f_1, c'_{f_1}] &= \{gf_1\} \times g[c'_{f_1}]_{H'} \\ &= \{e_1\} \times g[c'_f]_{H'} && \text{by the Zipping Lemma} \\ &= \{e_1\} \times [gc'_f]_H \\ &= \{e_1\} \times [c_e]_H && \text{since } [e, c_e] = g \cdot [f, c'_f] \\ &= \{e_1\} \times [c_{e_1}]_H && \text{by the Zipping Lemma} \\ &= [e_1, c_{e_1}]. \end{aligned}$$

This gives the $\supset$ inclusion in (6.7.3), and the $\subset$ inclusion follows similarly by considering $g^{-1} : P \to P'$. Conversely, (6.7.3) clearly implies compatibility of P and P'. □

Corollary 6.7.3. *Teleports form a groupoid on the set of portals, meaning that:*

- *Any portal is a 1-teleport of itself.*
- *If P is a g-teleport of P', then P' is a g^{-1}-teleport of P.*
- *If P is a g-teleport of P' and P' is a g'-teleport of P'', then P is a gg'-teleport of P''*

In particular, compatibility of portals is an equivalence relation, and we will refer to the equivalence classes as compatibility classes. *The compatibility class of a portal P is denoted $[P]$.*

For each triple $(Z, K, (c_x))$, K acts on the set of portals leading to Z, and these portals are disjoint by Lemma 6.6.6. This implies that the map $P/K_P \to Z/K$ is an embedding for each portal P. The aim is to glue together quotients Z/K and Z'/K' along portals P/K_P and $P'/K'_{P'}$ for compatible portals P and P'. Lemma 6.7.2 tells us that $P \cong P'$, which is a good start, but it doesn't imply that we get isomorphic quotients $P/K_P \cong P'/K'_{P'}$; the next step is to overcome this by taking finite covers of the complexes Z/K, or equivalently by replacing the groups K with finite index subgroups.

Lemma 6.7.4. *For each triple $(Z, K, (c_x))$ there is a finite index subgroup $\hat{K} \lhd K$, such that whenever portals P and P' lead to $(Z, K, (c_x)), (Z', K', (c'_x)) \in \mathcal{V}_j$, with P a g-teleport of P', then*

$$\hat{K}_P = g\hat{K}'_{P'}g^{-1}. \tag{6.7.4}$$

In particular $P/\hat{K}_P \cong P'/\hat{K}'_{P'}$.

Proof. To each group K we apply Theorem 6.1.5 with respect to the stabilisers of a set of K-orbit representatives of portals leading to Z. Note that these stabilisers form a malnormal collection by Lemma 6.6.7. This gives us finite index subgroups $\dot{K}_P < K_P$ for every portal P leading to Z, such that $k\dot{K}_P k^{-1} = \dot{K}_{kP}$ for $k \in K$.

Recall from Section 6.2 the group $N \triangleleft G$ with finite hyperplane quotients. For each portal P, define the group

$$\hat{K}_P := N_P \cap \bigcap_{g,P'} g\dot{K}'_{P'}g^{-1}, \tag{6.7.5}$$

where we range over all g, P' such that P is a g-teleport of P'. This is an intersection of subgroups of G_P that all act cocompactly on P (N_P acts cocompactly by Lemma 6.6.9), so they all have finite index in G_P. Moreover, there are finitely many triples in $\mathcal{V}_j$, so there is a uniform bound on the size of the quotients $P'/\dot{K}'_{P'}$, and hence a uniform bound on the indices $|G_P : g\dot{K}'_{P'}g^{-1}|$. This means that $\hat{K}_P$ has finite index in G_P. Also note that $1\dot{K}_P 1^{-1} = \dot{K}_P$ is one of the terms in the intersection, so $\hat{K}_P$ is a finite index subgroup of $\dot{K}_P$. Equation (6.7.4) follows from (6.7.5), Corollary 6.7.3 and the normality of N in G; in particular this implies $\hat{K}_P \triangleleft K_P$.

Finally, we apply Theorem 6.1.5 to obtain finite index subgroups $\hat{K} \triangleleft K$ for each triple in $(Z, K, (c_x)) \in \mathcal{V}_j$, such that the $\hat{K}$-stabilisers of portals leading to Z are indeed the subgroups $\hat{K}_P$ defined by (6.7.5). □

Replacing each $(Z, K, (c_x))$ by $(Z, \hat{K}, (c_x))$ would preserve all the properties of $\mathcal{V}_j$ except the Gluing Equations. $(Z, \hat{K}, (c_x))$ would contribute $|K : \hat{K}|$ times more to each set $\mathcal{V}_j^{\pm}(f, c)$ than $(Z, K, (c_x))$ does. But making $|K : \hat{K}|$ copies of $(Z, K, (c_x))$ in $\mathcal{V}_j$ would have the same effect. Therefore, for each triple $(Z, K, (c_x)) \in \mathcal{V}_j$, we can replace $(Z, K, (c_x))$ by some number of copies of $(Z, \hat{K}, (c_x))$ such that the Gluing Equations are preserved. To simplify notation we will not write $\hat{K}$ for the rest of the section, we will instead assume that the triples in $\mathcal{V}_j$ already satisfy

$$K_P = gK'_{P'}g^{-1} \tag{6.7.6}$$

whenever portals P and P' lead to $(Z, K, (c_x)), (Z', K', (c'_x)) \in \mathcal{V}_j$, with P a g-teleport of P'.

6.7.2 Matching up portal quotients

Definition 6.7.5. ($[P]^+$ and $[P]^-$ -compatible portals)
Let P be a portal leading to $Z \in \mathcal{V}_j$, and suppose it lies in a hyperplane H. For P' a portal leading to $Z' \in \mathcal{V}_j$ that is compatible with P, choose $g \in G$ such that P is a g-teleport of P'. We say that P' is a $[P]^+$*-compatible portal* if $gZ' \cap H^+ \neq \emptyset$ and a $[P]^-$*-compatible portal* if $gZ' \cap H^- \neq \emptyset$ (recall from Definition 2.4.8 and Remark 2.4.19 that $H^{\pm}$ denote the sets of vertices in $N(H)$ on either side of H). By Remark 6.2.1 the

labelling $H^{\pm}$ is G-equivariant, so the definition of $[P]^{\pm}$-compatible portal is independent of the choice of teleport g and of the choice of portal P within the compatibility class.

If P is a $[P]^+$-compatible portal leading to $(Z, K, (c_x))$, P' is a $[P]^-$-compatible portal leading to $(Z', K', (c'_x))$, and P is a g-teleport of P', then Z and gZ' lie on opposite sides of P (i.e. $Z \cap gZ' = P$), so we can glue them together to form a larger (convex) subspace of X, as illustrated below. Moreover, we can define local colouring data for this subspace by combining (c_x) with the g-translate of (c'_x), and these colourings will be compatible on the edges dual to P (in the sense of property (1) of $\mathcal{V}_j$) by Lemma 6.7.2.

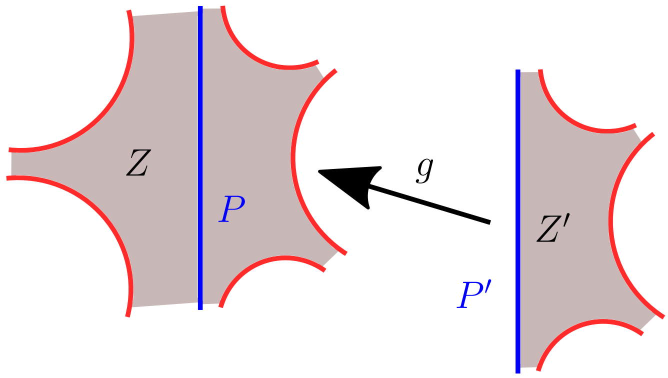

Figure 6.7: Gluing along portals.

By (6.7.6) this gluing descends to the quotients via the following commutative diagram.

$$\begin{array}{ccc} P & \xleftarrow{\;g\;} & P' \\ \downarrow & & \downarrow \\ P/K_P & \xleftarrow{\;\bar{g}\;} & P'/K'_{P'} \end{array} \tag{6.7.7}$$

Thus Z/K is glued to Z'/K' along $P/K_P \overset{\bar{g}}{\cong} P'/K'_{P'}$. We want to glue up all the portal quotients in pairs, so we need the following definition and lemma.

Definition 6.7.6. (Portal quotients)
For a portal P leading to $(Z, K, (c_x)) \in \mathcal{V}_j$ we have a *portal quotient* $P/K_P \hookrightarrow Z/K$ (we consider this as a subspace of Z/K, so portals in the same K-orbit define the same portal quotient). We define the *size* of P/K_P as

$$\text{size}(P/K_P) := |\{K_P \cdot e \,|\, e \in X^1 \text{ is dual to } P\}|.$$

By (6.7.6) compatible portals have quotients of the same size. We say that a portal quotient $P'/K'_{P'}$ is a *$[P]^{\pm}$-compatible portal quotient* if P' is a $[P]^{\pm}$-compatible portal.

Lemma 6.7.7. *For each compatibility class $[P]$, the number of $[P]^+$-compatible portal quotients equals the number of $[P]^-$-compatible portal quotients.*

Proof. Fix a portal P_0 leading to a triple $(Z_0, K^0, (c_x^0))$, and contained in a hyperplane H, and let $[c]_H$ be the associated colouring equivalence class (i.e. $c \in C_{L+1}(\Gamma)$ satisfies $[c]_H = [c_f^0]_H$ for any edge f dual to P_0, possible by the Zipping Lemma). The idea is to show that the total size of $[P_0]^+$-compatible portal quotients equals the total size of $[P_0]^-$-compatible portal quotients - the lemma then follows since all portals in the class $[P_0]$ have quotients of the same size. The key is to prove the following equality:

$$\left\{ (K \cdot e, Z) \;\middle|\; \begin{array}{c} e \in X^1 \text{ dual to a } [P_0]^{\pm}\text{-compatible} \\ \text{portal } P \text{ leading to } (Z, K, (c_x)) \in \mathcal{V}_j \end{array} \right\} = \bigcup_{f \text{ dual to } P_0} \mathcal{V}_j^{\pm}(f, c) \qquad (6.7.8)$$

- Firstly let's check the inclusion $\subset$. Given $(K \cdot e, Z)$ on the LHS of (6.7.8), if P_0 is a g-teleport of P, then $f := ge$ is dual to P_0 and $[ge, gc_e] = [f, c]$ by Lemma 6.7.2. Since P is a $[P_0]^{\pm}$-compatible portal, we also have $gZ \cap H(f)^{\pm} \neq \emptyset$, so $(K \cdot e, Z) \in \mathcal{V}_j^{\pm}(f, c)$ lies on the RHS of (6.7.8).

- For the other inclusion, suppose $(K \cdot e, Z) \in \mathcal{V}_j^{\pm}(f, c)$ for some edge f dual to P_0. Then by definition there is $g \in G$ with $gZ \cap H(f)^{\pm} \neq \emptyset$ and $g[e, c_e] = [f, c]$. We have $c_e(e) = c(f) = j$, so e is dual to some portal P, and P_0 is compatible with P by Lemma 6.7.2. This implies that $(K \cdot e, Z)$ is on the LHS of (6.7.8) as required.

Note that $\mathcal{V}_j^{\pm}(f, c)$ only depends on $\pm$ and the G-orbit of the equivalence class $[f, c]$. It also follows straight from the definition of $\mathcal{V}_j^{\pm}(f, c)$ that if edges f_1, f_2 are dual to P_0 with $[f_1, c] \notin G \cdot [f_2, c]$, then $\mathcal{V}_j^{\pm}(f_1, c) \cap \mathcal{V}_j^{\pm}(f_2, c) = \emptyset$. Thus we can write the RHS of (6.7.8) as a disjoint union by indexing over a finite set of edges f with distinct orbits $G \cdot [f, c]$. Whether we use $+$ or $-$, the Gluing Equations tell us that each set $\mathcal{V}_j^{\pm}(f, c)$ will be the same size, hence the RHS of (6.7.8) will also be the same size. In turn, the LHS of (6.7.8) has the same size for $+$ and $-$, which shows that the total size of $[P_0]^+$-compatible portal quotients equals the total size of $[P_0]^-$-compatible portal quotients, as required. □

Definition 6.7.8. (Graph of spaces $\mathcal{Y}$)
For each compatibility class $[P]$ we now choose a matching between the $[P]^+$-compatible portal quotients and $[P]^-$-compatible portal quotients (possible by Lemma 6.7.7). Define a graph of spaces $\mathcal{Y}$ as follows:

- The vertex spaces in $\mathcal{Y}$ are the spaces Z/K for $(Z, K, (c_x)) \in \mathcal{V}_j$.

- The edge spaces are the portal quotients $P/K_P \hookrightarrow Z/K$ (and we make a choice of portal P from each K-orbit of portals leading to Z).

- For each pair of portal quotients $(P/K_P, P'/K'_{P'})$ from the matching (which we will refer to as an *edge space pair*) we choose $g_P \in G$ such that P is a g_P-teleport of P', and glue the portal quotients together using the map $\bar{g}_P : P'/K'_{P'} \to P/K_P$ from (6.7.7) (and we choose $g_{P'} = g_P^{-1}$).

Now let Y be a connected component of $\mathcal{Y}$. We want to show that $\pi_1(Y) \in \mathcal{QVH}$, but first we must show that it is hyperbolic.

Lemma 6.7.9. *$\pi_1(Y)$ is hyperbolic.*

Proof. We just need to check that the conditions of Theorem 2.8.7 are satisfied for the graph of groups corresponding to Y.

- The vertex groups are hyperbolic because they come from actions of K on Z, which are free cocompact actions on hyperbolic cube complexes.
- The map $P/K_P \hookrightarrow Z/K$ from edge space to vertex space lifts to the inclusion $P \hookrightarrow Z$ of universal covers, which is an isometric embedding, hence the corresponding inclusion of groups $K_P \hookrightarrow K$ will be a quasi-isometric embedding.
- For $(Z, K, (c_x)) \in \mathcal{V}_j$, the collection of portal stabilisers (K_P), indexed over some set of K-orbit representatives of portals leading to Z, is malnormal in K because of Lemma 6.6.7. □

Lemma 6.7.10. *$\pi_1(Y) \in \mathcal{QVH}$.*

Proof. We have already shown that $\pi_1(Y)$ is hyperbolic; and it has vertex groups $K \in \mathcal{QVH}$ by property (4) of $\mathcal{V}_j$. Finally, the edge groups K_P are quasiconvex in $\pi_1(Y)$ by [54, Theorem 1.2]. □

6.7.3 Embedding in X

The next step is to embed the universal cover $\tilde{Y}$ of Y into X, such that the group of deck transformations of $\tilde{Y} \to Y$ is a subgroup of G. We can then complete the inductive step of the gluing construction by promoting $\tilde{Y}$ to a triple in $\mathcal{V}_{j-1}$. The definition of the vertex and edge spaces of Y means that we already have an immersion $Y \to X/G$, so one might think that we can use standard covering space theory to lift this to a map of universal covers $\tilde{Y} \to X$ that is equivariant with respect to a certain subgroup of G. Indeed this would follow if the action of G on X were free, but since G might contain torsion the quotient X/G is actually an orbi-complex, which makes the required covering space arguments more subtle. In the following definition and lemma we give an explicit construction of the map $\tilde{Y} \to X$ that does not assume any knowledge of orbi-complex covering space theory. We do this by first constructing $\tilde{Y}$ as a subspace of X, and then proving that we have a universal covering $\mu : \tilde{Y} \to Y$.

Definition 6.7.11. ($\tilde{Y}$ as a subspace of X)
Fix Z_0/K^0 a vertex space of Y to act as a base vertex space. For $0 \leq i \leq n$, suppose

Z_i/K^i are vertex spaces of Y, suppose P_i, P'_i are portals leading to Z_i, Z_{i+1} respectively such that $(P_i/K^i_{P_i}, P'_i/K^{i+1}_{P'_i})$ is an edge space pair, and let $k_i \in K^i$. Then define

$$g_i := k_0 g_{P_0} k_1 g_{P_1} \cdots g_{P_{i-1}} k_i \in G. \tag{6.7.9}$$

$\tilde{Y} \subset X$ will be the union of all spaces $g_i Z_i \subset X$ for all choices Z_i, P_i, k_i as above, and the restriction of the covering map $\mu : \tilde{Y} \to Y$ will be given by

$$\mu : g_i Z_i \xrightarrow{g_i^{-1}} Z_i \to Z_i/K^i \to Y.$$

Lemma 6.7.12. $\mu : \tilde{Y} \to Y$ *as constructed above is a well-defined universal covering.*

Proof. Consider Z_i, P_i, k_i for $0 \le i \le n$ as in Definition 6.7.11. Firstly we will check that the map μ agrees on the intersection of $g_i Z_i$ and $g_{i+1} Z_{i+1}$. Indeed $g_{P_i} : P'_i \subset Z_{i+1} \to P_i \subset Z_i$ is a teleport between a pair of $[P_i]^+-$ and $[P_i]^--$ compatible portals, so we know that

$$\begin{aligned} g_i Z_i \cap g_{i+1} Z_{i+1} &= g_i(Z_i \cap g_{P_i} k_{i+1} Z_{i+1}) \\ &= g_i P_i. \end{aligned} \tag{6.7.10}$$

Then the two ways of defining the map $\mu : g_i P_i \to Y$ are given by the following commutative diagram, hence they agree.

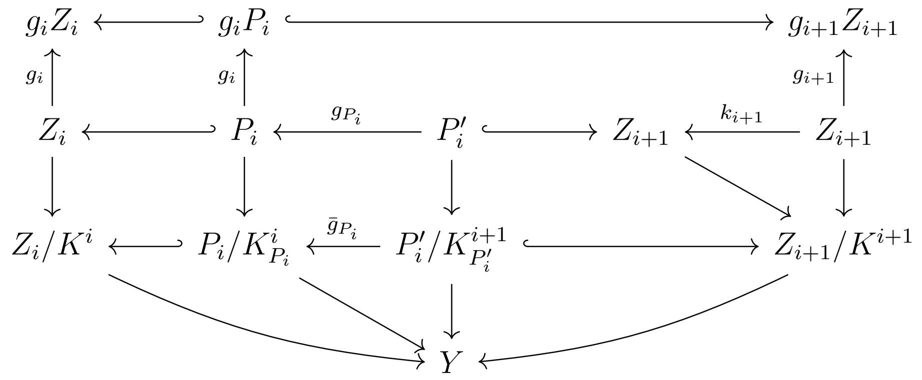

Next we want to investigate the effect of changing k_i to some other element $k'_i \in K$. Suppose this changes g_i to g'_i and g_{i+1} to g'_{i+1}. First note that $g_i Z_i = g'_i Z_i$, and that the map $\mu : g_i Z_i \to Y$ is the same whether defined with g_i or g'_i. For g'_{i+1} there are two cases to consider. If k'_i is in the same $K^i_{P_i}$-coset as k_i, then (6.7.6) implies that we can write

$$k'_i g_{P_i} k_{i+1} = k_i g_{P_i} k'_{i+1},$$

for some $k'_{i+1} \in K^{i+1}$, so as before we have $g_{i+1} Z_{i+1} = g'_{i+1} Z_{i+1}$ and the map $\mu : g_{i+1} Z_{i+1} \to Y$ is the same either way. On the other hand, if k'_i lies in a different $K^i_{P_i}$-coset than k_i, then $g_i P_i$ and $g'_i P_i$ will be distinct portals, and they will lie in disjoint j-boundary hyperplanes by Lemma 6.6.6, hence $g_{i+1} Z_{i+1}$ and $g'_{i+1} Z_{i+1}$ will lie in disjoint half-spaces of X as illustrated below (j-boundary hyperplane translates in red). Similarly, if we change

P_i to a different K^i-orbit of portal leading to Z_i, then $g_{i+1}Z_{i+1}$ will be shifted into a third half-space, disjoint from the other two.

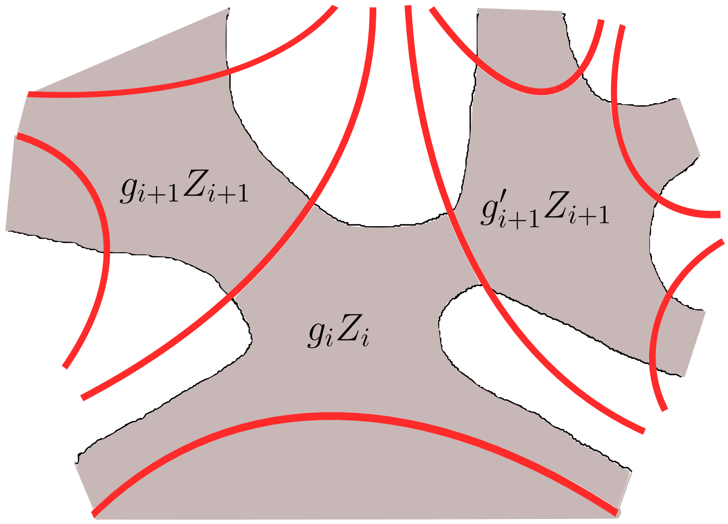

Figure 6.8: The tree structure of $\tilde{Y}$.

The conclusion of this discussion is that the collection of all possible g_iZ_i are arranged in a tree structure, with different g_iZ_i intersecting only along portals as in (6.7.10), and that each map $\mu : g_iZ_i \to Y$ only depends on the space g_iZ_i. $\tilde{Y}$ is a tree of simply connected spaces, so is itself simply connected; and for $g_iP_i \subset g_iZ_i \cap g_{i+1}Z_{i+1}$ as above, we have a translate of Z_{i+1} glued to g_iZ_i along every coset $g_ikK^i_{P_i}$ for $k \in K^i$, so $\mu : \tilde{Y} \to Y$ really is a universal covering map. □

6.7.4 Completing the induction

We now promote $\tilde{Y}$ to a triple in $\mathcal{V}_{j-1}$ by adding the data of a group action and colourings. We define these now, and afterwards we will verify the numbered properties of $\mathcal{V}_{j-1}$ given in Section 6.5.

Definition 6.7.13. (Group action and colourings for $\tilde{Y}$)

Define the group

$$K(Y) := \{g_i \,|\, g_i \text{ is an element of the form (6.7.9) with } Z_i = Z_0 \in \mathcal{V}_j\} < G.$$

If g_l is another element of the form (6.7.9), then g_ig_l is also of the form (6.7.9), and so $g_i : g_lZ_l \to g_ig_lZ_l \subset \tilde{Y}$, hence we get $g_i\tilde{Y} \subset \tilde{Y}$. It is also clear that $\mu g_i = \mu$, so $K(Y)$ is a subgroup of the deck transformation group of $\mu : \tilde{Y} \to Y$. What's more it follows from the construction of μ and $\tilde{Y}$ that $K(Y) \cdot x = \mu^{-1}(\mu(x))$ for all $x \in \tilde{Y}$, hence $K(Y)$ is the full group of deck transformations, and acts freely cocompactly on $\tilde{Y}$.

To complete the triple we need to provide colourings (c^Y_x) for every vertex in $\tilde{Y}$: if $x \in Z$ is a vertex and $(Z, K, (c_x)) \in \mathcal{V}_j$ then endow each translate $g_i x$ with the colouring $c^Y_{g_i x} := g_i c_x$. This defines a triple $(\tilde{Y}, K(Y), (c^Y_x)) \in \mathcal{V}_{j-1}$. Doing this for all connected components Y of $\mathcal{Y}$ defines the entire collection $\mathcal{V}_{j-1}$.

We have already seen that $K(Y)$ acts freely cocompactly on $\tilde{Y}$, and it is clear from the definition that the colourings (c^Y_x) are invariant under the action of $K(Y)$ - so in particular are well-defined. There is just one other thing we need to check before going on to the numbered properties of $\mathcal{V}_{j-1}$, and this is the following lemma.

Lemma 6.7.14. *$\tilde{Y}$ is an intersection of half-spaces in X.*

Proof. For each $g_i Z_i$ in $\tilde{Y}$ as in Definition 6.7.11, and boundary hyperplane H of Z_i of colour $< j$, consider the hyperplane $g_i H$. We claim that $\tilde{Y}$ is an intersection of half-spaces corresponding to these hyperplanes. Any edge leaving $\tilde{Y}$ must cross one of these hyperplanes, so it suffices to show that each of these hyperplanes has a half-space containing $\tilde{Y}$. Indeed consider $g_i Z_i$ and $g_i H$ as above, and consider P_i and $g_{i+1} Z_{i+1}$ as from Definition 6.7.11, so that we have $g_i Z_i$ glued to $g_{i+1} Z_{i+1}$ along $g_i P_i$. Let P_i lie in a hyperplane H_i and let M be the part of $\tilde{Y}$ on the opposite side of $g_i H_i$ to $g_i Z_i$. We will show that M and $g_i Z_i$ lie entirely on the same side of $g_i H$.

If $g_i H \cap g_i H_i = \emptyset$ then we are done. Now suppose $g_i H$ and $g_i H_i$ intersect, then apply Lemma 6.6.5 to find a vertex $x \in Z_i$ incident at edges e, e_i which are dual to H and H_i respectively, and form the corner of a square in X. Let $g_{P_i} y$ be the vertex at the other end of e_i (so $y \in Z_{i+1}$), then $g_{P_i} y$ is incident at an edge dual to H. P_i is a g_{P_i}-teleport of some portal $P'_i \subset Z_{i+1}$, so

$$[e_i, c^i_{e_i}] = [e_i, g_{P_i} c^{i+1}_{e_{i+1}}]$$

where $e_{i+1} := g^{-1}_{P_i} e_i$ and c^i, c^{i+1} are the colourings for Z_i, Z_{i+1}. So

$$j > c_{e_i}(\overline{H}) = g_{P_i} c^{i+1}_{e_{i+1}}(\overline{H}) = g_{P_i} c^{i+1}_y(\overline{H}).$$

Thus $g^{-1}_{P_i} H$ is a boundary hyperplane of Z_{i+1} of colour $< j$, so Z_{i+1} and y lie entirely on the same side of it. In turn $g_{P_i} Z_{i+1}$ and Z_i lie entirely on the same side of H, and $g_{i+1} Z_{i+1}$ and $g_i Z_i$ lie entirely on the same side of $g_i H$. Iterating this argument along branches of M shows that M and $g_i Z_i$ lie entirely on the same side of $g_i H$ as required.

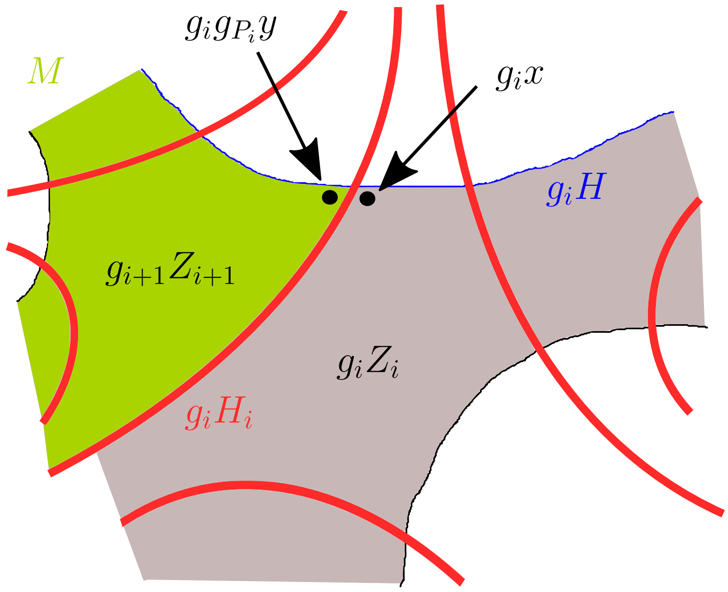

Figure 6.9: $\tilde{Y}$ as an intersection of half-spaces.

□

Lastly we check the numbered properties of $\mathcal{V}_{j-1}$ from Section 6.5:

(1) If $e \in X^1$ joins vertices $x, y \in \tilde{Y}$, we need $[e, c_x^Y] = [e, c_y^Y]$ to hold. This is clearly true if e is contained within a translate $g_i Z_i$ because c^Y is a translate of the colourings for Z_i. From the proof of Lemma 6.7.12, we know that the different $g_i Z_i$ translates are separated by hyperplanes in a tree structure, and that two translates are adjacent only if they are of the form $g_i Z_i$ and $g_{i+1} Z_{i+1}$ as in Definition 6.7.11. We may assume that $k_{i+1} = 1$, as changing k_{i+1} doesn't change the translate $g_{i+1} Z_{i+1}$ - so $g_{i+1} = g_i g_{P_i}$. If e crosses from $g_i Z_i$ to $g_{i+1} Z_{i+1}$, then it crosses the portal translate $g_i P_i = g_{i+1} P_i'$, and can be written $e = g_i e_i = g_{i+1} e_i'$ for e_i dual to P_i and e_i' dual to P_i'. Suppose c^i, c^{i+1} are the colourings for Z_i, Z_{i+1}, and that x is the end of e in $g_i Z_i$. Then putting $x = g_i x_i$ and $y = g_{i+1} y_i'$ we have

$$\begin{aligned}
[e, c_x^Y] &= [e, g_i c^i_{x_i}] \\
&= g_i [e_i, c^i_{x_i}] \\
&= g_i [e_i, c^i_{e_i}] \\
&= g_i g_{P_i} [e_i', c^{i+1}_{e_i'}] && \text{since } P_i \text{ is a } g_{P_i}\text{-teleport of } P_i' \\
&= g_{i+1} [e_i', c^{i+1}_{y_i'}] \\
&= [e, g_{i+1} c^{i+1}_{y_i'}] \\
&= [e, c_y^Y].
\end{aligned}$$

(2) Given $e \in X^1$ joining vertices x, y with $x \in \tilde{Y}$, $c_x^Y(e) > j$ if and only if e is contained within some translate $g_i Z_i$; $c_x^Y(e) = j$ if and only if e crosses a portal translate into a different $g_i Z_i$; so $y \in \tilde{Y}$ if and only if $c_x^Y(e) > j - 1$ as required.

(3) Let $(K \cdot e, Z) \in \mathcal{V}_j^{\pm}(f,c)$, with $(Z, K, (c_x)) \in \mathcal{V}_j$ and Z/K in a component Y of $\mathcal{Y}$. We have $Z = Z_i$ and $g_i Z_i \subset \tilde{Y}$ for some g_i as in Definition 6.7.11, so we get $(K(Y) \cdot g_i e, \tilde{Y}) \in \mathcal{V}_{j-1}^{\pm}(f,c)$. Moreover, $\mu : \tilde{Y} \to Y$ is a universal covering with deck transformation group $K(Y)$, so all choices of g_i put $g_i e$ in the same $K(Y)$-orbit. This means that we have a bijection $\mathcal{V}_j^{\pm}(f,c) \to \mathcal{V}_{j-1}^{\pm}(f,c)$ for all (f,c), and so the Gluing Equations hold in $\mathcal{V}_{j-1}$.

(4) Finally, $K(Y) \cong \pi_1(Y) \in \mathcal{QVH}$ by Lemma 6.7.10.

Chapter 7

Quasi-isometric rigidity for graphs of virtually free groups

The aim of this chapter is to prove a quasi-isometric rigidity theorem for graphs of virtually free groups (Theorem 1.3.3), which in particular applies to certain "generic" HNN extensions of a free group over cyclic subgroups (Theorem 1.3.1). Along the way we prove Theorem 1.3.2, characterising when a graph of virtually free groups is subgroup separable; plus we show that the relative hyperbolicity assumption in Theorem 1.3.3 is necessary (Section 7.6).

7.1 Rigid line patterns

7.1.1 Rigid line patterns and rigid trees

Given a vertex group of $G \in \mathscr{C}^\bullet$ (Definition 2.6.6), we need to understand the structure of its incident edge groups from a coarse geometry perspective. For this we review the notion of rigid line pattern.

Definition 7.1.1. (Line pattern)
A *line pattern* $\mathcal{L}$ on a metric space X is a collection of bi-infinite quasi-geodesics in distinct $\sim$-equivalence classes. If $(X, \mathcal{L}_X)$ and $(Y, \mathcal{L}_Y)$ are spaces with line patterns, then we say that a quasi-isometry $f : X \to Y$ *respects line patterns* if there is an associated bijection $f_* : \mathcal{L}_X \to \mathcal{L}_Y$ such that $f(l) \sim f_*(l)$ for all $l \in \mathcal{L}_X$. In this case we write $f : (X, \mathcal{L}_X) \to (Y, \mathcal{L}_Y)$. Observe that a composition of quasi-isometries respecting line patterns is itself a quasi-isometry respecting line patterns.

Definition 7.1.2. (Free group with line pattern, [21])
Consider a finitely generated free group F of rank greater than one. Let $\mathcal{H}$ be a finite collection of cyclic subgroups of F. The *line pattern* $\mathcal{L} = \mathcal{L}_\mathcal{H}$ generated by $\mathcal{H}$ is the collection of quasi-geodesics corresponding to left cosets of the subgroups in $\mathcal{H}$. Note

that the $(F, \mathcal{H})$ depends on a choice of finite generating sets, but all such choices are equivalent up to quasi-isometry respecting line patterns.

Definition 7.1.3. (Vertex group with induced line pattern)
If the graph of groups decomposition (G, Γ) contains a non-abelian free vertex group G_u whose incident edge groups are all cyclic, then the collection of G_u-conjugates of incident edge groups forms a line pattern $\mathcal{L}_u$ for G_u. If u lifts to a vertex $\tilde{u}$ in the Bass–Serre covering tree T for (G, Γ), then $G_{\tilde{u}}$ is non-abelian free and has cyclic incident edge groups, and the collection of these incident edge groups forms a line pattern $\mathcal{L}_{\tilde{u}}$ for $\tilde{u}$ (this time the collection of incident edge groups is already closed under conjugation in $G_{\tilde{u}}$).

Definition 7.1.4. (Rigid line pattern, [21])
If X is a space with line pattern $\mathcal{L}_X$, let $\mathcal{QI}(X, \mathcal{L}_X)$ denote the group of quasi-isometries from X to itself that respect the line pattern $\mathcal{L}_X$ (formally an element of $\mathcal{QI}(X, \mathcal{L}_X)$ is an $\approx$-equivalence class of quasi-isometries, but when we write $f \in \mathcal{QI}(X, \mathcal{L}_X)$ we will mean f to be a particular choice of quasi-isometry). Similarly, let $\mathrm{Isom}(X, \mathcal{L}_X)$ denote the group of isometries of X that respect $\mathcal{L}_X$. We say that $(X, \mathcal{L}_X)$ is a *rigid model space* if the natural map $\iota : \mathrm{Isom}(X, \mathcal{L}_X) \to \mathcal{QI}(X, \mathcal{L}_X)$ is an isomorphism.

A free group with line pattern $(F, \mathcal{L})$ is *rigid* if there is a quasi-isometry $\phi : (F, \mathcal{L}) \to (X, \mathcal{L}_X)$ to a rigid model space. If the group F is clear, then we will simply say that $\mathcal{L}$ is *rigid*.

Building on the work of [21], Cashen proved the following characterization of rigidity for a free group with line patterns:

Theorem 7.1.5. *(Cashen [20, Theorem 4.29])*
Let F be a finitely generated free group and $\mathcal{H}$ a finite set of cyclic subgroups in F. Then we have three mutually exclusive cases:

(1) $(F, \mathcal{L}_{\mathcal{H}})$ is rigid,

(2) F is the fundamental group of a hyperbolic surface with boundary, with the boundary components corresponding to the subgroups $\mathcal{H}$,

(3) F is not of type (2), and admits a non-trivial free or cyclic splitting relative to $\mathcal{H}$.

Definition 7.1.6. (ϕ-conjugacy action)
If $\phi : (F, \mathcal{L}) \to (X, \mathcal{L}_X)$ is a quasi-isometry to a rigid model space, then we get an isomorphism $\mathcal{QI}(F, \mathcal{L}) \to \mathrm{Isom}(X, \mathcal{L}_X)$ given by $f \mapsto \iota^{-1}(\phi f \phi^{-1})$, where $\iota : \mathrm{Isom}(X, \mathcal{L}_X) \to \mathcal{QI}(X, \mathcal{L}_X)$ is the isomorphism as above. We call the corresponding isometric action of $\mathcal{QI}(F, \mathcal{L})$ on $(X, \mathcal{L}_X)$ the *ϕ-conjugacy action.*

Remark 7.1.7. The ϕ-conjugacy action is independent of ϕ in the sense that if $\phi_1, \phi_2 : (F, \mathcal{L}) \to (X, \mathcal{L}_X)$ are two different quasi-isometries, then the isometry $\iota^{-1}(\phi_2\phi_1^{-1}) : (X, \mathcal{L}_X) \to (X, \mathcal{L}_X)$ is equivariant with respect to the ϕ_1-conjugacy action on the left hand side and the ϕ_2-conjugacy action on the right hand side. In particular, the translation length of an element $f \in \mathcal{QI}(F, \mathcal{L})$ with respect to the ϕ-conjugacy action is independent of ϕ. Sometimes we will just say the action of $\mathcal{QI}(F, \mathcal{L})$ on $(X, \mathcal{L}_X)$ if we do not wish to refer to a particular ϕ.

Remark 7.1.8. If $\mathcal{L} \subset \mathcal{L}'$ are two line patterns on F, and $\mathcal{L}$ is rigid, then $\mathcal{L}'$ must also be rigid. This is because if $\phi : (F, \mathcal{L}) \to (X, \mathcal{L}_X)$ is a quasi-isometry to a rigid model space, then $\phi : (F, \mathcal{L}') \to (X, \phi(\mathcal{L}'))$ is also a quasi-isometry to a rigid model space - as $\mathcal{L}_X \subset \phi(\mathcal{L}')$ and $\mathcal{QI}(X, \phi(\mathcal{L}')) < \mathcal{QI}(X, \mathcal{L}_X) \cong \mathrm{Isom}(X, \mathcal{L}_X)$.

The main theorem we will need about rigid line patterns is the following. Part (2) is due to Cashen–Macura, and part (3) is due to Hagen–Touikan (which also relies on the construction of Cashen–Macura).

Theorem 7.1.9. *(Cashen–Macura [21, Main Theorem], Hagen–Touikan [40, Theorem C])*
Let $(F, \mathcal{L})$ be a free group with line pattern. The following are equivalent:

(1) $\mathcal{L}$ is rigid.

(2) The decomposition space $\mathcal{D}_\mathcal{L}$*, obtained from ∂F by identifying the two limit points of each line $l \in \mathcal{L}$ and taking the quotient topology, is connected, has no cut points and no cut pairs.*

(3) There is a quasi-isometry $\alpha : (F, \mathcal{L}) \to (Y, \mathcal{L}_Y)$ to a rigid model space, where Y is a locally finite tree with no leaves and $\mathcal{L}_Y$ is a collection of bi-infinite geodesics.

We will call $(Y, \mathcal{L}_Y)$ from Theorem 7.1.9(3) a *rigid tree* for $(F, \mathcal{L})$. Note that distinct bi-infinite geodesics in Y cannot be at finite Hausdorff distance, so the α-conjugacy action of $\mathcal{QI}(F, \mathcal{L})$ on Y isometrically maps each geodesic in $\mathcal{L}_Y$ onto another geodesic in $\mathcal{L}_Y$. F acts on itself by left multiplication, preserving $\mathcal{L}$, and so we can view it as a subgroup of $\mathcal{QI}(F, \mathcal{L})$. The corresponding action of F on a rigid tree Y satisfies the following lemma.

Lemma 7.1.10. *Let $(F, \mathcal{L})$ be a rigid line pattern with a quasi-isometry $\alpha : (F, \mathcal{L}) \to (Y, \mathcal{L}_Y)$ to a rigid tree. Then the action of F on Y is free and cocompact, and α is at bounded distance from any orbit map of F.*

Proof. By definition of the α-conjugacy action, for $g \in F$ the diagram of quasi-isometries

$$\begin{array}{ccc} F & \xrightarrow{\alpha} & Y \\ \downarrow{g} & & \downarrow{g} \\ F & \xrightarrow{\alpha} & Y \end{array} \tag{7.1.1}$$

commutes up to bounded distance. The two g maps are actually isometries, so this diagram defines two quasi-isometries $F \to F$ at bounded distance from each other, with quasi-isometry constants only depending on α. As F has Cayley graph a regular tree, one can easily deduce that the distance between these two quasi-isometries $F \to F$ also just depends on α. This implies that α is at bounded distance from any orbit map of F. It immediately follows that the action of F on Y is cocompact, and it must also be free because F is torsion-free. □

Remark 7.1.11. "Random line patterns" are rigid line patterns in the following sense: working in the Cayley graph of F with respect to a given free basis $\mathcal{B}$, and taking geodesic representatives for the lines in $\mathcal{L}_\mathcal{H}$, if some $l \in \mathcal{L}_\mathcal{H}$ contains every reduced word of length 3 as a subsegment, then $\mathcal{L}_\mathcal{H}$ is rigid. This follows from [22, Corollary 5.5] and Theorem 7.1.5 (note that the only possibility of being in case (2) but not case (3) of Theorem 7.1.5 is if the hyperbolic surface is a pair of pants, but then F will admit a free splitting relative to each subgroup in $\mathcal{H}$ individually). In particular, if $w \in F$ is a random word of length n with respect to $\mathcal{B}$, then the probability that $\mathcal{L}_{\{\langle w \rangle\}}$ is rigid tends to 1 exponentially quickly as $n \to \infty$.

7.1.2 Rigid decompositions are JSJ decompositions

In this section we explore the close relation between rigid line patterns and vertex groups of $G \in \mathscr{C}_{tf}^\bullet$ (Definition 2.6.6).

Lemma 7.1.12. *Let $G \in \mathscr{C}_{tf}^\bullet$ with a JSJ tree T. Then for each $u \in V_0T_c$ the group G_u is a non-abelian free group and the induced line pattern $(G_u, \mathcal{L}_u)$ is rigid.*

Proof. The line pattern $(G_u, \mathcal{L}_u)$ must be in one of the three cases of Theorem 7.1.5. We cannot be in case (2) because the splitting of G has no QH vertex groups. G_u cannot split freely relative to its incident edge groups because G is one-ended. G_u cannot admit a cyclic splitting relative to its incident edge groups by Remark 2.6.7, so case (3) can't happen either. Therefore we must be in case (1), which means that $(G_u, \mathcal{L}_u)$ is rigid. □

We know from Section 2.7.2 that the group $\mathscr{G}$ of quasi-isometries $G \to G$ acts on the tree of cylinders T_c, and that each quasi-isometry restricts to maps between the vertex groups in G; we record here that these maps also respect the line patterns.

Lemma 7.1.13. *Let $G \in \mathscr{C}_{tf}^\bullet$ and $u \in V_0T_c$. Then $[f] \in \mathscr{G}$ induces a $\approx$-class of quasi-isometries $[f]_u : (G_u, \mathcal{L}_u) \to (G_{\hat{f}(u)}, \mathcal{L}_{\hat{f}(u)})$ that respect line patterns.*

Proof. This follows immediately from Theorem 2.7.9(3). □

We also have a converse to Lemma 7.1.12 as follows.

Proposition 7.1.14. *Let G be a finitely generated group that splits over two-ended subgroups by acting minimally on a tree T. Suppose that the vertex stabilisers are all either*

(1) virtually non-abelian free with incident edge stabilisers inducing rigid line patterns,

(2) virtually infinite cyclic,

with at least one vertex stabiliser of the first type. Then G is one ended and T is a JSJ tree for G with no QH vertex groups.

Proof. First suppose that G is not one-ended. Let T_{DS} be a G-tree with finite edge stabilisers and one ended vertex stabilisers (the *Dunwoody-Stallings decomposition*). Each vertex stabiliser G_v for T acts on a minimal subtree $S_v \subset T_{DS}$. If $u, v \in VT$ are the endpoints of an edge e, then S_u and S_v must intersect, else G_e would stabilise the arc between them, contradicting the finiteness of edge stabilisers in T_{DS}. The union of all S_v is then a G-invariant subtree of T_{DS}, and so by minimality it is the whole of T_{DS}. In particular, at least one of the S_v is non-trivial.

If all edge stabilisers for T are elliptic in T_{DS}, then the type (2) vertex stabilisers are also elliptic in T_{DS}, and so there must be a type (1) vertex stabiliser G_v that acts non-trivially on S_v relative to its incident edge stabilisers, and the same is true of any finite index subgroup of G_v. But G_v has a finite index subgroup with incident edge stabilisers inducing a rigid line pattern, contradicting Theorem 7.1.5. Hence at least some edge stabilisers for T are not elliptic in T_{DS}, but such an edge stabiliser G_e is two-ended, so must stabilise a unique axis $\ell_e \subset T_{DS}$, and moreover any finite index $\mathbb{Z} < G_e$ will act on ℓ_e by translations. Also note that $\ell_e \subset S_v$ for a vertex v incident at e. We now have the following claim.

<u>Claim:</u> There exists an edge $e_{DS} \in ET_{DS}$ and a type (1) vertex stabiliser G_v such that $e_{DS} \subset \ell_e$ for a unique edge $e \in \mathrm{lk}(v)$.

<u>Proof:</u> Suppose not. Let $e_{DS} \in ET_{DS}$ be contained in at least one axis ℓ_e. Given an axis ℓ_e, if G_e is incident at a vertex stabiliser G_v, if G_v is type (1) then by assumption there is another $e' \in \mathrm{lk}(v)$ with $e_{DS} \subset \ell_{e'}$, while if G_v is type (2) then all edge stabilisers incident at G_v will be commensurable in G and have the same axis ℓ_e, so again there is another $e' \in \mathrm{lk}(v)$ with $e_{DS} \subset \ell_{e'}$. Therefore, for any axis ℓ_e containing e_{DS}, there are two more edges incident at either end of e whose stabilisers have axes that also contain e_{DS}. Thus e_{DS} is contained in infinitely many axes ℓ_e. There are finitely many G-orbits of edges in T, so there exists $e \in ET$ with $e_{DS} \subset \ell_e$ and an infinite sequence (g_n) in G such that the edges $g_n(e)$ are all distinct and $e_{DS} \subset \ell_{g_n(e)}$ for all n. Noting that $g_n(\ell_e) = \ell_{g_n(e)}$, we can precompose the g_n by elements of G_e that translate along ℓ_e and assume that the edges $g_n(e_{DS})$ lie at bounded distance from e_{DS}. Passing to a subsequence of (g_n), we

can assume that the edges $g_n(e_{DS})$ are all at distance d from e_{DS} and all lie in the same component of $T_{DS} - e_{DS}$. But then there is an edge $e'_{DS} \subset \ell_e$ at distance d from e_{DS} such that $g_n(e'_{DS}) = e_{DS}$ for all n, and so e_{DS} has infinite stabiliser, a contradiction. ■

Taking e_{DS}, e and G_v as from the claim, we will now convert the action of G_v on S_v into an action on a different tree S that gives a splitting of G_v over finite subgroups relative to its incident edge stabilisers, contradicting Theorem 7.1.5 as before. S will be bipartite with respect to vertex sets $VS = V_0S \sqcup V_1S$, and is defined as follows:

- V_0S is the collection of components of $S_v - G_v \cdot e_{DS}$.
- V_1S is the collection of axes $\ell_{g(e)}$ for $g \in G_v$.
- $U \in V_0S$ and $\ell_{g(e)} \in V_1S$ form an edge if they intersect.

S_v has a tree of spaces decomposition formed by the components $U \in V_0S$ and edges $g(e_{DS})$ for $g \in G_v$, and each edge $g(e_{DS})$ is contained in the unique axis $\ell_{g(e)} \in V_1S$, therefore S is indeed a tree. The action of G_v on S_v induces an action on S. Each edge group $G_{g(e)}$ for $g \in G_v$ stabilises the axis $\ell_{g(e)} \in V_1S$, while each edge $e' \in \mathrm{lk}(v) - G_v \cdot e$ has axis $\ell_{e'}$ contained in some component $U \in V_0S$, and so $G_{e'}$ stabilises U. On the other hand, the G_v-stabiliser of an edge $(U, \ell_{g(e)}) \in ES$ must stabilise (setwise) the two G_v-translates of e_{DS} contained in $\ell_{g(e)}$ that touch U, and so this stabiliser must be finite. Therefore S gives a splitting of G_v over finite subgroups relative to its incident edge stabilisers, as required.

We now show that T is a JSJ tree for G with no QH vertex groups. Let T_J be a JSJ tree for G over two-ended subgroups. By [38, Lemma 2.6(3)], the edge stabilisers of T are all elliptic in T_J, and hence so are the vertex stabilisers of type (2). For a type (1) vertex stabiliser G_v, we can apply Theorem 7.1.5 to a finite index free subgroup of G_v whose incident edge stabilisers induce a rigid line pattern, and deduce that G_v is elliptic in T_J. Therefore each edge stabiliser of T is either contained in an edge stabiliser of T_J, or has both its adjacent vertex stabilisers contained in the same vertex stabiliser G^J_x of T_J. The second case can't happen, as then G^J_x would be flexible, and hence QH (Remark 2.6.7); one can then argue that the vertex stabilisers of T contained in G^J_x would have line patterns coming from compact hyperbolic surfaces with boundary, contradicting rigidity of the line patterns by Theorem 7.1.5. We conclude that every edge stabiliser of T is contained in an edge stabiliser of T_J, making T universally elliptic. We already showed that the vertex stabilisers of T are elliptic in T_J, hence they are elliptic in every universally elliptic tree for G, and so T is a JSJ tree for G. Finally, there are no QH vertex stabilisers of T by Theorem 7.1.5. □

Example 7.1.15. Proposition 7.1.14 allows us to construct explicit examples of groups satisfying the assumptions of Theorem 1.3.3, especially when combined with Remark 7.1.11. For example if $\mathbb{F}_m$ and $\mathbb{F}_n$ are finitely generated free groups, and $1 \neq w_1 \in \mathbb{F}_m$, $1 \neq w_2 \in \mathbb{F}_n$ are not proper powers, and w_1, w_2 can each be represented by cyclically reduced words that contain every possible length three subword, then the following amalgam satisfies the assumptions of Theorem 1.3.3.

$$G = \mathbb{F}_m *_{\mathbb{Z}} \mathbb{F}_n := \langle \mathbb{F}_m, \mathbb{F}_n \mid w_1 = w_2 \rangle$$

The assumption that w_1 and w_2 are not proper powers ensures that G is hyperbolic.

If instead we have $1 \neq w_1, w_2 \in \mathbb{F}_n$, but otherwise with the same properties, then the following HNN extension also satisfies the assumptions of Theorem 1.3.3.

$$G = \mathbb{F}_n *_{\mathbb{Z}} := \langle \mathbb{F}_n, t \mid t w_1 t^{-1} = w_2 \rangle$$

For the HNN extension, if $g^{-1} w_1 g = w_2$ or w_2^{-1} for $g \in \mathbb{F}_n$, then G is hyperbolic relative to $\langle w_1, gt \rangle$ - which is isomorphic to either $\mathbb{Z}^2$ or the Klein bottle group (and the latter has an index two $\mathbb{Z}^2$ subgroup). Otherwise G is hyperbolic.

7.2 Balanced groups, separability, and torsion

In this section we prove Theorem 1.3.2. In [97] Wise characterised subgroup separable groups in $\mathscr{C}_{tf}$ as being balanced. We generalise the notion of balanced in the obvious way to all groups in $\mathscr{C}$, and in Theorem 7.2.6 we prove that being balanced is equivalent to subgroup separability. The other implications of Theorem 1.3.2 are dealt with in Section 7.2.3.

7.2.1 Balanced graphs of groups

Definition 7.2.1. (Balanced graph of groups)
A finite graph of groups (G, Γ) with two-ended edge groups is *balanced* if the following equation holds for any loop in Γ given by edges $e_0, e_1, ..., e_n = e_0$, where $\iota(e_i) = v_i$ and $\tau(e_i) = v_{i+1}$, and $\zeta_{e_{i-1}}(G_{e_{i-1}})$ is commensurable to $g_i \zeta_{\bar{e}_i}(G_{e_i}) g_i^{-1}$ in G_{v_i} for some $g_i \in G_{v_i}$.

$$1 = \prod_{i=1}^{n} \frac{\left[g_i \zeta_{\bar{e}_i}(G_{e_i}) g_i^{-1} : \zeta_{e_{i-1}}(G_{e_{i-1}}) \cap g_i \zeta_{\bar{e}_i}(G_{e_i}) g_i^{-1}\right]}{\left[\zeta_{e_{i-1}}(G_{e_{i-1}}) : \zeta_{e_{i-1}}(G_{e_{i-1}}) \cap g_i \zeta_{\bar{e}_i}(G_{e_i}) g_i^{-1}\right]}. \tag{7.2.1}$$

Lemma 7.2.2. *(G, Γ) is balanced if and only if there is no relation $g h^p g^{-1} = h^q$ for h an infinite order element of an edge group and $|p| \neq |q|$.*

Proof. Let T be the Bass–Serre tree corresponding to (G, Γ). The edge loop of Definition 7.2.1 corresponds to an edge path $e_0, e_1, ..., e_n$ in T such that the edge stabilisers G_{e_i} are all commensurable in G and there exists $g \in G$ with $g(e_0) = e_n$. The product (7.2.1) becomes:

$$\prod_{i=1}^{n} \frac{[G_{e_i} : G_{e_i} \cap G_{e_{i-1}}]}{[G_{e_{i-1}} : G_{e_i} \cap G_{e_{i-1}}]} = \frac{[G_{e_n} : G_{e_0} \cap G_{e_n}]}{[G_{e_0} : G_{e_0} \cap G_{e_n}]} \tag{7.2.2}$$

Let $h \in G_{e_0}$ be infinite order, so $ghg^{-1} \in G_{e_n}$, and $\langle h \rangle < G_{e_0}$ and $\langle ghg^{-1} \rangle < G_{e_n}$ are finite index subgroups. Suppose $\langle h \rangle \cap \langle ghg^{-1} \rangle$ is generated by $h^q = gh^pg^{-1}$. Then (7.2.2) is equal to

$$\begin{aligned} \frac{[G_{e_n} : \langle gh^pg^{-1} \rangle]}{[G_{e_0} : \langle h^q \rangle]} &= \frac{|p|\,[G_{e_n} : \langle ghg^{-1} \rangle]}{|q|\,[G_{e_0} : \langle h \rangle]} \\ &= \frac{|p|\,[gG_{e_0}g^{-1} : g\langle h \rangle g^{-1}]}{|q|\,[G_{e_0} : \langle h \rangle]} \\ &= \frac{|p|}{|q|}, \end{aligned}$$

thus completing the proof of the lemma. □

Remark 7.2.3. Hyperbolic groups and CAT(0) groups are always balanced as they cannot contain a relation $gh^pg^{-1} = h^q$ for h an infinite order element and $|p| \neq |q|$. This is Corollary 2.3.7 for CAT(0) groups, and for hyperbolic groups the relation $g^n h^{p^n} g^{-n} = h^{q^n}$ combined with the fact that $\langle h \rangle$ is undistorted [15, Corollary III.Γ.3.10(1)] implies $|p| = |q|$.

Remark 7.2.4. It follows from Lemma 7.2.2 that, given $(\hat{G}, \hat{\Gamma}) \to (G, \Gamma)$ a finite cover of graphs of groups (or equivalently $\hat{G} < G$ finite index with the restricted action on the Bass–Serre tree T), $(\hat{G}, \hat{\Gamma})$ is balanced if and only if (G, Γ) is balanced.

Theorem 7.2.5. *(Wise [97, Theorem 5.1])*
Suppose a finitely generated group G splits as a finite graph of groups (G, Γ), where the edge groups are cyclic and the vertex groups are free. Then G is subgroup separable if and only if (G, Γ) is balanced.

We generalise Theorem 7.2.5 to the following, which gives us the equivalence of (1) and (2) in Theorem 1.3.2.

Theorem 7.2.6. *Let $G \in \mathscr{C}$ split as a finite graph of groups (G, Γ), where the edge groups are two-ended and the vertex groups are virtually free. Then G is subgroup separable if and only if (G, Γ) is balanced, and in this case G is virtually torsion-free.*

7.2.2 Removing torsion

In this section we prove Theorem 7.2.6. A key ingredient in Wise's proof of Theorem 7.2.5 is the omnipotence of free groups. The omnipotence of free groups can be viewed as a special case of Theorem 1.2.5, about commanding subgroups, and it is this theorem that we make use of here.

The direction of Theorem 7.2.6 where we assume that G is subgroup separable is straightforward. Indeed if (G,Γ) is not balanced then by Lemma 7.2.2 we have a relation $gh^pg^{-1} = h^q$ for h an infinite order element of an edge group and $|p| \neq |q|$. $\langle h^{|pq|}\rangle$ is separable in G, so there is a homomorphism $\rho : G \to \bar{G}$ to a finite group such that $\rho(h^i) \notin \rho(\langle h^{|pq|}\rangle)$ for $1 \le i < |pq|$, which implies that $\rho(h)$ has order $k|pq|$ for some integer k. But then $\rho(h^p)$ and $\rho(h^q) = \rho(gh^pg^{-1})$ are conjugate elements in $\bar{G}$ with distinct orders $k|q|$ and $k|p|$ respectively, a contradiction.

In the rest of this section we prove the other direction of Theorem 7.2.6, so suppose G has a balanced graph of groups decomposition (G,Γ) with virtually free vertex groups and two-ended edge groups. We will show that G is virtually torsion-free, subgroup separability then follows from Remark 7.2.4 and Theorem 7.2.5. Note that some vertex groups in (G,Γ) might be two-ended, and others infinite-ended, but this does not matter to us, as our arguments in this section will work for both.

Lemma 7.2.7. *There exist finite-index torsion-free subgroups $G'_{\bar{e}} = G'_e \lhd G_e$ and $G'_v \lhd G_v$ for $e \in E\Gamma$ and $v \in V\Gamma$, such that*

$$\zeta_e(G_e) \cap G'_v = \zeta_e(G'_e) \tag{7.2.3}$$

whenever $\tau(e) = v$.

Proof. Let $v \in V\Gamma$. We can assume that incident edge groups in G_v that are commensurable up to conjugacy in G_v are actually commensurable in G_v (composing the edge maps ζ_e by conjugations in G_v preserves equation 7.2.3 and (G,Γ) remains a balanced graph of groups decomposition). For each $e \in \mathrm{lk}(v)$, let $P_e < G_v$ be the unique maximal subgroup commensurable to $\zeta(G_e)$ (P_e will be the G_v-stabiliser of the limit set of $\zeta_e(G_e)$ in the Gromov boundary of G_v). Note that $P_{e_1} = P_{e_2}$ if and only if $\zeta_{e_1}(G_{e_1})$ and $\zeta_{e_2}(G_{e_2})$ are commensurable in G_v. The family

$$\mathbb{P}_v := \{P_e \mid e \in \mathrm{lk}(v)\}$$

is an almost malnormal collection of quasiconvex subgroups in G_v (Definition 2.8.5). G_v is virtually free, so it is hyperbolic and virtually special, and it commands $\mathbb{P}_v$ by Theorem 1.2.5. Hence there are finite-index subgroups $\dot{P} < P$ for each $P \in \mathbb{P}_v$ such that, for any further finite-index subgroups $P' < \dot{P}$ with $P' \lhd P$, there exists a finite-index normal subgroup $G'_v \lhd G_v$ with $P \cap G'_v = P'$. Let $\dot{G}_v \lhd G_v$ be a finite-index free normal subgroup.

Shrinking the $\dot{P}$ if necessary, we can assume that they lie inside $\dot{G}_v$. In particular $\dot{P} \cong \mathbb{Z}$. We may also assume $\dot{P} \lhd P$. We do all this for every $v \in V\Gamma$.

Now, for each $e \in E\Gamma$, choose a finite-index normal subgroup $G'_e = G'_{\bar{e}} \lhd G_e$ with $G'_e < \zeta_e^{-1}(\dot{P}_e) \cap \zeta_{\bar{e}}^{-1}(\dot{P}_{\bar{e}})$. Note that $G'_e \cong \mathbb{Z}$.

The fact that (G, Γ) is balanced implies there exist positive integers K_e for $e \in E\Gamma$, with $K_e = K_{\bar{e}}$, such that

$$\frac{K_{e_1}}{[\zeta_{e_1}(G_{e_1}) : \zeta_{e_1}(G_{e_1}) \cap \zeta_{e_2}(G_{e_2})]} = \frac{K_{e_2}}{[\zeta_{e_2}(G_{e_2}) : \zeta_{e_1}(G_{e_1}) \cap \zeta_{e_2}(G_{e_2})]} \in \mathbb{N} \tag{7.2.4}$$

whenever $\tau(e_1) = \tau(e_2) = v \in V\Gamma$ with $\zeta_{e_1}(G_{e_1})$ and $\zeta_{e_2}(G_{e_2})$ commensurable in G_v. If we choose the G'_e such that $[G_e : G'_e] = NK_e$ for some fixed N, then

$$\frac{[\zeta_{e_1}(G'_{e_1}) : \zeta_{e_1}(G'_{e_1}) \cap \zeta_{e_2}(G'_{e_2})]}{[\zeta_{e_2}(G'_{e_2}) : \zeta_{e_1}(G'_{e_1}) \cap \zeta_{e_2}(G'_{e_2})]} = \frac{K_{e_2}[\zeta_{e_1}(G_{e_1}) : \zeta_{e_1}(G'_{e_1}) \cap \zeta_{e_2}(G'_{e_2})]}{K_{e_1}[\zeta_{e_2}(G_{e_2}) : \zeta_{e_1}(G'_{e_1}) \cap \zeta_{e_2}(G'_{e_2})]} = 1, \tag{7.2.5}$$

and as $\zeta_{e_1}(G'_{e_1}), \zeta_{e_2}(G'_{e_2}) < \dot{P}_{e_1} = \dot{P}_{e_2} \cong \mathbb{Z}$, we deduce that $\zeta_{e_1}(G'_{e_1}) = \zeta_{e_2}(G'_{e_2})$. Hence for each $e \in E\Gamma$ we can define $P'_e := \zeta_e(G'_e) \lhd P_e$, and this only depends on P_e (the normality follows because P'_e is characteristic in $\dot{P}_e$ and $\dot{P}_e$ is normal in P_e).

As each G_v commands $\mathbb{P}_v$, there exist finite-index normal subgroups $G'_v \lhd G_v$ with $P \cap G'_v = P'$ for all $P \in \mathbb{P}_v$, which yields equation (7.2.3). And intersecting with $\dot{G}_v$ if necessary, we can assume that each G'_v is torsion-free. □

Proposition 7.2.8. *Let $G \in \mathscr{C}$. If G is balanced, then G is virtually torsion-free and therefore subgroup separable.*

Proof. We define a finite cover of graphs of groups $(\hat{G}, \hat{\Gamma}) \to (G, \Gamma)$, so that $\hat{G} < G$ is a finite index subgroup. The edge and vertex groups of $(\hat{G}, \hat{\Gamma})$ will be copies of the G'_e and G'_v constructed in Lemma 7.2.7, which are torsion-free, so $\hat{G}$ will be torsion-free.

The data for constructing the cover $(\hat{G}, \hat{\Gamma}) \to (G, \Gamma)$ is as follows.

- Have a surjective graph morphism $p : \hat{\Gamma} \to \Gamma$.
- For $\hat{v} \in V\hat{\Gamma}$ and $p(\hat{v}) = v$, have an inclusion $\iota_{\hat{v}} : \hat{G}_{\hat{v}} \hookrightarrow G_v$ with image G'_v. For $\hat{e} \in E\hat{\Gamma}$ and $p(\hat{e}) = e$, have an inclusion $\iota_{\hat{e}} : \hat{G}_{\hat{e}} \hookrightarrow G_e$ with image G'_e.
- If $\tau(\hat{e}) = \hat{v} \in V\hat{\Gamma}$, $p(\hat{e}) = e$ and $p(\hat{v}) = v$, then there is $h_{\hat{e}} \in G_v$ such that the following diagram commutes

$$\begin{array}{ccccc} \hat{G}_{\hat{e}} & & \xrightarrow{\quad\zeta_{\hat{e}}\quad} & & \hat{G}_{\hat{v}} \\ \downarrow{\scriptstyle \iota_{\hat{e}}} & & & & \downarrow{\scriptstyle \iota_{\hat{v}}} \\ G_e & \xrightarrow{\zeta_e} & G_v & \xrightarrow{h_{\hat{e}}(-)h_{\hat{e}}^{-1}} & G_v. \end{array} \tag{7.2.6}$$

  Moreover, the elements $h_{\hat{e}}$ provide a complete set of double coset representatives $G'_v h_{\hat{e}} \zeta_e(G_e)$ as $\hat{e}$ ranges over edges in $p^{-1}(e)$ with $\tau(\hat{e}) = \hat{v}$.

One can check that this is indeed the correct data by thinking in terms of graphs of spaces and considering elevations of the various edge maps (we omit an explanation of this), or alternatively one can compare this data with [5, Definitions 2.1 and 2.6].

An alternative characterisation of the $h_{\hat{e}}$ (again with fixed e and $\hat{v}$) is that they provide a complete set of coset representatives for the subgroup $G'_v\zeta_e(G_e)/G'_v$ in the finite quotient G_v/G'_v. Now

$$\begin{aligned} \left|\frac{G'_v\zeta_e(G_e)}{G'_v}\right| &= [\zeta_e(G_e) : \zeta_e(G_e) \cap G'_v] && (7.2.7) \\ &= [\zeta_e(G_e) : \zeta_e(G'_e)] \qquad \text{by (7.2.3)} \\ &= [G_e : G'_e], \end{aligned}$$

so there will be $[G_v : G'_v]/[G_e : G'_e]$ such cosets, and hence the same number of $\hat{e}$.

As a result, we must satisfy the gluing equation

$$|p^{-1}(v)|\frac{[G_v : G'_v]}{[G_e : G'_e]} = |p^{-1}(e)| \tag{7.2.8}$$

whenever $\tau(e) = v \in V\Gamma$; and conversely, if we have numbers $|p^{-1}(v)|$ and $|p^{-1}(e)|$ that solve the equations (7.2.8), then such a finite cover $(\hat{G}, \hat{\Gamma})$ can be constructed. But such a solution is easy, just set

$$|p^{-1}(v)| = \frac{M}{[G_v : G'_v]}, \qquad |p^{-1}(e)| = \frac{M}{[G_e : G'_e]}, \tag{7.2.9}$$

where M is a common multiple of the $[G_v : G'_v]$ and $[G_e : G'_e]$. □

7.2.3 Relative hyperbolicity and virtual specialness

In this section we prove the other implications of Theorem 1.3.2.

Lemma 7.2.9. *Let $G \in \mathscr{C}$ split as a finite balanced graph of groups (G, Γ), where the edge groups are two-ended and the vertex groups are virtually free. Then, replacing G by a finite index torsion-free subgroup, we can arrange that each cylinder Y in the corresponding Bass–Serre tree T has stabiliser G_Y which admits a product splitting $G_Y = \mathbb{Z} \times F_n$ $(n \geq 0)$ such that the $\mathbb{Z}$ factor fixes Y pointwise and the F_n factor acts freely cocompactly on Y.*

Proof. G_Y acts cocompactly on Y, and all edge and vertex stabilisers are two-ended by Lemma 2.7.12. By Theorem 7.2.5 we can assume that G is torsion-free, so then G_Y splits as a graph of groups with all edge and vertex stabilisers isomorphic to $\mathbb{Z}$ - such groups are called *generalised Baumslag–Solitar groups* (or GBS groups). It follows from the proof of [58, Proposition 2.6] that G_Y contains a finite index subgroup $\dot{G}_Y$ which admits a product splitting $\dot{G}_Y = \mathbb{Z} \times F_n$ $(n \geq 0)$ such that the $\mathbb{Z}$ factor fixes Y pointwise and the F_n factor acts freely cocompactly on Y. Alternatively, we can apply Proposition 2.5.3 to (G_Y, Y)

to produce a finite index subgroup $\dot{G}_Y < G_Y$ such that each vertex stabiliser of $\dot{G}_Y$ in Y is equal to its incident edge stabilisers (recall that subgroup separability of G implies subgroup separability of G_Y). This $\dot{G}_Y$ will thus admit a product splitting $G_Y = \mathbb{Z} \times F_n$ where the $\mathbb{Z}$ factor is equal to any vertex or edge stabiliser.

Any finite index subgroup of $\dot{G}_Y$ will admit a similar product splitting, so we may apply Proposition 2.5.3 to the action of G on the tree of cylinders T_c and a set of G-orbit representatives of cylinder vertices, and this will produce a finite index subgroup of G satisfying the conclusions of the lemma. □

Proof of Theorem 1.3.2. The equivalence of (1) and (2) is Theorem 7.2.6. It remains to show the equivalence of (2), (3) and (4). Fix an action of G on a tree T with two-ended edge stabilisers and virtually free vertex stabilisers and let (G, Γ) be the quotient graph of groups decomposition.

Let's start by showing the equivalence of (2), that (G, Γ) is balanced, and (3), that G is hyperbolic relative to peripheral subgroups that are virtually $\mathbb{Z} \times \mathbb{F}_n$ ($n \geq 0$). (2) implies (3) by combining Lemma 7.2.9 and Proposition 2.8.8. Conversely, suppose for contradiction we have (3) but not (2), then Lemma 7.2.2 gives us infinite order elements h, g such that $gh^pg^{-1} = h^q$ with $|p| \neq |q|$. By [68, Corollary 4.21] the element h must lie in a (conjugate of a) peripheral subgroup, call it P. Moreover, g will also belong to P, otherwise $\langle gh^pg^{-1} \rangle = \langle h^q \rangle < P \cap gPg^{-1}$, contradicting the almost malnormality of the peripheral subgroups. But then we contradict P being virtually $\mathbb{Z} \times \mathbb{F}_n$.

Next we'll show the equivalence of (2), that (G, Γ) is balanced, and (4), that G is virtually special. Firstly suppose that (G, Γ) is balanced, and let T be the corresponding Bass–Serre tree; by Lemma 7.2.9 we may assume that G is torsion-free and that its cylinder stabilisers are isomorphic to $\mathbb{Z} \times \mathbb{F}_n$, with all stabilisers of edges in the cylinder being equal to the $\mathbb{Z}$ factor. By [51], G is the fundamental group of a non-positively curved cube complex X; moreover, the v-arcs from [51, Definition 10.1] are hyperplanes that correspond to the edge groups in (G, Γ), so X decomposes as a graph of cube complexes in the sense of [53] corresponding to (G, Γ). We want to show that X is virtually special. By [53, Theorem 1.4] it is enough to show that G has finite stature with respect to its vertex stabilisers in T (finite stature is defined in [53, Definition 1.2]). It suffices to show that for any $e_1, e_2 \in ET$ either $G_{e_1} \cap G_{e_2} = G_{e_1}$ or $G_{e_1} \cap G_{e_2} = \{1\}$. Indeed if e_1 and e_2 belong to the same cylinder then $G_{e_1} \cap G_{e_2} = G_{e_1}$ by our assumption on the cylinders, and otherwise the edge groups are not commensurable so intersect trivially.

Finally, suppose that (G, Γ) is not balanced. Again, by Lemma 7.2.2, G contains infinite order elements g, h with $gh^pg^{-1} = h^q$ and $|p| \neq |q|$, hence so will any finite index subgroup of G. This implies that G is not virtually cubulated - as this would contradict the conjugation invariance of the combinatorial translation length of isometries of a CAT(0) cube complex (see [42, 102]). □

Remark 7.2.10. We observe that if we know G is hyperbolic relative to a family $\mathcal{P}$ of virtually abelian peripheral subgroups (where $\mathcal{P}$ might not be the family of cylinder stabilisers), then the cylinder stabilisers will also be virtually abelian. Indeed by Theorem 1.3.2 we know that the cylinder stabilisers are virtually $\mathbb{Z} \times \mathbb{F}_n$, so we just need to show that $n \leq 1$. Each cylinder stabiliser is undistorted (because G is hyperbolic relative to its cylinder stabilisers) and unconstricted (meaning one of its asymptotic cones has no global cut point; each cylinder stabiliser is non-virtually cyclic with an infinite order element in its centre, hence unconstricted [29, p965]); so we may apply [29, Theorem 1.7] to conclude that each cylinder stabiliser is contained in a neighbourhood of a conjugate of some $P \in \mathcal{P}$ (note that [29, Theorem 1.7] has a typo, $G' \to G$ should be a quasi-isometric embedding rather than a quasi-isometry). The observation then follows because there is no quasi-isometric embedding $\mathbb{Z} \times \mathbb{F}_n \to P$ if $n \geq 2$ (for example because $\mathbb{Z} \times \mathbb{F}_n$ has exponential growth and P has polynomial growth).

7.3 Leighton's theorem for graphs with coloured fins

Leighton's Theorem for graphs with fins was proven by Woodhouse [101, Theorem 0.1]; in this section we build on this result by adding colours and orientations to the fins and arranging for the common finite cover to satisfy a symmetry property. The orientations of the fins are particularly important. In Sections 7.4 and 7.5 we will construct graphs of spaces by taking graphs with fins and gluing the ends of certain fins together by homeomorphisms. The homotopy type of such a graph of spaces will not only depend on which fins you glue together, but on the orientations of the fins that get matched up by the gluing.

7.3.1 Definitions

Definition 7.3.1. (Graph with coloured fins)
Let X be a graph, which we now consider to be a 1-dimensional cube complex. Let Δ be a collection of combinatorial immersions $\gamma : S \to X$, where each S is a circle or a bi-infinite line subdivided into $\ell(S)$ edges ($\ell(S) = \infty$ if S is a bi-infinite line). A *graph with fins* $\mathbf{X}$ is a non-positively curved square complex obtained by taking the mapping cylinder of

$$\cup_\Delta \gamma : \bigsqcup_\Delta S \to X.$$

A graph with fins $\mathbf{X}$ is finite if it is a finite cube complex. The subset

$$\bigsqcup_\Delta S \times \{1\} \subseteq \mathbf{X}$$

is the *boundary* of the graph with fins. Each component of the boundary, $S \times \{1\}$, is called a *fin* - for ease of notation we will always write S instead of $S \times \{1\}$. The collection of fins is denoted $\partial\mathbf{X}$. Meanwhile, the subsets $S \times \{0\}$ lie in X. The natural retraction $r : \mathbf{X} \to X$ restricted to the boundary allows us to recover the collection Δ - i.e. for each fin $S \in \partial\mathbf{X}$ the restriction $r : S \to X$ is exactly the original immersion $\gamma : S \to X$.

A fin $S \in \partial\mathbf{X}$ is a 1-manifold, so can be given an orientation o. The pair (S, o) is an *oriented fin*, and will often be written as $\mathbb{S}$. If $\mathbb{S} = (S, \mathsf{o})$ then we write $\bar{\mathbb{S}} = (S, \bar{\mathsf{o}})$ for the fin with opposite orientation. The *length* of $\mathbb{S}$ is $\ell(\mathbb{S}) := \ell(S)$. The collection of oriented fins is denoted $\partial_{\mathsf{o}}\mathbf{X}$. If we have a colouring $\lambda : \partial_{\mathsf{o}}\mathbf{X} \to \mathcal{C}$ (where $\mathcal{C}$ is just a set), then we say that $\mathbf{X}$ is a *graph with coloured fins*.

Definition 7.3.2. (Coverings and automorphisms of graphs with coloured fins)
A *covering of graphs with fins* $\Phi : \widehat{\mathbf{X}} \to \mathbf{X}$ is a covering of square complexes that restricts to a graph covering $\widehat{X} \to X$. We require $\mathbf{X}$ to be connected but $\widehat{\mathbf{X}}$ doesn't need to be.

The restriction of Φ to a fin $\hat{S} \in \partial\widehat{\mathbf{X}}$ is a covering $\hat{S} \to S$ of a fin $S \in \partial\mathbf{X}$. If $\hat{\mathbb{S}} = (\hat{S}, \hat{\mathsf{o}})$ and $\mathbb{S} = (S, \mathsf{o})$ are orientations respected by the covering, then we say that $\hat{\mathbb{S}} \to \mathbb{S}$ is a *covering of oriented fins* (we will usually just say that $\hat{\mathbb{S}} \to \mathbb{S}$ is a covering). Thus we get a map $\Phi : \partial_{\mathsf{o}}\widehat{\mathbf{X}} \to \partial_{\mathsf{o}}\mathbf{X}$ where each $\hat{\mathbb{S}} \to \Phi(\hat{\mathbb{S}})$ is a covering. We call $\Phi : \widehat{\mathbf{X}} \to \mathbf{X}$ a *covering of graphs with coloured fins* if the induced map $\Phi : \partial_{\mathsf{o}}\widehat{\mathbf{X}} \to \partial_{\mathsf{o}}\mathbf{X}$ preserves colours (both $\mathbf{X}$ and $\widehat{\mathbf{X}}$ must use the same set of colours $\mathcal{C}$).

A covering $\widehat{\mathbf{X}} \to \mathbf{X}$ is an *isomorphism* if it is an isomorphism of square complexes. An isomorphism $\mathbf{X} \to \mathbf{X}$ is an *automorphism*. Let $\mathrm{Aut}(\mathbf{X})$ denote the group of automorphisms of $\mathbf{X}$. We note that any automorphism of $\mathbf{X}$ also induces an automorphism on $\partial_{\mathsf{o}}\mathbf{X}$. A covering $\widetilde{\mathbf{X}} \to \mathbf{X}$ is a *universal covering* if $\widetilde{X}$ is a tree, or equivalently if $\widetilde{\mathbf{X}} \to \mathbf{X}$ is a universal covering of square complexes. In this case, the deck transformations of $\widetilde{\mathbf{X}} \to \mathbf{X}$ induce a subgroup of $\mathrm{Aut}(\widetilde{\mathbf{X}})$.

Example 7.3.3. Let X be the bouquet of two circles - the graph given by a single vertex and two edges. We fix a generating set $\pi_1 X = \langle x, y\rangle$ so that the generators x and y correspond to the two edges. Let $\mathbf{X}$ be the graph with fins determined by the geodesic paths given by the set $\{x, y, xy\}$. In this example the oriented fins can be written out as $\partial_{\mathsf{o}}\mathbf{X} = \{x, x^{-1}, y, y^{-1}, xy, y^{-1}x^{-1}\}$. See Figure 7.1 for an illustration of $\mathbf{X}$.

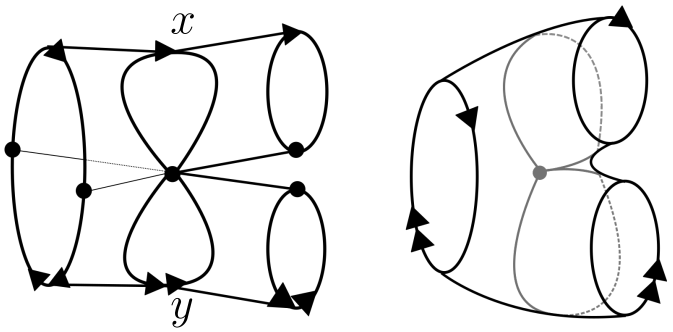

Figure 7.1: A graph with fins - drawn again on the right to emphasize in this case it is homeomorphic to a surface with boundary.

Remark 7.3.4. A graph with coloured fins $\mathbf{X}$ and a graph covering $\widehat{X} \to X$ uniquely determine a coloured fin structure $\widehat{\mathbf{X}}$ on $\widehat{X}$ and a covering $\widehat{\mathbf{X}} \to \mathbf{X}$.

Definition 7.3.5. If $\Phi_i : \widehat{\mathbf{X}} \to \mathbf{X}_i$ are coverings for $i = 1, 2$, and $\mathbb{S}_i \in \partial_\circ \mathbb{X}_i$ are oriented fins, then we write

$$\partial_\circ \widehat{\mathbf{X}}(\mathbb{S}_1, \mathbb{S}_2) := \Phi_1^{-1}(\mathbb{S}_1) \cap \Phi_2^{-1}(\mathbb{S}_2)$$

for the collection of oriented fins in $\widehat{\mathbf{X}}$ that cover both $\mathbb{S}_1$ and $\mathbb{S}_2$.

Definition 7.3.6. (Density)
For $\mathbf{X}$ a finite graph with coloured fins and $c \in \mathcal{C}$ a colour, define the *density* ρ_c by

$$\rho_c := \sum_{\lambda(\mathbb{S})=c} \ell(\mathbb{S})/|X|, \tag{7.3.1}$$

where $|X|$ is the number of vertices in X. Note that densities ρ_c are preserved by finite coverings, and are therefore invariants of the commensurability class of $\mathbf{X}$.

7.3.2 The theorem

Theorem 7.3.7. *(Leighton's Theorem for graphs with coloured fins)*
Let $\mathbf{X}_1$ and $\mathbf{X}_2$ be graphs with coloured fins that have a common universal cover $\widetilde{\mathbf{X}}$. Denote the covering maps by $\Psi_i : \widetilde{\mathbf{X}} \to \mathbf{X}_i$ and let $\Gamma_1, \Gamma_2 < \mathrm{Aut}(\widetilde{\mathbf{X}})$ be the corresponding deck transformation groups. Suppose that $\mathrm{Aut}(\widetilde{\mathbf{X}})$ acts transitively on the oriented fins of each colour in $\widetilde{\mathbf{X}}$. Then $\mathbf{X}_1$ and $\mathbf{X}_2$ have a common finite cover $\widehat{\mathbf{X}}$ such that

$$\sum_{\hat{\mathbb{S}} \in \partial_\circ \widehat{\mathbf{X}}(\mathbb{S}_1, \mathbb{S}_2)} \ell(\hat{\mathbb{S}}) = \left(\frac{|\widehat{X}|}{\rho_c |X_1||X_2|} \right) \ell(\mathbb{S}_1)\ell(\mathbb{S}_2), \tag{7.3.2}$$

for any $\mathbb{S}_i \in \partial_\circ \mathbf{X}_i$ of the same colour c.

The rest of this section is devoted to proving this theorem, so fix $\mathbf{X}_1$, $\mathbf{X}_2$, $\widetilde{\mathbf{X}}$, and $\Gamma_1, \Gamma_2 < \mathrm{Aut}(\widetilde{\mathbf{X}})$ as above. For brevity we will write $H = \mathrm{Aut}(\widetilde{\mathbf{X}})$ for the rest of this section. We will assume that the graphs X_i are simplicial, and that H doesn't invert edges in $\widetilde{X}$. We can achieve these properties by subdividing the edges of the graphs X_i. Note that equation (7.3.2) is preserved by subdividing edges in underlying graphs; indeed $|\widehat{X}|/|X_2|$ is the degree of $\widehat{X} \to X_2$, so is unchanged, and the quantities $\ell(\hat{\mathbb{S}})$, $\ell(\mathbb{S}_1)$, $\ell(\mathbb{S}_2)$ and $\rho_c|X_1|$ all increase by a factor of two.

Definition 7.3.8. (Polyhedra and faces)
Let $\mathbf{X}$ be a graph with coloured fins. A hyperplane in $\mathbf{X}$ is *vertical* if it is dual to an edge in X - let $\mathscr{H}$ denote the set of vertical hyperplanes. Let $\dot{\mathbf{X}}$ denote the square complex obtained from $\mathbf{X}$ by subdividing along the vertical hyperplanes. A *polyhedron* (P, ϕ) is a square complex P equipped with a cubical embedding $\phi : P \to \dot{\mathbf{X}}$ such that $\phi(P)$ is the cubical neighbourhood in $\dot{\mathbf{X}}$ of a vertex $x \in X$. Alternatively, we can think of $\phi(P)$ as the closure of the component of $\mathbf{X} - \mathscr{H}$ containing x. We call x the *centre* of $\phi(P)$. A *face* (F, φ) is a finite tree F equipped with a cubical embedding $\varphi : F \to \dot{\mathbf{X}}$ such that $\varphi(F)$ is a vertical hyperplane in $\mathbf{X}$ (which is a subcomplex in $\dot{\mathbf{X}}$). We say that (F, φ) is a *face* of (P, ϕ) if there is a commutative diagram of cubical embeddings

$$\begin{array}{ccc} F & \longrightarrow & P \\ & \searrow^{\varphi} & \downarrow{\scriptstyle \phi} \\ & & \dot{\mathbf{X}}. \end{array} \tag{7.3.3}$$

Fixing an orientation on each edge in X, we have a notion of being on the *left* or *right* of a vertical hyperplane in $\mathbf{X}$. We say that (P, ϕ) is on the *left* (resp. *right*) of a face (F, φ) if $\phi(P)$ is on the left (resp. right) of $\varphi(F)$ (there is no ambiguity as we have assumed X is simplicial). Up to isomorphism there is a unique polyhedron on the left and right of each face.

If (P, ϕ) and (P', ϕ') are polyhedra on the left and right of a face (F, φ), then the polyhedra can be glued together along the embeddings of F to make a new complex $P \cup P'$ that maps into $\mathbf{X}$ via $\phi \cup \phi'$.

Definition 7.3.9. (Polyhedral pairs and face pairs)
A *polyhedral pair* is a triple $\mathbf{P} = (P, \phi_1, \phi_2)$ where each pair (P, ϕ_i) is a polyhedron for $\mathbf{X}_i$. We say that $\mathbf{P}$ is *H-admissible* if there is a commutative diagram as follows, which we will refer to as the *admissibility diagram*,

$$\begin{array}{ccccc} \widetilde{\mathbf{X}} & & \xrightarrow{\quad h \quad} & & \widetilde{\mathbf{X}} \\ & \nwarrow^{\tilde{\phi}_1} & & \nearrow^{\tilde{\phi}_2} & \\ \Psi_1 \downarrow & & P & & \downarrow \Psi_2 \\ & \swarrow_{\phi_1} & & \searrow_{\phi_2} & \\ \mathbf{X}_1 & & & & \mathbf{X}_2, \end{array} \tag{7.3.4}$$

where $\tilde{\phi}_i$ are lifts of the maps ϕ_i and $h \in H$. Note that the lifts $\tilde{\phi}_i$ are unique up to post-composition by $g_i \in \Gamma_i$, so if $\mathbf{P}$ is admissible then the diagram (7.3.4) can be constructed for any lifts $\tilde{\phi}_i$.

Similarly, a *face pair* is a triple $\mathbf{F} = (F, \varphi_1, \varphi_2)$ where each pair (F, φ_i) is a face for $\mathbf{X}_i$. We say that $\mathbf{F}$ is H-*admissible* if there is a commutative diagram

$$\begin{array}{ccccc} \widetilde{\mathbf{X}} & & \xrightarrow{\quad h \quad} & & \widetilde{\mathbf{X}} \\ & \nwarrow^{\tilde{\varphi}_1} & & \nearrow^{\tilde{\varphi}_2} & \\ \Psi_1 \big\downarrow & & F & & \big\downarrow \Psi_2 \\ & \swarrow_{\varphi_1} & & \searrow_{\varphi_2} & \\ \mathbf{X}_1 & & & & \mathbf{X}_2, \end{array} \tag{7.3.5}$$

where $\tilde{\varphi}_i$ are lifts of the maps φ_i and $h \in H$. We say that a polyhedral pair $\mathbf{P} = (P, \phi_1, \phi_2)$ is on the *left* (resp. *right*) of a face pair $\mathbf{F} = (F, \varphi_1, \varphi_2)$ if (P, ϕ_i) is on the left (resp. right) of (F, φ_i) for $i = 1, 2$ with respect to the same embedding $F \to P$. Note that it is impossible for (P, ϕ_1) to be on the left of (F, φ_1) and for (P, ϕ_2) to be on the right of (F, φ_2) with respect to the same embedding $F \to P$ because H has no edge-inversions. Let $\overleftarrow{\mathbf{F}}$ (resp. $\overrightarrow{\mathbf{F}}$) denote the set of admissible polyhedral pairs on the left (resp. right) of $\mathbf{F}$. Note that $\overleftarrow{\mathbf{F}}$ and $\overrightarrow{\mathbf{F}}$ are finite since $\mathbf{X}_1$ and $\mathbf{X}_2$ are. If $\mathbf{P} \in \overleftarrow{\mathbf{F}}$ and $\mathbf{P}' \in \overrightarrow{\mathbf{F}}$ then we can glue together P and P' along the embeddings of F to obtain a complex $P \cup P'$ with maps $\phi_1 \cup \phi_1'$ and $\phi_2 \cup \phi_2'$ to $\mathbf{X}_1$ and $\mathbf{X}_2$.

Given a polyhedron (P, ϕ_1) for $\mathbf{X}_1$, we will be interested in counting the ways it can be extended to an admissible polyhedral pair $\mathbf{P} = (P, \phi_1, \phi_2)$, subject to forcing $\mathbf{P} \in \overleftarrow{\mathbf{F}}$ for a fixed face pair $\mathbf{F}$.

Lemma 7.3.10. *Let (P, ϕ_1) be a polyhedron for $\mathbf{X}_1$ and choose a lift $\tilde{\phi}_1 : P \to \widetilde{\mathbf{X}}$ with image $\widetilde{P}$. Let (P, ϕ_1) be on the left (resp. right) of a face (F, φ_1), and let $\tilde{\phi}_1(F) = \widetilde{F}$ (viewing F as a subset of P). Suppose (F, φ_1) extends to an admissible face pair $\mathbf{F} = (F, \varphi_1, \varphi_2)$. Then the choices ϕ_2 such that $(P, \phi_1, \phi_2) \in \overleftarrow{\mathbf{F}}$ (resp. $\overrightarrow{\mathbf{F}}$) are in one to one correspondence with the quotient $H_{(\widetilde{F})}/H_{(\widetilde{P})}$ - where $H_{(\widetilde{F})}$ and $H_{(\widetilde{P})}$ are the pointwise stabilisers of $\widetilde{F}$ and $\widetilde{P}$ respectively.*

Proof. Assume (P, ϕ_1) is on the left of (F, φ_1). Now $(F, \varphi_1, \varphi_2)$ fits into a commutative diagram (7.3.5) for some $h \in H$ and lifts $\tilde{\varphi}_i$, and we can choose $\varphi_1 = \phi_1|_F$. Then any ϕ_2 such that $(P, \phi_1, \phi_2) \in \overleftarrow{\mathbf{F}}$ will fit into an admissibility diagram

$$\begin{array}{ccccc} \widetilde{\mathbf{X}} & & \xrightarrow{\quad h' \quad} & & \widetilde{\mathbf{X}} \\ & \nwarrow^{\tilde{\phi}_1} & & \nearrow^{\tilde{\phi}_2} & \\ \Psi_1 \big\downarrow & & P & & \big\downarrow \Psi_2 \\ & \swarrow_{\phi_1} & & \searrow_{\phi_2} & \\ \mathbf{X}_1 & & & & \mathbf{X}_2, \end{array} \tag{7.3.6}$$

for some $h' \in H$. As $\phi_2|_F = \varphi_2$, we know that $\tilde{\phi}_2|_F$ and $\tilde{\varphi}_2$ differ by an element of Γ_2; so by composing h' with an element of Γ_2, we may assume that $\tilde{\phi}_2|_F = \tilde{\varphi}_2$. Then $h'|_{\widetilde{F}} = h|_{\widetilde{F}}$, hence $h' \in hH_{(\widetilde{F})}$. Conversely, any $h' \in hH_{(\widetilde{F})}$ defines a polyhedral pair $(P, \phi_1, \phi_2) \in \overleftarrow{\mathbf{F}}$ via (7.3.6). Finally, the map ϕ_2 only depends on the coset $h'H_{(\widetilde{P})}$, again because of (7.3.6). This establishes the desired bijection between the choices ϕ_2 and the quotient $H_{(\widetilde{F})}/H_{(\widetilde{P})}$. □

We want to take appropriate numbers of copies of each admissible polyhedral pair so that we can glue them all together along face pairs (as we described at the end of Definition 7.3.9) to form a common finite cover of $\mathbf{X}_1$ and $\mathbf{X}_2$. We formalise this with the following definition.

Definition 7.3.11. (Gluing Equations)
Let $\mathcal{P}$ be the (finite) collection of all admissible polyhedral pairs, and let $\omega : \mathcal{P} \to \mathbb{Z}_{>0}$ denote a weight function on $\mathcal{P}$. For each admissible face pair $\mathbf{F}$ we have the following *Gluing Equation*:

$$\sum_{\mathbf{P} \in \overleftarrow{\mathbf{F}}} \omega(\mathbf{P}) = \sum_{\mathbf{P} \in \overrightarrow{\mathbf{F}}} \omega(\mathbf{P}). \tag{7.3.7}$$

Given a solution, we can take $\omega(\mathbf{P})$ copies of each $\mathbf{P}$, and glue them together along faces according to (arbitrary) bijections

$$\{(\mathbf{P}, i) \mid \mathbf{P} \in \overleftarrow{\mathbf{F}},\ 1 \leq i \leq \omega(\mathbf{P})\} \leftrightarrow \{(\mathbf{P}, i) \mid \mathbf{P} \in \overrightarrow{\mathbf{F}},\ 1 \leq i \leq \omega(\mathbf{P})\}, \tag{7.3.8}$$

and this will give us a common finite cover of $\mathbf{X}_1$ and $\mathbf{X}_2$. For the moment we won't worry about colouring the oriented fins in this finite cover.

To solve the gluing equations, we will consider the Haar measure μ for the group H. As H contains a uniform lattice - for example Γ_1 - H is unimodular and μ is both left and right H-invariant. Note that μ is positive on every open set and finite on every compact set, both of which apply to the stabilisers $H_{(\widetilde{P})}$ and $H_{(\widetilde{F})}$. There are finitely many H-orbits of images of polyhedra $\widetilde{P}$ in $\widetilde{\mathbf{X}}$, and so, by H-invariance of μ, there are finitely many values $\mu(H_{(\widetilde{P})})$; furthermore, the stabilisers $H_{(\widetilde{P})}$ are all commensurable in H, so by rescaling we can assume that all $\mu(H_{(\widetilde{P})})$ are positive integers. For each $\mathbf{P} = (P, \phi_1, \phi_2) \in \mathcal{P}$, choose a lift $\tilde{\phi}_1 : P \to \widetilde{\mathbf{X}}$ with image $\widetilde{P}$, and set

$$\omega(\mathbf{P}) = \mu(H_{(\widetilde{P})}). \tag{7.3.9}$$

Observe that $\omega(\mathbf{P})$ is independent of the choice of lift $\tilde{\phi}_1$ because of the left and right H-invariance of μ.

Lemma 7.3.12. *The Haar measure weight function (7.3.9) solves the Gluing Equations (7.3.7).*

Proof. Given an admissible face pair $\mathbf{F} = (F, \varphi_1, \varphi_2)$, let (P, ϕ_1) be the polyhedron on the left of (F, φ_1). All $\mathbf{P} \in \overleftarrow{\mathbf{F}}$ can be obtained by choosing a map ϕ_2 such that $(P, \phi_1, \phi_2) \in \overleftarrow{\mathbf{F}}$, and by Lemma 7.3.10 there are $H_{(\widetilde{F})}/H_{(\widetilde{P})}$ such choices, where $\widetilde{F} \subset \widetilde{P} \subset \widetilde{\mathbf{X}}$ comes from a lift of (P, ϕ_1). Substituting (7.3.9) into the left hand side of (7.3.7) then gives us

$$\begin{aligned}\sum_{\mathbf{P}\in\overleftarrow{\mathbf{F}}} \omega(\mathbf{P}) &= \sum_{\mathbf{P}\in\overleftarrow{\mathbf{F}}} \mu(H_{(\widetilde{P})}) \\ &= |H_{(\widetilde{F})} : H_{(\widetilde{P})}|\mu(H_{(\widetilde{P})}) \\ &= \mu(H_{(\widetilde{F})}).\end{aligned}$$

Observe that this only depends on $\mathbf{F}$, and so by a symmetric argument we get the same value if we substitute (7.3.9) into the right hand side of (7.3.7). □

We have now constructed a common finite cover of $\mathbf{X}_1$ and $\mathbf{X}_2$, call it $\widehat{\mathbf{X}}$ say. Denote the covering maps by $\Phi_i : \widehat{\mathbf{X}} \to \mathbf{X}_i$. We colour the oriented fins of $\widehat{\mathbf{X}}$ by pulling back the colours from $\mathbf{X}_1$ and $\mathbf{X}_2$, which is well-defined by the following lemma. This makes the Φ_i coverings of graphs with coloured fins.

Lemma 7.3.13. *If we pull back the colours on $\partial_\circ \mathbf{X}_1$ to $\partial_\circ \widehat{\mathbf{X}}$, then the covering $\Phi_2 : \widehat{\mathbf{X}} \to \mathbf{X}_2$ preserves colours.*

Proof. Take an oriented fin $\hat{\mathbb{S}} = (\hat{S}, \hat{\mathsf{o}}) \in \partial_\circ \widehat{\mathbf{X}}$ with $\Phi_i(\hat{\mathbb{S}}) = \mathbb{S}_i$. We must show that $\mathbb{S}_1$ and $\mathbb{S}_2$ have the same colour. Let $\tilde{\mathbb{S}} \in \partial_\circ \widetilde{\mathbf{X}}$ be a lift of $\mathbb{S}_1$ to $\widetilde{\mathbf{X}}$. If $\hat{\mathbb{S}}$ crosses a polyhedral pair $\mathbf{P}$ in $\widehat{\mathbf{X}}$, then we have an admissibility diagram (7.3.4) with some $h \in H$. Restricting to the fins, we get the following commutative diagram.

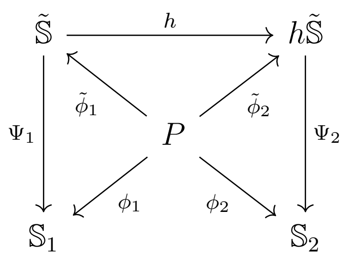

(7.3.10)

As $h : \widetilde{\mathbf{X}} \to \widetilde{\mathbf{X}}$ preserves colours of oriented fins, we see that $\mathbb{S}_1$, $\tilde{\mathbb{S}}$, $h\tilde{\mathbb{S}}$ and $\mathbb{S}_2$ all have the same colour, as required. □

We now turn to proving equation (7.3.2) from Theorem 7.3.7. We will need the following definition.

Definition 7.3.14. (Arcs and oriented arcs)
Let (P, ϕ) be a polyhedron for a graph with fins $\mathbf{X}$. An *arc* in (P, ϕ) is a component of $\phi^{-1}(\partial \mathbf{X})$, and its image in $\mathbf{X}$ is also called an *arc*. If $A \subset \mathbf{X}$ is an arc, then it is contained in a unique fin $S \in \partial \mathbf{X}$, it is homeomorphic to an interval, and it contains exactly one vertex of S. An arc A can be given an orientation o as a 1-manifold, making

it an *oriented arc* $\mathbb{A} = (A, \mathsf{o})$. If $\mathbb{A}$ is an oriented arc contained in an oriented fin $\mathbb{S}$ such that the orientations agree, then we say that $\mathbb{A}$ is an *oriented subarc of* $\mathbb{S}$.

If $\Phi : \widehat{\mathbf{X}} \to \mathbf{X}$ is a covering of graphs with fins, then it maps each arc in $\widehat{\mathbf{X}}$ homeomorphically to an arc in $\mathbf{X}$. Moreover, if $\hat{A}$ is an arc in $\widehat{\mathbf{X}}$ with orientation $\hat{\mathbb{A}}$, and $A = \Phi(\hat{A})$ is its image in $\mathbf{X}$, then we get an induced orientation $\mathbb{A}$ on A such that $\hat{\mathbb{A}} \to \mathbb{A}$ is an orientation preserving homeomorphism, and we write $\mathbb{A} = \Phi(\hat{\mathbb{A}})$. Note that if $\hat{\mathbb{A}}$ is an oriented subarc of $\hat{\mathbb{S}}$, then $\Phi(\hat{\mathbb{A}})$ is an oriented subarc of $\Phi(\hat{\mathbb{S}})$.

Similarly, an *arc* in a polyhedral pair $\mathbf{P} = (P, \phi_1, \phi_2)$ is a component of $\phi_1^{-1}(\partial \mathbf{X}_1) = \phi_2^{-1}(\partial \mathbf{X}_2)$. As $\widehat{\mathbf{X}}$ is built from the polyhedral pairs in $\mathcal{P}$, we see that the arcs in $\widehat{\mathbf{X}}$ correspond exactly to the arcs in elements of $\mathcal{P}$. We let $\partial_{\mathsf{o}} \mathbf{P}$ denote the set of oriented arcs in $\mathbf{P}$.

Fix oriented fins $\mathbb{S}_1 \in \partial_{\mathsf{o}} \mathbf{X}_1$ and $\mathbb{S}_2 \in \partial_{\mathsf{o}} \mathbf{X}_2$ of the same colour c. Our goal is to sum lengths of fins in $\partial_{\mathsf{o}} \widehat{\mathbf{X}}(\mathbb{S}_1, \mathbb{S}_2)$; we will do this by counting admissible polyhedral pairs whose images contain certain oriented subarcs of $\mathbb{S}_1$ and $\mathbb{S}_2$.

Definition 7.3.15. Let $\mathbb{A}_i$ be oriented subarcs of the oriented fins $\mathbb{S}_i$, and define

$$P(\mathbb{A}_1, \mathbb{A}_2) := \big\{ \mathbf{P} = (P, \phi_1, \phi_2) \in \mathcal{P} \mid \exists \mathbb{A} \in \partial_{\mathsf{o}} \mathbf{P}, \;\; \phi_i(\mathbb{A}) = \mathbb{A}_i \big\}, \tag{7.3.11}$$

the collection of polyhedral pairs containing oriented arcs that map to the $\mathbb{A}_i$. Then define

$$A(\mathbb{A}_1, \mathbb{S}_2) := \{ \mathbb{A}_2 \mid \mathbb{A}_2 \text{ is an oriented subarc of } \mathbb{S}_2 \text{ with } P(\mathbb{A}_1, \mathbb{A}_2) \neq \emptyset \}. \tag{7.3.12}$$

In order to enumerate the elements of $P(\mathbb{A}_1, \mathbb{A}_2)$, we fix a polyhedron (P, ϕ_1) for $\mathbf{X}_1$ that contains $\mathbb{A}_1$ in its image $\phi_1(P)$ (this polyhedron will be unique up to isomorphism). Let $\mathbb{A} = (A, \mathsf{o})$ be the (unique) oriented arc in P with $\phi_1(\mathbb{A}) = \mathbb{A}_1$. Enumerating the elements of $P(\mathbb{A}_1, \mathbb{A}_2)$ is now equivalent to enumerating maps $\phi_2 : P \to \mathbf{X}_2$ such that $(P, \phi_1, \phi_2) \in P(\mathbb{A}_1, \mathbb{A}_2)$. For the following two lemmas it will also be helpful to fix a lift $\tilde{\phi}_1 : P \to \widetilde{\mathbf{X}}$ of ϕ_1, and letting $\widetilde{P} := \tilde{\phi}_1(P)$, $\tilde{A} := \tilde{\phi}_1(A)$ and $\tilde{\mathbb{A}} := \tilde{\phi}_1(\mathbb{A})$.

Lemma 7.3.16. *If $P(\mathbb{A}_1, \mathbb{A}_2)$ is non-empty, then it is in bijection with $H_{(\tilde{A})}/H_{(\widetilde{P})}$ - where $H_{(\tilde{A})}$ is the pointwise stabiliser of $\tilde{A}$.*

Proof. The proof is very similar to Lemma 7.3.10. Suppose $(P, \phi_1, \phi_2) \in P(\mathbb{A}_1, \mathbb{A}_2)$. We then get an admissibility diagram (7.3.4) for some $h \in H$. Any other $(P, \phi_1, \phi_2') \in P(\mathbb{A}_1, \mathbb{A}_2)$ also gives an admissibility diagram, but with h replaced by h' and $\tilde{\phi}_2$ replaced by $\tilde{\phi}_2'$. Since $\phi_2(\mathbb{A}) = \phi_2'(\mathbb{A}) = \mathbb{A}_2$, we know that $h(\tilde{\mathbb{A}})$ and $h'(\tilde{\mathbb{A}})$ differ by an element of Γ_2, so by composing h' with an element of Γ_2 we may assume that $h(\tilde{\mathbb{A}}) = h'(\tilde{\mathbb{A}})$. This implies that $h' \in hH_{(\tilde{A})}$, and conversely any $h' \in hH_{(\tilde{A})}$ defines a map ϕ_2' with $(P, \phi_1, \phi_2') \in P(\mathbb{A}_1, \mathbb{A}_2)$. Finally, the map ϕ_2' only depends on the coset $h'G_{(\widetilde{P})}$, and so we obtain a bijection between the choices ϕ_2' and the quotient $H_{(\tilde{A})}/H_{(\widetilde{P})}$. □

Lemma 7.3.17. *The ratio* $|A(\mathbb{A}_1, \mathbb{S}_2)|/\ell(\mathbb{S}_2)$ *only depends on the colour* c.

Proof. Let $\tilde{\mathbb{S}}_2 \in \partial\widetilde{\mathbf{X}}$ be an oriented fin that covers $\mathbb{S}_2$, and let $\tilde{\mathbb{A}}_2$ be an oriented subarc of $\tilde{\mathbb{S}}_2$.

We claim that $\Psi_2(\tilde{\mathbb{A}}_2) \in A(\mathbb{A}_1, \mathbb{S}_2)$ if and only if $\tilde{\mathbb{A}}_2$ is a H-translate of $\tilde{\mathbb{A}}$. Indeed, if $\mathbb{A}_2 := \Psi_2(\tilde{\mathbb{A}}_2) \in A(\mathbb{A}_1, \mathbb{S}_2)$, then there is an admissible (P, ϕ_1, ϕ_2) with $\phi_2(\mathbb{A}) = \mathbb{A}_2$, and it has an associated admissibility diagram (7.3.4) for some $h \in H$. Then $h(\tilde{\mathbb{A}})$ and $\tilde{\mathbb{A}}_2$ will both be a lifts of $\mathbb{A}_2$, so by composing h with an element of Γ_2 we may assume that $h(\tilde{\mathbb{A}}) = \tilde{\mathbb{A}}_2$. Conversely, if $h(\tilde{\mathbb{A}}) = \tilde{\mathbb{A}}_2$ for some $h \in H$, then we get an admissibility diagram (7.3.4), and $\Psi_2(\tilde{\mathbb{A}}_2) = \phi_2(\mathbb{A}) \in A(\mathbb{A}_1, \mathbb{S}_2)$.

Thus the proportion of oriented subarcs of $\mathbb{S}_2$ that lie in $A(\mathbb{A}_1, \mathbb{S}_2)$ is equal to the proportion of oriented subarcs of $\tilde{S}_2$ that lie in the H-orbit of $\tilde{\mathbb{S}}$. In turn this is equal to the smallest positive translation length of elements of $H_{\tilde{\mathbb{S}}}$. It follows that it is independent of the choice of $\tilde{\mathbb{S}}_2$ that covers $\mathbb{S}_2$, in fact it only depends on the H-orbit of $\tilde{\mathbb{S}}_2$, thus only depends on the colour of the oriented fin. □

For $\hat{\mathbb{S}} \in \partial_\circ\widehat{\mathbf{X}}(\mathbb{S}_1, \mathbb{S}_2)$, the proportion of oriented subarcs $\mathbb{A}$ of $\hat{\mathbb{S}}$ that descend to $\mathbb{A}_1$ is $1/\ell(\mathbb{S}_1)$, and any such $\mathbb{A}$ must lie in some $(P, \phi_1, \phi_2) \in \mathcal{P}$ that forms a piece of $\widehat{\mathbf{X}}$, with $\phi_1(\mathbb{A}) = \mathbb{A}_1$ and $\phi_2(\mathbb{A}) = \mathbb{A}_2$ some oriented subarc of $\mathbb{S}_2$. There are $\omega(P, \phi_1, \phi_2)$ copies of (P, ϕ_1, ϕ_2) in $\widehat{\mathbf{X}}$, thus we can make the following computation:

$$\begin{aligned}
\sum_{\hat{S} \in \partial_\circ \widehat{\mathbf{X}}(\mathbb{S}_1, \mathbb{S}_2)} \ell(\hat{\mathbb{S}}) &= \ell(\mathbb{S}_1) \sum_{\substack{\mathbb{A}_2 \subset \mathbb{S}_2 \\ (P,\phi_1,\phi_2) \in P(\mathbb{A}_1, \mathbb{A}_2)}} \omega(P, \phi_1, \phi_2) \\
&= \ell(\mathbb{S}_1) \sum_{\mathbb{A}_2 \subset \mathbb{S}_2} |P(\mathbb{A}_1, \mathbb{A}_2)| \mu(H_{(\tilde{P})}) \\
&= \ell(\mathbb{S}_1) |A(\mathbb{A}_1, \mathbb{S}_2)| |H_{(\tilde{A})} : H_{(\tilde{P})}| \mu(H_{(\tilde{P})}) \qquad \text{by Lemma 7.3.16} \\
&= \ell(\mathbb{S}_1) \ell(\mathbb{S}_2) \left(\frac{|A(\mathbb{A}_1, \mathbb{S}_2)|}{\ell(\mathbb{S}_2)} \right) \mu(H_{(\tilde{A})}) \\
&= K_{\tilde{\mathbb{A}}}\, \ell(\mathbb{S}_1) \ell(\mathbb{S}_2), \qquad (7.3.13)
\end{aligned}$$

where $K_{\tilde{\mathbb{A}}}$ only depends on the H-orbit of $\tilde{\mathbb{A}}$ by Lemma 7.3.17. The key point now is that the oriented fins in $\widetilde{\mathbf{X}}$ of colour c are all in the same H-orbit. As a result, if we had chosen different $\mathbb{S}_1$ and $\mathbb{S}_2$ of colour c, then by choosing a suitable oriented subarc $\mathbb{A}_1$ in $\mathbb{S}_1$ we can arrange for the oriented subarc $\tilde{\mathbb{A}}$ to be in the same H-orbit as before. Thus $K_{\tilde{\mathbb{A}}}$ in fact only depends on c, and we will write it as K_c.

To complete the proof of equation (7.3.2) it remains to compute a formula for K_c. We do this by summing (7.3.13) over all $\mathbb{S}_1 \in \partial_\circ \mathbf{X}_1$ and $\mathbb{S}_2 \in \partial_\circ \mathbf{X}_2$ of colour c:

$$\sum_{\substack{\lambda_1(\mathbb{S}_1)=\lambda_2(\mathbb{S}_2)=c\\ \hat{\mathbb{S}}\in\partial_\circ\widehat{\mathbf{X}}(\mathbb{S}_1,\mathbb{S}_2)}} \ell(\hat{\mathbb{S}}) = K_c\left(\sum_{\lambda_1(\mathbb{S}_1)=c}\ell(\mathbb{S}_1)\right)\left(\sum_{\lambda_2(\mathbb{S}_2)=c}\ell(\mathbb{S}_2)\right)$$

We can then substitute in the definition of the density ρ_c (Definition 7.3.1), which will be the same for $\mathbf{X}_1$, $\mathbf{X}_2$ and $\widehat{\mathbf{X}}$ since they are all commensurable, to obtain:

$$\rho_c|\widehat{X}| = K_c\rho_c|X_1|\rho_c|X_2|$$

This gives the required formula for K_c:

$$K_c = \frac{|\widehat{X}|}{\rho_c|X_1||X_2|}$$

7.4 Building graphs of spaces

In this section and the next we prove Theorem 1.3.3. In this section we build graphs of spaces for the two groups that share a number of properties, while in the next section we construct a common finite cover.

From now on let G be the group from Theorem 1.3.3, i.e. a group in $\mathscr{C}^\bullet$ that is hyperbolic relative to virtually abelian subgroups, and let $\psi : G \to G'$ be a fixed quasi-isometry to another finitely generated group. Being finitely presented is a quasi-isometry invariant, and by Theorem 2.7.8 so is having a JSJ decompositions over two-ended subgroups containing only virtually free vertex groups and no QH vertex groups. Being hyperbolic relative to virtually abelian subgroups is also preserved by quasi-isometry because of [29, Theorem 1.6] and the quasi-isometric rigidity of abelian groups [14, 70]. Hence G' is also a group in $\mathscr{C}^\bullet$ that is hyperbolic relative to virtually abelian subgroups.

Let T and T' be JSJ trees for G and G' and let T_c and T'_c be their associated trees of cylinders. By Theorem 2.7.9, there is an isomorphism $\hat{\psi} : T_c \to T'_c$ such that ψ restricts to quasi-isometries $G_v \to G'_{\hat{\psi}(v)}$ and $G_e \to G'_{\hat{\psi}(e)}$ for $v \in VT_c$ and $e \in ET_c$. We know from Remark 7.2.10 that the cylinder stabilisers for G are virtually abelian, so if T_c is just a single vertex then G is either virtually free or virtually abelian, and such groups are already known to be quasi-isometrically rigid. So we may assume that T_c is not a single vertex.

Notation 7.4.1. Recall that the group $\mathscr{G}$ of Hausdorff equivalence classes of quasi-isometries $G \to G$ acts on the tree of cylinders T_c by Corollary 2.7.10, and similarly $\mathscr{G}'$ acts on T'_c. From now on it will be convenient to identify $\mathscr{G}$ with $\mathscr{G}'$ via the isomorphism $[f] \mapsto [\psi f \psi^{-1}]$, and to identify T_c with T'_c via the isomorphism $\hat{\psi}$ given by Theorem 2.7.9.

G acts on itself by left multiplication, so we have a homomorphism $G \to \mathscr{G}$. G has edge stabilisers in T_c that intersect trivially, so G acts on T_c faithfully and the homomorphism $G \to \mathscr{G}$ is injective, thus we can think of G as a subgroup of $\mathscr{G}$. Similarly, G' is a subgroup of $\mathscr{G}'$, so also a subgroup of $\mathscr{G}$ by the above identification (i.e. when we say G' is a subgroup of $\mathscr{G}$ we mean the subgroup $\psi^{-1}G'\psi$). This means that G' acts on the tree T_c; since this action is conjugate to the action of G' on T'_c, we know that T_c is a tree of cylinders for G'.

Notation 7.4.2. Recall that the tree of cylinders has a partition of the vertex set $VT_c = V_0T_c \sqcup V_1T_c$. V_0T_c corresponds to vertices of T that lie in more than one cylinder; in our case all vertex groups of the JSJ decomposition are rigid, so we will refer to vertices $u \in V_0T_c$ as *rigid vertices.* Note that the stabilisers of rigid vertices will all be virtually non-abelian free. The vertices V_1T_c correspond to cylinders in T; in Section 2.7.2 we denoted cylinders by $Y \subset T$ or $Y \in V_1T_c$, but as we no longer need to work with them as subtrees of T we will instead denote them by $v \in V_1T_c$, and refer to them as *cylindrical vertices.* By Lemma 7.2.9 and Remark 7.2.10 the stabilisers of cylindrical vertices will be virtually $\mathbb{Z}$ or $\mathbb{Z}^2$.

7.4.1 Cylindrical factors and orientations

By Theorem 1.3.2, we know that G and G' are both balanced.

Lemma 7.4.3. *By passing to a finite index subgroup of G, we can assume that G is torsion-free, and that for each $v \in V_1T_c$ there is subgroup $\mathbb{Z} \cong \mathbb{Z}_v < G_v$, such that either $G_v = \mathbb{Z}_v$ or $G_v = \mathbb{Z}_v \times \mathbb{Z}$, and $G_e = \mathbb{Z}_v$ for any $e \in \mathrm{lk}(v)$.*

Proof. By Lemma 7.2.9 we can pass to a finite index torsion-free subgroup of G such that, for each $v \in V_1T_c$ representing a cylinder $Y \subset T$, we get a product splitting $G_v = \mathbb{Z}_v$ or $G_v = \mathbb{Z}_v \times \mathbb{Z}$, where the $\mathbb{Z}_v$ factor pointwise fixes Y, and in the second case the second factor acts freely cocompactly on Y. For any $e \in \mathrm{lk}(v)$, G_e will act elliptically on T, hence it will be a subgroup of $\mathbb{Z}_v$, but we also know that $\mathbb{Z}_v$ fixes $e \in EY$, so $G_e = \mathbb{Z}_v$. □

The subgroup $\mathbb{Z}_v$ from Lemma 7.4.3 is called the *cylindrical factor* of G_v. Similarly, we apply Lemma 7.4.3 to G' to make it torsion-free and get cylindrical factors $\mathbb{Z}'_v$ for the vertex stabilisers G'_v, where $v \in V_1T_c$.

Definition 7.4.4. (Oriented cylinders and edge groups)
An *orientation* $\mathcal{O}$ on a cylindrical factor $\mathbb{Z}_v$ is a choice of one of its two ends. The pair $(v, \mathcal{O})$ is called an *oriented cylinder.* If $e \in \mathrm{lk}(v)$ then $G_e = \mathbb{Z}_v$, so $\mathcal{O}$ is also a choice of end of the subgroup G_e, and we call the pair $(e, \mathcal{O})$ an *oriented edge group.* Let $\bar{\mathcal{O}}$ denote the opposite end of $\mathbb{Z}_v$.

The group $\mathscr{G}$ of Hausdorff equivalence classes of quasi-isometries $G \to G$ acts on the tree of cylinders T_c by Corollary 2.7.10. For $[f] \in \mathscr{G}$ that acts by $\hat{f} \in \mathrm{Aut}(T_c)$ we also have $f(G_e) \sim G_{\hat{f}(e)}$ for $e \in ET_c$ by Theorem 2.7.9(3), hence $\mathscr{G}$ acts on the set of oriented edge groups, and also on the set of oriented cylinders. To avoid over-counting we will always consider the edges $E_1T_c \subset ET_c$ with terminus in a cylindrical vertex. We denote $\mathscr{G}$-orbits using square brackets, so for $e \in E_1T_c$ and $v \in V_1T_c$ we have:

$$[e, \mathcal{O}] := \mathscr{G} \cdot (e, \mathcal{O})$$
$$[v, \mathcal{O}] := \mathscr{G} \cdot (v, \mathcal{O})$$

Let $\mathcal{C}$ denote the set of $\mathscr{G}$-orbits $[e, \mathcal{O}]$ for $e \in E_1T_c$, which we will think of as a colouring of the oriented edge groups. If $[f] \cdot (e_1, \mathcal{O}_1) = (e_2, \mathcal{O}_2)$ then

$$[f] \cdot (e_1, \bar{\mathcal{O}}_1) = (e_2, \bar{\mathcal{O}}_2), \tag{7.4.1}$$

so for $c = [e, \mathcal{O}]$ we can define $\bar{c} := [e, \bar{\mathcal{O}}]$.

Notation 7.4.5. We will also use square brackets to denote $\mathscr{G}$-orbits in E_1T_c and V_0T_c: $[e] := \mathscr{G} \cdot e$ for $e \in E_1T_c$ and $[u] := \mathscr{G} \cdot u$ for $u \in V_0T_c$. Note that $[e, \mathcal{O}]$ determines $[e]$.

7.4.2 Cylinder numbers and ratios

Definition 7.4.6. (Cylinder numbers and ratios)
Let $v \in V_1T_c$ be a cylindrical vertex. For an orientation $\mathcal{O}$ on $\mathbb{Z}_v$ define

$$\mathrm{lk}(v, \mathcal{O}) := \{(e, \mathcal{O}) \mid e \in \mathrm{lk}(v)\}.$$

For a colour $c \in \mathcal{C}$, define

$$\mathrm{lk}(v, \mathcal{O}, c) := \{(e, \mathcal{O}) \mid e \in \mathrm{lk}(v), [e, \mathcal{O}] = c\}.$$

The *cylinder number* $t_c(v, \mathcal{O})$ is the number of G_v-orbits of oriented edge groups in $\mathrm{lk}(v, \mathcal{O}, c)$. The *cylinder ratio* of $(v, \mathcal{O})$ is

$$t(v, \mathcal{O}) = [c \mapsto t_c(v, \mathcal{O})],$$

where the brackets indicate that we only define the function up to rescaling. Similarly, let $t'_c(v, \mathcal{O})$ be the number of G'_v-orbits of oriented edge groups in $\mathrm{lk}(v, \mathcal{O}, c)$, and $t'(v, \mathcal{O}) = [c \mapsto t'_c(v, \mathcal{O})]$.

The motivation for cylinder numbers is the following, which will be made more precise later on. In a graph of spaces for the splitting of G induced by T_c, we can take the vertex space for G_v to be a circle or a torus, and the edge spaces for incident edge stabilisers to be circles; the orientation $\mathcal{O}$ induces orientations on the edge spaces as 1-manifolds, so $\mathcal{C}$ gives a colouring of oriented edge spaces. The cylinder number $t_c(v, \mathcal{O})$ is just the number of edge spaces incident at the vertex space for G_v, with orientation induced by $\mathcal{O}$, of colour c.

We note that v is a finite valence vertex in the case $G_v \cong \mathbb{Z}$, and by Lemma 7.4.3 the stabiliser G_v fixes $\mathrm{lk}(v)$, so the cylinder number $t_c(v, \mathcal{O})$ is just the size of $\mathrm{lk}(v, \mathcal{O}, c)$. So in this case not only is $t(v, \mathcal{O}) = t'(v, \mathcal{O})$, but $t_c(v, \mathcal{O}) = t'_c(v, \mathcal{O})$. In the case that $G_v \cong \mathbb{Z}^2$, although $t_c(v, \mathcal{O})$ is not in general equal to $t'_c(v, \mathcal{O})$, we will show that the cylinder ratios are in fact equal and that we can pass to finite index subgroups of G and G' such that the cylinder numbers are equal.

Remark 7.4.7. Consider a cylindrical vertex stabiliser $G_v = \mathbb{Z}_v \times \mathbb{Z}$ and fix an orientation $\mathcal{O}$ on $\mathbb{Z}_v$. For an edge $e \in \mathrm{lk}(v)$, any $g \in G_v$ maps $G_e = \mathbb{Z}_v \times \{0\}$ to some coset $\mathbb{Z}_v \times \{n\}$ by a translation of $\mathbb{Z}_v \times \mathbb{Z}$, so the induced quasi-isometry $G_e \to G_{ge} = G_e$ is at bounded distance from the identity, and the orientation $\mathcal{O}$ is preserved

$$\begin{aligned} g \cdot (e, \mathcal{O}) &= (ge, \mathcal{O}) \\ g \cdot (v, \mathcal{O}) &= (v, \mathcal{O}). \end{aligned} \tag{7.4.2}$$

Equations (7.4.2) also hold for $g \in G'_v$ by considering its action on $G'_v = \mathbb{Z}'_v \times \mathbb{Z}$. So G_v- and G'_v-orbits in $\mathrm{lk}(v, \mathcal{O})$ just correspond to G_v- and G'_v-orbits in $\mathrm{lk}(v)$.

Lemma 7.4.8. *(1) $t(v, \mathcal{O}) = t'(v, \mathcal{O})$ for all $\mathbb{Z}^2$ cylinders $v \in V_1 T_c$ with orientation $\mathcal{O}$.*

(2) $t(v, \mathcal{O})$ only depends on the $\mathscr{G}$-orbit $[v, \mathcal{O}]$.

Proof. We will only give a proof of (1), but (2) can be proven by the same argument applied to a quasi-isometry $[f] \in \mathscr{G}$ instead of the quasi-isometry $\psi : G \to G'$.

Let $v \in V_1 T_c$ with $G_v \cong \mathbb{Z}^2$. Since ψ induces a quasi-isometry from G_v to G'_v it follows $G'_v \cong \mathbb{Z}^2$ also. If $e \in \mathrm{lk}(v)$, then G_e is equal to the cylindrical factor $\mathbb{Z}_v < G_v$. Let $e_1, \ldots, e_N \in \mathrm{lk}(v)$ be G_v-orbit representatives of the edges. There exists $g \in G_v$ that corresponds to the generator of the second factor in the decomposition $G_v \cong \mathbb{Z}_v \times \mathbb{Z}$. It follows that $G_v = \bigcup_k g^k \mathbb{Z}_v$ and $\mathrm{lk}(v) = \bigcup_{i,k} g^k e_i$. Then we have a function $n : \mathrm{lk}(v) \to \mathbb{Z}$ given by $n(g^k e_i) = k$. Similarly for G'_v we let $e'_1, \ldots, e'_{N'} \in \mathrm{lk}(v)$ be G'_v-orbit representatives, $g' \in G'_v$ be an element generating the second factor in the decomposition $G'_v = \mathbb{Z}'_v \times \mathbb{Z}$, to obtain $n' : \mathrm{lk}(v) \to \mathbb{Z}$ given by $n((g')^k e'_i) = k$.

We now observe that there exists some $L > 0$ such that $G(e_i) \sim_L \mathbb{Z}_v \times \{0\}$ for all i, where $G(e_i)$ is the coset corresponding to e_i from Notation 2.7.6. By G-invariance of the metric, for all $e \in \mathrm{lk}(v)$ we have

$$G(e) \sim_L \mathbb{Z}_v \times \{n(e)\}. \tag{7.4.3}$$

We now consider the following five pseudo-metrics on $\mathrm{lk}(v)$.

$$d(e_1, e_2) = \begin{cases} |n(e_1) - n(e_2)| & \text{(a)} \\ d_H(G(e_1), G(e_2)) & \text{(b)} \\ d_H(\psi(G(e_1)), \psi(G(e_2))) & \text{(c)} \\ d_H(G'(e_1), G'(e_2)) & \text{(d)} \\ |n'(e_1) - n'(e_2)| & \text{(e)} \end{cases} \tag{7.4.4}$$

Claim: Pseudo metrics (a)–(e) are all equivalent up to quasi-isometry.

Proof: With respect to the standard generators of $\mathbb{Z}^2$, the Hausdorff distance $d_H(\mathbb{Z} \times \{n_1\}, \mathbb{Z} \times \{n_2\})$ is simply $|n_1 - n_2|$. Cylinder stabilisers are quasi-isometrically embedded in G (peripheral subgroups are always quasiconvex [29, Lemma 4.15]), and quasi-isometries coarsely preserve Hausdorff distance between subsets, so (7.4.3) implies that metrics (a) and (b) are equivalent. Similarly, (d) and (e) are equivalent. (b) and (c) are equivalent because ψ is a quasi-isometry, and finally (c) and (d) are equivalent precisely because of Theorem 2.7.9(2). ■

The maps $n, n' : \mathrm{lk}(v) \to \mathbb{Z}$ are both surjective, so the equivalence of metrics (a) and (e) gives us a quasi-isometry $\nu : \mathbb{Z} \to \mathbb{Z}$ such that

$$n' \approx \nu \circ n. \tag{7.4.5}$$

After perturbing ν by bounded distance, we can assume that it is monotonic; indeed, if $\lim_{i \to \pm\infty} \nu(i) = \pm\infty$ then $i \mapsto \max_{j \leq i} \nu(j)$ is increasing and at bounded distance from ν.

For each $a \in \mathbb{Z}$, $n^{-1}(a)$ has one edge from each G_v-orbit, and so by Remark 7.4.7 we have that $|\{(e, \mathcal{O}) \in \mathrm{lk}(v, \mathcal{O}, c) \mid n(e) = a\}|$ is equal to $t_c(v, \mathcal{O})$ from Definition 7.4.6. For $c_1, c_2 \in \mathcal{C}$ (with $\mathrm{lk}(v, \mathcal{O}, c_i) \neq \emptyset$) we have

$$\frac{t_{c_1}(v, \mathcal{O})}{t_{c_2}(v, \mathcal{O})} = \lim_{b-a \to \infty} \frac{|\{(e, \mathcal{O}) \in \mathrm{lk}(v, \mathcal{O}, c_1) \mid n(e) \in [a, b]\}|}{|\{(e, \mathcal{O}) \in \mathrm{lk}(v, \mathcal{O}, c_2) \mid n(e) \in [a, b]\}|}, \tag{7.4.6}$$

and a similar equation holds for t'_{c_i}. By (7.4.5), and the fact that ν is monotonic, we see that

$$\lim_{b-a \to \infty} \frac{|\{(e, \mathcal{O}) \in \mathrm{lk}(v, \mathcal{O}, c) \mid n(e) \in [a, b]\}|}{|\{(e, \mathcal{O}) \in \mathrm{lk}(v, \mathcal{O}, c) \mid n'(e) \in [\nu(a), \nu(b)]\}|} = 1 \tag{7.4.7}$$

for any $c \in \mathcal{C}$. Combining (7.4.6) with (7.4.7) we deduce that $t_{c_1}(v, \mathcal{O})/t_{c_2}(v, \mathcal{O}) = t'_{c_1}(v, \mathcal{O})/t'_{c_2}(v, \mathcal{O})$. Hence $t(v, \mathcal{O}) = t'(v, \mathcal{O})$ as required. □

The next task is to prove the following lemma.

Lemma 7.4.9. *There exist finite index subgroups* $\hat{G} \triangleleft G$ *and* $\hat{G}' \triangleleft G'$ *and integers* $N_c[v, \mathcal{O}]$ *such that* $N_c[v, \mathcal{O}] = N_{\bar{c}}[v, \bar{\mathcal{O}}]$*, and for each oriented cylinder* $(v, \mathcal{O})$ *the cylinder numbers of* $\hat{G}$ *and* $\hat{G}'$ *both equal the numbers* $N_c[v, \mathcal{O}]$*:*

$$\hat{t}_c(v, \mathcal{O}) = \hat{t}'_c(v, \mathcal{O}) = N_c[v, \mathcal{O}].$$

We will need the following remark.

Remark 7.4.10. For an oriented cylinder $(v, \mathcal{O})$ and a colour $c \in \mathcal{C}$, we have a bijection

$$\begin{aligned} \mathrm{lk}(v, \mathcal{O}, c) &\to \mathrm{lk}(v, \bar{\mathcal{O}}, \bar{c}) \\ (e, \mathcal{O}) &\mapsto (e, \bar{\mathcal{O}}). \end{aligned} \tag{7.4.8}$$

Remark 7.4.7 tells us that G_v acts on both $\mathrm{lk}(v, \mathcal{O}, c)$ and $\mathrm{lk}(v, \bar{\mathcal{O}}, \bar{c})$, so by (7.4.1) we know that the map (7.4.8) is G_v-equivariant. It follows that

$$t_c(v, \mathcal{O}) = t_{\bar{c}}(v, \bar{\mathcal{O}}). \tag{7.4.9}$$

As discussed earlier, if a cylindrical vertex stabiliser is cyclic, then the cylinder numbers are already equal, and will be stable under passing to finite index subgroups, so Lemma 7.4.9 is all about modifying the $\mathbb{Z}^2$ cylinders. If $v \in V_1 T_c$ is a cylindrical vertex with stabiliser $G_v \cong \mathbb{Z}_v \times \mathbb{Z}$, then let $\pi_v : G_v \to \mathbb{Z}$ denote the projection onto the second factor in the product decomposition. Any finite index subgroup $\hat{G} \triangleleft G$ will have stabiliser $\hat{G}_v$ finite index in G_v. Suppose for a moment we have $\hat{G} \triangleleft G$ finite index such that $\pi_v(\hat{G}_v) = N\mathbb{Z}$. Then each G_v-orbit of edges in $\mathrm{lk}(v)$ would split into N many $\hat{G}_v$-orbits, so by Remark 7.4.7 the cylinder numbers for $\hat{G}$ and G would be related by

$$\hat{t}_c(v, \mathcal{O}) = N t_c(v, \mathcal{O}). \tag{7.4.10}$$

It follows readily from Definition 7.4.6 that $t_c(v, \mathcal{O})$ only depends on c and the G-orbit of $(v, \mathcal{O})$, hence there are only finitely many cylinder numbers. Furthermore, Lemma 7.4.8 says that the cylinder ratio $t(v, \mathcal{O})$ only depends on the $\mathscr{G}$-orbit $[v, \mathcal{O}]$. Therefore, for each $[v, \mathcal{O}]$ we can pick numbers $N_c[v, \mathcal{O}]$ that are in the ratio $t(v, \mathcal{O})$, and that are common multiples of all cylinder numbers. By (7.4.9) we can assume that $N_c[v, \mathcal{O}] = N_{\bar{c}}[v, \bar{\mathcal{O}}]$. Again by (7.4.9), we deduce that

$$N_v := \frac{N_c[v, \mathcal{O}]}{t_c(v, \mathcal{O})} = \frac{N_{\bar{c}}[v, \bar{\mathcal{O}}]}{t_{\bar{c}}(v, \bar{\mathcal{O}})} \tag{7.4.11}$$

only depends on v, in fact it only depends on the G-orbit of v because $t_c(v, \mathcal{O})$ only depends on c and the G-orbit of $(v, \mathcal{O})$. We define integers N'_v similarly, such that N'_v only depends on the G'-orbit of v.

Proof of Lemma 7.4.9. By (7.4.10) and (7.4.11), it suffices to construct finite-index normal subgroups $\hat{G} \lhd G$ and $\hat{G}' \lhd G'$ such that

$$\begin{aligned} \pi_v(\hat{G}_v) &= N_v\mathbb{Z} \\ \pi_v(\hat{G}'_v) &= N'_v\mathbb{Z}, \end{aligned} \tag{7.4.12}$$

for each v. Note that it is enough to have (7.4.12) hold for a set of G-orbit representatives of $v \in V_1T_c$ with $G_v \cong \mathbb{Z}^2$ (and similarly for G') because $\hat{G}$ is normal in G (and because each map $\pi_v : G_v \to \mathbb{Z}$ is determined by the edge stabilisers incident to G_v, up to a factor of ± 1, so they are preserved by conjugation in G). During this construction, we are allowed to multiply all the $N_v, N'_v, N_c[v, \mathcal{O}]$ by some fixed constant, as this preserves equation (7.4.11), or in other words we are allowed to assume that they are multiples of any given finite set of integers.

Let $\{P_1, \ldots, P_m\}$ be the G-stabilisers of a set of G-orbit representatives $\{v_1, \ldots, v_m\}$ of cylindrical vertices in T_c. Note that G is virtually special by Theorem 1.3.2 and it is hyperbolic relative to $\{P_1, \ldots, P_m\}$ by Proposition 2.8.8. And by Theorem 2.8.9 we may disregard any cyclic peripheral subgroups and assume that each P_i is isomorphic to $\mathbb{Z}^2$.

By Theorem 5.4.6, G commands $(P_1, ..., P_m)$, so there exist finite-index subgroups $\dot{P}_i < P_i$ such that, for any further finite-index subgroups $\hat{P}_i < \dot{P}_i$, there exists a finite-index normal subgroup $\hat{G} \lhd G$ with $\hat{G}_{v_i} = P_i \cap \hat{G} = \hat{P}_i$. Suppose $\pi_{v_i}(\dot{P}_i) = N_i\mathbb{Z}$. As discussed above, we may assume that N_{v_i} is a multiple of N_i for each i; so we can arrange $\pi_{v_i}(\hat{P}_i) = N_{v_i}\mathbb{Z}$ by setting $\hat{P}_i = \dot{P}_i \cap \pi_{v_i}^{-1}(N_{v_i}\mathbb{Z})$. Hence the $\hat{G} \lhd G$ provided above satisfies (7.4.12). By exactly the same argument we can construct a finite-index $\hat{G}' \lhd G'$ that satisfies (7.4.12). □

By Lemma 7.4.9, we can assume going forward that for each oriented cylinder $(v, \mathcal{O})$ the cylinder numbers of G and G' both equal the numbers $N_c[v, \mathcal{O}]$:

$$t_c(v, \mathcal{O}) = t'_c(v, \mathcal{O}) = N_c[v, \mathcal{O}]. \tag{7.4.13}$$

7.4.3 A tree of trees with fins

For each rigid vertex $u \in V_0T_c$, recall that the incident edge groups for the stabiliser G_u induce a line pattern $\mathcal{L}_u$ (Definition 7.1.3). By Lemma 7.1.12 this line pattern will be rigid, and so by Theorem 7.1.9 there is a quasi-isometry to a tree with line pattern that is a rigid model space:

$$\alpha_u : (G_u, \mathcal{L}_u) \to (Y_u, \mathcal{L}_u).$$

Recall Lemma 7.1.13, which says that any $[f] \in \mathscr{G}$ induces a $\approx$-class of quasi-isometries

$$[f]_u : (G_u, \mathcal{L}_u) \to (G_{\hat{f}(u)}, \mathcal{L}_{\hat{f}(u)}) \tag{7.4.14}$$

that respect line patterns. So for each $\mathscr{G}$-orbit of vertices u, the free groups with line patterns $(G_u, \mathcal{L}_u)$ are all quasi-isometric, and hence we may choose the rigid model spaces $(Y_u, \mathcal{L}_u)$ to be isometric. We can encode the line pattern $\mathcal{L}_u$ in the tree Y_u as a set of fins to obtain a quasi-isometry to a graph with fins (see Definition 7.3.1):

$$\beta_u : (G_u, \mathcal{L}_u) \to (\mathbf{Y}_u, \partial\mathbf{Y}_u).$$

Since the underlying graph Y_u is a tree, we will refer to $(\mathbf{Y}_u, \partial\mathbf{Y}_u)$ as a *tree with fins.* Note that $(\mathbf{Y}_u, \partial\mathbf{Y}_u)$ also serves as a rigid model space, and its group of isometries is precisely its automorphism group in the sense of Definition 7.3.2. Moreover, the isometry type of the rigid model space $(Y_u, \mathcal{L}_u)$ only depends on the $\mathscr{G}$-orbit $[u]$, and so the isomorphism type of the tree with fins $(\mathbf{Y}_u, \partial\mathbf{Y}_u)$ also just depends on $[u]$. Combining these two facts with (7.4.14) yields the following lemma.

Lemma 7.4.11. *For each $[f] \in \mathscr{G}$ and $u \in V_0T_c$, there is a unique isomorphism*

$$[\boldsymbol{f}]_u : (\mathbf{Y}_u, \partial\mathbf{Y}_u) \to (\mathbf{Y}_{\hat{f}(u)}, \partial\mathbf{Y}_{\hat{f}(u)})$$

such that $[\boldsymbol{f}]_u \approx \beta_{\hat{f}(u)} \circ [f]_u \circ \beta_u^{-1}$.

Proof. We know that u and $\hat{f}(u)$ are in the same $\mathscr{G}$-orbit, so $(\mathbf{Y}_u, \partial\mathbf{Y}_u)$ and $(\mathbf{Y}_{\hat{f}(u)}, \partial\mathbf{Y}_{\hat{f}(u)})$ are isomorphic. As these trees with fins are rigid model spaces, the line-pattern-preserving quasi-isometry (or more precisely $\approx$-class of quasi-isometries) $\beta_{\hat{f}(u)}\circ[f]_u\circ\beta_u^{-1} : (\mathbf{Y}_u, \partial\mathbf{Y}_u) \to (\mathbf{Y}_{\hat{f}(u)}, \partial\mathbf{Y}_{\hat{f}(u)})$ between them is finite Hausdorff distance from a unique isometry $[\boldsymbol{f}]_u$. □

This gives us the data to define an action of $\mathscr{G}$ on the disjoint union of the $\mathbf{Y}_u$ - which we think of as a "tree of trees with fins".

Lemma 7.4.12. *The maps $[\boldsymbol{f}]_u$ define an action of $\mathscr{G}$ on the graph with fins $\mathbf{Y} := \sqcup_{u\in V_0T_c}\mathbf{Y}_u$.*

Proof. We know from Corollary 2.7.10 that $\mathscr{G}$ acts on the rigid vertices V_0T_c. It follows from Lemma 7.4.11 that we have a well-defined map $\mathscr{G} \to \mathrm{Aut}(\mathbf{Y})$, we must show that this is a homomorphism. It is clear that id_G maps to the identity, so it remains to show that this map respects composition. Let $[f_1], [f_2] \in \mathscr{G}$ with $\hat{f}_1(u_1) = u_2$ and $\hat{f}_2(u_2) = u_3$. We know that $[f_2 \circ f_1]_{u_1} \approx [f_2]_{u_2} \circ [f_1]_{u_1}$ as these maps come from restricting the quasi-isometries to the vertex groups, so it follows from Lemma 7.4.11 that $[\boldsymbol{f_2} \circ \boldsymbol{f_1}]_{u_1} \approx [\boldsymbol{f_2}]_{u_2} \circ [\boldsymbol{f_1}]_{u_1}$, but this second $\approx$ must be an equality since both sides are isometries between rigid model spaces. □

For $u \in V_0T_c$ we know that the lines in $\mathcal{L}_u$ correspond to the incident edge stabilisers G_e, and this is a one-to-one correspondence because no two incident edge stabilisers are commensurable in G_u (as they come from different cylinders). In turn these lines correspond via β_u to the fins of $\mathbf{Y}_u$. Let $S_e \in \partial\mathbf{Y}_u$ be the fin corresponding to G_e (with $\iota(e) = u$). A choice of end $\mathcal{O}$ of the edge stabiliser G_e defines an oriented edge group $(e, \mathcal{O})$, which will correspond via β_u to a choice of end of the fin S_e, or equivalently a choice of orientation $\mathbb{S}_e = (S_e, \circ)$ of the fin as a 1-manifold, as in Definition 7.3.1. It follows from the way we defined the $\mathscr{G}$-action on $\mathbf{Y}$ that the action of $\mathscr{G}$ on oriented edge groups is conjugate to the action of $\mathscr{G}$ on oriented fins in $\mathbf{Y}$. We defined $\mathcal{C}$ to be the set of $\mathscr{G}$-orbits of oriented edge groups, so this also corresponds to $\mathscr{G}$-orbits of oriented fins, which we will think of as a colouring of the oriented fins $\lambda : \partial_\circ\mathbf{Y} \to \mathcal{C}$. This makes $\mathbf{Y}$ and each of the $\mathbf{Y}_u$ into graphs with coloured fins, and the $\mathscr{G}$-action preserves colours.

Remark 7.4.13. For a rigid vertex $u \in V_0T_c$ the action of $\mathscr{G}$ on Y restricts to an action of G_u on $\mathbf{Y}_u$, where $g \in G_u$ acts by $[\![g]\!]_u$. It follows from the definition of $[\![g]\!]_u$ that this action of G_u on $\mathbf{Y}_u$ is the β_u-conjugacy action in the sense of Definition 7.1.6. Similarly, the quasi-isometry $\psi : G \to G'$ restricts to a quasi-isometry $\psi : G_u \to G'_u$, and the action of G'_u on $\mathbf{Y}_u$ is the $\beta_u\psi^{-1}$-conjugacy action. It then follows from Lemma 7.1.10 that the actions of G_u and G'_u on $\mathbf{Y}$ are free and cocompact, that $\beta_u : G_u \to \mathbf{Y}_u$ is Hausdorff equivalent to any orbit map of G_u, and that $\beta_u\psi^{-1} : G'_u \to \mathbf{Y}_u$ is Hausdorff equivalent to any orbit map of G'_u.

Remark 7.4.14. The space $\mathbf{Y}$ is disconnected, so it is tempting to try and connect it up into some simply connected metric space that's quasi-isometric to G, and that admits an action of $\mathscr{G}$ quasi-conjugate to its action on G. The natural way to try and do this is to take a copy of $\mathbb{R}$ or $\mathbb{R}^2$ for each cylindrical vertex $v \in V_1T$, and glue them to the appropriate fins in $\mathbf{Y}$ according to how the edge stabilisers G_e embed in the vertex stabilisers G_v. There is no real advantage in doing this however, because the action of $\mathscr{G}$ would not be isometric - it would induce isometries between the vertex spaces as it does for $\mathbf{Y}$, but in general it would act via "shearing" maps between the edge spaces. Such a construction was used however in [9].

7.4.4 Stretch ratios

Definition 7.4.15. (Stretch ratio)
Let $v \in V_1T_c$ be a cylindrical vertex and let $g \in \mathbb{Z}_v$ be a non-trivial element. Let $e \in E_1T_c$ be an edge with $\tau(e) = v$ and $\iota(e) = u \in V_0T_c$, then the automorphism

$$[\![g]\!]_u : (\mathbf{Y}_u, \partial\mathbf{Y}_u) \to (\mathbf{Y}_u, \partial\mathbf{Y}_u),$$

acts by translation on the fin S_e.

Let r_e be the translation length of $[\![g]\!]_u$, which is equal to the distance that it translates along the fin S_e. Note that $r_e \neq 0$.

The *stretch ratio* of $v \in V_1T_c$ is the function $\mathrm{lk}(v) \to \mathbb{Q}$ given by $e \mapsto r_e$ determined by $g \in \mathbb{Z}_v$, but as we are only interested in the ratio between the r_e terms we will only consider this function to be defined up to scaling. We will denote this equivalence class of functions by

$$\mathrm{Str}(v) = [e \mapsto r_e].$$

The stretch ratio does not depend on the choice of non-trivial element $g \in \mathbb{Z}_v$, since each element is a power of a fixed generator, and the translation lengths scale linearly by the power.

We can also define the stretch ratio for $v \in V_1T_c$ with respect to G' by using elements $g' \in \mathbb{Z}'_v$. It is a result of Cashen–Martin [23] that the stretch ratios defined using G and G' will coincide. Their result is more general, but the two consequences that will be relevant to us are the following. We also include a proof because the result is slightly simpler in our setting, and it highlights how we make use of rigid model spaces.

Lemma 7.4.16. *(Cashen–Martin [23, Proposition 5.14])*

(1) The stretch ratio $\mathrm{Str}(v)$ *is the same for* G *and* G'*.*

(2) There exist integers $r_{[e]}$ *for* $e \in E_1T_c$*, where* $[e]$ *denotes the* $\mathscr{G}$*-orbit of* e*, such that* $Str(v) = [e \mapsto r_{[e]}]$ *for all* $v \in V_1T_c$*.*

We recall that a *coarse M-similitude* is a function $f : X \to Y$ between metric spaces such that

$$Md_X(x_1, x_2) - \epsilon \leq d_Y(f(x_1), f(x_2)) < Md_X(x_1, x_2) + \epsilon$$

for all $x_1, x_2 \in X$ and some fixed $\epsilon \geq 0$. We make four remarks about such an f:

- Any map Hausdorff equivalent to f will also be a coarse M-similitude.
- If $f : X \to Y$ is a quasi-isometry, then its quasi-inverse f^{-1} will be a coarse M^{-1}-similitude.
- If $g : Y \to Z$ is a coarse N-similitude, then $g \circ f$ is a coarse MN-similitude.
- An equivariant quasi-isometry of $\mathbb{Z}$ into a tree will be a coarse M-similitude, where M is determined by the translation length along the axis.

Proof of Lemma 7.4.16. Let $e \in E_1T_c$ be an edge with $\iota(e) = u \in V_0T_c$ and $\tau(e) = v \in V_1T_c$. We know from Remark 7.4.13 that $\beta_u : G_u \to \mathbf{Y}_u$ is Hausdorff equivalent to any orbit map of G_u. We also know that $\mathbb{Z}_v = G_e < G_u$ acts on $\mathbf{Y}_u$ by translating along the fin S_e, say the translation length of a generator is r_e, so it follows that (up to Hausdorff equivalence) β_u restricts to a coarse r_e-similitude $\mathbb{Z}_v \to S_e$.

Similarly, we know from Remark 7.4.13 that $\beta_u \circ \psi^{-1} : G'_u \to \mathbf{Y}_u$ is Hausdorff equivalent to any orbit map of G'_u, and that $\mathbb{Z}'_v = G'_e < G'_u$ acts on $\mathbf{Y}_u$ by translating along the fin S_e, with translation length of a generator being r'_e say. So it follows that (up to

Hausdorff equivalence) $\beta_u \psi^{-1}$ restricts to a coarse r'_e-similitude $\mathbb{Z}'_v \to S_e$. Composing the two coarse similitudes tells us that $\psi : \mathbb{Z}_v \to \mathbb{Z}'_v$ is a coarse r_e/r'_e-similitude. But the map $\psi : \mathbb{Z}_v \to \mathbb{Z}'_v$ doesn't depend on the choice of e, so the ratio r_e/r'_e is the same for all edges $e \in \text{lk}(v)$ - thus proving (1).

For (2), we must show that the action of $\mathscr{G}$ preserves stretch ratio. More precisely, if $[f] \in \mathscr{G}$ and $\text{Str}(v) = [e \mapsto r_e]$, then we must show that

$$\text{Str}(\hat{f}(v)) = [\hat{f}(e) \mapsto r_e \mid e \in \text{lk}(v)]. \tag{7.4.15}$$

Observe that, for $e \in \text{lk}(v)$ with $\iota(e) = u$, we have the following diagram that commutes up to Hausdorff equivalence.

$$\begin{array}{ccccc} \mathbb{Z}_v & \hookrightarrow & G_u & \xrightarrow{\beta_u} & \mathbf{Y}_u \\ \downarrow f & & \downarrow f & & \downarrow [\boldsymbol{f}]_u \\ \mathbb{Z}_{\hat{f}(v)} & \hookrightarrow & G_{\hat{f}(u)} & \xrightarrow{\beta_{\hat{f}(u)}} & \mathbf{Y}_{\hat{f}(v)} \end{array} \tag{7.4.16}$$

We know that $\mathbb{Z}_{\hat{f}(v)}$ acts on $\mathbf{Y}_{\hat{f}(v)}$ by translating along the fin $S_{\hat{f}(e)}$, with the translation length of a generator being $r_{\hat{f}(e)}$ say. And as before $\beta_{\hat{f}(u)}$ restricts to a coarse $r_{\hat{f}(e)}$-similitude $\mathbb{Z}_{\hat{f}(v)} \to S_{\hat{f}(e)}$. But we know that $[\boldsymbol{f}]_u$ restricts to an isometry $S_e \to S_{\hat{f}(e)}$, so composing coarse similitudes implies that $f : \mathbb{Z}_v \to \mathbb{Z}_{\hat{f}(v)}$ is a coarse $r_e/r_{\hat{f}(e)}$-similitude. As before we note that the ratio $r_e/r_{\hat{f}(e)}$ must be the same for all edges $e \in \text{lk}(v)$, which completes the proof of (7.4.15). □

7.4.5 Constructing graphs of spaces from graphs with fins

Consider the graphs of groups decompositions (G, Γ) and (G', Γ') for G and G' given by their respective actions on $T_c \cong T'_c$. The vertices in Γ and Γ' are either *rigid* or *cylindrical* according to their lifts in T_c, and have corresponding vertex partitions $V\Gamma = V_0\Gamma \sqcup V_1\Gamma$ and $V\Gamma' = V_0\Gamma' \sqcup V_1\Gamma'$. As for T_c, we always consider edges with terminus a cylindrical vertex, and we write $E_1\Gamma$ and $E_1\Gamma'$ for the sets of these edges. We also colour edges and rigid vertices according to the $\mathscr{G}$-orbits of their lifts in T_c, so we write $[e] := [\tilde{e}]$ for $e \in E_1\Gamma \sqcup E_1\Gamma'$ with lift $\tilde{e} \in E_1T_c$, and $[u] := [\tilde{u}]$ for $u \in V_0\Gamma \sqcup V_0\Gamma'$ with lift $\tilde{u} \in V_0T_c$ (using Notation 7.4.5).

We now build graphs of spaces $(\mathcal{X}, \Gamma)$ and $(\mathcal{X}', \Gamma')$, for (G, Γ) and (G', Γ') respectively, following the conventions given in Section 2.6.1.

Definition 7.4.17. (Graphs of spaces $(\mathcal{X}, \Gamma)$ and $(\mathcal{X}', \Gamma')$)

For each rigid vertex $u \in V_0\Gamma$, take a lift $\tilde{u} \in V_0T_c$, and consider the action of $G_{\tilde{u}}$ on its corresponding tree with coloured fins $\mathbf{Y}_{\tilde{u}}$ as described in Remark 7.4.13. This action is free and cocompact, so the quotient $\mathbf{X}_u := \mathbf{Y}_{\tilde{u}}/G_{\tilde{u}}$ is a finite graph with coloured fins.

The colouring $\lambda : \partial_\circ \mathbf{Y}_{\tilde{u}} \to \mathcal{C}$ descends to a colouring $\lambda : \partial_\circ \mathbf{X}_u \to \mathcal{C}$. The fundamental group $\pi_1 \mathbf{X}_u$ is identified with the deck transformations $G_{\tilde{u}}$ of the covering $\mathbf{Y}_{\tilde{u}} \to \mathbf{X}_u$, which in turn is identified with the vertex group G_u of (G, Γ). We let $\mathcal{X}_u = \mathbf{X}_u$.

This is independent of the choice of lift $\tilde{u}$, because if $\tilde{u}_1$ and $\tilde{u}_2$ are two lifts of u with $g(\tilde{u}_1) = \tilde{u}_2$ ($g \in G$), then $[\boldsymbol{g}]_{\tilde{u}_1} : \mathbf{Y}_{\tilde{u}_1} \to \mathbf{Y}_{\tilde{u}_2}$ is an isomorphism that is equivariant with respect to the actions of $G_{\tilde{u}_1}$ and $G_{\tilde{u}_2}$ respectively via the conjugation $h \in G_{\tilde{u}_1} \mapsto ghg^{-1} \in G_{\tilde{u}_2}$.

For each cylindrical vertex $v \in V_1\Gamma$ we let $\mathcal{X}_v$ be homeomorphic to a circle S^1 if $G_v \cong \mathbb{Z}$ or a torus $S^1 \times S^1$ if $G_v \cong \mathbb{Z}^2$ and identify $\pi_1 \mathcal{X}_v$ with G_v. We have $G_v \cong G_{\tilde{v}} = \mathbb{Z}_{\tilde{v}} \times \mathbb{Z}$ or $\mathbb{Z}_{\tilde{v}}$ for any lift $\tilde{v} \in V_1 T_c$ of v, and since the cylindrical factor $\mathbb{Z}_{\tilde{v}}$ is preserved by G-conjugation we can define the *cylindrical factor* $\mathbb{Z}_v < G_v$. We then fix a *cylindrical fibre* $S_v \subseteq \mathcal{X}_v$, a subspace homeomorphic to a circle whose embedding gives the embedding of the cylindrical factor. Note that in the case $G_v \cong \mathbb{Z}$ we have $S_v = \mathcal{X}_v$.

Let $e \in E_1\Gamma$ be an edge such that $\iota(e) = u \in V_0\Gamma$ and $\tau(e) = v \in V_1\Gamma$. By construction, the fins in $\mathbf{X}_u$ correspond to $G_{\tilde{u}}$-orbits of fins in $\mathbf{Y}_{\tilde{u}}$, which in turn correspond to $G_{\tilde{u}}$-orbits of edges in $\mathrm{lk}(\tilde{u})$. Hence we get one fin $S_e \in \partial \mathbf{X}_u$ for each edge e with $\iota(e) = u$, and for each lift $\tilde{e}$ with $\iota(\tilde{e}) = \tilde{u}$ the covering $\mathbf{Y}_{\tilde{u}} \to \mathbf{X}_u$ restricts to a covering $S_{\tilde{e}} \to S_e$ of fins. On the level of fundamental groups, the fin S_e corresponds to the G_u-conjugacy class of the image $\zeta_{\bar{e}}(G_e) < G_u$. Having an orientation $\mathbb{S}_e$ of the fin S_e corresponds to choosing an orientation of the fin $S_{\tilde{e}}$, which corresponds to a choice of end $\mathcal{O}$ of $G_{\tilde{e}}$. Then the colour of the oriented fin is $\lambda(\mathbb{S}_e) = [\tilde{e}, \mathcal{O}]$, while the colour of the edge is $[e] = [\tilde{e}]$ - in particular $\lambda(\mathbb{S}_e)$ determines $[e]$. Let $\mathcal{X}_e$ be homeomorphic to the circle and let $\phi_{\bar{e}} : \mathcal{X}_e \to \mathcal{X}_u$ be the homeomorphism onto $S_e \subseteq \mathbf{X}_u$ that induces $\zeta_{\bar{e}}$, and let $\phi_e : \mathcal{X}_e \to \mathcal{X}_v$ be the homeomorphism onto the cylindrical fibre $S_v \subseteq \mathcal{X}_v$ that induces ζ_e. Having determined the vertex spaces $\{\mathcal{X}_v \mid v \in V\Gamma\}$, edge spaces $\{\mathcal{X}_e \mid e \in E\Gamma\}$, and attaching maps $\{\phi_e \mid e \in E\Gamma\}$, we obtain the graph of spaces $(\mathcal{X}, \Gamma)$.

We construct $(\mathcal{X}', \Gamma')$ similarly. So we have a vertex space $\mathcal{X}'_{u'} = \mathbf{X}_{u'} := \mathbf{Y}_{\tilde{u}'}/G_{\tilde{u}'}$ for a rigid vertex $u' \in V_0\Gamma'$ with a lift $\tilde{u}' \in V_0 T_c$, and for e' with $\iota(e') = u'$ we have a fin $S_{e'} \in \partial \mathbf{X}_{u'}$. For a cylindrical vertex $v' \in V_1\Gamma'$ we let $\mathcal{X}'_{v'}$ be a torus containing a cylindrical fibre $S_{v'} \cong S^1$ corresponding to the cylindrical factor $\mathbb{Z}'_{v'} < G_{v'}$. For $e' \in E_1\Gamma'$ an edge with $\iota(e') = u' \in V_0\Gamma'$ and $\tau(e') = v' \in V_1\Gamma'$, we let $\mathcal{X}'_{e'}$ be a circle, and $\phi'_{\bar{e}'} : \mathcal{X}'_{e'} \to \mathcal{X}'_{u'}$, $\phi'_{e'} : \mathcal{X}'_{e'} \to \mathcal{X}'_{v'}$ maps that are homeomorphisms onto $S_{e'}$ and $S_{v'}$ respectively.

Definition 7.4.18. (Orientations)
Let $v \in V_1\Gamma$ be a cylindrical vertex with lift $\tilde{v} \in V_1 T_c$. Because we have identified $\pi_1 S_v$ with $\mathbb{Z}_{\tilde{v}}$, a choice of end $\mathcal{O}$ on $\mathbb{Z}_{\tilde{v}}$ induces an orientation o on S_v as a 1-manifold. In keeping with Definition 7.3.1, we use the notation $\mathbb{S}_v = (S_v, \mathsf{o})$, and we call this an *oriented cylindrical fibre*. We colour oriented cylindrical fibres according to the $\mathscr{G}$-orbit of the corresponding oriented cylinders, and denote these colours with square brackets,

so $[\mathbb{S}_v] := [\tilde{v}, \mathcal{O}]$ for any lift $\tilde{v}$ of v and choice of end $\mathcal{O}$ of $\mathbb{Z}_{\tilde{v}}$ that induces the orientation $\mathbb{S}_v$. Note that different lifts $\tilde{v}$ will give oriented cylinders in the same G-orbit, so the colouring on $\mathbb{S}_v$ is well-defined.

Similarly, we can put orientations on the edge spaces $\mathbb{X}_e = (\mathcal{X}_e, \mathsf{o})$, and of course we already have the notion of oriented fin $\mathbb{S}_e = (S_e, \mathsf{o})$. We colour oriented edge spaces according to the $\mathscr{G}$-orbit of the corresponding oriented edge groups, and denote these colours with square brackets, so $[\mathbb{X}_e] := [\tilde{e}, \mathcal{O}] \in \mathcal{C}$ for any lift $\tilde{e}$ of e and choice of end $\mathcal{O}$ of $G_{\tilde{e}}$ that induces the orientation $\mathbb{X}_e$. As for oriented fins we use bars to denote the opposite orientation, so $\bar{\mathbb{S}}_v$ is the opposite orientation to $\mathbb{S}_v$ and $\bar{\mathbb{X}}_e$ is the opposite orientation to $\mathbb{X}_e$. When $\phi_e : \mathbb{X}_e \to \mathbb{S}_v$ is orientation preserving we write $\phi_e(\mathbb{X}_e) = \mathbb{S}_e$, and when $\phi_{\tilde{e}} : \mathbb{X}_e \to \mathbb{S}_e$ is orientation preserving we write $\phi_{\tilde{e}}(\mathbb{X}_e) = \mathbb{S}_e$. We make analogous definitions for $\mathbb{S}_{v'} = (S_{v'}, \mathsf{o})$ and $\mathbb{X}_{e'} = (\mathcal{X}'_{e'}, \mathsf{o})$ in $\mathcal{X}'$.

At this point we have ways of defining orientations on several different objects, so we should take a moment to check that these orientations are compatible by chasing the definitions. Suppose $\tilde{e}$ is a lift of an edge $e \in E_1\Gamma$ with $\tau(e) = v \in V_1\Gamma$, $\tau(\tilde{e}) = \tilde{v}$, $\iota(e) = u \in V_0\Gamma$ and $\iota(\tilde{e}) = \tilde{u}$. If $\mathcal{O}$ is a choice of end of $\mathbb{Z}_{\tilde{v}} = G_{\tilde{e}}$ then we get an oriented cylindrical fibre $\mathbb{S}_v$ as above, but also an oriented fin $\mathbb{S}_{\tilde{e}}$ as in Section 7.4.3, which descends to an oriented fin $\mathbb{S}_e \in \partial_{\mathsf{o}}\mathbf{X}_u$. So we have a diagram

$$\begin{array}{ccccc} (\tilde{v}, \mathcal{O}) & \dashleftarrow & (\tilde{e}, \mathcal{O}) & \dashrightarrow & \mathbb{S}_{\tilde{e}} \\ \downarrow & & & & \downarrow \\ \mathbb{S}_v & \xleftarrow{\phi_e} & \mathbb{X}_e & \xrightarrow{\phi_{\tilde{e}}} & \mathbb{S}_e, \end{array} \tag{7.4.17}$$

where the dotted arrows represent one orientation inducing another, and the solid arrows are orientation preserving maps of 1-manifolds. The colours are also compatible, so $[\mathbb{S}_v] = [\tilde{v}, \mathcal{O}]$ and $[\mathbb{X}_e] = [\tilde{e}, \mathcal{O}] = \lambda(\mathbb{S}_{\tilde{e}}) = \lambda(S_e)$.

Definition 7.4.19. (Stretch ratio)
For a rigid vertex $\tilde{u} \in V_0T_c$ and $\iota(\tilde{e}) = \tilde{u}$, in Definition 7.4.15 we set $r_{\tilde{e}}$ to be the translation length of a generator of $g \in G_{\tilde{e}}$ acting on $\mathbf{Y}_{\tilde{u}}$. We know that $G_{\tilde{e}}$ is the $G_{\tilde{u}}$-stabiliser of the fin $S_{\tilde{e}}$, and that the quotient of $S_{\tilde{e}}$ is the fin $S_e \in \partial\mathbf{X}_u$, where $\tilde{u}$ and $\tilde{e}$ descend to u and e in Γ, and so $r_{\tilde{e}} = \ell(S_e)$. For $\tilde{v} \in V_1T_c$, we defined the stretch ratio $\mathrm{Str}(\tilde{v})$ to be the ratio of the numbers $r_{\tilde{e}}$ for $\tilde{e} \in \mathrm{lk}(\tilde{v})$, thus it makes sense to define the *stretch ratio of* $v \in V_1\Gamma \sqcup V_1\Gamma'$ to be the class of functions $\mathrm{Str}(v) := [e \in \mathrm{lk}(v) \mapsto \ell(S_e)]$.

Lemma 7.4.16 tells us that the stretch ratio depends only on the $\mathscr{G}$-orbits of the edges. More precisely, there are numbers $r_{[\tilde{e}]}$ such that $\mathrm{Str}(\tilde{v}) = [\tilde{e} \mapsto r_{[\tilde{e}]}]$ for $\tilde{v} \in V_1T_c$, which implies that

$$\mathrm{Str}(v) = [e \mapsto r_{[e]}] \tag{7.4.18}$$

for $v \in V_1\Gamma \sqcup V_1\Gamma'$.

7.4.6 Density coefficients

Definition 7.4.20. Given our graph of spaces $(\mathcal{X}, \Gamma)$ we define the *volume* of $\mathcal{X}$ to be the following sum (recall that $|X_u|$ is the number of vertices in the graph X_u):

$$|\mathcal{X}| := \sum_{u \in V_0\Gamma} |X_u|.$$

Given a rigid vertex $u \in V_0\Gamma$, we define the *density* of the colour $[u]$ in $(\mathcal{X}, \Gamma)$, denoted $\rho_{[u]}$, to be the value

$$\rho_{[u]} := \sum_{u_* \in V_0\Gamma,\, [u_*]=[u]} |X_{u_*}|/|\mathcal{X}|. \tag{7.4.19}$$

Remark 7.4.21. We can also consider the density of $[u]$ in $(\mathcal{X}', \Gamma')$, since the vertices of $V_0\Gamma'$ are labelled with the same colours, but prima facie there is no reason to believe that they will be equal. However, because density is preserved by finite covers of graphs of spaces, after we have constructed a common finite cover $\widehat{\mathcal{X}}$ we will know that $\rho_{[u]}$ gives the same value whether defined with Γ or Γ'.

We recall Definition 7.3.6, the notion of the density ρ_c of a colour c given a graph with coloured fins. The following lemma relates the local notion of density ρ_c of a particular vertex space $\mathbf{X}_u$, with the global density $\rho_{[u]}$ of the vertex spaces of colour $[u]$.

Lemma 7.4.22. *Let $\mathbb{S}_e \in \partial_\circ \mathbf{X}_u$ be an oriented fin of colour c. Then*

$$\sum_{\lambda(\mathbb{S}_{e_*})=c,\, e_* \in E_1\Gamma} \ell(\mathbb{S}_{e_*}) = \rho_c \rho_{[u]} |\mathcal{X}| \tag{7.4.20}$$

Proof. All $\mathbf{X}_{u_*}$ containing an oriented fin of colour c have $[u_*] = [u]$ and are covered by $\mathbf{Y}_{\tilde{u}}$ for some $\tilde{u} \in V_0 T_c$ a lift of u. Hence by Theorem 7.3.7, all these $\mathbf{X}_{u_*}$ have a common finite cover, and so they all have the same density ρ_c. We can then make the following computation:

$$\begin{aligned}
\sum_{\lambda(\mathbb{S}_{e_*})=c,\, e_* \in E_1\Gamma} \ell(\mathbb{S}_{e_*}) &= \sum_{\substack{u_* \in V_0\Gamma, \\ [u_*]=[u]}} \left[\sum_{\substack{\mathbb{S}_{e_*} \in \partial_\circ \mathbf{X}_{u_*}, \\ \lambda(\mathbb{S}_{e_*})=c}} \ell(\mathbb{S}_{e_*}) \right] \\
&= \sum_{\substack{u_* \in V_0\Gamma, \\ [u_*]=[u]}} \rho_c |X_{u_*}| \\
&= \rho_c \rho_{[u]} |\mathcal{X}|.
\end{aligned}$$

□

7.5 A common finite cover

In this section we complete the proof of Theorem 1.3.3 by constructing a common finite cover of the graphs of spaces $\mathcal{X}$ and $\mathcal{X}'$ from the previous section.

7.5.1 A template for our desired common cover

More precisely, we will construct finite covers $\widehat{\mathcal{X}} \to \mathcal{X}$ and $\widehat{\mathcal{X}}' \to \mathcal{X}'$ such that $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{X}}'$ are homotopy equivalent. This will be achieved by constructing $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{X}}'$ such that their induced decompositions are over graphs $\widehat{\Gamma}$ and $\widehat{\Gamma}'$ that are type and colour preserving. Indeed if we identify $\widehat{\Gamma}$ and $\widehat{\Gamma}'$, then we will have homeomorphic vertex spaces $\widehat{\mathcal{X}}_v \cong \widehat{\mathcal{X}}'_v$ for all $v \in V\widehat{\Gamma}$ and homeomorphic edge spaces $\widehat{\mathcal{X}}_e \cong \widehat{\mathcal{X}}'_e$ for all $e \in E\widehat{\Gamma}$. The attaching maps $\hat{\phi}_e, \hat{\phi}'_e : \widehat{\mathcal{X}}_e \to \widehat{\mathcal{X}}_v$ will be homotopic for all $e \in E\widehat{\Gamma}$. By a standard result in topology the graphs of spaces will therefore be homotopic. Commensurability of G and G' will follow.

7.5.2 Common covers of vertex and edge spaces

In this section we define the vertex and edge spaces of $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{X}}'$.

Definition 7.5.1. (Common covers of rigid vertex spaces)
For rigid vertices $u \in V_0\Gamma$ and $u' \in V_0\Gamma'$ of the same colour $[u] = [u']$, we describe how to produce a common cover $\widehat{\mathbf{X}}_{u,u'}$ of the graphs with fins $\mathcal{X}_u = \mathbf{X}_u$ and $\mathcal{X}'_{u'} = \mathbf{X}_{u'}$. These two graphs with fins are defined by the quotients $\mathbf{Y}_{\tilde{u}}/G_{\tilde{u}}$ and $\mathbf{Y}_{\tilde{u}'}/G'_{\tilde{u}'}$, where $\tilde{u}$ and $\tilde{u}'$ are lifts of u and u' respectively to T_c. As $[\tilde{u}] = [u] = [u'] = [\tilde{u}']$, we know that there exists $[f] \in \mathscr{G}$ with $\hat{f}(\tilde{u}) = \tilde{u}'$ and $[\boldsymbol{f}]_{\tilde{u}} : \mathbf{Y}_{\tilde{u}} \to \mathbf{Y}_{\tilde{u}'}$ an isomorphism. We know that the action of $[f]^{-1}G'_{\tilde{u}'}[f] < \mathscr{G}_{\tilde{u}}$ on $\mathbf{Y}_{\tilde{u}}$ is conjugate to the action of $G'_{\tilde{u}'}$ on $\mathbf{Y}_{\tilde{u}'}$ via $[\boldsymbol{f}]_{\tilde{u}}$, so $\mathbf{X}_{u'} \cong \mathbf{Y}_{\tilde{u}}/[f]^{-1}G'_{\tilde{u}'}[f]$. We can then apply Theorem 7.3.7 to $\mathbf{Y}_{\tilde{u}}$, with $\Gamma_1, \Gamma_2 < \mathrm{Aut}(\mathbf{Y}_{\tilde{u}})$ the images of $G_{\tilde{u}}, [f]^{-1}G'_{\tilde{u}'}[f] < \mathscr{G}_{\tilde{u}}$ under the homomorphism $\mathscr{G}_{\tilde{u}} \to \mathrm{Aut}(\mathbf{Y}_{\tilde{u}})$, to produce a common finite cover $\widehat{\mathbf{X}}_{u,u'}$ of $\mathbf{X}_u$ and $\mathbf{X}_{u'}$ that satisfies equation (7.3.2). Note that the colours of oriented fins in $\mathbf{Y}_{\tilde{u}}$ were defined to correspond to $\mathscr{G}$-orbits (Section 7.4.3), so $\mathscr{G}_{\tilde{u}}$ does indeed act transitively on the oriented fins of each colour in $\mathbf{Y}_{\tilde{u}}$. Additionally note that, while the definitions of $\mathbf{X}_u$ and $\mathbf{X}_{u'}$ did not depend on the choice of lifts $\tilde{u}$ and $\tilde{u}'$, the definition of $\widehat{\mathbf{X}}_{u,u'}$ does depend on these choices, and also on the choice of $[f] \in \mathscr{G}$.

The following lemma is a direct application of omnipotence of free groups.

Lemma 7.5.2. *We can choose integers $\ell_{[e]}$ for $e \in E_1\Gamma \sqcup E_1\Gamma'$ and replace each $\widehat{\mathbf{X}}_{u,u'}$ with a finite cover, such that the length of a fin $\hat{S} \in \partial\widehat{\mathbf{X}}_{u,u'}$ that covers a fin $S_e \in \partial\mathbf{X}_u \sqcup \partial\mathbf{X}_{u'}$ is $\ell_{[e]}$. Moreover, for a vertex $v \in V_1\Gamma \sqcup V_1\Gamma'$ we have $\mathrm{Str}(v) = [e \mapsto \ell_{[e]}]$, or equivalently there is an integer d_v such that*

$$\ell_{[e]} = d_v \ell(S_e), \tag{7.5.1}$$

for all $e \in \mathrm{lk}(v)$ - so the degree of the covering $\hat{S} \to S_e$ is d_v and depends only on v.

Proof. By omnipotence of free groups [97, Theorem 3.5], there exists $N > 0$ such that for any $k : \partial\widehat{\mathbf{X}}_{u,u'} \to \mathbb{N}$ there exists a normal cover $\Phi : \overline{\mathbf{X}} \to \widehat{\mathbf{X}}_{u,u'}$ such that the length of any fin in $\Phi^{-1}(S)$ is $Nk(S)$. If $\hat{S} \in \partial\widehat{\mathbf{X}}_{u,u'}$ covers fins $S_e \in \partial\mathbf{X}_u$ and $S_{e'} \in \partial\mathbf{X}_{u'}$ then

$[e] = [e']$, because S_e and $S_{e'}$ will have orientations of the same colour, and the colour of an oriented fin determines the colour of the corresponding edge (see Definition 7.4.17). Therefore, we can replace the $\widehat{\mathbf{X}}_{u,u'}$ with further finite covers and assume that the length of a fin covering S_e is $\ell_{[e]}$. We know from (7.4.18) that $\mathrm{Str}(v) = [e \mapsto r_{[e]}]$, so if we set $\ell_{[e]} = N r_{[e]}$, then we have that

$$\mathrm{Str}(v) = [e \mapsto \ell_{[e]}] \tag{7.5.2}$$

for a vertex $v \in V_1\Gamma \sqcup V_1\Gamma'$. Note that equation (7.3.2) from Theorem 7.3.7 is preserved by passing to a further finite cover. □

We also need common finite covers for the cylindrical vertex spaces.

Definition 7.5.3. (Common covers of cylindrical vertex spaces)
Given cylindrical vertices $v \in V_1\Gamma$ and $v' \in V_1\Gamma'$ and oriented cylindrical fibres $\mathbb{S}_v$ and $\mathbb{S}_{v'}$ of the same colour (see Definition 7.4.18), we let $\hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'})$ be an oriented circle equipped with orientation preserving covering maps to $\mathbb{S}_v$ and $\mathbb{S}_{v'}$ of degrees d_v and $d_{v'}$ respectively (where d_v and $d_{v'}$ come from (7.5.1)). We extend each $\hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'})$ to a common cover $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ of the vertex spaces $\mathcal{X}_v$ and $\mathcal{X}'_{v'}$. If $G_v \cong G'_{v'} \cong \mathbb{Z}$ then no extension is necessary, while if $G_v \cong G'_{v'} \cong \mathbb{Z}^2$ then we make $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ a torus containing $\hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'})$ as an embedded circle, so that $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ is the cover corresponding to the subgroups $d_v\mathbb{Z}_v \times \mathbb{Z} < \mathbb{Z}_v \times \mathbb{Z} = G_v = \pi_1(\mathcal{X}_v)$ and $d_{v'}\mathbb{Z}_{v'} \times \mathbb{Z} < \mathbb{Z}_{v'} \times \mathbb{Z} = G_{v'} = \pi_1(\mathcal{X}'_{v'})$. We consider $\hat{\mathbb{S}}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ to be the same embedded circle as $\hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'})$ but with orientation reversed, while $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'}) = \widehat{\mathcal{X}}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ is just a space with no orientation. Thus we obtain a pair of common covers for each pair of vertices v and v'.

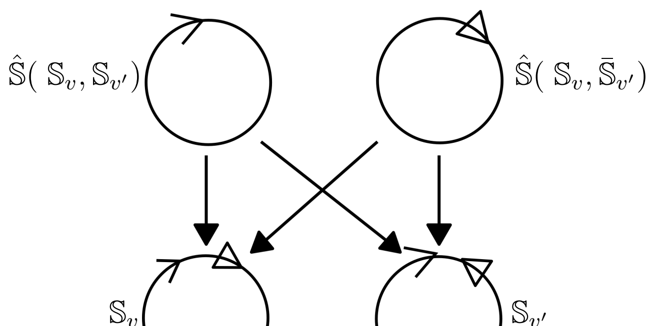

Figure 7.2: Each cylindrical fibre has the clockwise orientation and the covering maps are determined by the arrows. Note that if we take the anticlockwise orientations we obtain $\hat{\mathbb{S}}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ and $\hat{\mathbb{S}}(\bar{\mathbb{S}}_v, \mathbb{S}_{v'})$. Thus $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'}) = \widehat{\mathcal{X}}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ and $\widehat{\mathcal{X}}(\mathbb{S}_v, \bar{\mathbb{S}}_{v'}) = \widehat{\mathcal{X}}(\bar{\mathbb{S}}_v, \mathbb{S}_{v'})$.

Definition 7.5.4. (Common covers of edge spaces)
If $e \in E_1\Gamma$ and $e' \in E_1\Gamma'$ are edges with $\tau(e) = v \in V_1\Gamma$ and $\tau(e') = v' \in V_1\Gamma'$, and $\mathbb{X}_e$ and $\mathbb{X}_{e'}$ are orientations of the same colour, then we define $\widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ to be an oriented circle equipped with orientation preserving covering maps to $\mathbb{X}_e$ and $\mathbb{X}_{e'}$ of degrees d_v and

$d_{v'}$ respectively. We identify $\widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ and $\widehat{\mathbb{X}}(\bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'})$ as two orientations of the same space $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'}) = \widehat{\mathcal{X}}(\bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'})$. So again we obtain a pair of common covers for each pair of edges.

7.5.3 Link maps

Having defined common covers of the edge and vertex spaces, we now need to glue them together, or rather enumerate the possible ways of gluing them together. The following definition will be used to describe the ways of gluing a cylindrical vertex space $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ to edge spaces $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$.

Definition 7.5.5. (Link maps)
Let $v \in V_1\Gamma$ and $v' \in V_1\Gamma'$ be cylindrical vertices and consider oriented cylindrical fibres $\mathbb{S}_v$ and $\mathbb{S}_{v'}$ of the same colour. This induces orientations $\mathbb{X}_e$ and $\mathbb{X}_{e'}$ on the incident edge spaces. A *link map* from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$ is a colour preserving bijection between the incident oriented edge spaces, so in symbols it is a bijection

$$\sigma : \mathrm{lk}(v) \to \mathrm{lk}(v')$$

such that $[\mathbb{X}_e] = [\mathbb{X}_{\sigma(e)}]$ for all $e \in \mathrm{lk}(v)$. We let $\mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ be the set of all link maps from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$.

Lemma 7.5.6. *Let $c \in \mathcal{C}$. The number of $e \in \mathrm{lk}(v)$ with $[\mathbb{X}_e] = c$ is equal to the number of $e' \in \mathrm{lk}(v')$ with $[\mathbb{X}_{e'}] = c$ is equal to $N_c[\mathbb{S}_v]$. In particular,* $\mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ *is non-empty.*

Proof. The oriented cylindrical fibre $\mathbb{S}_v$ corresponds to a choice of end $\mathcal{O}$ of a cylindrical factor $\mathbb{Z}_{\tilde{v}}$ for $\tilde{v}$ a lift of v. The incident oriented edge spaces $\mathbb{X}_e$ correspond to $G_{\tilde{v}}$-orbits of oriented edge groups $(\tilde{e}, \mathcal{O})$ with $\tilde{e} \in \mathrm{lk}(\tilde{v})$. Moreover, for $c \in \mathcal{C}$, the number of $G_{\tilde{v}}$-orbits of oriented edge groups $(\tilde{e}, \mathcal{O})$ of colour c is equal to the cylinder number $t_c(v, \mathcal{O})$ by Definition 7.4.6. Hence the number of incident oriented edge spaces $\mathbb{X}_e$ of colour c is also equal to $t_c(v, \mathcal{O})$, and by (7.4.13) we have

$$t_c(v, \mathcal{O}) = N_c[v, \mathcal{O}] = N_c[\mathbb{S}_v],$$

so it only depends on the colours c and $[\mathbb{S}_v]$. Again by (7.4.13), we know that the number of oriented edge spaces incident to $\mathbb{S}_{v'}$ of colour c is equal to $N_c[\mathbb{S}_{v'}] = N_c[\mathbb{S}_v]$. □

Remark 7.5.7. $\sigma : \mathrm{lk}(v) \to \mathrm{lk}(v')$ defines a link map from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$ if and only if it defines a link map from $\bar{\mathbb{S}}_v$ to $\bar{\mathbb{S}}_{v'}$. This is because $\bar{\mathbb{S}}_v$ and $\bar{\mathbb{S}}_{v'}$ induce the orientations $\bar{\mathbb{X}}_e$ and $\bar{\mathbb{X}}_{e'}$ on the incident edge spaces, so if σ defines a link map from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$ then $[\bar{\mathbb{X}}_e] = \overline{[\mathbb{X}_e]} = \overline{[\mathbb{X}_{\sigma(e)}]} = [\bar{\mathbb{S}}_{\sigma(e)}]$ for each $e \in \mathrm{lk}(v)$.

Given a link map $\sigma : \text{lk}(v) \to \text{lk}(v')$ from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$ and $e \in \text{lk}(v)$ with $\sigma(e) = e'$, suppose $\iota(e) = u$ and $\iota(e') = u'$. Let $\phi_{\bar{e}}(\mathbb{X}_e) = \mathbb{S}_e \in \partial_{\circ}\mathbf{X}_u$ and $\phi_{\bar{e}'}(\mathbb{X}_{e'}) = \mathbb{S}_{e'} \in \partial_{\circ}\mathbf{X}_{u'}$. The fins $\mathbb{S}_e$ and $\mathbb{S}_{e'}$ both have colours equal to $[\mathbb{X}_e] = [\mathbb{X}_{e'}]$, so equation (7.3.2) implies that there exists $\hat{\mathbb{S}} \in \partial_{\circ}\widehat{\mathbf{X}}_{u,u'}(\mathbb{S}_e, \mathbb{S}_{e'})$ that covers both of them. Equation (7.5.1) tells us that these covering maps of fins have degrees d_v and $d_{v'}$ respectively, so we get two commutative diagrams as follows.

$$\begin{array}{ccccccccc}
\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'}) & \hookleftarrow & \hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'}) & \xleftarrow{\sim} & \widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'}) & \xrightarrow{\sim} & \hat{\mathbb{S}} & \hookrightarrow & \widehat{\mathbf{X}}_{u,u'} \\
\downarrow & & \downarrow{\scriptstyle d_v} & & \downarrow{\scriptstyle d_v} & & \downarrow{\scriptstyle d_v} & & \downarrow \\
\mathcal{X}_v & \hookleftarrow & \mathbb{S}_v & \xleftarrow[\sim]{\phi_e} & \mathbb{X}_e & \xrightarrow[\sim]{\phi_{\bar{e}}} & \mathbb{S}_e & \hookrightarrow & \mathbf{X}_u
\end{array} \tag{7.5.3}$$

$$\begin{array}{ccccccccc}
\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'}) & \hookleftarrow & \hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'}) & \xleftarrow{\sim} & \widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'}) & \xrightarrow{\sim} & \hat{\mathbb{S}} & \hookrightarrow & \widehat{\mathbf{X}}_{u,u'} \\
\downarrow & & \downarrow{\scriptstyle d_{v'}} & & \downarrow{\scriptstyle d_{v'}} & & \downarrow{\scriptstyle d_{v'}} & & \downarrow \\
\mathcal{X}'_{v'} & \hookleftarrow & \mathbb{S}_{v'} & \xleftarrow[\sim]{\phi'_{e'}} & \mathbb{X}_{e'} & \xrightarrow[\sim]{\phi'_{\bar{e}'}} & \mathbb{S}_{e'} & \hookrightarrow & \mathbf{X}_{u'}
\end{array} \tag{7.5.4}$$

In these diagrams a homeomorphism is indicated by $\sim$. Also note that the middle six spaces in each diagram have associated orientations, which are preserved by the maps between them.

These diagrams give us the right local data to define edge maps in the covers $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{X}}'$. The vertical maps are the coverings from vertex and edge spaces of $\widehat{\mathcal{X}}$ and $\widehat{\mathcal{X}}'$ to vertex and edge spaces of $\mathcal{X}$ and $\mathcal{X}'$, as defined in Section 7.5.2. Then the top row of (7.5.3) can be used to define edge maps of $\widehat{\mathcal{X}}$ while the top row of (7.5.4) can be used to define edge maps of $\widehat{\mathcal{X}}'$. The two maps from $\widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ to $\hat{\mathbb{S}}$ are both orientation preserving homeomorphisms of circles, hence they are homotopic, similarly the two maps from $\widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ to $\hat{\mathbb{S}}(\mathbb{S}_v, \mathbb{S}_{v'})$ are homotopic. See Figure 7.3. From now on we will only care about these edge maps up to homotopy, so we will just talk about a single cover $\widehat{\mathcal{X}}$ of $\mathcal{X}$ and $\mathcal{X}'$.

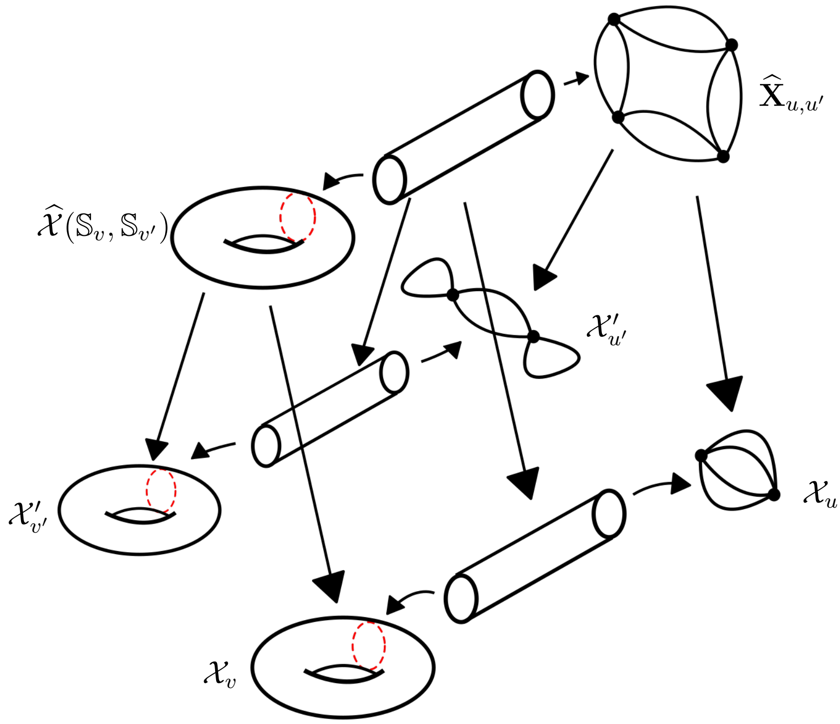

Figure 7.3: An illustration of how the common cover is constructed. The arrows in the diagram commute, and the dashed lines in the tori denote the cylindrical fibres.

Remark 7.5.8. Under the replacement $\mathbb{S}_v, \mathbb{S}_{v'}, \mathbb{X}_e, \mathbb{X}_{e'}, \mathbb{S}_e, \mathbb{S}_{e'}, \hat{\mathbb{S}} \mapsto \bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'}, \bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'}, \bar{\mathbb{S}}_e, \bar{\mathbb{S}}_{e'}, \bar{\hat{\mathbb{S}}}$, diagrams (7.5.3) and (7.5.4) will consist of the same spaces and maps, the orientations of the spaces will just reverse. So when using σ to construct the local data of edge maps in $\widehat{\mathcal{X}}$, it doesn't matter whether we regard σ as a link map from $\mathbb{S}_v$ to $\mathbb{S}_{v'}$ or as a link map from $\bar{\mathbb{S}}_v$ to $\bar{\mathbb{S}}_{v'}$.

7.5.4 From local common covers to global

A finite common cover $\widehat{\mathcal{X}}$ of $\mathcal{X}$ and $\mathcal{X}'$ will be constructed by taking as vertex spaces $\omega(u, u')$ copies of each $\widehat{\mathbf{X}}_{u,u'}$ and $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$ copies of each $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$, and as edge spaces $\omega(\mathbb{X}_e, \mathbb{X}_{e'})$ copies of each $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$. We require $\omega(\mathbb{S}_v, \mathbb{S}_{v'}) = \omega(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ and $\omega(\mathbb{X}_e, \mathbb{X}_{e'}) = \omega(\bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'})$ because $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'}) = \widehat{\mathcal{X}}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ and $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'}) = \widehat{\mathcal{X}}(\bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'})$. To each copy of $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ we associate a link map $\sigma \in \mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$, and then for each $e \in \mathrm{lk}(v)$ we glue an edge space $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ to $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ and also to a vertex space $\widehat{\mathbf{X}}_{u,u'}$, all according to the diagrams (7.5.3) and (7.5.4) (so $e' = \sigma(e)$, $u = \iota(e)$ and $u' = \iota(e')$). By Remark 7.5.8 it doesn't matter whether we regard σ as lying in $\mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ or $\mathrm{LkMap}(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$. The different $\sigma \in \mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ will be evenly distributed across the $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$ copies of $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$ (so in particular $|\mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})|$ will divide $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$).

For this to form a cover of $\mathcal{X}$ and $\mathcal{X}'$, we must ensure that each edge space $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ gets used exactly once, and that each fin in each vertex space $\widehat{\mathbf{X}}_{u,u'}$ has exactly one edge space glued to it. This requirement can be captured by a set of Gluing Equations, which we describe in the following lemma.

Lemma 7.5.9. *(Gluing Equations)*

We can form a common finite cover $\hat{\mathcal{X}}$ of $\mathcal{X}$ and $\mathcal{X}'$ by the above gluing instructions if the following Gluing Equations have a positive solution:

$$\frac{\omega(\mathbb{S}_v, \mathbb{S}_{v'})}{N_c[\mathbb{S}_v]} = \omega(\mathbb{X}_e, \mathbb{X}_{e'}) = \omega(u, u')|\partial_\circ \widehat{\mathbf{X}}_{u,u'}(\mathbb{S}_e, \mathbb{S}_{e'})| \tag{7.5.5}$$

Here $\mathbb{S}_v$ and $\mathbb{S}_{v'}$ are oriented cylindrical fibres from $\mathcal{X}$ and $\mathcal{X}'$ of the same colour; $e \in \mathrm{lk}(v)$ and $e' \in \mathrm{lk}(v')$ are edges such that the edge spaces with induced orientations $\mathbb{X}_e$ and $\mathbb{X}_{e'}$ have the same colour $c \in \mathcal{C}$; $\iota(e) = u$ and $\iota(e') = u'$; and $\phi_{\bar{e}}(\mathbb{X}_e) = \mathbb{S}_e \in \partial_\circ \mathbf{X}_u$ and $\phi_{\bar{e}'}(\mathbb{X}_{e'}) = \mathbb{S}_{e'} \in \partial_\circ \mathbf{X}_{u'}$ are the oriented fins corresponding to $\mathbb{X}_e$ and $\mathbb{X}_{e'}$.

Proof. By Lemma 7.5.6, there are $N_c[\mathbb{S}_v]$ edges $e'_* \in \mathrm{lk}(v')$ whose oriented edge spaces $\mathbb{X}_{e'_*}$ have colour c, and any choice $e \mapsto e'_*$ can be extended to a link map $\sigma : \mathrm{lk}(v) \to \mathrm{lk}(v')$. Moreover, the number of possible extensions is independent of e'_*, thus the proportion of link maps $\sigma \in \mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ with $\sigma(e) = e'$ is $1/N_c[\mathbb{S}_v]$. By the local gluing data of (7.5.3) and (7.5.4), a copy of $\widehat{\mathcal{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ is used in the construction of $\widehat{\mathcal{X}}$ precisely when a link map $\sigma \in \mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})$ with $\sigma(e) = e'$ is associated to a vertex space $\widehat{\mathcal{X}}(\mathbb{S}_v, \mathbb{S}_{v'})$. This explains the first equality in (7.5.5).

For the second equality in (7.5.5), note that the local gluing data of (7.5.3) and (7.5.4) glues each copy of an oriented edge space $\widehat{\mathbb{X}}(\mathbb{X}_e, \mathbb{X}_{e'})$ to an oriented fin $\hat{\mathbb{S}} \in \partial_\circ \widehat{\mathbf{X}}_{u,u'}(\mathbb{S}_e, \mathbb{S}_{e'})$, for one of the $\omega(u, u')$ copies of $\widehat{\mathbf{X}}_{u,u'}$; and these are the only edge spaces that could be glued to $\hat{\mathbb{S}}$ because $\mathbb{X}_e$ and $\mathbb{X}_{e'}$ are the unique oriented edge spaces that attach to the oriented fins $\mathbb{S}_e$ and $\mathbb{S}_{e'}$.

Of course we also need $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$, $\omega(\mathbb{X}_e, \mathbb{X}_{e'})$ and $\omega(u, u')$ to be positive integers, and for $|\mathrm{LkMap}(\mathbb{S}_v, \mathbb{S}_{v'})|$ to divide $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$, but this can be achieved by scaling the solution suitably. □

Lemma 7.5.2 tells us that all fins in $\widehat{\mathbf{X}}_{u,u'}$ that cover $S_e \in \partial \mathbf{X}_u$ have length $\ell_{[e]}$, so Theorem 7.3.7 tells us that we can substitute

$$|\partial_\circ \widehat{\mathbf{X}}_{u,u'}(\mathbb{S}_e, \mathbb{S}_{e'})| = \left(\frac{|\widehat{X}_{u,u'}|}{\rho_c |X_u||X_{u'}|} \right) \frac{\ell(\mathbb{S}_e)\ell(\mathbb{S}_{e'})}{\ell_{[e]}}$$

into equations (7.5.5). Thus we can solve the gluing equations by taking

$$\omega(u, u') = \frac{|X_u||X_{u'}|}{\rho_{[u]}|\widehat{X}_{u,u'}|}, \text{ and } \frac{\omega(\mathbb{S}_v, \mathbb{S}_{v'})}{N_c[\mathbb{S}_v]} = \omega(\mathbb{X}_e, \mathbb{X}_{e'}) = \frac{\ell(\mathbb{S}_e)\ell(\mathbb{S}_{e'})}{\ell_{[e]}\rho_c\rho_{[u]}}. \tag{7.5.6}$$

It remains to show that this solution is well-defined. Note that the replacement $\mathbb{S}_v, \mathbb{S}_{v'} \mapsto \bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'}$ will flip the orientations on the fins $\mathbb{S}_e$ and $\mathbb{S}_{e'}$, and the colour c will turn to $\bar{c}$; but this will not change the lengths of the fins, and $N_{\bar{c}}[\bar{\mathbb{S}}_v] = N_c[\mathbb{S}_v]$ by Lemma 7.4.9; and $\rho_{\bar{c}} = \rho_c$ because by definition this is proportional to the sum of lengths of oriented fins of colour $\bar{c}$, which equals the sum of lengths of oriented fins of colour c

since these are different orientations of the same fins. Hence $\omega(\mathbb{S}_v, \mathbb{S}_{v'}) = \omega(\bar{\mathbb{S}}_v, \bar{\mathbb{S}}_{v'})$ and $\omega(\mathbb{X}_e, \mathbb{X}_{e'}) = \omega(\bar{\mathbb{X}}_e, \bar{\mathbb{X}}_{e'})$ as required.

It is easy to see that the formula for $\omega(u, u')$ depends only on u and u', and that the formula for $\omega(\mathbb{X}_e, \mathbb{X}_{e'})$ depends only on $\mathbb{X}_e$ and $\mathbb{X}_{e'}$; but the reason that the formula for $\omega(\mathbb{S}_v, \mathbb{S}_{v'})$ depends only on $\mathbb{S}_v$ and $\mathbb{S}_{v'}$ is more subtle, which is the task of our final lemma.

Lemma 7.5.10. *The expression*

$$\frac{N_c[\mathbb{S}_v]\ell(\mathbb{S}_e)\ell(\mathbb{S}_{e'})}{\ell_{[e]}\rho_c\rho_{[u]}}$$

depends only on $\mathbb{S}_v$ *and* $\mathbb{S}_{v'}$.

Proof.

$$\begin{aligned}
\ell_{[e]}\rho_c\rho_{[u]}|\mathcal{X}| &= \ell_{[e]} \sum_{\lambda(\mathbb{S}_{e_*})=c,\, e_*\in E_1\Gamma} \ell(\mathbb{S}_{e_*}) && \text{by Lemma 7.4.22,}\\
&= \ell_{[e]} \sum_{\substack{[\mathbb{S}_{v_*}]=[\mathbb{S}_v],\, v_*\in V_1\Gamma \\ \phi_{e_*}\phi_{\bar{e}_*}^{-1}(\mathbb{S}_{e_*})=\mathbb{S}_{v_*},\, \lambda(\mathbb{S}_{e_*})=c,\, e_*\in \mathrm{lk}(v)}} \ell(\mathbb{S}_{e_*})\\
\ell_{[e]} &\sum_{[\mathbb{S}_{v_*}]=[\mathbb{S}_v],\, v_*\in V_1\Gamma} N_c[\mathbb{S}_v]\ell(\mathbb{S}_{e_*}) && \text{by Lemma 7.5.6,}\\
&= \sum_{[\mathbb{S}_{v_*}]=[\mathbb{S}_v],\, v_*\in V_1\Gamma} \frac{N_c[\mathbb{S}_v]\ell_{[e]}^2}{d_{v_*}} && \text{by (7.5.1),}\\
&= \sum_{[\mathbb{S}_{v_*}]=[\mathbb{S}_v],\, v_*\in V_1\Gamma} \frac{N_c[\mathbb{S}_v]d_v d_{v'}\ell(\mathbb{S}_e)\ell(\mathbb{S}_{e'})}{d_{v_*}} && \text{again by (7.5.1).}
\end{aligned}$$

And so our required expression

$$\frac{N_c[\mathbb{S}_v]\ell(\mathbb{S}_e)\ell(\mathbb{S}_{e'})}{\ell_{[e]}\rho_c\rho_{[u]}} = |\mathcal{X}| \left(\sum_{[\mathbb{S}_{v_*}]=[\mathbb{S}_v],\, v_*\in V_1\Gamma} \frac{d_v d_{v'}}{d_{v_*}} \right)^{-1},$$

only depends on $\mathbb{S}_v$ and $\mathbb{S}_{v'}$. □

We conclude that (7.5.6) gives a well-defined solution to the Gluing Equations, and so by Lemma 7.5.9 we can form a common finite cover $\hat{\mathcal{X}}$ of $\mathcal{X}$ and $\mathcal{X}'$. Thus G and G' are commensurable, completing the proof of Theorem 1.3.3.

7.6 Counter example for higher rank cylinders

We now consider the wider class of groups $\mathscr{C}^\bullet$ of all subgroup separable, one-ended, finitely presented groups with JSJ decomposition consisting of virtually free vertex groups, and no QH vertex groups. By Theorem 1.3.2 such groups are hyperbolic relative to virtually free-by-cyclic vertex groups.

We present the following pair of groups which we assert are quasi-isometric, but not commensurable.

Let $w \in \mathbb{F}_2 = \langle x, y \rangle$ be a word that induces a rigid line pattern in $\mathbb{F}_2$. We consider the following groups:

$$G = \mathbb{F}_2 *_{\mathbb{Z}} (\mathbb{F}_2 \times \mathbb{Z}) \ = \ \langle x, y, a, b, z \mid w = z, \ [a, z] = [b, z] = 1 \rangle,$$

and

$$G' = \mathbb{F}_2 *_{\mathbb{Z}} (\mathbb{F}_3 \times \mathbb{Z}) \ = \ \langle x, y, a, b, c, z \mid w = z, \ [a, z] = [b, z] = [c, z] = 1 \rangle.$$

We note that $\mathcal{X}(G) = \mathcal{X}(G') = -1$ since the free-by-cyclic factors contribute nothing. These groups are torsion-free, and in the language of Guiradel and Levitt [38], the given splitting corresponds to the canonical tree of cylinders with respect to a JSJ decomposition. We also note that these groups are virtually special.

Lemma 7.6.1. *G and G' are quasi-isometric.*

Proof. Let $f : \mathbb{F}_2 = \langle a, b \rangle \to \mathbb{F}_3 = \langle a, b, c \rangle$ be a bi-Lipschitz bijection with bi-Lipschitz constant $C \geq 1$ - this exists by [71].

Write an element of G as $g = \alpha_1\beta_1\alpha_2\beta_2 \cdots \alpha_k\beta_k$, where $\alpha_i \in \langle x, y \rangle$, $\beta_i \in \langle a, b \rangle$, and $\alpha_i, \beta_i \neq 1$, $\alpha_i \notin \langle w \rangle$ (except possibly α_1 and β_k). Define a map $\psi : G \to G'$ by $\psi(g) := \alpha_1 f(\beta_1)\alpha_2 f(\beta_2) \cdots \alpha_k f(\beta_k)$ (viewing $\langle x, y \rangle$ as a common subgroup of G and G'). We check that ψ is well-defined: indeed the β_i are uniquely determined by g, and the only ambiguity in the α_i comes from making replacements $(\alpha_i, \alpha_{i+1}) \mapsto (\alpha_i w^j, w^{-j}\alpha_{i+1})$, which does not change $\psi(g)$.

We claim that ψ is a quasi-isometry. Take elements $g, \bar{g} \in G$ written in the above normal form, and make replacements as above so that they agree on an initial subword of maximum possible length. If the first term where they differ is an α_i term then we can write $g = \alpha_1\beta_1\alpha_2\beta_2 \cdots \alpha_k\beta_k$ and $\bar{g} = \alpha_1\beta_1 \cdots \alpha_{l-1}\beta_{l-1}\bar{\alpha}_l\bar{\beta}_l \cdots \bar{\alpha}_m\bar{\beta}_m$ with $\bar{\alpha}_l \notin \alpha_l\langle w \rangle$. Working with respect to the given generators for G and G', we use d to denote the metrics on G and G' and $|\cdot|$ to denote the distance to the identity. Then for appropriate choices of the α_i and $\bar{\alpha}_i$ we have

$$\begin{aligned} d(g, g') &= |\beta_k^{-1}\alpha_k^{-1} \cdots \beta_l^{-1}\alpha_l^{-1}\bar{\alpha}_l\bar{\beta}_l \cdots \bar{\alpha}_m\bar{\beta}_m| \\ &= |\beta_k^{-1}| + |\alpha_k^{-1}| + \cdots + |\beta_l^{-1}| + |\alpha_l^{-1}\bar{\alpha}_l| + |\bar{\beta}_l| + \cdots + |\bar{\alpha}_m| + |\bar{\beta}_m| \\ &= |\beta_k| + |\alpha_k| + \cdots + |\beta_l| + |\alpha_l^{-1}\bar{\alpha}_l| + |\bar{\beta}_l| + \cdots + |\bar{\alpha}_m| + |\bar{\beta}_m|. \end{aligned} \quad (7.6.1)$$

On the other hand

$$\begin{aligned} d(\psi(g), \psi(g')) &= |f(\beta_k)^{-1}\alpha_k^{-1} \cdots f(\beta_l)^{-1}\alpha_l^{-1}\bar{\alpha}_l f(\bar{\beta}_l) \cdots \bar{\alpha}_m f(\bar{\beta}_m)| \\ &\leq |f(\beta_k)^{-1}| + |\alpha_k^{-1}| + \cdots + |f(\beta_l)^{-1}| + |\alpha_l^{-1}\bar{\alpha}_l| + |f(\bar{\beta}_l)| + \cdots + |\bar{\alpha}_m| + |f(\bar{\beta}_m)| \\ &= |f(\beta_k)| + |\alpha_k| + \cdots + |f(\beta_l)| + |\alpha_l^{-1}\bar{\alpha}_l| + |f(\bar{\beta}_l)| + \cdots + |\bar{\alpha}_m| + |f(\bar{\beta}_m)| \\ &\leq C(|\beta_k| + \ldots |\beta_l| + |\bar{\beta}_l| + \cdots + |\bar{\beta}_m|) + |\alpha_k| + \cdots + |\alpha_l^{-1}\bar{\alpha}_l| + \ldots |\bar{\alpha}_m| \\ &\leq Cd(g, g'). \end{aligned} \quad (7.6.2)$$

A similar argument works if the first term where g and $\bar{g}$ differ is a β_i term rather than an α_i term. Using f^{-1} we can define an inverse to ψ (so in particular ψ is a bijection), and by the same argument as above we get $d(g, g') \leq Cd(\psi(g), \psi(g'))$ for any $g, \bar{g} \in G$. □

Lemma 7.6.2. *G and G' are not commensurable.*

Proof. Suppose that there exist finite index subgroups $\hat{G} < G$ and $\hat{G}' < G'$, such that $\hat{G} \cong \hat{G}'$.

There is an induced graph of groups decomposition of $\hat{G}$ from the decomposition of G. Let $(\hat{G}, \hat{\Gamma})$ denote that decomposition. There is also an induced decomposition of $\hat{G}'$, that we can denote by $(\hat{G}', \hat{\Gamma}')$, but at this point we argue from the uniqueness of these tree of cylinders decompositions ([38][Corollary 7.4]) that they are the same decomposition.

We now consider the vertex groups in $(\hat{G}, \hat{\Gamma})$ that cover the free-by-cyclic vertex group $\mathbb{F}_2 \times \mathbb{Z} = \langle a, b \rangle \times \langle z \rangle$ in G. If $\hat{G}_v$ is such a vertex group, then we have an embedding $\hat{G}_v \hookrightarrow \langle a, b \rangle \times \langle z \rangle$ as a finite index subgroup. We know that $\langle z \rangle$ is the edge group incident at $\langle a, b \rangle \times \langle z \rangle$ in G, so the edges incident at v correspond to double cosets $\hat{G}_v g \langle z \rangle$ for $g \in \langle a, b \rangle \times \langle z \rangle$.

Let $\pi : \langle a, b \rangle \times \langle z \rangle \to \langle a, b \rangle$ be the projection map, and consider the short exact sequence

$$1 \to \hat{G}_v \cap \langle z \rangle \hookrightarrow \hat{G}_v \xrightarrow{\pi} \pi(\hat{G}_v) \to 1. \tag{7.6.3}$$

$\pi(\hat{G}_v)$ is free, so there is a section $\sigma : \pi(\hat{G}_v) \to \hat{G}_v$, with image F say. As $\langle z \rangle$ is central in $\langle a, b \rangle \times \langle z \rangle$, we see that $\hat{G}_v$ splits as a product $\hat{G}_v = F \times (\hat{G}_v \cap \langle z \rangle)$. Note that the rank $n(v)$ of $F \cong \pi(\hat{G}_v)$ is an invariant of $\hat{G}_v$ (it is one less than the rank of the abelianisation of $\hat{G}_v$ for example). A double coset $\hat{G}_v g \langle z \rangle$ must equal $\pi(\hat{G}_v)\pi(g) \times \langle z \rangle < \langle a, b \rangle \times \langle z \rangle$, so the number of such double cosets is equal to the index $|\langle a, b \rangle : \pi(\hat{G}_v)|$. But we know $\langle a, b \rangle$ and $\pi(\hat{G}_v)$ are free groups of rank 2 and $n(v)$ respectively, so this index must equal $n(v) - 1$, and as discussed above this is the degree of the vertex v in $\hat{\Gamma}$.

We can run exactly the same arguments for $\hat{G}'_v \cong \hat{G}_v$ embedded in $\mathbb{F}_3 \times \mathbb{Z} = \langle a, b, c \rangle \times \langle z \rangle < G'$, and we get the same rank $n(v)$; the only difference is that we compute the degree of v in Γ as the index $|\langle a, b, c \rangle : \pi(\hat{G}_v)|$, which is the index between free groups of rank 3 and $n(v)$, and hence equals $(n(v) - 1)/2$, a contradiction. □

Chapter 8

Future directions

8.1 Special cube complexes

Section 1.1 contains results that generalise Leighton's Theorem and Section 1.2 contains results about controlling finite covers of special cube complexes, so this all points tantalizingly towards the following conjecture of Haglund [41, Problem 2.4].

Conjecture 1.2.9. *(Haglund)*
If X_1 and X_2 are finite special cube complexes with a common universal cover $\tilde{X}$, then they have a common finite cover.

We note that Conjecture 1.2.9 does not hold for arbitrary non-positively curved cube complexes, with the square complexes of Wise [100] and Burger–Mozes [19] providing counter-examples. However, there are a few cases where Conjecture 1.2.9 is known to hold, for example if $\tilde{X}$ is the CAT(0) cube complex obtained by subdividing a certain right-angled Fuchsian building [41], or if $\tilde{X}$ is the universal cover of a certain Salvetti complex [52] (including those with defining graph an n-gon with $n \geq 5$).

The author plans to extend Huang's result to other Salvetti complexes. Of particular interest are the cases where the defining graph is a tree and where the Salvetti complex is two-dimensional. Similarly, one could consider the case where X_1 and X_2 are special $\mathcal{VH}$-complexes [44], although there is less symmetry to exploit here compared with Salvetti complexes. The techniques involved would likely be similar to those of Section 7.3, where we proved Leighton's Theorem for graphs with coloured fins, although there would probably be multiple stages of gluing. In particular, if an argument involved replacing pieces of a cube complex by finite covers before gluing them together along circular hyperplanes, then Theorem 1.2.3 would be very useful.

The author is also interested in proving combination theorems for special cube complexes in non-hyperbolic settings, extending the results of Haglund–Wise [45] and Huang–Wise [53]. Theorem 1.2.3 lends itself to proving combination theorems where the edge groups are cyclic. For more complicated edge groups one would ideally like the vertex groups to command the incident edge groups, so it would be good to prove Conjecture 1.2.6, or at least a stronger form of Theorem 5.6.4.

8.2 Action rigidity

In Section 1.3 we discussed how Leighton-type theorems are useful for proving that pairs of groups are commensurable, or that certain groups are quasi-isometrically rigid. We now introduce another relevant notion of group rigidity: a group G is *action rigid* if, for any group G', the existence of proper cocompact actions of G and G' on the same proper geodesic space X implies that G and G' are abstractly commensurable. We note that action rigidity is a weaker notion than quasi-isometric rigidity; to prove action rigidity it is enough to find a common finite/compact cover for the quotients X/G and X/G', whereas for quasi-isometric rigidity we're only given that G and G' are quasi-isometric and the first step is usually to find actions on a common space X.

Stark and Woodhouse used the Symmetry-restricted Leighton's Theorem (Theorem 1.1.5) to prove action rigidity for free products of closed hyperbolic surface groups [87], and the author is currently in a joint project with Margolis, Stark and Woodhouse to extend this to many more graphs of groups. We now describe some of the main ideas, but without giving a precise statement of what we're trying to prove. As above, we start with groups G and G' acting geometrically on a proper geodesic space X, and we wish to show that they are abstractly commensurable. Our strategy has two main steps. Firstly, we modify the space X so that it becomes a tree of simplicial complexes, but still with geometric actions of G and G'. Secondly, for each vertex space X_v, we construct a common finite cover $\hat{X}_v$ of X_v/G_v and X_v/G'_v, and then glue these all together using the Graph of Objects Leighton's Theorem (Theorem 1.1.6) to obtain a common finite cover of X/G and X/G'. For the first step, the tree decomposition comes from the JSJ decomposition of G (either over finite edge groups or two-ended edge groups). For the second step we will usually assume that G_v and G'_v have finite index in $\mathrm{Aut}(X_v)$, as then we can take $\hat{X}_v = X_v/G_v \cap G'_v$. This happens for example if X_v is a triangulated manifold. Another source of examples comes from a theorem of Trofimov [90]: if G is a group of polynomial growth that acts geometrically on a graph X, then there exists a graph Y and an equivariant quasi-isometry $X \to Y$ with associated homomorphism $\mathrm{Aut}(X) \to \mathrm{Aut}(Y)$ such that $\mathrm{Aut}(Y)$ acts properly on Y.

Looking further ahead, the author would like to work on action rigidity for hierarchies of groups, including the following fundamental conjecture.

Conjecture 8.2.1. *Every group in the hierarchy $\mathcal{H} = (\mathcal{H}_n)$ is action rigid, where $\mathcal{H}_0$ is the class of finite groups, and $\mathcal{H}_{n+1}$ consists of those hyperbolic groups that split over finite and two-ended subgroups with vertex groups in $\mathcal{H}_n$.*

References

[1] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[2] Ian Agol. Virtual properties of 3-manifolds. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. 1*, pages 141–170. Kyung Moon Sa, Seoul, 2014.

[3] Ian Agol, Daniel Groves, and Jason Fox Manning. An alternate proof of Wise's malnormal special quotient theorem. *Forum Math. Pi*, 4:e1, 54, 2016.

[4] Jitendra Bajpai. Omnipotence of surface groups. Master's thesis, McGill University, 2007.

[5] Hyman Bass. Covering theory for graphs of groups. *J. Pure Appl. Algebra*, 89(1-2):3–47, 1993.

[6] Hyman Bass and Ravi Kulkarni. Uniform tree lattices. *J. Amer. Math. Soc.*, 3(4):843–902, 1990.

[7] Jason Behrstock and Ruth Charney. Divergence and quasimorphisms of right-angled Artin groups. *Math. Ann.*, 352(2):339–356, 2012.

[8] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher. Geometry and rigidity of mapping class groups. *Geom. Topol.*, 16(2):781–888, 2012.

[9] Jason A. Behrstock and Walter D. Neumann. Quasi-isometric classification of graph manifold groups. *Duke Math. J.*, 141(2):217–240, 2008.

[10] Jason A. Behrstock and Walter D. Neumann. Quasi-isometric classification of non-geometric 3-manifold groups. *J. Reine Angew. Math.*, 669:101–120, 2012.

[11] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[12] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

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

[14] M. R. Bridson and S. M. Gersten. The optimal isoperimetric inequality for torus bundles over the circle. *Quart. J. Math. Oxford Ser. (2)*, 47(185):1–23, 1996.

[15] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[16] Martin R. Bridson, Alan W. Reid, and Henry Wilton. Profinite rigidity and surface bundles over the circle. *Bull. Lond. Math. Soc.*, 49(5):831–841, 2017.

[17] Martin R. Bridson and Sam Shepherd. Leighton's theorem: extensions, limitations, and quasitrees.

[18] Martin R. Bridson and Henry Wilton. The triviality problem for profinite completions. *Invent. Math.*, 202(2):839–874, 2015.

[19] Marc Burger and Shahar Mozes. Finitely presented simple groups and products of trees. *C. R. Acad. Sci. Paris Sér. I Math.*, 324(7):747–752, 1997.

[20] Christopher H. Cashen. Splitting line patterns in free groups. *Algebr. Geom. Topol.*, 16(2):621–673, 2016.

[21] Christopher H. Cashen and Nataša Macura. Line patterns in free groups. *Geom. Topol.*, 15(3):1419–1475, 2011.

[22] Christopher H. Cashen and Jason F. Manning. Virtual geometricity is rare. *LMS J. Comput. Math.*, 18(1):444–455, 2015.

[23] Christopher H. Cashen and Alexandre Martin. Quasi-isometries between groups with two-ended splittings. *Math. Proc. Cambridge Philos. Soc.*, 162(2):249–291, 2017.

[24] Andrew Casson and Douglas Jungreis. Convergence groups and Seifert fibered 3-manifolds. *Invent. Math.*, 118(3):441–456, 1994.

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

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

[27] Pallavi Dani, Emily Stark, and Anne Thomas. Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups. *Groups Geom. Dyn.*, 12(4):1273–1341, 2018.

[28] Andreas W. M. Dress and Rudolf Scharlau. Gated sets in metric spaces. *Aequationes Math.*, 34(1):112–120, 1987.

[29] Cornelia Druţu and Mark Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by Denis Osin and Mark Sapir.

[30] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.

[31] Eduard Einstein. Hierarchies for relatively hyperbolic virtually special groups. arXiv:1903.12284 (2019).

[32] Alex Eskin. Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces. *J. Amer. Math. Soc.*, 11(2):321–361, 1998.

[33] Benson Farb and Lee Mosher. On the asymptotic geometry of abelian-by-cyclic groups. *Acta Math.*, 184(2):145–202, 2000.

[34] David Gabai. Convergence groups are Fuchsian groups. *Ann. of Math. (2)*, 136(3):447–510, 1992.

[35] Rita Gitik. Ping-pong on negatively curved groups. *J. Algebra*, 217(1):65–72, 1999.

[36] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[37] D. Groves and J. Manning. Specializing cubulated relatively hyperbolic groups. *arXiv: Group Theory*, 2020.

[38] Vincent Guirardel and Gilbert Levitt. JSJ decompositions of groups. *Astérisque*, (395):vii+165, 2017.

[39] Mark Hagen. Lecture notes on CAT(0) cube complexes, median graphs, and cubulating groups. https://www.wescac.net/into_the_forest.pdf, 2019.

[40] Mark F. Hagen and Nicholas W.M. Touikan. Panel collapse and its applications. arXiv:1712.06553v2 (2018).

[41] Frédéric Haglund. Commensurability and separability of quasiconvex subgroups. *Algebr. Geom. Topol.*, 6:949–1024, 2006.

[42] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. pages 1–17, 2007.

[43] Frédéric Haglund and Frédéric Paulin. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 181–248. Geom. Topol. Publ., Coventry, 1998.

[44] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[45] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Ann. of Math. (2)*, 176(3):1427–1482, 2012.

[46] Wolfgang Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961.

[47] U. Hamenstädt. Geometry of the mapping class groups iii: Quasi-isometric rigidity. *arXiv: Geometric Topology*, 2005.

[48] John Hempel. Residual finiteness for 3-manifolds. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 379–396. Princeton Univ. Press, Princeton, NJ, 1987.

[49] Susan Hermiller and John Meier. Algorithms and geometry for graph products of groups. *J. Algebra*, 171(1):230–257, 1995.

[50] Tim Hsu and Daniel T. Wise. On linear and residual properties of graph products. *Michigan Math. J.*, 46(2):251–259, 1999.

[51] Tim Hsu and Daniel T. Wise. Cubulating graphs of free groups with cyclic edge groups. *Amer. J. Math.*, 132(5):1153–1188, 2010.

[52] Jingyin Huang. Commensurability of groups quasi-isometric to RAAGs. *Invent. Math.*, 213(3):1179–1247, 2018.

[53] Jingyin Huang and Daniel T. Wise. Virtual specialness of certain graphs of special cube complexes, 2019.

[54] Ilya Kapovich. The combination theorem and quasiconvexity. *Internat. J. Algebra Comput.*, 11(2):185–216, 2001.

[55] Michael Kapovich. Lectures on quasi-isometric rigidity. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 127–172. Amer. Math. Soc., Providence, RI, 2014.

[56] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.

[57] Frank Thomson Leighton. Finite common coverings of graphs. *J. Combin. Theory Ser. B*, 33(3):231–238, 1982.

[58] Gilbert Levitt. On the automorphism group of generalized baumslag–solitar groups. *Geom. Topol.*, 11(1):473–515, 2007.

[59] Gilbert Levitt. Generalized Baumslag-Solitar groups: rank and finite index subgroups. *Ann. Inst. Fourier (Grenoble)*, 65(2):725–762, 2015.

[60] William Malone. *Topics in geometric group theory*. ProQuest LLC, Ann Arbor, MI, 2010. Thesis (Ph.D.)–The University of Utah.

[61] Alex Margolis. The geometry of groups containing almost normal subgroups. *arXiv:1905.03062*, 2019.

[62] Alexander Margolis. Quasi-isometry classification of raags that split over cyclic subgroups. *arXiv: Group Theory*, 2018.

[63] Ashot Minasyan. Hereditary conjugacy separability of right-angled Artin groups and its applications. *Groups Geom. Dyn.*, 6(2):335–388, 2012.

[64] Lee Mosher, Michah Sageev, and Kevin Whyte. Quasi-actions on trees. I. Bounded valence. *Ann. of Math. (2)*, 158(1):115–164, 2003.

[65] Walter D. Neumann. On Leighton's graph covering theorem. *Groups Geom. Dyn.*, 4(4):863–872, 2010.

[66] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.

[67] Eduardo Oregón-Reyes. On cubulated relatively hyperbolic groups, 2020.

[68] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.

[69] Denis V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.

[70] Pierre Pansu. Croissance des boules et des géodésiques fermées dans les nilvariétés. *Ergodic Theory Dynam. Systems*, 3(3):415–445, 1983.

[71] P. Papasoglu. Homogeneous trees are bi-Lipschitz equivalent. *Geom. Dedicata*, 54(3):301–306, 1995.

[72] Panos Papasoglu. Quasi-isometry invariance of group splittings. *Ann. of Math. (2)*, 161(2):759–830, 2005.

[73] Panos Papasoglu and Kevin Whyte. Quasi-isometries between groups with infinitely many ends. *Comment. Math. Helv.*, 77(1):133–144, 2002.

[74] Yu. V. Prokhorov. Convergence of random processes and limit theorems in probability theory. *Teor. Veroyatnost. i Primenen.*, 1:177–238, 1956.

[75] Luis Ribes and Pavel Zalesskii. *Profinite groups*, volume 40 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]*. Springer-Verlag, Berlin, second edition, 2010.

[76] Martin Roller. Poc sets, median algebras and group actions, 2016.

[77] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[78] Michah Sageev and Daniel T. Wise. The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.*, 37(5):706–710, 2005.

[79] Richard Evan Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.*, (82):133–168 (1996), 1995.

[80] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.

[81] Jean-Pierre Serre. *Arbres, amalgames,* SL_2. Société Mathématique de France, Paris, 1977. Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46.

[82] Herman Servatius. Automorphisms of graph groups. *J. Algebra*, 126(1):34–60, 1989.

[83] Sam Shepherd. Agol's theorem on hyperbolic cubulations. arXiv:1905.06199 (2019), *Rocky Mountain J. Math.*, to appear.

[84] Sam Shepherd. Two generalisations of Leighton's Theorem. With an appendix by Giles Gardam and Daniel J. Woodhouse. arXiv:1908.00830 (2019).

[85] Sam Shepherd and Daniel J. Woodhouse. Quasi-isometric rigidity for graphs of virtually free groups with two-ended edge groups. arXiv:2007.10034 (2020).

[86] Emily Stark. Abstract commensurability and quasi-isometry classification of hyperbolic surface group amalgams. *Geom. Dedicata*, 186:39–74, 2017.

[87] Emily Stark and Daniel J. Woodhouse. Action rigidity for free products of hyperbolic manifold groups, 2019.

[88] Jacek Świątkowski. Regular path systems and (bi)automatic groups. *Geom. Dedicata*, 118:23–48, 2006.

[89] Alexander Taam and Nicholas W.M. Touikan. On the quasi-isometric rigidity of graphs of surface groups. *arXiv:1904.10482*, 2019.

[90] V. I. Trofimov. Graphs with polynomial growth. *Mat. Sb. (N.S.)*, 123(165)(3):407–421, 1984.

[91] M. van de Vel. Matching binary convexities. *Topology Appl.*, 16(3):207–235, 1983.

[92] Kevin Whyte. Amenability, bi-Lipschitz equivalence, and the von Neumann conjecture. *Duke Math. J.*, 99(1):93–112, 1999.

[93] Gareth Wilkes. Profinite rigidity of graph manifolds and JSJ decompositions of 3-manifolds. *J. Algebra*, 502:538–587, 2018.

[94] Gareth Wilkes. Profinite rigidity of graph manifolds, II: Knots and mapping classes. *Israel J. Math.*, 233(1):351–378, 2019.

[95] Henry Wilton. Virtual retractions, conjugacy separability and omnipotence. *J. Algebra*, 323(2):323–335, 2010.

[96] Daniel T. Wise. The structure of groups with a quasiconvex hierarchy.

[97] Daniel T. Wise. Subgroup separability of graphs of free groups with cyclic edge groups. *Q. J. Math.*, 51(1):107–129, 2000.

[98] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.

[99] Daniel T. Wise and Daniel J. Woodhouse. A cubical flat torus theorem and the bounded packing property. *Israel J. Math.*, 217(1):263–281, 2017.

[100] D.T. Wise. *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups.* PhD thesis, Princeton University, 1996.

[101] Daniel J. Woodhouse. Revisiting leighton's theorem with the haar measure. arXiv:1712.06553v3 (2019).

[102] Daniel J. Woodhouse. A generalized axis theorem for cube complexes. *Algebr. Geom. Topol.*, 17(5):2737–2751, 2017.

[103] P. A. Zalesskiĭ and O. I. Tavgen'. Closure of orbits and residual finiteness with respect to the conjugacy of free amalgamated products. *Mat. Zametki*, 58(4):525–535, 639, 1995.