

Leighton's Theorem: extensions, limitations, and quasitrees

Martin R. Bridson and Sam Shepherd *

February 2, 2021

Abstract

Leighton's Theorem states that if there is a tree T that covers two finite graphs G_1 and G_2 , then there is a finite graph \hat{G} that is covered by T and covers both G_1 and G_2 . We prove that this result does not extend to regular covers by graphs other than trees. Nor does it extend to non-regular covers by a quasitree, even if the automorphism group of the quasitree contains a uniform lattice. But it does extend to regular coverings by quasitrees.

1 Introduction

Leighton's Theorem [18] states that if a pair of finite graphs G_1, G_2 have a common covering space, then they have a common finite covering. This serves as a prototype for a type of theorem that one might hope to prove in many different categories. In particular, for different categories of (orbi)spaces one might hope to prove that if two compact spaces have a common covering of a specified kind, then there is a common intermediate covering that is compact. The search for such theorems is related to the study of rigidity, particularly quasi-isometric rigidity: if one is in a situation where one can promote the existence of a quasi-isometry of groups $\Gamma_1 \sim \Gamma_2$ into the existence of proper cocompact actions on a common space X , then proving that Γ_1 and Γ_2 are abstractly commensurable is equivalent to proving that the orbispaces X/Γ_1 and X/Γ_2 have a common finite covering. Beyond quasi-isometric rigidity, the study of Leighton-type theorems fits into a rich and active area of geometry and group theory: in geometry, particularly low-dimensional topology, understanding the existence and nature of finite-sheeted coverings has been a dominant theme in recent years (e.g. [16], [1]); in parallel, there have been significant advances in group theory concerning the difficulty of recognising finite-index subgroups and the extent to which groups of geometric interest are determined by their lattice of finite-index subgroups (e.g. [8], [9]).

Against this background, the second author and others have extended Leighton's Theorem in various ways [17], [20], [27], [28], [33] with applications to rigidity [5], [29], [30]. Our main purpose in this article is to understand the extent to which Leighton's Theorem applies to coverings of graphs by quasitrees, i.e. locally finite graphs that are quasi-isometric to trees. We shall see that this is a subtle question, the answer to which depends crucially on whether the coverings considered are regular or not and whether the automorphism group of the quasitree in question contains uniform lattices.

2010 *Mathematics subject classification*. 20F65, 20F67, 05C25

Key words and phrases. Quasitrees, covering spaces, Leighton's Theorem.

*Shepherd was funded by the Engineering and Physical Sciences Research Council.

Moving beyond quasitrees and virtually free groups, we shall also describe situations where finite graphs with a common regular covering do not share a common intermediate covering because one of the Galois groups involved does not have sufficiently many (or indeed any) finite-index subgroups (Theorem 1.1).

In the classical setting, it is easy to reduce Leighton's Theorem to the case where the graphs G_1 and G_2 are connected. Standard arguments about covering spaces show that in this setting the conclusion of the theorem is that if the universal coverings of G_1 and G_2 are isomorphic to the same tree T , then after composing with automorphisms of T , the covering maps $T \rightarrow G_i$ both factor through a covering $T \rightarrow \hat{G}$ where \hat{G} is a finite graph. In particular, there is a diagram of covering maps

$$\begin{array}{ccc} & T & \\ & \downarrow & \\ & \hat{G} & \\ \swarrow & & \searrow \\ G_1 & & G_2 \end{array} \tag{1.1}$$

Moreover, one can arrange for all of these coverings to be regular: this is a useful refinement made explicit in Bass and Kuhlmann's proof [3]. If one thinks in terms of actions rather than quotient spaces, then Leighton's Theorem is the torsion-free case of the following theorem: if $\Gamma_1, \Gamma_2 < \text{Aut}(T)$ act properly and cocompactly on T , then there exists $g \in \text{Aut}(T)$ such that $\Gamma_1^g := g^{-1}\Gamma_1g$ and Γ_2 are commensurable, i.e. $\Gamma_1^g \cap \Gamma_2$ has finite index in both Γ_1^g and Γ_2 . We shall rely on a strengthening of this formulation proved in [27] that constrains the conjugator g .

Our first result shows that one cannot extend Leighton's Theorem to arbitrary regular coverings $G \rightarrow G_i$ of finite connected graphs.

Theorem 1.1. *There exist connected graphs G, G_1, G_2 and regular covering maps $G \rightarrow G_1$ and $G \rightarrow G_2$, such that G_1 and G_2 are finite but there does not exist a covering $G \rightarrow \hat{G}$ of any finite graph \hat{G} such that \hat{G} covers both G_1 and G_2 .*

Theorem 1.1 illustrates the need to impose conditions on the common covering of G_1 and G_2 as one moves away from trees. Our proof of this theorem encodes structures from another setting in which the natural analogue of Leighton's Theorem fails, namely that of compact, non-positively curved squared complexes. Key examples in this setting were constructed by Wise [31] and Burger-Mozes [10].

Henceforth, we shall restrict our attention to the case where the common covering of G_1 and G_2 is quasi-isometric to a tree, i.e. is a *quasitree*. In the positive direction, relying on the Symmetry-Restricted Leighton Theorem from [27], we shall prove the following result, for which we need to establish some notation. For a locally finite graph X , we write $\text{Aut}(X)$ to denote the group of automorphisms of X (i.e. the isometries that preserve the vertex set). We equip $\text{Aut}(X)$ with the usual (compact-open) topology, making it a locally-compact Hausdorff group. A *uniform lattice* $\Gamma < \text{Aut}(X)$ is a discrete subgroup that acts properly and cocompactly on X . For $H < \text{Aut}(X)$, the set of automorphisms that coincide on balls of radius R with elements of H will be denoted by $\mathcal{S}_R(H, X)$; see definition 3.4.

Theorem 1.2. *Let X be a locally finite quasitree, let $H < \text{Aut}(X)$, and let $\Gamma_1, \Gamma_2 < H$ be uniform lattices in $\text{Aut}(X)$. Then for all $R \in \mathbb{N}$ there exists $g \in \mathcal{S}_R(H, X)$ such that Γ_1^g is commensurable to Γ_2 in $\text{Aut}(X)$.*

As a corollary of Theorem 1.2, we deduce that any two finite simplicial complexes with free fundamental groups and a common universal cover must have a common finite cover (Corollary 3.6).

Let us concentrate now on the case where $H = \text{Aut}(X)$ and Γ_1 and Γ_2 are torsion-free. In this case Theorem 1.2 tells us that, given finite graphs X_1 and X_2 , if there is a quasitree X and *regular* coverings $X \rightarrow X_i$, then there is a covering $X \rightarrow \hat{X}$ of a finite graph \hat{X} such that \hat{X} covers both X_1 and X_2 . However, our next result shows that this does not remain true if one drops the requirement that both of the coverings $X \rightarrow X_i$ are regular, even if one retains the condition that $\text{Aut}(X)$ contains a uniform lattice (which implies that $\text{Aut}(X)$ is unimodular). Note that, just like Theorem 1.1, this theorem excludes the existence of arbitrary coverings $X \rightarrow \hat{X} \rightarrow X_i$, not just those for which the compositions are the original coverings $X \rightarrow X_i$.

Theorem 1.3. *There exists a quasitree X , finite graphs X_1 and X_2 , and covering maps $X \rightarrow X_1$ and $X \rightarrow X_2$, with $X \rightarrow X_2$ regular, for which there is no finite graph \hat{X} fitting into the following diagram of covers.*

$$\begin{array}{ccc} & X & \\ & \downarrow & \\ & \hat{X} & \\ \swarrow & & \searrow \\ X_1 & & X_2 \end{array} \tag{1.2}$$

In Section 4 we shall prove a version of Theorem 1.3 in which the regularity requirement on $X \rightarrow X_2$ is replaced by the weaker requirement that $\text{Aut}(X)$ acts cocompactly on X , which serves to fix the main ideas in the proof and expose the key phenomena. In Section 5 we develop the technical refinements needed to prove Theorem 1.3.

We have already alluded to the importance of the fact that when a tree covers a finite graph, the covering is regular. When one moves away from trees, regularity is a more subtle issue. We explore this theme in the final section of this paper, where we prove the following theorem, in which $\mathcal{C}(X, Y)$ denotes the compact space of local-isometries $X \rightarrow Y$ that are covering maps.

Theorem 1.4. *Let X be a locally finite quasitree such that $\text{Aut}(X)$ contains a uniform lattice. Suppose that X covers a finite graph Y . Then there exists an $\text{Aut}(X)$ -invariant probability measure on $\mathcal{C}(X, Y)$ if and only if there exists a finite cover $\hat{Y} \rightarrow Y$ that admits a regular covering map $X \rightarrow \hat{Y}$.*

The remainder of this paper is structured in accordance with the foregoing discussion. In Section 2 we prove Theorem 1.1. In Section 3 we prove Theorem 1.2; we do so by reducing to the case where X is a tree, which is the Symmetry-Restricted Leighton Theorem proved by the second author and Gardam-Woodhouse in [27]. This reduction is based on the existence of an $\text{Aut}(X)$ -equivariant quasi-isometry from X to a locally finite tree (Theorem 3.1), which is established using elements of the theory of CAT(0) cube complexes. (An alternative proof can be based on [19]; see Remark 3.3). We then prove Theorem 1.3 in two steps over Sections 4 and 5, as described above, before closing with a proof of Theorem 1.4 in Section 6.

We thank the referee for their careful reading and helpful comments.

2 The need to control the geometry of the common cover

We begin with a weak form of Theorem 1.1 that exemplifies the main idea with very simple graphs G_1, G_2 . This result does not correspond to diagram (1.1) but rather to the stronger conclusion of Leighton's Theorem that, given coverings $p_i : T \rightarrow G_i$ of finite graphs G_i by a tree T , there is an automorphism g of T and a commutative diagram of coverings with \hat{G} finite

$$\begin{array}{ccccc}
 T & \xrightarrow{g} & T & & \\
 \downarrow p_1 & \searrow \tau_1 & \swarrow \tau_2 & \downarrow p_2 & \\
 & \hat{G} & & & \\
 \swarrow \lambda_1 & & \searrow \lambda_2 & & \\
 G_1 & & G_2 & &
 \end{array} \tag{2.1}$$

Proposition 2.1. *There exist connected graphs G, G_1, G_2 and covering maps $p_i : G \rightarrow G_i$ ($i = 1, 2$) such that G_1 and G_2 are finite but there does **not** exist a finite graph \hat{G} with covering maps $\lambda_i : \hat{G} \rightarrow G_i$ ($i = 1, 2$) and $\tau_i : G \rightarrow \hat{G}$ such that $p_i = \lambda_i \circ \tau_i$ ($i = 1, 2$).*

Proof. Let $\Pi = F_k/N$ be an infinite group that has no proper subgroups of finite index, where F_k is the free group on $\{a_1, \dots, a_k\}$. Let G_1 be the k -petal rose – i.e. the graph with one vertex and $E(G_1) = \{a_1, \dots, a_k\}$. Let G be the Cayley graph of Π with respect to the generators a_i and consider the regular covering $p_1 : G \rightarrow G_1$ with Galois group Π . The intermediate covers $G \xrightarrow{\tau_1} \hat{G} \xrightarrow{\lambda_1} G_1$ correspond to the subgroups $H < F_k$ that contain N , and $\hat{G} \rightarrow G_1$ is a finite covering if and only if H has finite index in F_k . By assumption, there are no proper subgroups of finite index in F_k/N and therefore no proper subgroups of finite index in F_k that contain N . Thus $\hat{G} = G_1$ if \hat{G} is finite.

This general construction will provide us with examples as described in the proposition provided that we can exhibit pairs of groups Π_1 and Π_2 with the following properties:

1. there are finite generating sets \mathcal{A}_i for Π_i (of the same cardinality k) such that the corresponding Cayley graphs $\mathcal{C}(\Pi_i, \mathcal{A}_i)$ (considered as unlabelled graphs) are isomorphic;
2. Π_2 has a proper subgroup of finite index, but Π_1 does not.

With such a pair of groups in hand, we can complete the proof by arguing as follows. Take G_1 to be the k -petalled rose with covering $p_1 : \mathcal{C}(\Pi_1, \mathcal{A}_1) \rightarrow G_1$ and take G_2^\dagger to be the k -petalled rose with covering $p_2^\dagger : \mathcal{C}(\Pi_2, \mathcal{A}_2) \rightarrow G_2^\dagger$. Then take G_2 to be a proper finite-sheeted cover of G_2^\dagger corresponding to a proper finite-index subgroup of F_k that contains $\ker(F_k \rightarrow \Pi_2)$ and lift p_2^\dagger to $p_2 : \mathcal{C}(\Pi_2, \mathcal{A}_2) \rightarrow G_2$. G_2 is larger than G_1 , so any \hat{G} as in the proposition must be a proper cover of G_1 , which would contradict the fact that Π_1 has no proper finite-index subgroups.

A rich source of examples (Π_1, Π_2) with the required properties are the BMW groups, which act freely on a product of two regular trees, with a single orbit of vertices; the 1-skeleton of the product of trees will be the Cayley graph of Π_i for a suitable choice of generators. The direct product of two finitely generated free groups is such a group and this will serve us as Π_2 . But there are many other fascinating examples – see [11] for a recent survey. Of particular note is the example constructed by Wise in his thesis [31] (see below), which has no proper subgroups of finite index, and the examples of Burger and Mozes [10], which are simple. \square

Example 2.2. The following group was constructed by Wise [31]; it has no proper subgroups of finite index. The universal cover of the standard 2-complex of this presentation is the product of a regular tree of valence 4 and a regular tree of valence 6. Thus the Cayley graph (as an unlabelled graph) is isomorphic to the standard Cayley graph of $F_2 \times F_3$.

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

2.1 Terminology for graphs, #-subdivision and square-completion

For the most part, we regard graphs as geometric objects. We write $V(X)$ for the vertex set of a graph X and $E(X)$ for its set of (closed, unoriented) edges, and we metrize X as a geodesic metric space in which each edge has length 1. But it is also convenient to refer to the combinatorial structure of X ; for example, we say that X is *finite* if $V(X)$ and $E(X)$ are finite. To work with edge-paths, it is convenient to replace $E(X)$ by the set of oriented edges $E^\pm(X)$, which for each $e \in E(X)$ contains the pair of local isometries $e^\pm : [0, 1] \rightarrow X$ with image e and $e^+(t) = e^-(1 - t)$. If $\varepsilon = e^\pm$ then $\bar{\varepsilon} := e^\mp$. A *reduced circuit* in X is the loop determined by a finite cyclically-ordered set $(\varepsilon_1, \dots, \varepsilon_n)$ with $\varepsilon_i \in E^\pm(X)$ and $\varepsilon_i(1) = \varepsilon_{i+1}(0)$, $\varepsilon_{i+1} \neq \bar{\varepsilon}_i$ for $i = 1, \dots, n$, indices mod n . (We specify *cyclic* ordering so that $(\varepsilon_1, \dots, \varepsilon_n) = (\varepsilon_2, \dots, \varepsilon_n, \varepsilon_1)$ etc.) If the vertices $\varepsilon_i(1)$ are all distinct, then the subgraph consisting of these vertices and the unoriented edges corresponding to the ε_i is called an *n-cycle*.

Definition 2.3. The *square completion* $\square(G)$ of a graph G is the combinatorial 2-complex obtained from G by attaching a 2-cell to each reduced circuit of length 4.

Definition 2.4. 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. $K_\#$ is metrized as a piecewise-Euclidean complex where each (new) edge has length 1 and each (new) 2-cell is a square.

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. We use the standard notation $K^{(1)}$ for the 1-skeleton of a combinatorial complex K .

Lemma 2.5. *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 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 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.

We stated the preceding lemma for non-positively curved complexes because that is our focus in this section, but the proof shows that it is enough to assume that K is *simple* in the sense that links of vertices in K are simplicial graphs.

2.2 Proof of Theorem 1.1

In the proof of Proposition 2.1 we used the requirement that $p_i = \lambda_i \circ \tau_i$ to restrict our attention to intermediate covers of $p_i : G \rightarrow G_i$, but that is not good enough now because we do not require that the leftmost and rightmost triangles in the analogue of diagram (2.1) commute. We shall remedy this by describing a more complicated graph G_1 which has the property that no finite graph can properly cover G_1 if it is itself covered by G .

To this end, we consider the quotient K of a product of regular trees (the universal cover \tilde{K} of K) by a BMW group; to be definite we take K to be the standard 2-complex of the presentation in Example 2.2. We subdivide K to form $K_{\#}$, and subdivide its universal cover $\tilde{K}_{\#}$ in the same canonical manner. Lemma 2.5 assures us that $K_{\#}$ and $\tilde{K}_{\#}$ can be recovered from their 1-skeleta by taking the square-completion.

The vertices of $G_1 := K_{\#}^{(1)}$ 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. A key point about the use of the integer 5 in the definition of $K_{\#}$ is that the only reduced circuits of length 4 in G_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 $G := \tilde{K}_{\#}^{(1)}$.

Let \hat{G} be a graph and let $\pi : \hat{G} \rightarrow G_1$ be a covering. As G_1 contains no reduced circuits of length less than 4, neither does \hat{G} . So the number of 4-cycles passing through a vertex $v \in \hat{G}$ 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 $q : G \rightarrow \hat{G}$ is a covering map then the number of reduced circuits of length 4 based at $w \in V(G)$ 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 $G \rightarrow \hat{G} \rightarrow G$ extend to coverings of combinatorial 2-complexes $\tilde{K}_{\#} = \square(G) \rightarrow \square(\hat{G}) \rightarrow \square(G_1) = K_{\#}$. It follows that if \hat{G} is finite, then $\hat{G} = G_1$, since $K_{\#}$ has no proper finite coverings.

To complete the proof, it suffices to take G_2 to be the 1-skeleton of the quotient Q of $\tilde{K}_{\#}$ by a

direct product of free groups, acting freely and transitively on the vertex set of \tilde{K} . Lemma 2.5 tells us that $Q = \square(G_2)$, so in particular $G_1 \neq G_2$. As G_1 has only one vertex of valence 10, we see that G_1 does not cover any smaller graph, hence G_1 does not cover G_2 . Any finite graph \hat{G} that admits coverings $G \rightarrow \hat{G}$ and $\hat{G} \rightarrow G_1, G_2$ must equal G_1 as we explained above, which then contradicts the existence of the covering $\hat{G} \rightarrow G_2$, so such a \hat{G} cannot exist. \square

3 Leighton's Theorem for coverings by quasitrees

As we explained in the introduction, our strategy for proving Theorem 1.2 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 [19] (see Remark 3.3). In order to understand the proof that we shall present, the reader should be familiar with the basic properties of CAT(0) cube complexes [7], in particular the definition and basic properties of hyperplanes and the duality between cube complexes and spaces with walls. Key ideas in this theory originate in the work of Gromov [13], Sageev [24] and Haglund-Paulin [15]. There are several variants on the basic construction, adapted to different settings (e.g. Roller [23], Nica [21], Chatterji-Niblo [12]). See [6] and [25] for surveys. Wall spaces and group actions on cube complexes play central roles in major recent advances in group theory and low-dimensional topology (most dramatically [1]), following a programme of Wise [32].

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 $V(X)$ of a quasitree) into two disjoint subspaces (half-spaces); the wall is said to *separate* $x, y \in V(X)$ 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 $\{H, H'\}$ is a wall, then exactly one of H, H' lies in ω , and if $\{H_1, H'_1\}$ and $\{H_2, H'_2\}$ are walls with $H_1 \subset H_2$ and $H_1 \in \omega$ then $H_2 \in \omega$. 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 $V(X)$ 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 3.1. *There is a canonical process that, given a locally finite quasitree X on which $G = \text{Aut}(X)$ acts cocompactly, will construct a locally finite tree T with an action of G on T and a continuous G -equivariant quasi-isometry $f : X \rightarrow T$ that restricts to an injection $V(X) \hookrightarrow V(T)$.*

Proof. The idea of the proof is to define walls on X that enable us embed it in a CAT(0) cube complex Ψ , then use the panel-collapse procedure of Hagen and Touikan [14] to canonically retract Ψ onto a tree.

To this end, we define a constant C_X by considering all continuous quasi-isometries $h : X \rightarrow Y$

from X to locally finite trees Y that send vertices to vertices, and set

$$C_X := 3 + \inf_{h: X \rightarrow Y} \sup_{e \in E(Y)} \text{diam}(h^{-1}(e))$$

For each set of edges $S \subset E(X)$, we let \hat{S} denote the set of midpoints, and we equip the vertex set $V(X)$ with the structure of a wall space by defining \mathcal{W} to be the set of all \hat{S} such that $X \setminus \hat{S}$ 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 $V(X)$ into the intersections of $V(X)$ with the two connected components of $X \setminus \hat{S}$.) It is clear that any two points of $V(X)$ are separated by only finitely many walls. 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 $\theta : V(X) \rightarrow \Psi$ is injective.

To verify this basic property, we must argue that each pair of distinct vertices $x_1, x_2 \in V(X)$ 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 $\hat{S}' \in \mathcal{W}$. This proves the claim.

The canonical nature of the construction ensures that the action of $G = \text{Aut}(X)$ on $V(X)$ 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: $\theta : V(X) \hookrightarrow \Psi$ is a quasi-isometry with respect to the combinatorial metric on Ψ .

As G acts cocompactly on Ψ , every point is within a bounded distance of the image of θ . By definition, for $x_1, x_2 \in V(X)$, 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 $d(x_1, x_2)$.

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

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

for all $x_1, x_2 \in V(X)$, where $K > 1$ is constant. For each $e \in E(Y)$, the union of the edges in $S(e) := \{e' \in E(X) \mid h(e') \cap e \neq \emptyset\}$ has diameter less than C_X (if non-empty). Moreover, if $X \setminus \widehat{S(e)}$ 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 $\text{diam}(h(e)) \leq 2K$ for all edges $e \in E(X)$, the sets $S(e_1)$ and $S(e_2)$ are disjoint for $e_1, e_2 \in E(Y)$ 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 .

To obtain an upper bound on the number of walls separating x_1 and x_2 , we 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 $S \subset E(X)$

of diameter less than C_X that contain any given edge $e \in E(X)$. 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 .

Claim: Ψ is locally finite.

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 $(\hat{S}, x) \in \mathcal{W} \times V(X)$ 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 [14] 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 [14], but it is obvious from the proof. By composing the inverse of this bijection with $\theta : V(X) \hookrightarrow \Psi$ we get a G -equivariant quasi-isometry $f : V(X) \hookrightarrow T$, and by extending linearly along edges we obtain a continuous G -equivariant quasi-isometry $f : X \rightarrow T$ that proves the theorem. \square

Remark 3.2. 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 3.1 will produce a G -equivariant map $X \rightarrow T$ 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 $\Omega < \text{Aut}(X)$ that have finite support.

Remark 3.3. Theorem 3.1 is reminiscent of the work of Mosher, Sageev and Whyte [19] 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 3.1 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 $h : X \rightarrow Y$ a quasi-isometry to a tree, we modify Y

to make it bushy as follows. (The choice of $h : X \rightarrow Y$ 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 $Y \rightarrow Y'$, we obtain a quasi-isometry from X to a bushy tree, and [19, Theorem 1] promotes the resulting quasi-action of $G = \text{Aut}(X)$ 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 $f : X \rightarrow T$ such that $d(f(gx), gf(x))$ is uniformly bounded independent of $g \in \text{Aut}(X)$ 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 $V(X)$ 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 $f' : X \rightarrow T$.

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

3.1 Proof of Theorem 1.2

In what follows $B_R(x)$ will denote the R -ball centred on a vertex x of a graph.

Definition 3.4. Let X be a graph, let $H < \text{Aut}(X)$, and let R be an integer. We define the *R -symmetry restricted closure* of H to be:

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

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

We recall the statement of Theorem 1.2 for the convenience of the reader. In the case where X is a tree, this is the *Symmetry-Restricted Leighton Theorem* proved in [27], 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 3.5. *Let X be a locally finite quasitree graph, let $H < \text{Aut}(X)$, and let $\Gamma_1, \Gamma_2 < H$ be uniform lattices in $\text{Aut}(X)$. Then for all $R \in \mathbb{N}$ there exists $g \in \mathcal{S}_R(H, X)$ such that Γ_1^g is commensurable to Γ_2 in $\text{Aut}(X)$.*

Proof. By subdividing we may assume that X is simplicial. By Theorem 3.1 there exists a locally finite tree T , an injective homomorphism $\iota : \text{Aut}(X) \rightarrow \text{Aut}(T)$ and a continuous ι -equivariant quasi-isometry $f : X \rightarrow T$ which is injective on $V(X)$. 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 $\text{Aut}(T)$. Let M be large enough so that for all $x \in V(X)$

$$f(B_R(x)) \subset B_M(f(x)). \quad (3.1)$$

We then have a homomorphism

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

satisfying $\alpha \circ \iota = \text{id}$ on restriction to $\mathcal{S}_R(H, X)$, defined as follows. Take $\phi \in \mathcal{S}_M(\iota(H), T)$. For $x \in V(X)$ there exists $h \in H$ such that $\iota(h)$ agrees with ϕ on $B_M(f(x))$. By ι -equivariance of f and (3.1), we get the following commutative diagram.

$$\begin{array}{ccc} B_R(x) & \xrightarrow{h} & X \\ \downarrow f & & \downarrow f \\ B_M(f(x)) & \xrightarrow{\iota(h)} & T \end{array} \quad (3.2)$$

$\curvearrowright \phi$

We define the restriction of $\alpha(\phi)$ to $B_R(x)$ to agree with h . This gives a well-defined description of an automorphism $\alpha(\phi) \in \text{Aut}(X)$, because for each vertex $y \in B_R(x)$ the diagram (3.2) 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 $V(X)$. And since $\alpha(\phi)$ agrees on R -balls with elements of H , we get that $\alpha(\phi) \in \mathcal{S}_R(H, X)$. It is also clear that $\alpha \circ \iota = \text{id}$ on restriction to $\mathcal{S}_R(H, X)$.

We know from [27] that the theorem is true when $X = T$, so there exists $g \in \mathcal{S}_M(\iota(H), T)$ such that $\iota(\Gamma_1)^g$ is commensurable to $\iota(\Gamma_2)$ in $\text{Aut}(T)$. Therefore $\alpha \circ \iota(\Gamma_1)^{\alpha(g)} = \Gamma_1^{\alpha(g)}$ is commensurable to $\alpha \circ \iota(\Gamma_2) = \Gamma_2$ in $\text{Aut}(X)$. \square

As a consequence of Theorem 3.5, we obtain a version of Leighton's Theorem for simplicial complexes with free fundamental group.

Corollary 3.6. *Let X_1 and X_2 be finite simplicial complexes with 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 $\Gamma_1, \Gamma_2 < \text{Aut}(\tilde{X})$ be the corresponding deck-groups. To obtain a common finite cover of X_1 and X_2 , we must find $g \in \text{Aut}(\tilde{X})$ such that Γ_1^g is commensurable to Γ_2 . Since Γ_1 and Γ_2 are 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 $\text{Aut}(Y) \cong \text{Aut}(\tilde{X})$ of topological groups, and so the result follows by applying Theorem 3.5 to Y . \square

4 Leighton's Theorem fails for cocompact quasitrees

Our purpose in this section is to prove the following special case of Theorem 1.3. Our reasons for proving this special case first were discussed in the introduction. A feature of the construction in this

proof is that the quasitree X is cocompact (i.e. there are only finitely many $\text{Aut}(X)$ -orbits of edges and vertices) but $\text{Aut}(X)$ does not contain a uniform lattice. The proof ends with a counting argument that draws inspiration from the following elementary argument based on [3] Example 4.12(1).

Example 4.1. Let T be the simplicial tree whose vertex set $V(T)$ 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 $\text{Aut}(T)$ 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 4.2. *There exists a cocompact quasitree graph X and finite graphs X_1 and X_2 such that X covers X_1 and X_2 , but there does not exist a finite graph \hat{X} fitting into the following diagram of covers.*

$$\begin{array}{ccc}
 & X & \\
 & \downarrow & \\
 & \hat{X} & \\
 \swarrow & & \searrow \\
 X_1 & & X_2
 \end{array} \tag{4.1}$$

Proof. We will build X_1 and X_2 by augmenting a pair of finite graphs Y_1 and Y_2 for which there is a finite graph Y and covering maps $p_i : Y \rightarrow Y_i$ for $i = 1, 2$. We will then build X by assembling infinitely many copies of Y in a tree-like manner. The following figure portrays the covering maps p_i , which are determined by giving three labels at each vertex of Y_1 and Y_2 describing the three preimages of that vertex in Y .




Figure 2: The coverings $p_i : Y \rightarrow Y_i$.

The graphs X_1 and X_2 are built from Y_1 and Y_2 by attaching both ends of a new subdivided edge to each vertex (i.e. each vertex of Y_i has two edges emanating from it and they meet at a valence-2 vertex of $X_i \setminus Y_i$). For ease of depiction in figures we will draw these subdivided edges as dashed edges, and refer to them as dashed edges for the rest of the proof. We will refer to the other edges as solid edges. X is constructed by joining copies of the graph Y by dashed edges in a tree-like structure as suggested by the following figure.




Figure 3: Tree-like structure of X .

In more detail, we take an 18-valent regular tree T , index copies of Y by the vertices of this tree, and at each vertex of each copy of Y we add two dashed edges connecting to neighbouring copies of Y . To complete the definition of X (up to graph isomorphism) we must specify the labels of the vertices at the endpoints of each dashed edge. We do this so that for each bi-infinite line of dashed edges in X , the labels of the vertices will have one of the following six repeating patterns:

$$\begin{aligned}
 &1, 1, 6, 1, 1, 6 & (4.2) \\
 &2a, 2a, 2a, 2a, 2a, 2a \\
 &2b, 2b, 2b, 2b, 2b, 2b \\
 &3a, 4a, 3a, 4a, 3a, 4a \\
 &3b, 4b, 3b, 4b, 3b, 4b \\
 &5, 5, 5, 5, 5, 5
 \end{aligned}$$

Working out radially from a basepoint $* \in T$, we attach the copies of Y in the next sphere by dashed edges so as to continue the 6 specified patterns in the forwards and backwards directions. No pair of copies of Y is joined by more than one edge. At the first stage, the copy Y_* of Y at $*$ has its vertex labelled 1 attached to an adjacent copy of Y by a dashed edge ending at a vertex with label 1, and to another adjacent copy with an edge ending at a vertex labelled 6; the vertex of Y_* labelled $2a$ is connected to vertices with the same label in two other neighbouring copies of Y ; likewise the vertices with label $2b$ and 5 ; while the vertex of Y_* with label $3a$ is connected to two neighbours at their vertices labelled $4a$, and the vertex of Y_* with label $3b$ is connected to the two remaining neighbours at their vertices labelled $4b$. At subsequent stages, as one works away from the basepoint $* \in T$, at $v \in V(T)$ one first deals with the vertex of Y_v that is already connected to a copy of Y one-step closer to $*$ than v , connecting that vertex to a copy of Y in the next generation by a dashed edge ending at a vertex

that continues the pattern of the word labelling the dashed path from $*$ to v ; the remaining vertices of Y_v are then connected to the remaining 16 neighbouring copies of Y by the same prescription used at the basepoint.

It follows from this constructive description of X that any copy of Y in X can be sent to any other copy by a suitable automorphism of X : send the first copy of Y to the second by the unique label-preserving map, then proceed radially outwards through the template tree T to extend the map to dashed edges and neighbouring copies of Y in the unique label-preserving manner. In particular the quasitree X is cocompact.

Our next task is to describe covering maps from X to X_1 and X_2 . To this end, observe that each of our six repeating patterns of labels (4.2) stays within the set of labels of one of the vertices of Y_1 . This enables us to define a covering $X \rightarrow X_1$ by requiring that it restrict to $p_1 : Y \rightarrow Y_1$ on each copy of Y in X . (There is an arbitrary choice to be made in orientation of how each dashed edge then maps to X_1 , but this is of no consequence.) In the case of Y_2 , the middle vertex has all the labels for the second, third and sixth patterns in (4.2), but the other three patterns have labels that switch between the top and bottom vertices of Y_2 , so we have to twist p_2 appropriately to account for this. We fix an automorphism $\tau : Y_2 \rightarrow Y_2$ of order 2 that interchanges the top and bottom vertices, then define a covering $X \rightarrow X_2$ that restricts to either $p_2 : Y \rightarrow Y_2$ or $\tau p_2 : Y \rightarrow Y_2$ on each copy of Y in X ; to determine whether p_2 or τp_2 is applied to a particular copy of Y , we again work out radially from the basepoint $*$ in the template tree T ; first we map Y_* by p_2 ; subsequently, each copy Y_λ of Y is attached to a unique copy $Y_{\lambda'}$ of Y closer to $*$, and the restriction $\pi_\lambda : Y_\lambda \rightarrow Y_2$ of $X \rightarrow X_2$ is defined to be $\pi_{\lambda'}$ or $\tau \pi_{\lambda'}$ according to whether the vertices of Y_λ and $Y_{\lambda'}$ that are connected by a dashed edge have labels at the same vertex of Y_2 or not. Having defined $X \rightarrow X_2$ on the disjoint union of the copies of Y , we again complete the description of the covering map by arbitrarily choosing an orientation for the image of each dashed edge.

To complete the proof of the theorem, we must argue that there is no finite graph \hat{X} that fits into a diagram of covers (4.1). If \hat{X} did exist, it would consist of solid and dashed edges in the same way as X , and each connected component \hat{Y} of the solid subgraph in \hat{X} would fit into the following diagram of covers:

$$\begin{array}{ccc}
 & Y & \\
 & \downarrow & \\
 & \hat{Y} & \\
 \swarrow & & \searrow \\
 Y_1 & & Y_2
 \end{array} \tag{4.3}$$

But the covers from Y to Y_1 and Y_2 are degree three, and Y_1 is not isomorphic to Y_2 , so this forces $Y \rightarrow \hat{Y}$ to be an isomorphism. Thus \hat{X} could be constructed from copies of Y connected by dashed edges. One can easily verify that the numbers in the labels of vertices of Y correspond to orbits of $\text{Aut}(Y)$ (for example the vertex 1 is fixed by any automorphism because it is the base of the only loop of length one). It follows that the vertices of \hat{X} inherit a labelling by numbers from Y (ignoring the letter suffix of labels) such that the putative covering $X \rightarrow \hat{X}$ preserves the numbering of vertices. Therefore, the bi-infinite lines of dashed edges in X , which have itineraries described by the patterns (4.2) must cover circles of dashed edges in \hat{X} whose itineraries (cyclic orderings of numerical vertex labels) follow the same pattern. This leads to a contradiction as follows: each vertex of \hat{X} labelled 1 or 6 lies on a unique dashed circle (since there is a unique dashed line through each of its pre-images in

X), and the itinerary of this dashed circle is a repetition of the cyclic pattern 1, 1, 6, 1, 1, 6; this forces \hat{X} to contain twice as many vertices labelled 1 versus vertices labelled 6; this contradicts the fact that each copy of $Y = \hat{Y}$ in \hat{X} has one vertex labelled 1 and one labelled 6. \square

Remark 4.3. The above proof in fact shows that $\text{Aut}(X)$ does not contain a cocompact lattice; equivalently, there is no *regular* covering $X \rightarrow Z$ with Z a finite graph. In more detail, what the argument shows is that X does not cover any finite graph constructed from (solid) copies of Y by adding dashed circles so that one such circle passes through each vertex. Any uniform lattice in $\text{Aut}(X)$ would be virtually free, and the quotient of X by a free subgroup Λ of finite index would be a finite graph assembled from copies of Y in this manner. Indeed the action of $\text{Aut}(X)$ permutes the obvious copies of Y , and the non-trivial elements of Λ move each copy of Y off itself.

5 Leighton's Theorem can fail for quasitrees even if $\text{Aut}(X)$ contains a uniform lattice

Theorem 1.2 concerns regular coverings by a quasitree, while in the contrasting result Theorem 4.2 the coverings are not regular and the automorphism group of the quasitree does not contain a uniform lattice (see Remark 4.3). This leaves open the question of whether one can exploit the irregularity of covers to prove a version of Theorem 4.2 in which the automorphism group of the quasitree does contain a uniform lattice. Our purpose in this section is to settle this question by proving Theorem 1.3 from the introduction. We repeat the statement for the reader's convenience.

Theorem 5.1. *There exists a locally finite quasitree X , finite graphs X_1 and X_2 , and covering maps $X \rightarrow X_1$ and $X \rightarrow X_2$, with $X \rightarrow X_2$ regular, for which there is no finite graph \hat{X} fitting into the following diagram of covers.*

$$\begin{array}{ccc}
 & X & \\
 & \downarrow & \\
 & \hat{X} & \\
 \swarrow & & \searrow \\
 X_1 & & X_2
 \end{array} \tag{5.1}$$

5.1 The graphs X_1, X_2 and X

As in the proof of Theorem 4.2, we will build X by assembling infinitely many copies of a finite graph Y in a tree-like pattern. We shall again 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. Once again, X_1 and X_2 will be obtained from finite graphs Y_1 and Y_2 that are covered by Y . This time we take $Y_2 = Y$, so there is only one non-trivial covering map $p_1 : Y \rightarrow Y_1$ to be considered; this is portrayed in the following figure, where we again label vertices to encode the map.




Figure 4: The covering $p_1 : Y \rightarrow Y_1$.

Let \mathcal{P} be the collection of all covering maps $Y \rightarrow Y_1$ (not just those two that respect the labelling of Figure 4). For each vertex $v \in V(Y)$ 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 $Y \rightarrow Y_1$ maps all of the vertices of either the left or right triangle onto a . (5.2)

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 5: 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 $V(X_1) = V(Y_1)$ and $V(X_2) = V(Y)$.




Figure 6: The graphs X_1 and X_2 .

We define the graph X by taking the regular covering $q_2 : X \rightarrow X_2$ corresponding to the normal subgroup $\langle\langle \pi_1 Y_2 \rangle\rangle$ of $\pi_1 X_2$; the Galois group Γ of this covering is free of rank 18. Note that $\Gamma < \text{Aut}(X)$ 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 according to a tree-like template, as in the proof of Theorem 4.2, and we will again use this template to build maps by working outwards from a basepoint in a radial manner.

5.2 Extendible edges and maps

We want to prove that there is a covering map $q_1 : X \rightarrow X_1$. This covering map will restrict to a covering $p : Y \rightarrow Y_1$ 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 $Y \rightarrow Y_1$ on adjacent copies of Y . With this in mind we make the following definition, where we work with oriented edges, writing $\varepsilon(0)$ and $\varepsilon(1)$ to denote the initial and terminal vertices of an edge $\varepsilon \in E^\pm(X_i)$. This definition is made more awkward by the need to exclude an *exceptional set* \mathcal{E} of pairs of edges; by definition \mathcal{E} is any pair of edges $(\varepsilon_1, \varepsilon_2) \in E^\pm(X_1) \times E^\pm(X_2)$ that are shown as red and thick-dashed in Figure 7 (with any orientations).




Figure 7: The red thick-dashed pairs in the exceptional set

We define a pair of oriented dashed edges $(\varepsilon_1, \varepsilon_2) \in E^\pm(X_1) \times E^\pm(X_2)$ to be *extendible* if $(\varepsilon_1, \varepsilon_2) \notin \mathcal{E}$ and there exist $p, p' \in \mathcal{P}$ such that $p(\varepsilon_2(0)) = \varepsilon_1(0)$ and $p'(\varepsilon_2(1)) = \varepsilon_1(1)$. Extendible pairs can easily be read off from Figure 7, because the condition just requires that the label of $\varepsilon_1(0)$ is an element of the set of labels of $\varepsilon_2(0)$ and that the label of $\varepsilon_1(1)$ is an element of the set of labels at $\varepsilon_2(1)$. Note that $(\varepsilon_1, \varepsilon_2)$ is extendible if and only if the pair with reversed orientations $(\bar{\varepsilon}_1, \bar{\varepsilon}_2)$ is extendible.

Given a vertex v let $\text{lk}_d(v)$ denote the set of oriented dashed edges with initial vertex v . If $v \in V(X_2)$ and $p \in \mathcal{P}$, we say that a bijection $\sigma : \text{lk}_d(v) \rightarrow \text{lk}_d(p(v))$ is *extendible* if every pair $(\sigma(\varepsilon), \varepsilon)$ is extendible.

Lemma 5.2. (1) For any $p \in \mathcal{P}$ and vertex $v \in V(X_2)$, there exists an extendible bijection $\sigma : \text{lk}_d(v) \rightarrow \text{lk}_d(p(v))$.

(2) For any $p \in \mathcal{P}$, vertex $v \in V(X_2)$ and extendible pair $(\varepsilon_1, \varepsilon_2) \in \text{lk}_d(p(v)) \times \text{lk}_d(v)$, there exists an extendible bijection $\sigma : \text{lk}_d(v) \rightarrow \text{lk}_d(p(v))$ such that $\sigma(\varepsilon_2) = \varepsilon_1$.

Proof. To prove these claims one must consider every pair of vertices $(p(v), v)$ with $v \in V(X_2)$ 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 8 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 8 so that we can refer to them accurately without getting overrun by notation.

To begin, note that each yellow edge in $\text{lk}_d(v)$ 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 .

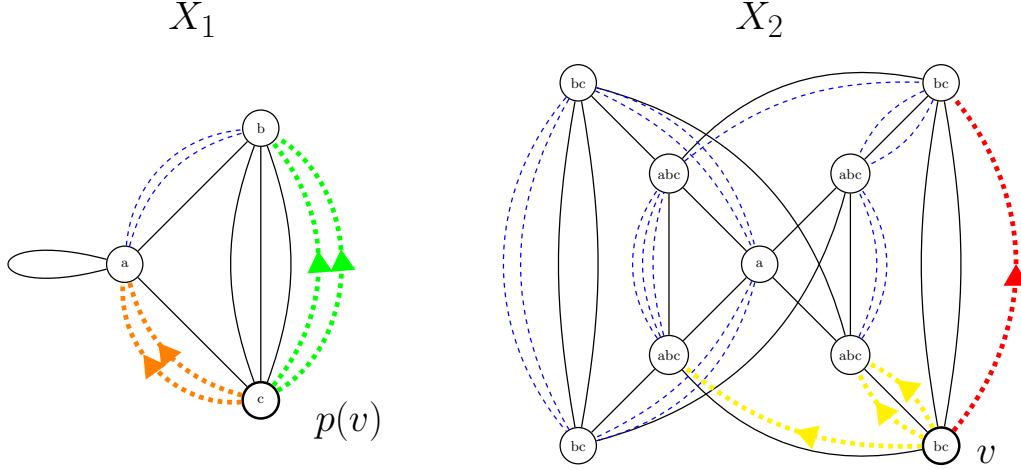


Figure 8: Analysis of maps $\sigma : \text{lk}_d(v) \rightarrow \text{lk}_d(p(v))$ for a particular pair $(p(v), v)$.

For the other pairs $(p(v), v)$ the proof is equally straightforward, having excluded the troublesome edges from Figure 7. 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 $\sigma : \text{lk}_d(v) \rightarrow \text{lk}_d(p(v))$ mapping a red edge to a red edge, since then one of the blue edges in $\text{lk}_d(v)$ would have to map to a blue edge in $\text{lk}_d(p(v))$. We removed this problem by fiat. \square

5.3 The covering $X \rightarrow X_1$

For Z a subgraph of X , we say that an immersion (locally injective graph morphism) $q : Z \rightarrow X_1$ is *extendible* if the pair $(q(\varepsilon), q_2(\varepsilon))$ is extendible for every dashed edge ε in Z . We construct the cover $q_1 : X \rightarrow X_1$ 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 5.3. *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 $q : Z \rightarrow X_1$ is extendible. If $Y' \subset X$ is a copy of Y adjacent to Z , then we can extend q to an extendible map $q' : Z' \rightarrow X_1$, 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 $p_2 : X \rightarrow X_2$. 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 $v \in V(Y')$ we set $q' = \sigma_v q_2$ on $\text{lk}_d(v)$ for some extendible bijection $\sigma_v : \text{lk}_d(q_2(v)) \rightarrow \text{lk}_d(pq_2(v))$. The fact that the σ_v are extendible ensures that q' is extendible. However, q' must agree with q on the unique dashed edge $\tilde{\varepsilon} \in E^\pm(Z)$ with $\tilde{v} := \tilde{\varepsilon}(0) \in V(Y')$. This gives us the following two constraints:

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

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

The idea is then to meet these constraints by using the fact that q is extendible. Indeed, extendibility of q implies that the pair $(q(\tilde{\varepsilon}), q_2(\tilde{\varepsilon}))$ is extendible, so there exists $p \in \mathcal{P}$ satisfying (1). And Lemma 5.2(2) ensures that there exists an extendible bijection $\sigma_{\tilde{v}} : \text{lk}_d(q_2(\tilde{v})) \rightarrow \text{lk}_d(pq_2(\tilde{v}))$ satisfying (2). Finally, the existence of extendible bijections σ_v for the other vertices v follows from Lemma 5.2(1). \square

5.4 Final step in the proof of Theorem 1.3

To complete the proof of the theorem, we must prove that there does not exist a finite graph \hat{X} fitting into the diagram (5.1). Let $\hat{p}_i : \hat{X} \rightarrow X_i$ denote the coverings of X_1 and X_2 . Just as in Theorem 4.2, each connected component \hat{Y} of the solid subgraph of \hat{X} would fit into the following diagram of covers:

$$\begin{array}{ccc} & Y & \\ & \downarrow & \\ & \hat{Y} & \\ \hat{p}_1 \swarrow & & \searrow \hat{p}_2 \\ Y_1 & & Y \end{array} \tag{5.3}$$

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 $p = \hat{p}_1 \circ \hat{p}_2^{-1} : Y \rightarrow Y_1$ induced from the above diagram. Recall that Y contains two triangles, called the left and right triangles (highlighted in bold below). We saw in (5.2) 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 oriented dashed edge ε in \hat{X} , the existence of the maps \hat{p}_1 and \hat{p}_2 on the adjacent pieces tells us that the pair $(\hat{p}_1(\varepsilon), \hat{p}_2(\varepsilon))$ 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 4.1). 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 9: 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. \square

Remark 5.4. We believe that our example for Theorem 1.3 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 $Y \rightarrow Y_1$ 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 $Y \rightarrow Y_1$ 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 $X \rightarrow X_1$ would then extend to a covering $Z \rightarrow Z_1$, and this would be a regular covering since Z is simply connected. We could then deduce that $X \rightarrow X_1$ was a regular covering, and find a finite graph \hat{X} fitting into diagram (5.1) by applying Theorem 1.2.

6 Space of covers

We conclude with a technical result that extends our previous comments on the interplay between the regularity of coverings, cocompactness and unimodularity. Bass and Kulkarni [4] prove that if X is a locally finite tree on which $\text{Aut}(X)$ acts cocompactly, then $\text{Aut}(X)$ is unimodular (in the sense that it supports a Haar measure that is both left and right invariant) if and only if it contains a uniform lattice. We shall not appeal to this fact, but we do need the more elementary fact that a locally compact group that admits a lattice is unimodular; see [22][Chapter I, Remark 1.9].

In the following definition, as elsewhere in this article, we work in the category of metric graphs in which all edges have length 1. We allow only morphisms that map vertices to vertices and restrict to local isometries on edges. In particular, this is true of the covering maps that we allow.

Definition 6.1. Let X be an infinite connected graph that covers a finite graph Y . We define $\mathcal{C}(X, Y)$ to be the space of covering maps $X \rightarrow Y$; this is a closed subspace of the product space $E^\pm(Y)^{E^\pm(X)}$, endowed with the Tychonoff topology.

If a group G acts on X then it acts continuously on the compact Hausdorff space $\mathcal{C}(X, Y)$. Thus it makes sense to talk about a G -invariant probability measure on $\mathcal{C}(X, Y)$.

Theorem 6.2. *Let X be a locally finite infinite quasitree such that $\text{Aut}(X)$ contains a uniform lattice, and suppose that X covers a finite graph Y . Then there exists an $\text{Aut}(X)$ -invariant probability measure on $\mathcal{C}(X, Y)$ if and only if there exists a finite cover $\hat{Y} \rightarrow Y$ and a regular covering map $X \rightarrow \hat{Y}$.*

Remark 6.3. Theorem 1.3 gives an example of a quasitree X that covers finite graphs X_1 and X_2 , with the covering $X \rightarrow X_2$ being regular, such that there is no covering $X \rightarrow \hat{X}$ of a finite graph \hat{X} that covers X_1 and X_2 . There cannot exist a regular covering $X \rightarrow \hat{Y}$ of a finite graph \hat{Y} that covers X_1 , as then Theorem 1.2 would provide the aforementioned prohibited graph \hat{X} . Hence Theorem 6.2 implies that $\mathcal{C}(X, X_1)$ admits no $\text{Aut}(X)$ -invariant probability measure.

Remark 6.4. Theorem 6.2 is false if we remove the assumption that $\text{Aut}(X)$ contains a uniform lattice. For instance if $\text{Aut}(X)$ is compact then the right-invariant Haar measure on $\text{Aut}(X)$ pushes forward to give an $\text{Aut}(X)$ -invariant probability measure on $\mathcal{C}(X, Y)$ via any orbit map $\text{Aut}(X) \rightarrow \mathcal{C}(X, Y)$. But there cannot exist a regular covering of a finite graph $X \rightarrow \hat{Y}$, because its Galois group would be

an infinite lattice in $\text{Aut}(X)$, contradicting compactness of $\text{Aut}(X)$. An example of such a covering $X \rightarrow Y$ is depicted in Figure 10 - $\text{Aut}(X)$ is compact because all automorphisms of X must fix the unique loop.




Figure 10: A covering $X \rightarrow Y$ with $\text{Aut}(X)$ compact.

Proof of Theorem 6.2. First suppose that there is a covering $p : X \rightarrow Y$ that factors through a regular covering $X \rightarrow \hat{Y}$ with \hat{Y} finite, and let $\Gamma < \text{Aut}(X)$ be the Galois group of $X \rightarrow \hat{Y}$. The existence of Γ shows that $\text{Aut}(X)$ is unimodular, so has a Haar measure μ which is both left and right invariant. As $\Gamma < \text{Aut}(X)$ is a uniform lattice, the Haar measure on $\text{Aut}(X)$ induces a finite right invariant measure $\bar{\mu}$ on the right coset space $(\text{Aut}(X) : \Gamma)$, which we may assume is a probability measure by rescaling. We then have a continuous right $\text{Aut}(X)$ -invariant map

$$\begin{aligned} \Phi : (\text{Aut}(X) : \Gamma) &\rightarrow \mathcal{C}(X, Y) \\ \Gamma\phi &\mapsto p \circ \phi. \end{aligned}$$

This is well-defined because the action of Γ on X preserves fibres of the covering map $X \rightarrow \hat{Y}$ (and hence also of $p : X \rightarrow Y$). Thus we can push forward the probability measure $\bar{\mu}$ on $(\text{Aut}(X) : \Gamma)$ to an $\text{Aut}(X)$ -invariant probability measure λ on $\mathcal{C}(X, Y)$.

Conversely, suppose that we have an $\text{Aut}(X)$ -invariant probability measure λ on $\mathcal{C}(X, Y)$. We apply Theorem 3.1 to get an $\text{Aut}(X)$ -invariant quasi-isometry $f : X \rightarrow T$ to a tree T . $\text{Aut}(X)$ contains a uniform lattice by assumption, and since X is a quasitree we may pass to a free uniform finite-index sublattice $\Gamma < \text{Aut}(X)$, and Γ also acts freely and cocompactly on T . By passing to a further finite-index sublattice if necessary, we may assume that the quotient T/Γ has no 1-cycles.

We now define polyhedral pairs and face pairs, which we will eventually glue together to build the intermediate cover \hat{Y} ; this construction is closely analogous to those in [33] and in the appendix of

[27]. For $v \in V(T)$, let $\text{star}(v) \subset T$ be the union of v and the interiors of the edges incident to it, and let $P(v) := f^{-1}(\text{star}(v))$, which we call the *polyhedron at v* . (The term “polyhedron” is suggestive but this set might be disconnected and need not be a subgraph of X ; it will be bounded however.) By boundedness, we know that there are only finitely many maps $\eta : P(v) \rightarrow Y$ that extend to a cover $X \rightarrow Y$, and we denote this set of maps by $\text{Map}(v)$. For $e \in E(T)$, let \mathring{e} be the interior of e , and define $F(e) := f^{-1}(\mathring{e})$ to be the *face at e* , and let $\text{Map}(e)$ be the (finite) set of maps $\eta : F(e) \rightarrow Y$ that extend to covers $X \rightarrow Y$.

We use square brackets to denote Γ -orbits of edges and vertices in T . Pick $v \in V(T)$, let $P[v]$ be a space isometric to $P(v)$ via an isometry $\alpha_v : P[v] \rightarrow P(v)$, and define $\text{Map}[v] := \text{Map}(v) \circ \alpha_v$. For any $u \in V(T)$ with $[u] = [v]$ there is a unique $g \in \Gamma$ with $g(v) = u$; this means that $gP(v) = P(u)$, and we get a bijection

$$\begin{aligned} \text{Map}(u) &\rightarrow \text{Map}(v) \\ \eta &\mapsto \eta \circ g, \end{aligned}$$

because $\eta \in \text{Map}(u)$ extending to $\bar{\eta} : X \rightarrow Y$ implies that $\eta \circ g$ extends to $\bar{\eta} \circ g$. Thus we can define $\alpha_u := g \circ \alpha_v$ and get $\text{Map}[v] = \text{Map}(u) \circ \alpha_u$. Similarly, we have spaces $F[e]$ equipped with sets of maps $\text{Map}[e]$, such that for any $e \in E(T)$ there is an isometry $\alpha_e : F[e] \rightarrow F(e)$, and $\text{Map}[e] = \text{Map}(e) \circ \alpha_e$, and for $g \in \Gamma$ we have $\alpha_{g(e)} = g \circ \alpha_e$. For $v \in V(T)$ and $\phi \in \text{Map}[v]$, we say that $\mathbf{P} = (P[v], \phi)$ is a *polyhedral pair*, while for $e \in E(T)$ and $\varphi \in \text{Map}[e]$ we say that $\mathbf{F} = (F[e], \varphi)$ is a *face pair*. We write \mathcal{P} and \mathcal{F} for the collections of polyhedral pairs and face pairs respectively.

If e is an edge incident at a vertex v in T , then we have an inclusion $F(e) \hookrightarrow P(v)$. We can then define an inclusion $\iota_{e,v} : F[e] \rightarrow P[v]$ via the following commutative diagram:

$$\begin{array}{ccc} F[e] & \xhookrightarrow{\iota_{e,v}} & P[v] \\ \downarrow \alpha_e & & \downarrow \alpha_v \\ F(e) & \hookrightarrow & P(v) \end{array} \tag{6.1}$$

Moreover, for $g \in \Gamma$ we have that $\iota_{g(e),g(v)} = \iota_{e,v}$. Since no Γ -translate of e is incident at v (by freeness of the action and the fact that T/Γ has no 1-cycles), we conclude that $\iota_{e,v}$ only depends on $[e]$ and $[v]$; we write it as $\iota_{[e],[v]}$.

As Γ acts freely on T , we can choose a Γ -invariant orientation on the edges of T , and for each $e \in E(T)$ we will refer to its two endpoints as the vertices on the *left* and *right* of e . If v is on the left (resp. right) of e , and $\mathbf{P} = (P[v], \phi)$ and $\mathbf{F} = (F[e], \varphi)$ are such that $\phi \circ \iota_{[e],[v]} = \varphi$, then we say that \mathbf{P} is the *left (right)* of \mathbf{F} , and denote this by $\mathbf{P} \in \overleftarrow{\mathbf{F}}$ ($\mathbf{P} \in \overrightarrow{\mathbf{F}}$). If $\mathbf{P}_1 = (P[v_1], \phi_1) \in \overleftarrow{\mathbf{F}}$ and $\mathbf{P}_2 = (P[v_2], \phi_2) \in \overrightarrow{\mathbf{F}}$, then we can glue \mathbf{P}_1 and \mathbf{P}_2 together along \mathbf{F} and the maps $\iota_{[e],[v_1]}$ and $\iota_{[e],[v_2]}$, and we can define a map

$$\phi_1 \cup \phi_2 : P[v_1] \cup_{F[e]} P[v_2] \rightarrow Y.$$

To construct the intermediate cover $X \rightarrow \hat{Y} \rightarrow Y$, we will take $\omega(\mathbf{P})$ copies of each polyhedral pair \mathbf{P} and $\omega(\mathbf{F})$ copies of each face pair \mathbf{F} , and each face pair will be glued to one polyhedral pair on the left and one polyhedral pair on the right. We will discuss the covering maps $X \rightarrow \hat{Y}$ and $\hat{Y} \rightarrow Y$ later, but for the gluing to even be possible the numbers $\omega(\mathbf{P})$ and $\omega(\mathbf{F})$ must satisfy the following Gluing

Equations for each face pair \mathbf{F} :

$$\sum_{\mathbf{P} \in \overleftarrow{\mathbf{F}}} \omega(\mathbf{P}) = \omega(\mathbf{F}) = \sum_{\mathbf{P} \in \overrightarrow{\mathbf{F}}} \omega(\mathbf{P}) \quad (6.2)$$

We note that by [26, claim on p19], a non-zero solution $\omega : \mathcal{P} \cup \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$ to the Gluing Equations implies the existence of a non-zero integer solution $\omega : \mathcal{P} \cup \mathcal{F} \rightarrow \mathbb{Z}_{\geq 0}$. We now show how the $\text{Aut}(X)$ -invariant probability measure λ on $\mathcal{C}(X, Y)$ can be used to solve the Gluing Equations.

Claim: The weights

$$\begin{aligned} \omega(P[v], \phi) &= \lambda(\{\theta \mid \theta|_{P(v)} = \phi \circ \alpha_v^{-1}\}) \\ \omega(F[e], \varphi) &= \lambda(\{\theta \mid \theta|_{F(v)} = \varphi \circ \alpha_e^{-1}\}) \end{aligned}$$

solve the Gluing Equations (6.2).

Firstly, note that the above formulae for ω are independent of the choice of vertex and edge in the orbits $[v]$ and $[e]$ because λ is Γ -invariant. If v is on the left of e and $\mathbf{F} = (F[e], \varphi)$ is a face pair, then $(P[v], \phi) \in \overleftarrow{\mathbf{F}}$ if and only if $\phi \circ \alpha_v^{-1} : P(v) \rightarrow Y$ is an extension of the map $\varphi \circ \alpha_e^{-1} : F(e) \rightarrow Y$; and so

$$\{\theta \mid \theta|_{F(v)} = \varphi \circ \alpha_e^{-1}\} = \bigsqcup_{(P[v], \phi) \in \overleftarrow{\mathbf{F}}} \{\theta \mid \theta|_{P(v)} = \phi \circ \alpha_v^{-1}\}.$$

Applying λ to both sides yields

$$\omega(\mathbf{F}) = \sum_{\mathbf{P} \in \overleftarrow{\mathbf{F}}} \omega(\mathbf{P}),$$

as required. The argument works similarly if v is on the right of e . Finally, the map $\omega : \mathcal{P} \cup \mathcal{F} \rightarrow \mathbb{Z}_{\geq 0}$ is non-zero because

$$1 = \lambda(\mathcal{C}(X, Y)) = \sum_{\phi \in \text{Map}[v]} \lambda(\{\theta \mid \theta|_{P(v)} = \phi \circ \alpha_v^{-1}\}) = \sum_{\phi \in \text{Map}[v]} \omega(P[v], \phi)$$

for $v \in V(T)$.

Now assume that $\omega : \mathcal{P} \cup \mathcal{F} \rightarrow \mathbb{Z}_{\geq 0}$ is a non-zero integer solution to the Gluing Equations. We can then take $\omega(\mathbf{P})$ copies of each polyhedral pair \mathbf{P} and $\omega(\mathbf{F})$ copies of each face pair \mathbf{F} , and glue them all together as described earlier to form a space \hat{Y} . This gluing can be described by a finite graph Z : take a vertex for each copy of a polyhedral pair and an edge for each copy of a face pair, then connect edges to vertices according to which face pairs get glued to which polyhedral pairs. We can then define a covering $Z \rightarrow T/\Gamma$, where a vertex associated to $(P[v], \phi)$ maps to $[v]$ and an edge associated to $(F[e], \varphi)$ maps to $[e]$. The covering $T \rightarrow T/\Gamma$ then factors as $T \rightarrow Z \rightarrow T/\Gamma$ for some covering $T \rightarrow Z$, which on vertices and edges we denote by $v \mapsto \bar{v}$ and $e \mapsto \bar{e}$. It will be convenient to always refer to vertices and edges of Z as images \bar{v} and \bar{e} of vertices and edges of T . We can then write $\mathbf{P}_{\bar{v}} = (P[v], \phi_{\bar{v}})$ for the polyhedral pair associated to a vertex $\bar{v} \in V(Z)$ and $\mathbf{F}_{\bar{e}} = (F[e], \varphi_{\bar{e}})$ for the face pair associated to an edge $\bar{e} \in E(Z)$. If \bar{v} is on the left (right) of \bar{e} , then $\mathbf{P}_{\bar{v}}$ will be glued to

$\mathbf{F}_{\bar{e}}$ on the left (right) to form \hat{Y} - in particular $\mathbf{P}_{\bar{v}} \in \overleftarrow{\mathbf{F}}_{\bar{e}}$ (resp. $\mathbf{P}_{\bar{v}} \in \overrightarrow{\mathbf{F}}_{\bar{e}}$). Therefore, we can write \hat{Y} as

$$\hat{Y} = \bigsqcup_{\bar{v} \in V(Z)} P[v] \times \{\bar{v}\} \sqcup \bigsqcup_{\bar{e} \in E(Z)} F[e] \times \{\bar{e}\} / (\iota_{[e],[v]}(a), \bar{v}) \sim (a, \bar{e}) \text{ for } \bar{e} \text{ incident at } \bar{v}.$$

Note that the images of the maps $\iota_{[e],[v]}$ in a polyhedron $P[v]$ are disjoint, so each point in \hat{Y} either lies in a single polyhedron $P[v] \times \{\bar{v}\}$, or it lies in a face $F[e] \times \{\bar{e}\}$ as well as in the two polyhedra $P[v] \times \{\bar{v}\}$ with \bar{e} incident at \bar{v} , but not in any other polyhedra. Since T/Γ has no 1-cycles, each polyhedron $P[v] \times \{\bar{v}\}$ is embedded in \hat{Y} .

We define the map $\chi : X \rightarrow \hat{Y}$ by

$$\begin{aligned} x \in P(v) &\mapsto (\alpha_v^{-1}(x), \bar{v}) \\ x \in F(e) &\mapsto (\alpha_e^{-1}(x), \bar{e}), \end{aligned}$$

for $v \in V(T)$ and $e \in E(T)$. Note that every point $x \in X$ lies in some polyhedron $P(v)$ because $f(x) \subset \text{star}(v)$ for some $v \in V(T)$. If e is incident at v then the two maps above agree on $x \in F(e) \subset P(v)$ because $(\alpha_v^{-1}(x), \bar{v}) = (\iota_{[e],[v]}\alpha_e^{-1}(x), \bar{v}) \sim (\alpha_e^{-1}(x), \bar{e})$. Moreover, if $\chi(x) \in F[e] \times \{\bar{e}\}$, then x is contained in a face in X , and the polyhedra on the left and right of this face will map down through χ to the two polyhedra $P[v] \times \{\bar{v}\}$ with \bar{e} incident at \bar{v} .

We define the map $\psi : \hat{Y} \rightarrow Y$ by

$$\begin{aligned} (a, \bar{v}) &\mapsto \phi_{\bar{v}}(a) \\ (a, \bar{e}) &\mapsto \varphi_{\bar{e}}(a), \end{aligned}$$

for $\bar{v} \in V(Z)$ and $\bar{e} \in E(Z)$. This is well defined, because if \bar{e} is incident at \bar{v} then $(a, \bar{e}) \sim (\iota_{[e],[v]}(a), \bar{v}) \mapsto \phi_{\bar{v}}\iota_{[e],[v]}(a) = \varphi_{\bar{e}}(a)$ since $\mathbf{P}_{\bar{v}}$ is on the left or right of $\mathbf{F}_{\bar{e}}$. It remains to show that these maps are coverings, and that $X \rightarrow \hat{Y}$ is a regular covering. Note it is not obvious from the above description that \hat{Y} is even a graph, but this follows from the existence of the covering maps.

Claim: $\chi : X \rightarrow \hat{Y}$ is a regular covering map.

We have a factoring of covers $T \rightarrow Z \rightarrow T/\Gamma$, so $Z = T/\hat{\Gamma}$ for $\hat{\Gamma} < \Gamma$ of finite index. If $g \in \hat{\Gamma}$ and $v \in V(T)$, then $\bar{v} = \overline{g(v)}$ and we get the following commutative diagram.

$$\begin{array}{ccc} P(v) & \xrightarrow{g} & P(g(v)) \\ & \swarrow \alpha_v \quad \searrow \alpha_{g(v)} & \\ & P[v] & \\ & \downarrow & \\ & P[v] \times \{\bar{v}\} & \end{array} \quad \begin{array}{c} \chi \swarrow \quad \searrow \chi \end{array} \quad (6.3)$$

Hence $\chi \circ g = \chi$.

On the other hand, if $x_1 \in P(v_1)$ and $x_2 \in X$ satisfy $\chi(x_1) = \chi(x_2)$, then $x_2 \in P(v_2)$ for some v_2 with $\bar{v}_1 = \bar{v}_2$, and so there exists $g \in \hat{\Gamma}$ with $g(v_1) = v_2$. This also tells us that $g(x_1), x_2 \in P(v_2)$ have the same image under χ , and since the restriction of χ to $P(v_2)$ is injective we deduce that $g(x_1) = x_2$. $\hat{\Gamma}$ acts properly on X , so we conclude that χ is a covering map with deck transformation group $\hat{\Gamma}$.

Claim: $\psi : \hat{Y} \rightarrow Y$ is a covering map.

As \hat{Y} is compact (by the previous claim), it is enough to show that each point in \hat{Y} has an open neighbourhood that maps homeomorphically onto an open subspace of Y . Given $y \in \hat{Y}$, take a lift $x \in X$ and consider $f(x) \in T$. We know that $f(x) \in \text{star}(v)$ for some vertex v , so $x \in P(v)$. We also know that $y = \chi(x) = (\alpha_v^{-1}(x), \bar{v})$, and that we have a commutative diagram

$$\begin{array}{ccc}
 P(v) & \xrightarrow[\sim]{\chi} & P[v] \times \{\bar{v}\} \\
 \nwarrow \alpha_v \sim & & \nearrow \sim \\
 & P[v] & \\
 \downarrow \phi_{\bar{v}} & & \\
 \eta \nearrow & Y & \nwarrow \psi
 \end{array} \tag{6.4}$$

where \sim indicates homeomorphisms, and η is some map in $\text{Map}(v)$. The map $\eta : P(v) \rightarrow Y$ extends to a cover $X \rightarrow Y$, so there is a neighbourhood of x in $P(v)$ that maps homeomorphically to an open subspace of Y . Since $\chi : X \rightarrow \hat{Y}$ is a cover, we know that this neighbourhood of x maps homeomorphically to a neighbourhood of y , which in turn is mapped homeomorphically by ψ to the same open subspace of Y .

This completes the proof of the theorem. □

References

- [1] I. Agol, *The virtual Haken conjecture* (with appendix by I. Agol, D. Groves and J. Manning), Documenta Math. **18**, 1045–1087, 2013.
- [2] H. Bass, *Covering theory for graphs of groups*, J. Pure Appl. Algebra 89, 3–47, 1993.
- [3] H. Bass and R. Kulkarni, *Uniform tree lattices*, J. Amer. Math. Soc. 3, 843–902, 1990.
- [4] H. Bass and A. Lubotzky, *Tree lattices*. Progress in Mathematics, 176. Birkhäuser Boston, Inc., Boston, MA, 2001.
- [5] Jason A. Behrstock and Walter D. Neumann, *Quasi-isometric classification of non-geometric 3-manifold groups*, J. Reine Angew. Math., 669:101–120, 2012.
- [6] M. Bestvina, *Geometric group theory and 3-manifolds hand in hand: the fulfillment of Thurston's vision*, Bull. Amer. Math. Soc. 51, 53–70, 2014.
- [7] M.R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der Mathematischen Wissenschaften, Vol. 319, Springer-Verlag, Heidelberg-Berlin, 1999.
- [8] M.R. Bridson and H. Wilton, *The triviality problem for profinite completions*, Invent. Math. 202, 839–874, 2015.

- [9] M.R. Bridson, B. McReynolds, A.W. Reid, R. Spitler, *Absolute profinite rigidity and hyperbolic geometry*, Ann. of Math. (2) 192 (2020), no. 3, 679–719.
- [10] Marc Burger and Shahar Mozes, *Finitely presented simple groups and products of trees*, C. R. Acad. Sci. Paris Sér. I Math. 324, no. 7, 747–752, 1997
- [11] Pierre-Emmanuel Caprace, *Finite and infinite quotients of discrete and indiscrete groups*, London Math. Soc. Lecture Note Ser., 455, Cambridge Univ. Press, Cambridge, 2019.
- [12] I. Chatterji and G. Niblo, *From wall spaces to $CAT(0)$ cube complexes*, Internat. J. Algebra Comput. 15, 875–885, 2005.
- [13] M. Gromov, *Hyperbolic groups*, in “Essays in Group Theory”, pp. 75–263, Math. Sci. Res. Inst. Publ. vol. 8, Springer, New York, 1987.
- [14] Mark F. Hagen and Nicholas W.M. Touikan, *Panel collapse and its applications*, Groups Geom. Dyn. 13 (2019), no. 4, 1285–1334.
- [15] Frédéric Haglund and Frédéric Paulin, *Simplicité de groupes d’automorphismes d’espaces à courbure négative*, in The Epstein birthday schrift, 181–248, Geom. Topol. Monogr. 1, Coventry, 1998.
- [16] Frédéric Haglund and Daniel T. Wise, *Special cube complexes*, Geom. Funct. Anal. 17, 1551–1620, 2008.
- [17] Jingyin Huang, *Commensurability of groups quasi-isometric to RAAGs*, Invent. Math. 213, no. 3, 1179–1247, 2018.
- [18] Frank T. Leighton, *Finite common coverings of graphs*, J. Combin. Theory Ser. B 33, 231–238, 1982.
- [19] Lee Mosher, Michah Sageev, and Kevin Whyte, *Quasi-actions on trees I. Bounded valence*, Annals of Mathematics, 158, 115–164, 2003.
- [20] Walter D. Neumann, *On Leighton’s graph covering theorem*, Groups Geom. Dyn. 4, 863–872, 2010.
- [21] Bogdan Nica, *Cubulating spaces with walls*, Algebr. Geom. Topol. 4, 297–309, 2004.
- [22] M. S. Raghunathan, *Discrete subgroups of Lie groups*, Ergebnisse der Math., vol. 68, Springer-Verlag, 1972.
- [23] Martin Roller, *Poc-sets, median algebras and group actions. An extended study of Dunwoody’s construction and Sageev’s theorem*, arXiv:1607.07747, 2016.
- [24] Michah Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71, 585–617, 1995.
- [25] Michah Sageev, *$CAT(0)$ cube complexes and groups*, in “Geometric group theory”, 7–54, IAS/Park City Math. Ser., 21, Amer. Math. Soc., Providence, RI, 2014.
- [26] Sam Shepherd, *Agol’s theorem on hyperbolic cubulations*, Rocky Mountain J. Math., to appear.
- [27] Sam Shepherd, (with appendix by Giles Gardam and Daniel J. Woodhouse), *Two generalisations of Leighton’s Theorem*, arXiv:1908.00830, 2019.

- [28] Sam Shepherd, *Finite covers of graphs and cube complexes*, PhD Thesis Oxford University, in preparation
- [29] Sam Shepherd and Daniel J. Woodhouse, *Quasi-isometric rigidity for graphs of virtually free groups with two-ended edge groups*, arXiv:2007.10034, 2020.
- [30] Emily Stark and Daniel J. Woodhouse, *Action rigidity for free products of hyperbolic manifold groups*, arXiv:1910.09609, 2019.
- [31] Daniel T. Wise, *Non-positively curved squared complexes: Aperiodic tilings and non- residually finite groups*, PhD Thesis Princeton University, ProQuest LLC, Ann Arbor, MI., 1996.
- [32] Daniel T. Wise, *From Riches to RAAGs: 3-manifolds, Right-angled Artin Groups, and Cubical Geometry*, CBMS Regional Conference Series in Mathematics 117, Amer. Math. Soc., Providence, RI, 2012.
- [33] Daniel J. Woodhouse, *Revisiting Leighton's Theorem with the Haar measure*, Math. Proc, Cambridge Philos. Soc., 1-9. doi:10.1017/S0305004119000550.