Algebraic & Geometric Topology 20 (2020) 1877–1938

Immersed cycles and the JSJ decomposition

SURAJ KRISHNA MEDA SATISH

We present an algorithm to construct the JSJ decomposition of one-ended hyperbolic groups which are fundamental groups of graphs of free groups with cyclic edge groups. Our algorithm runs in double exponential time and is the first algorithm on JSJ decompositions to have an explicit time bound. Our methods are combinatorial/geometric and rely on analysing properties of immersed cycles in certain CAT(0) square complexes.

20E06, 20E08, 20F65, 20F67

1 Introduction

In [34], Sela showed the existence of a canonical decomposition of a torsion-free one-ended hyperbolic group over its infinite cyclic subgroups. This decomposition, which Sela called a *JSJ decomposition* (see Definition 1.1), is a generalisation to group theory of JSJ decompositions of 3–manifolds (due to Jaco and Shalen [20] and Johannson [21]).

We present an algorithm (Theorem 1.3) to construct the JSJ decomposition of the fundamental group G of a graph of free groups with cyclic edge groups when G is one-ended and hyperbolic. We develop this algorithm by first obtaining one such algorithm in a special case (Theorem 1.4): when G is the fundamental group of a compact nonpositively curved square complex called a *tubular graph of graphs* (which we introduced in [37]).

A tubular graph of graphs (see Definition 2.8 for the precise definition) is a square complex obtained by attaching finitely many tubes (a *tube* is a Cartesian product of the unit interval and a circle) to a finite collection of finite graphs. Tubular graphs of graphs are thus nonpositively curved $\mathcal{VH}$–complexes (introduced by Wise in [40]) whose vertical hyperplanes are circles (see Section 2). A typical example of the fundamental group of a tubular graph of graphs is the amalgamated product of two free groups over cyclic subgroups.

As an application of Theorem 1.4, we obtain an algorithm that takes a finite-rank free group F and a finite family of cyclic subgroups $\mathcal{H}$ such that F is freely indecomposable

Published: 20 July 2020 DOI: 10.2140/agt.2020.20.1877

relative to $\mathcal{H}$ as input and constructs the JSJ decomposition of F relative to H in double exponential time (Theorem 13.3).

A consequence of our main result is a double-exponential time solution to the isomorphism problem for graphs of free groups with cyclic edge groups in the hyperbolic case (Theorem 15.1). Recall that the isomorphism problem is the algorithmic problem of deciding whether two finite presentations of groups present isomorphic groups; see Dehn [14].

1.1 JSJ decompositions

We adopt the terminology of Sela [34]. Let G be a torsion-free hyperbolic group. A *hanging surface subgroup* G' of G is a subgroup isomorphic to the fundamental group of a surface with boundary such that there exists a graph of groups decomposition of G in which G' is a vertex group whose incident edge groups are precisely the peripheral subgroups of G'. A *maximal hanging surface subgroup* is a hanging surface subgroup that is not properly contained in any hanging surface subgroup. A noncyclic vertex group G' of G is *rigid* if it is elliptic in every cyclic splitting of G. A subgroup is *full* (in the sense of Bowditch [4]) if it is not properly contained as a finite-index subgroup in any subgroup of G.

We can now define JSJ decompositions in the sense of Sela [34], as modified by Bowditch [4]; see [4, Theorem 0.1 and Theorem 5.28].

Definition 1.1 (JSJ decomposition) Let G be a torsion-free hyperbolic group. A *JSJ splitting* of G is a finite graph of groups decomposition of G where each edge group is cyclic and each vertex group is full and of one of the following three types:

(1) a cyclic subgroup,

(2) a maximal hanging surface subgroup, or

(3) a rigid subgroup.

If a vertex v of type (1) has valence one, then the incident edge group does not surject onto the vertex group G_v. Moreover, exactly one endpoint of any edge is of type (1) and the edge groups that connect to any vertex group of type (2) are precisely the peripheral subgroups of that group.

Theorem 1.2 [34] *Let G be a torsion-free one-ended hyperbolic group which is not the fundamental group of a closed surface. Then a JSJ decomposition of G exists and is unique.*

We are now ready to state our main result.

Theorem 1.3 (Theorem 14.1) *There exists an algorithm of double exponential time complexity that takes a graph of free groups with cyclic edge groups with one-ended hyperbolic fundamental group G as input and returns the JSJ decomposition of G.*

As mentioned earlier, the algorithm is a consequence of the theorem below, proving which takes up a major part of the current article.

Theorem 1.4 (Theorem 12.44) *There exists an algorithm of double exponential time complexity that takes a tubular graph of graphs with one-ended hyperbolic fundamental group G as input and returns a tubular graph of graphs whose graph of groups structure is the JSJ decomposition of G.*

Other authors have obtained algorithms to compute JSJ decompositions of groups under different conditions. In [12], Dahmani and Guirardel give an algorithm to compute JSJ decompositions of one-ended hyperbolic groups over maximal virtually cyclic subgroups with infinite centre. In [13], Dahmani and Touikan give an algorithm to compute JSJ decompositions of torsion-free hyperbolic groups over its cyclic subgroups. In [2], Barrett gives an algorithm to compute JSJ decompositions of one-ended hyperbolic groups over virtually cyclic subgroups, while Cashen and Manning [7; 9] develop an implementable algorithm to construct the relative JSJ of a free group relative to a family of cyclic subgroups. We remark that the time complexity of these algorithms is not known.

Our approach is combinatorial/geometric. We will now describe this approach briefly.

1.2 Coarse behaviour and Brady–Meier tubular graphs of graphs

Definition 1.5 A cube complex is *Brady–Meier* if all its vertex links are connected and moreover each vertex link remains connected after removing any simplex in the link.

Theorem 1.6 (Brady and Meier [5]) *The fundamental group of a finite connected Brady–Meier nonpositively curved cube complex is one-ended.*

Theorem 1.7 [37] *There exists an algorithm of polynomial time complexity that takes a tubular graph of graphs with one-ended fundamental group as input and returns a Brady–Meier tubular graph of graphs with isomorphic fundamental group.*

Thanks to the above result, we work with Brady–Meier tubular graphs of graphs throughout this article. Let X be a Brady–Meier tubular graph of graphs endowed with its $\mathcal{VH}$ structure. Each vertical hyperplane of X is a circle (Proposition 2.9). If $\widetilde{X}$ denotes the CAT(0) universal cover of X, then the vertical hyperplanes of $\widetilde{X}$ are lines. Let G denote the fundamental group of X. Adopting the terminology of Scott and Wall [33], X has a structure of a graph of spaces (see Section 2.2 for details), where each vertex space is itself a graph. Similarly, $\widetilde{X}$ has a structure of a tree of spaces, where each vertex space is a (vertical) tree.

In order to construct the JSJ decomposition of G, one has to first find cyclic subgroups over which G splits. We address this issue using the Brady–Meier structure of X:

A geodesic line L of $\widetilde{X}$ *separates* $\widetilde{X}$ if $\widetilde{X} \setminus L$ is not connected. Two separating geodesic lines L_1 and L_2 of $\widetilde{X}$ *cross* if L_1 meets two distinct components of $\widetilde{X} \setminus L_2$ and vice versa. An *axis* in $\widetilde{X}$ of an element $g \in G$ is a geodesic line in $\widetilde{X}$ that is invariant under the action of the cyclic subgroup $\langle g \rangle$. Given $g \in G$, an axis L of g always exists in $\widetilde{X}$; see Bridson and Haefliger [6].

Suppose that G splits over $\langle g \rangle$ and that L is contained in a vertical tree of $\widetilde{X}$. Then an application of a result of Papasoglu [27] to the fact that $\widetilde{X}$ is a Brady–Meier complex implies that L separates $\widetilde{X}$ and L does not cross any of its translates (Lemma 6.3). Conversely, if L is separating and does not cross any of its translates, then G splits over a subgroup of the stabiliser of L (Proposition 6.10).

The properties of separation and crossing have local characterisations in the Brady–Meier complex $\widetilde{X}$. Denote by $N_R(L)$ the set of all points in $\widetilde{X}$ at distance at most R from a point of L.

Lemma 1.8 (Lemma 4.7) *The line L separates $\widetilde{X}$ if and only if it separates $N_{\frac{1}{4}}(L)$.*

Proposition 1.9 (Proposition 5.2) *Two separating lines L_1 and L_2 cross if and only if*

(1) *$L_1 \cap L_2$ is nonempty and compact, and*

(2) *L_2 meets two components of $N_{\frac{1}{4}}(L_1 \cap L_2) \setminus L_1$.*

The quotient of L by $\langle g \rangle$ is an immersed circle C, which we call a *cycle*, in X. The *regular neighbourhood* of C is the quotient of $N_{\frac{1}{4}}(L)$ by the action of $\langle g \rangle$. The fact that L separates $N_{\frac{1}{4}}(L)$, along with a condition that is satisfied since G splits over $\langle g \rangle$ (Definition 7.18), implies the following result:

Lemma 1.10 (Lemma 7.24) *C separates its regular neighbourhood and $\widetilde{C}$ does not cross any of its translates in $\widetilde{X}$.*

We need another property to construct the JSJ decomposition. A cyclic subgroup over which G splits is said to be *universally elliptic* if it is elliptic in the Bass–Serre tree of any cyclic splitting of G; see Guirardel and Levitt [16]. The edge groups of the JSJ decomposition are universally elliptic.

Let L_1 (respectively L_2) be an axis of g_1 (respectively g_2) such that G splits over $\langle g_1\rangle$ and $\langle g_2\rangle$. Then $\langle g_1\rangle$ is elliptic in the Bass–Serre tree of the splitting over $\langle g_2\rangle$ only if L_1 and any translate of L_2 don't cross (Lemma 8.1).

1.3 Repetitive cycles and JSJ splittings

In Section 9, we introduce an important notion, namely repetitivity, that bounds the length of a cycle that induces a universally elliptic splitting. Let $\widetilde{C}$ denote a lift of a cycle C in $\widetilde{X}$.

Definition 1.11 (Definition 9.2; Lemma 9.4) A cycle C is k*–repetitive* if $\widetilde{C}$ is a separating line and there exists an edge e in $\widetilde{X}$ and elements $g_1, \dots, g_k \in G$ such that

(1) each translate $g_i\widetilde{C}$ contains e,

(2) the distance between e and $g_i e$ is strictly less than the length of C, and

(3) any two squares s and s' that contain e are either separated by all translates $g_i\widetilde{C}$ or by none of them.

There are two important reasons for introducing the notion of repetitive cycles. The first is that any cycle that is longer than a certain bound is k–repetitive (Proposition 9.9). Here, the bound depends only on k and the number of squares of X. The second reason is the following:

Proposition 1.12 (Proposition 10.1) *Let C be a k–repetitive cycle with $k \geq 3$. Suppose that $\pi_1(C)$ is a maximal cyclic subgroup of G. Then there exists a separating line L in $\widetilde{X}$ such that L and $\widetilde{C}$ coarsely cross.*

This implies that $\pi_1(C)$ conjugates into a hanging surface subgroup of the JSJ splitting of G, by Bowditch [4, Proposition 5.30]. Hence, $\pi_1(C)$ is not universally elliptic and the length of a cycle which induces a universally elliptic cyclic subgroup is bounded. This leads to the following result:

Theorem 1.13 (Theorem 11.1) *There exists an algorithm of double exponential time complexity that takes a Brady–Meier tubular graph of graphs with hyperbolic fundamental group G as input and returns a finite list of splitting cycles that contains all universally elliptic subgroups of G up to commensurability.*

We remark that a similar, but weaker, result is obtained by Cashen and Macura [8, Proposition 4.11] using a similar idea: given a free group F and a finite family of cyclic subgroups $\mathcal{H}$, the authors obtain a bound depending on F and $\mathcal{H}$ such that if F admits a cyclic splitting relative to $\mathcal{H}$, then there is a word of length less than the bound over which F splits relative to $\mathcal{H}$.

1.4 Obtaining a JSJ complex

In Section 12, we modify the given tubular graph of graphs X to a tubular graph of graphs X_{jsj} such that the fundamental group of the graph of spaces structure of X_{jsj} is the JSJ decomposition of G.

The first step involves a modification of the initial tubular graph of graphs X by cutting along the finite list of cycles supplied by Theorem 1.13. We do this cutting procedure using the machinery of *spaces with walls* (due to Haglund and Paulin [17]). The vertex set of $\widetilde{X}$ is a space with walls, with walls defined by its vertical and horizontal hyperplanes. We enrich the wall set by adding lifts of cycles supplied by Theorem 1.13 (see Section 12.3). We then remove tubes which are attached to cyclic vertex graphs on both sides. In Proposition 12.36, we show that each edge group of the JSJ decomposition of G is a conjugate of an edge group of the underlying graph of groups of the new tubular graph of graphs. Thus, an edge stabiliser of the Bass–Serre tree of the new tubular graph of graphs is either an edge stabiliser of the JSJ tree or a cyclic subgroup that conjugates into a maximal hanging surface subgroup of the JSJ splitting. It only remains to identify the maximal surface subgroups that appear as vertex groups in the JSJ decomposition. Once identified, removing tubes corresponding to edge stabilisers which conjugate into maximal hanging surface subgroups gives the JSJ decomposition, proving the main result (Theorem 1.4).

1.5 Identifying surfaces

We give a criterion to identify surfaces in the Brady–Meier setup. A vertex graph of a tubular graph of graphs is a *surface graph* if the fundamental group of the graph is a surface group whose peripheral subgroups are precisely the incident edge subgroups. Then:

Lemma 1.14 (Lemma 12.42) *A vertex graph of a Brady–Meier tubular graph of graphs is a surface graph if and only if every edge of its double is contained in exactly two squares.*

We refer the reader to Definition 12.41 for the definition of a double.

Acknowledgements I thank my advisors Frédéric Haglund and Thomas Delzant for their time and for the many useful discussions. I thank Henry Wilton for the many references. I am grateful to the referee for the careful reading and helpful comments which improved the article, and especially for motivating Theorem 14.1. Part of this work was done when I was visiting IRMA, Strasbourg. This work was supported by a French public grant for research, as part of the Investissement d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH.

2 The setup

2.1 $\mathcal{VH}$–complexes

The notion of $\mathcal{VH}$–complexes was first introduced in [40].

Definition 2.1 A *square complex* is a two-dimensional CW complex in which each 2–cell is attached to a combinatorial loop of length 4 and is isometric to the standard Euclidean unit square $I^2 = [0,1]^2$.

All our square complexes will be locally finite.

Definition 2.2 (vertex links) Let $v \in X$ be a vertex of a square complex. The link of v, denoted by $\mathrm{link}(v)$, is a graph whose vertex set is $\{e \mid e \text{ is a half-edge incident to } v\}$. The number of edges between two vertices e and f is the number of squares of X in which e and f are adjacent half-edges.

Definition 2.3 A square complex is *nonpositively curved* if the length of a closed path in the link of any of its vertices is at least four.

By a result of Gromov [15], a simply connected nonpositively curved square complex is $\mathrm{CAT}(0)$ in the metric sense.

Definition 2.4 [32] Let X be a square complex. A *mid-edge* of a square s in X is an edge (after subdivision of s) running through the centre of s and parallel to two of the edges of s. Declare two edges e and f to be equivalent if there exists a sequence $e = e_1, \ldots, e_n = f$ of edges such that e_i and e_{i+1} are opposite edges of some square of X. Given an equivalence class $[e]$ of edges, the *hyperplane dual to* e, denoted by h_e, is the collection of mid-edges which intersect edges in $[e]$.

Definition 2.5 [40] A $\mathcal{VH}$*–complex* is a square complex in which every 1–cell is labelled as either vertical or horizontal in such a way that each 2–cell is attached to a loop which alternates between horizontal and vertical 1–cells.

The labelling of the edges of a $\mathcal{VH}$–complex as horizontal and vertical induces a labelling of the vertices in the link of any vertex as horizontal and vertical, thus making the link a bipartite graph. Similarly, the hyperplanes of a $\mathcal{VH}$–complex are also labelled as vertical and horizontal, with a vertical hyperplane being dual to an equivalence class of horizontal edges and a horizontal hyperplane being dual to an equivalence class of vertical edges.

Remark 2.6 Since the link of any vertex of a $\mathcal{VH}$–complex is bipartite, the length of a closed path is even. Thus a $\mathcal{VH}$–complex is nonpositively curved if there exists no bigon in any vertex link.

2.2 Graphs of spaces

Graphs of groups are the basic objects of study in Bass–Serre theory [35]. They were studied from a topological perspective in [33] by looking at graphs of spaces instead of graphs of groups. We will adopt this point of view.

Definition 2.7 By a *graph of spaces*, we mean the following data: Γ is a connected graph, called the underlying graph. For each vertex s (respectively edge a) of Γ, there is a topological space X_s (respectively X_a). Further, whenever a is incident to s, $\partial_{a,s}\colon X_a \to X_s$ is a π_1–injective continuous map. The *geometric realisation* of the above graph of spaces is the space $X = \big(\bigsqcup_{s\in\Gamma^{(0)}} X_s \sqcup \bigsqcup_{a\in\Gamma^{(1)}} X_a \times [0,1]\big)/{\sim}$, where $(x,0)$ and $(x,1)$ are identified respectively with $\partial_{a,s}(x)$ and $\partial_{a,s'}(x)$. Here, s and s' are the two endpoints of a.

Note that the universal cover of X has the structure of a *tree of spaces*, a graph of spaces whose underlying graph is the Bass–Serre tree of the associated graph of groups structure of X [33].

2.3 Tubular graphs of graphs

Definition 2.8 A *tubular graph of graphs* is a finite graph of spaces in which each vertex space is a finite connected simplicial graph and each edge space is a simplicial graph homeomorphic to a circle. Further, the attaching maps are simplicial immersions. We will always assume that the underlying graph is connected.

As a consequence of the definition, unless the underlying graph is trivial, no vertical graph is a tree. We note that asking for each vertex graph to be simplicial is not a serious restriction as every one-dimensional CW complex is a simplicial graph after subdivision.

We remark that not all graphs of free groups with cyclic edge groups can be realised as tubular graphs of graphs since we require both the images of each edge space to have identical length. This can be achieved, however, if the underlying graph is a tree.

It is easy to see (compare with [40, Theorem 1.18]) that:

Proposition 2.9 *The geometric realisation of a tubular graph of graphs is a finite* (*hence compact*), *connected nonpositively curved $\mathcal{VH}$–complex whose vertical hyperplanes are circles.*

Indeed, the geometric realisation is $\mathcal{VH}$, where vertical edges are the edges of vertex graphs and horizontal edges are the edges induced by vertices of edge graphs. Note that the squares are obtained from the "tubes" (edge graph times the unit interval, for each edge graph). Then the link of any vertex of a vertex graph does not have bigons because attaching maps of edge spaces are simplicial immersions.

Convention Throughout this article, we will use the same notation for a tubular graph of graphs and the $\mathcal{VH}$–complex which is its geometric realisation. X will always denote a Brady–Meier tubular graph of graphs with fundamental group G, while X_s will denote a vertex graph (a component of the vertical 1–skeleton) in X and $\widetilde{X}$ the CAT(0) universal cover of X. Unless mentioned otherwise, we work with the CAT(0) metric in $\widetilde{X}$.

Definition 2.10 (thickness) For an edge e in X, the *thickness* of e is the number of squares of X which contain e.

Observe that a horizontal edge of X always has thickness equal to two. Since X is Brady–Meier, we have:

Lemma 2.11 [37] *Every edge of X has thickness at least two.*

Definition 2.12 (paths, lines) Recall that a path in a space Z is a continuous map from a closed interval to Z.

A *combinatorial path* (see [25] for instance) is a map of graphs $\rho\colon P \to \Gamma$, where P is a subdivided compact interval and Γ is a graph. Further, all our combinatorial paths will be assumed to be immersions of graphs. P is always assumed to be oriented. When there is no confusion about Γ, we will refer to $\rho\colon P \to \Gamma$ as the path P. Unless mentioned to the contrary, a path between two vertices of X or $\widetilde{X}$ is a combinatorial path.

A *segment* is an embedded combinatorial path. Note that any compact graph homeomorphic to an interval is the image of a segment. We will often refer to such graphs as segments.

A *cycle* is an immersion of graphs $\phi\colon C \to \Gamma$, where C is a subdivided circle. We will often denote it by C.

A *line* is an isometric embedding $\mathbb{R} \hookrightarrow \widetilde{X}$ (with the $\mathrm{CAT}(0)$ metric), while a ray is an isometric embedding of $[0,\infty)$. A *combinatorial line* is an isometric embedding of graphs $R \to \widetilde{X}^1$, where R is the real line subdivided at integer intervals. We will only consider combinatorial lines that are also lines in the $\mathrm{CAT}(0)$ metric (see Remark 2.15 below).

Since horizontal edges of $\widetilde{X}$ are of thickness two, vertical hyperplanes of $\widetilde{X}$ are lines. Further:

Fact 2.13 The first cubical neighbourhood in $\widetilde{X}$ of a vertical hyperplane h, or the set of all closed squares of $\widetilde{X}$ that meet h, is convex [32] and hence isometric to a Euclidean strip $[0,1]\times\mathbb{R}$ with $\mathsf{h} \cong \{\frac{1}{2}\}\times\mathbb{R}$. Thus maximal geodesics in such a strip are of the form either $\{t_0\}\times\mathbb{R}$ or geodesics from $(0,x)$ to $(1,y)$.

We next divide the set of lines in $\widetilde{X}$ into the following three types.

Definition 2.14 A *vertical* line is a combinatorial line contained in a vertical tree. A *tubular* line is one that is parallel to a vertical hyperplane in the first cubical neighbourhood of the hyperplane. A *transversal* line is a line that hits at least two vertical trees.

Observe that a given line can be both vertical and tubular. We note that a tubular line that is not vertical is disjoint from the vertical 1–skeleton and hits any horizontal edge at most at one point, while a transversal line hits at least one vertical hyperplane (in exactly one point).

Remark 2.15 As mentioned in the introduction, edge groups of the JSJ decomposition are universally elliptic. By Lemma 8.4, transversal lines are not stabilised by universally elliptic subgroups. Hence, only vertical and tubular lines play a role in the analysis that follows.

3 Regular neighbourhoods and regular spheres

Recall that a cell of a square complex is either a vertex, an edge or a square.

Definition 3.1 (cubical neighbourhoods) The *first cubical neighbourhood* Y^{+1} of a subset Y of a square complex Z is a subcomplex of Z given by the union of all cells of Z that meet the closure of Y. The *n^{th} cubical neighbourhood* Y^{+n} is defined inductively as $(Y^{+(n-1)})^{+1}$.

Definition 3.2 (cubical subdivisions) The *first cubical subdivision* $Z^{(1)}$ of a square complex Z is a square complex obtained by subdividing Z in the following way: Each edge of Z is subdivided into two edges with the midpoint of the initial edge forming a new vertex. Each square of Z is subdivided into four squares of equal area by taking the centre of the square as a new vertex and taking four new edges between the centre of the square and each of the midpoints of the edges of the square. The *n^{th} cubical subdivision* $Z^{(n)}$ of Z is the first cubical subdivision of $Z^{(n-1)}$.

We will now define an abstract neighbourhood for a combinatorial path in a square complex. The path may not embed in the square complex, but it will embed in its abstract neighbourhood. Fix a combinatorial path ρ from P to the 1–skeleton of a square complex Z. Here the path may or may not be a cycle. We allow P to be a combinatorial ray or a combinatorial line. We remind the reader that ρ is an immersion of graphs. We will consider ρ as a map from P to the 1–skeleton of $Z^{(2)}$, the second cubical subdivision of Z.

Definition 3.3 The *regular neighbourhood* $N(P)$ of P in a square complex Z is a square complex constructed as follows. Let c be a cell of $Z^{(2)}$. We take one copy of c for each component of $\rho^{-1}(\mathsf{c})$ (see Figure 1). The adjacency of cells is given by the adjacency of arcs of P, where each arc is a component of the preimage of a cell of $Z^{(2)}$.

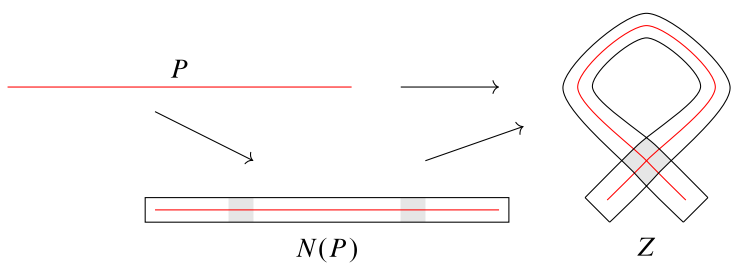

Figure 1: Two disjoint subpaths of P are mapped to the grey square.

Since ρ restricted to each arc of the preimage of a cell of $Z^{(2)}$ is an (isometric) embedding, we observe:

Fact 3.4 There is a natural embedding of P in $N(P)$ such that ρ factors through this embedding:

$$\begin{array}{ccc} P & \xrightarrow{\quad\rho\quad} & Z \\ & \searrow \quad \nearrow & \\ & N(P) & \end{array}$$

Note that the map $N(P) \to Z$ is an immersion. In particular, if ρ is an embedding, then $N(P)$ embeds in Z, since $\rho^{-1}(\mathsf{c})$ of any cell c contains a single component.

The reason for choosing the second cubical subdivision instead of the first in the definition of $N(P)$ is to make it easier to define certain operations (see Definition 3.7).

Definition 3.5 The *regular sphere* around P, denoted by $\partial N(P)$, is the union of all cells of $N(P)$ that are disjoint from P.

Fact 3.6 The regular sphere around a vertex is isomorphic as graphs to the first barycentric subdivision of the vertex link.

3.1 The regular sphere around an edge

The goal of this subsection is to show that the regular sphere around an edge of a square complex can be built from the regular spheres around its endpoints.

Let Y be a subset of a simplicial graph Γ. Recall that the *star of* Y, denoted by $\operatorname{star}(Y)$, is the subgraph of Γ consisting of all vertices and edges of Γ that meet Y. The *open star of* Y, denoted by $\mathring{\operatorname{star}}(Y)$, is the interior of $\operatorname{star}(Y)$.

Definition 3.7 [24] Let Γ_1 and Γ_2 be graphs. Let $v_1 \in \Gamma_1$ and $v_2 \in \Gamma_2$ be vertices of equal valence, say k. Let $\phi_i\colon \{1,\dots,k\} \to \operatorname{adj}(v_i)$ be a labelling of the vertices adjacent to v_i, for $i = 1,2$. Then the *spliced graph* $\Gamma_1 \,{}_{(v_1,\phi_1)}\!\oplus_{(v_2,\phi_2)} \Gamma_2$ is defined as a quotient of $\Gamma_1 \setminus \mathring{\operatorname{star}}(v_1) \sqcup \Gamma_2 \setminus \mathring{\operatorname{star}}(v_2)$, where $\phi_1(j)$ is glued to $\phi_2(j)$, for $1 \le j \le k$. If $v_1 \ne v_2$ are vertices in Γ_1 as above, then we define the *self-spliced graph* ${}_{(v_1,\phi_1)}\!\bigoplus_{(v_2,\phi_2)} \Gamma_1$ as a quotient of $\Gamma_1 \setminus (\mathring{\operatorname{star}}(v_1) \cup \mathring{\operatorname{star}}(v_2))$, where $\phi_1(j)$ is glued to $\phi_2(j)$, for $1 \le j \le k$.

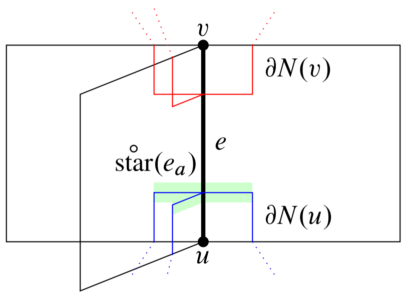

Figure 2: Regular spheres around two adjacent vertices. The star of e_a is highlighted in green.

Recall that the *dipole graph of order* d is a multigraph consisting of two vertices and d edges joining them.

Let e be an edge of a nonpositively curved square complex Z with endpoints u and v. We can now state the main result of this subsection.

Lemma 3.8 *The regular sphere around e is homeomorphic to the spliced graph of the regular spheres around u and v with the natural labelling induced by the squares containing e.*

Proof Let m be the midpoint of e. Then m is a vertex after a subdivision of the square complex. Observe that $\partial N(m)$ is homeomorphic to a dipole graph of order d, where d is the thickness of e. Let e_a be the initial half-edge of e and e_b its second half. Then e_a and e_b meet $\partial N(m)$ at distinct vertices of valence d, which we will also call e_a and e_b respectively (see Figure 2). Thus $\partial N(m) \setminus \mathring{\text{star}}(e_a) \cup \mathring{\text{star}}(e_b)$ is a disjoint union of d segments, one for each square that contains e.

Similarly, e_a (respectively e_b) meets $\partial N(u)$ (respectively $\partial N(v)$) at a vertex of valence d; see Figure 2. So $\partial N(u) \setminus \mathring{\text{star}}(e_a)$ (respectively $\partial N(v) \setminus \mathring{\text{star}}(e_b)$) is a graph with d "hanging" edges: edges with one of their endpoints having valence one. We thus see that $\partial N(e) \cong (\partial N(u) \setminus \mathring{\text{star}}(e_a) \sqcup \partial N(v) \setminus \mathring{\text{star}}(e_b))/\sim$ with the natural gluing. □

3.2 The regular sphere around a combinatorial path

Henceforth, till the end of Section 3, Z is either X or $\widetilde{X}$. Assume that P is not a vertex. Let e be an edge in P. If P is not a cycle, then P is a concatenation of paths P_1, e and P_2. If P is a cycle, we denote the connected complement of $\mathring{e}$ by P_1.

Lemma 3.9 *The regular sphere around P is homeomorphic to a*

(1) *spliced graph of the regular spheres around P_1 and P_2 (with labelling induced by the squares containing e) if P is not a cycle, and*

(2) *self-spliced graph of the regular sphere around P_1 (with labelling induced by the squares containing e) if P is a cycle.*

The proof is analogous to the proof of Lemma 3.8.

3.3 Connected regular spheres

For the rest of the section, P will always be a noncyclic path. We recall that Z is either X or $\widetilde{X}$ and P is either a compact interval, a combinatorial ray or a combinatorial line.

Recall that a point y of a topological space Y is said to be a *cut point* of Y if $Y \setminus \{y\}$ is not connected.

Lemma 3.10 *The regular sphere around any vertex or the midpoint of any edge of Z is connected and has no cut points if and only if Z is Brady–Meier.* □

We now state the main result of the section.

Proposition 3.11 *If P is compact, then the regular sphere around P is connected and has no cut points.*

The proof requires the following lemma. It was observed in [8], but without a proof.

Lemma 3.12 *Let Γ_1 and Γ_2 be connected graphs with no cut points. Suppose that Γ is the spliced graph $\Gamma_1 \,_{(v_1,\phi_1)}\!\oplus_{(v_2,\phi_2)} \Gamma_2$. Then Γ has no cut points.*

Proof First observe that $\Gamma_i \setminus \mathring{\mathrm{star}}(v_i)$ is connected by assumption. Let $v \in \Gamma$. We will show that v is not a cut point. Assume that $v \in \Gamma_1$. The important case to consider is of a point $x \neq v \in \Gamma_1$. Since $\Gamma_1 \setminus \{v\}$ is connected, there exists a path from x to v_1 in Γ_1 disjoint from v. Let u be a vertex adjacent to v_1 on this path. Then u is glued to a vertex of Γ_2 in Γ. Thus there exists a path from x to Γ_2 disjoint from v. □

Proof of Proposition 3.11 The proof is by induction on the length of P. If P is a vertex, then the result is obviously true. Suppose that P is of length at least one. Let e be an edge in P and P_1 and P_2 be subpaths such that P is the concatenation of P_1, e and P_2. By induction, the regular sphere around P_i has no cut points. Lemmas 3.9 and 3.12 then give the result for P. □

Lemma 3.13 *The regular sphere around a combinatorial ray P of $\widetilde{X}$ is connected.*

Proof Let $v \in \partial N(P)$. By Lemma 3.9, there exists $p \in P$ such that $v \in \partial N(p)$. If p_0 denotes the initial point of P, then P is a concatenation of the paths P_1 and P_2, where P_1 is the subpath of P from p_0 to p and P_2 is its complement. Since $\partial N(P_1)$ has no cut points (by Proposition 3.11, there exists a path in $\partial N(P_1)$ from v to a point v_0 in $\partial N(p_0)$ disjoint from P_2), the result follows. □

Corollary 3.14 (rays don't separate) *A combinatorial ray of $\widetilde{X}$ does not separate $\widetilde{X}$.*

The following powerful result for $\widetilde{X}$ will be used repeatedly in later sections.

Lemma 3.15 (path-abundance lemma) *Let P be a combinatorial geodesic in $\widetilde{X}$ and $x \in \widetilde{X} \setminus P$. Then, given $p \in P$, there exists a path α from x to p such that $\alpha \cap P = \{p\}$.*

Proof First note that $N(P)$ embeds in $\widetilde{X}$, by Fact 3.4. Let γ be a path from x to p. Let γ' be the maximal initial subpath of γ such that $\mathring{\gamma}' \cap P$ is empty. If γ' ends at p, then declare $\gamma' = \alpha$.

Suppose not. Let p' be the endpoint of γ'. Then P is a concatenation $P_1 \cdot [p', p] \cdot P_2$. By Proposition 3.11, the regular sphere around $[p', p]$ has no cut points. In particular, $\partial N([p', p]) \setminus P_1$ is connected. We recall that we denote the point at which P_i meets $\partial N([p', p])$ also as P_i.

Denoting $\gamma' \cap \partial N([p, p'])$ by γ', we note that there exists a path β between γ' and P_2 in $\partial N([p, p']) \setminus P_1$. Let h be a vertex adjacent to P_2 such that β meets h. Note that $h \in \partial N(p) \setminus P_2$. The required path α is a concatenation of γ', β and a path in $N(p)$ from h to p. □

4 Separating and coarsely separating lines

Recall that a subspace Y of a topological space Z *separates two points* z_1 and z_2 in Z if z_1 and z_2 lie in different components of $Z \setminus Y$. A subspace Y *separates* $Y' \subset Z$ if Y separates two points of Y'.

Definition 4.1 (separating lines) A *separating line* in $\widetilde{X}$ is a line that separates $\widetilde{X}$.

Given a subspace Y of a metric space Z, recall that $N_R(Y)$ denotes the set of points in Z at distance at most R from Y.

Definition 4.2 (coarsely separating lines [28]) A line L *coarsely separates* $\widetilde{X}$ if there exists $R > 0$ such that

(1) $N_R(L)$ separates $\widetilde{X}$, and

(2) there exist components $Y_1 \neq Y_2$ of $\widetilde{X} \setminus N_R(L)$ such that for any $R' \geq R$, $Y_i \nsubseteq N_{R'}(L)$.

Since a line is an embedding in $\widetilde{X}$, the regular sphere around a combinatorial line embeds in $\widetilde{X}$, by Fact 3.4.

Let h be a vertical hyperplane. Note that h is a combinatorial line in the first cubical subdivision of $\widetilde{X}$.

Definition 4.3 The regular sphere around a nonvertical tubular line L at distance at most $\frac{1}{2}$ from a vertical hyperplane h in $\widetilde{X}$ is defined to be the regular sphere around h in the first cubical subdivision of $\widetilde{X}$.

Lemma 4.4 *Let L be a combinatorial separating line in $\widetilde{X}$ and $P \subset L$ be a combinatorial subpath. Then L separates $\partial N(P)$.*

Proof Suppose the lemma is not true. Then note that $N(P) \setminus L$ is connected. Let $x, y \in \widetilde{X} \setminus L$. Fix $p \in P$. By Lemma 3.15, there exist paths α from x to p and β from y to p such that $\alpha \cap L = \beta \cap L = \{p\}$. Since $\partial N(P) \setminus L$ is connected, there exists a path in $\partial N(P) \setminus L$ between $\alpha \cap \partial N(P)$ and $\beta \cap \partial N(P)$. Thus x and y are not separated by L for any $x, y \in \widetilde{X}$, a contradiction. □

4.1 Separating implies coarsely separating

Throughout this subsection, L refers to a vertical or a tubular line.

Definition 4.5 (half-spaces of a line) A *half-space* of L is the closure in $\widetilde{X}$ of a component of $\widetilde{X} \setminus L$.

We warn the reader that there can be more than two half-spaces of a separating line in general.

An easy consequence of Lemma 3.15 is the following result:

Lemma 4.6 *Let Y be a half-space of L. Then $L \subset Y$.* □

In fact, we can read the number of half-spaces of L off its regular sphere:

Lemma 4.7 *There exists a natural map from the set of half-spaces of L to the set of components of the regular sphere around L. Further, this map is bijective.*

Proof Observe that each component of $\partial N(L)$ lies in a half-space of L. Let Y be a half-space of L, and $h_1, h_2 \in Y \cap \partial N(L)$. Then there exists a path between h_1 and h_2 in the component $Y \setminus L$. There also exists a path between h_1 and h_2 through L, since $h_i \in \partial N(L)$. These two paths between h_1 and h_2 bound a disk D, as $\widetilde{X}$ is simply connected, and $D \cap \partial N(L)$ gives a path between h_1 and h_2 in $\partial N(L)$. The required map is the one that sends a half-space Y of L to $Y \cap \partial N(L)$. □

Corollary 4.8 *Given an edge e in L, for each component K of $\partial N(L)$, there exists a square s containing e such that $\mathsf{s} \cap \partial N(L) \subset K$.*

Proof Let Y be the half-space of L corresponding to K, by Lemma 4.7. By Lemma 4.6, Y meets e. Let m be the midpoint of e. By Lemma 3.15, there exists a path between any point in the interior of Y to m that does not meet $L \setminus \{m\}$. Hence Y contains a square s that contains e and is as required. □

Fact 4.9 It is easy to see that L is a separating line whenever it is tubular. Clearly, if L is not vertical, then it separates the strip that contains it. Otherwise, any strip that contains L induces a component (line) of the regular sphere around L.

Lemma 4.10 *Let Y be a half-space of L. Then $Y \not\subseteq N_R(L)$ for any $R > 0$.*

Proof A hyperplane of a CAT(0) cube complex is, after subdivision, a CAT(0) subcomplex [32]. Thus every hyperplane of $\widetilde{X}$ is a tree. But since each edge of $\widetilde{X}$ is of thickness at least 2 (Lemma 2.11), every hyperplane is an unbounded tree. Observe that if L meets a hyperplane h at exactly one point, then h has points at arbitrarily large distances from L. It is easy to see that Y contains the interior of at least one square s. Choose s such that s meets L. Then the horizontal hyperplane through s meets L at a single point. □

Proposition 4.11 *L coarsely separates $\widetilde{X}$.* □

4.2 Coarsely separating periodic lines

Definition 4.12 The *translation length* of an element $g \in G$ is the infimum of $d(x, gx)$ over all $x \in \widetilde{X}$. An *axis of* $g \in G$ in $\widetilde{X}$ is a line L in $\widetilde{X}$ such that $gL \subset L$ and g moves an element of L by its translation length. A line L in $\widetilde{X}$ is *periodic* if it is an axis of some element of G.

Given $g \in G$, an axis in $\widetilde{X}$ of g always exists (see [6, Theorem II.6.8] for details).

Lemma 4.13 *Given a combinatorial periodic line L, either L is vertical or each vertical subpath of L is compact.*

Proof Let $g \in G$ be such that $gL \subset L$. Suppose that a vertical component of L is not compact, and hence contains a ray γ. Let e be an edge of L adjacent to γ. Then either g or g^{-1} sends e into γ. Since G sends vertical edges to vertical edges, e is vertical. Continuing this way, we conclude that L is vertical. □

The main result of this subsection is the following:

Proposition 4.14 *A periodic coarsely separating combinatorial line L of $\widetilde{X}$ separates $\widetilde{X}$.*

The proof uses the following lemma. Recall that h^{+1} is the first cubical neighbourhood of a hyperplane h.

Lemma 4.15 *Let Y be a half-space of a periodic combinatorial line L such that L is not contained in h^{+1} for any vertical hyperplane h in Y. Then $Y \setminus L^{+k}$ is connected for each $k \in \mathbb{N}$.*

Proof of Proposition 4.14 If L is contained in h^{+1} for some vertical hyperplane h, then L is tubular and hence separating (Fact 4.9). So assume that L is not contained in h^{+1} for any vertical hyperplane h. Suppose that L does not separate. Let $Y = \widetilde{X}$ be the unique half-space of L. By Lemma 4.15, $Y \setminus L^{+k}$ is connected for all k, implying that L does not coarsely separate. □

The proof of Lemma 4.15 requires some work. For the rest of the subsection, we fix a periodic combinatorial line L and a half-space Y of L such that L is not contained in h^{+1} for any vertical hyperplane in Y.

Remark 4.16 By [18, Lemma 13.15], L^{+k} is convex for any k.

Definition 4.17 A hyperplane h is *tangent to* a subcomplex Z of $\widetilde{X}$ if Z is disjoint from h but meets h^{+1}.

Fact 4.18 As L^{+K} is convex, any element of $L^{+(k+1)}$ is contained either in L^{+k} or in the first cubical neighbourhood of a hyperplane tangent to L^{+k}.

Lemma 4.19 *Given a vertical hyperplane h in Y and $k \in \mathbb{N}$, $\mathsf{h} \cap L^{+k}$ is compact.*

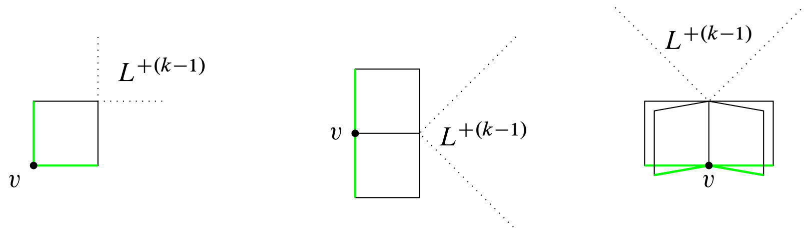

Figure 3: The edges of ∂L^{+k} are in green: no edge at v meeting $L^{+(k-1)}$ (left), a horizontal edge at v in L^{+k} (middle) and a vertical edge at v in L^{+k} (right).

Proof Suppose there exists a vertical hyperplane h such that $\mathsf{h} \cap L^{+k}$ is not compact. Let T be the underlying Bass–Serre tree of the tree of spaces structure of $\widetilde{X}$. By Lemma 4.13, the image of L in T is either a point or a line. We first claim that L is vertical. This follows from the fact that the image of h in T is a point and a ray of h is at finite Hausdorff distance from a ray of L.

Let α be the path in T between the image of h and the image of L. Let h' be the unique vertical hyperplane tangent to L such that its image in T lies in α. Then it is easy to see that $\mathsf{h}'^{+1} \cap L$ contains a ray. It then follows that $L \subset \mathsf{h}'^{+1}$ as both h' and L are periodic and G acts freely on $\widetilde{X}$. This is a contradiction. □

We will denote by ∂L^{+k} the set of all cells in L^{+k} disjoint from $L^{+(k-1)}$.

Lemma 4.20 *Let v be a vertex in ∂L^{+k}. Then exactly one of the following holds*:

(1) *One vertical and one horizontal edge incident to v lie in ∂L^{+k}.*

(2) *Two vertical edges* (*and no horizontal edge*) *incident to v lie in ∂L^{+k}.*

(3) *Finitely many horizontal edges* (*and no vertical edge*) *incident to v lie in ∂L^{+k}.*

Proof Since $v \notin L^{+(k-1)}$, observe that at most one edge incident to v meets $L^{+(k-1)}$. The three mutually exclusive cases to then consider are that of no edge incident to v meeting $L^{+(k-1)}$, a unique horizontal edge meeting $L^{+(k-1)}$ or a unique vertical edge meeting $L^{+(k-1)}$ (see Figure 3). □

Before proving Lemma 4.15, we will have to prove:

Lemma 4.21 *Let h_1 and h_2 be hyperplanes in Y tangent to L^{+k}. Suppose that $\mathsf{h}_1^{+1} \cap L^{+k}$ and $\mathsf{h}_2^{+1} \cap L^{+k}$ intersect. Then $\mathsf{h}_1^{+1} \cap \partial L^{+(k+1)}$ and $\mathsf{h}_2^{+1} \cap \partial L^{+(k+1)}$ lie in a component of $\partial L^{+(k+1)}$.*

We will denote $\mathsf{h}_i^{+1} \cap L^{+k}$ by σ_i.

Lemma 4.22 *Suppose that σ_1 and σ_2 are horizontal. Then $\sigma_1 \cap \sigma_2$ is a singleton.*

Proof Let v be a vertex in $\sigma_1 \cap \sigma_2$. If (1) or (2) of Lemma 4.20 holds at $\{v\}$, we are done. If (3) holds, observe that any horizontal edge f in σ_1 is in the first cubical neighbourhood of exactly two horizontal hyperplanes, h_1 and h, where h meets L^{+k}. □

By Lemma 4.19, σ_i is compact whenever it is vertical. Thus, $\sigma_1 \cap \sigma_2$ is always compact.

Lemma 4.23 *Let v be a terminal vertex in $\sigma_1 \cap \sigma_2$. Then either $\sigma_1 \cap \sigma_2 = \{v\}$ or v is a terminal vertex of σ_1 or σ_2.*

Proof If $\sigma_1 \cap \sigma_2$ contains an edge, then, by Lemma 4.22, σ_1 and σ_2 are both vertical. Then either (1) or (2) of Lemma 4.20 holds at the terminal vertex v. If (1) holds, then v is terminal in both σ_1 and σ_2. If (2) holds (see Figure 4), then it is easy to see that v is terminal in one of the two segments. □

Proof of Lemma 4.21 Let v be a terminal vertex of $\sigma_1 \cap \sigma_2$. Let e_i be the edge incident to v such that the hyperplane h_i passes through v. We have three cases given by Lemma 4.20:

Case 1 Only one vertical edge f incident to v lies in ∂L^{+k}. Since $\widetilde{X}$ is Brady–Meier, there exists a path β in $\operatorname{link}(v) \setminus \{f\}$ between e_1 and e_2. The projection of β to $\partial\{v\}^{+1}$ hits the other endpoints of e_1 and e_2, which lie in $\partial L^{+(k+1)}$. Further, β and thus its projection are disjoint from L^{+k}. Hence the result.

Case 2 Two vertical edges f_1 and f_2 incident to v lie in ∂L^{+k}. By Lemma 4.23, either σ_1, say, is horizontal or v is terminal in σ_1. Thus one of the edges, say f_2, does not lie in σ_1. Let β be a path in $\operatorname{link}(v) \setminus \{f_1\}$ between e_1 and e_2. Since f_2 does not lie in σ_1, the path β and its projection to $\partial\{v\}^{+1}$ are disjoint from σ_1. If f_2 does not lie in σ_2 or β is disjoint from f_2, then we are done as the projection of β gives the required path in $\partial L^{+(k+1)}$. If not, then we repeat the procedure at v', the other endpoint of f_2. We continue until the path no longer meets the compact σ_2. Hence the result.

Case 3 Only horizontal edges incident to v lie in ∂L^{+k}. Let f be the vertical edge incident to v and contained in L^{+k}. Let β be a path between e_1 and e_2 in $\operatorname{link}(v) \setminus \{f\}$. Then β is disjoint from L^{+k} and so is its projection to $\partial\{v\}^{+1}$. □

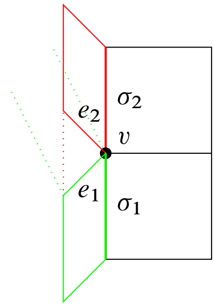

Figure 4: The edge e_i passes through the vertical hyperplane h_i.

We are now ready to prove Lemma 4.15.

Proof of Lemma 4.15 The proof is by induction. Note that $Y \setminus L^{+k}$ is connected whenever $Y \cap \partial L^{+k}$ is connected. Since Y is a half-space of L, $Y \cap \partial N(L)$ is connected, by Lemma 4.7. Thus $Y \cap \partial L^{+1}$ is connected.

Assume that $Y \cap \partial L^{+k}$ is connected for some k. We will now show that $Y \cap \partial L^{+(k+1)}$ is connected. Indeed, $L^{+(k+1)}$ is contained in the union of L^{+k} and the first cubical neighbourhoods of hyperplanes tangent to L^{+k}, by Fact 4.18. Thus, given two vertices u and u' in $Y \cap \partial L^{+(k+1)}$, there exist hyperplanes h and h' tangent to L^{+k} such that $u \in \mathsf{h}^{+1}$ and $u' \in \mathsf{h}'^{+1}$. Let $\sigma = \mathsf{h}^{+1} \cap L^{+k}$ and $\sigma' = \mathsf{h}'^{+1} \cap L^{+k}$. By the induction assumption, there exists a path between σ and σ' in ∂L^{+k}. This implies that there exists a finite sequence of tangent hyperplanes $\mathsf{h} = \mathsf{h}_1, \dots, \mathsf{h}_n = \mathsf{h}'$ such that if $\sigma_i = \mathsf{h}_i^{+1} \cap L^{+k}$, then $\sigma_i \cap \sigma_{i+1}$ is nonempty. Lemma 4.21 then implies that u and u' lie in a component of ∂L^{k+1}. Hence the result. □

5 A crossing criterion for lines

Definition 5.1 (crossing of lines) Let L and L' be two separating lines of $\widetilde{X}$. We say that L *crosses* L' if $L \not\subseteq Y'$ for any half-space Y' of L'. Two separating lines L and L' *don't cross* if neither L crosses L' nor L' crosses L.

Note that two disjoint lines don't cross. Thus a vertical line and a tubular line that is not vertical never cross. We will see later that in fact, no vertical line crosses a tubular line. Two intersecting lines may or may not cross. The main goal of this section is to obtain the following criterion for the crossing of two lines.

Proposition 5.2 (crossing criterion) *Let L and L' be two separating combinatorial lines in $\widetilde{X}$. Then L crosses L' if and only if*

(1) *$L \cap L' = P$ is nonempty and compact*, *and*

(2) *$L' \cap \partial N(P)$ separates $L \cap \partial N(P)$.*

Throughout this section, L and L' are two separating combinatorial lines and P denotes their intersection. The following three lemmas are consequences of the Brady–Meier property of $\widetilde{X}$ and the path-abundance lemma (Lemma 3.15), and we omit the proofs.

Lemma 5.3 *If P is either empty or noncompact*, *then L and L' don't cross.* □

Lemma 5.4 (crossing is symmetric) *L is contained in a half-space of L' if and only if L' is contained in a half-space of L.* □

Lemma 5.5 *L and L' don't cross if and only if for each half-space Y of L*, *there exists a half-space Y' of L' such that either $Y \subset Y'$ or $Y' \subset Y$*, *and similarly for each half-space Y' of L'*, *there exists a half-space Y of L such that either $Y \subset Y'$ or $Y' \subset Y$.* □

Before we go to the proof of Proposition 5.2, we will collect a couple of results about graphs without cut points as $\partial N(P)$ has no cut points whenever P is compact (Proposition 3.11).

5.1 Graphs with no cut points

We fix a connected graph Γ in this subsection such that Γ has no cut points. We further assume that Γ contains at least one edge. A *cut pair* is a pair of points that separates Γ.

We now draw the attention of the reader to certain similarities between cut pairs in Γ and separating lines in $\widetilde{X}$. If $\{a,b\}$ is a cut pair, then a *half-space* of $\{a,b\}$ is the closure of a component of $\Gamma \setminus \{a,b\}$. The first result is analogous to Lemma 4.6.

Lemma 5.6 *Let Y be a half-space of a cut pair $\{a,b\}$. Then $\{a,b\} \subset Y$.*

Proof Since Γ is connected, at least one of the two, say a, is contained in Y. If b is not contained in Y, then a is a cut point as a separates Y from b. □

The second result is analogous to Lemma 5.4.

Lemma 5.7 [7, Lemma 2.3(1)] *Let $\{a,b\}$ and $\{a',b'\}$ be cut pairs in Γ. Then $\{a',b'\}$ separates $\{a,b\}$ if and only if $\{a,b\}$ separates $\{a',b'\}$.*

Corollary 5.8 [7, Lemma 2.3(3)] *Let $\{a,b\}$ and $\{a',b'\}$ be cut pairs in Γ. If there exist at least three half-spaces of $\{a,b\}$, then $\{a',b'\}$ is not separated by $\{a,b\}$.*

5.2 The criterion

Proof of Proposition 5.2 By Lemma 4.4, both L and L' separate $\partial N(P)$. Suppose that P is compact and $L' \cap \partial N(P)$ separates $L \cap \partial N(P)$. Then, by Corollary 5.8, L' separates $\partial N(P)$ into exactly two components. This implies that L' has exactly two half-spaces as each half-space of L' meets P, by Lemma 4.6. This in turn implies that different components of $\partial N(P) \setminus L'$ are contained in different half-spaces of L'. Hence L crosses L'.

For the converse, note that if P is not compact then L and L' don't cross, by Lemma 5.3. Similarly, if $L' \cap \partial N(P)$ does not separate $L \cap \partial N(P)$, then clearly, L lies in a half-space of L'. □

6 Cyclic splittings and separating lines

Definition 6.1 A group G *splits over a subgroup* H if G decomposes as a nontrivial free product with amalgamation over H or as an HNN extension over H. In the sequel, a decomposition of G either as a free product with amalgamation or as an HNN extension (over H) will be called a *basic splitting* of G (over H).

Proposition 6.2 [35] *G splits over H if and only if G acts without edge inversions on an unbounded tree T such that H is the stabiliser of some edge of T and there exists no proper G–invariant subtree of T.*

As a consequence of Proposition 4.14, we have the following application of a result of Papasoglu:

Lemma 6.3 [27] *Let H be a cyclic subgroup over which G splits. Suppose that a vertical line L is an axis of H. Then L is a separating line that does not cross any of its translates.*

Lemma 6.4 *Let L be a periodic line that separates $\widetilde{X}$ and does not cross any of its translates. Then there exists an unbounded G–tree T_L and a vertex in T_L whose stabiliser is the stabiliser of L. Further, there exists no proper G–invariant subtree of T_L.*

The construction of such a *dual tree* when L has exactly two half-spaces and is disjoint from all its translates is standard. In that case, the dual tree is bipartite and is constructed as follows: Each component of $\widetilde{X} \setminus \bigsqcup_{g \in G} gL$ defines a black vertex while each translate of L defines a white vertex. The adjacency is given by containment: a white vertex is adjacent to a black vertex if it is contained in the closure of the black vertex.

In our case, L may not be disjoint from its translates and may have more than two half-spaces. This necessitates a more careful treatment, but the underlying idea is still the same. Our construction, in fact, coincides with the above standard construction when L is disjoint from its translates and has only two half-spaces.

We start with an observation that will be used in the proof.

Lemma 6.5 *Let L_1 and L_2 be separating lines that don't cross. Given half-spaces Y_1 of L_1 and Y_2 of L_2 such that $(Y_1 \setminus L_1) \cap (Y_2 \setminus L_2)$ is nonempty, then either L_1 is contained in Y_2 or L_2 is contained in Y_1.*

Proof Since L_1 and L_2 don't cross, there exist half-spaces Y_1' of L_1 and Y_2' of L_2 such that $L_1 \subset Y_2'$ and $L_2 \subset Y_1'$, by definition. We claim that either $Y_1' = Y_1$ or $Y_2' = Y_2$. Suppose not. Since $L_1 \subset Y_2'$, the separating line L_1 is disjoint from $\widetilde{X} \setminus Y_2' \supset Y_2 \setminus L_2$. Thus $Y_2 \setminus L_2$ is contained in a half-space Y_1'' of L_1. But the fact that L_2 is contained in Y_1' implies $Y_1'' = Y_1'$, and hence $Y_2 \setminus L_2$ is disjoint from $Y_1 \setminus L_1$, a contradiction. □

The required tree T_L will be the CAT(0) cube complex dual to a space with walls. Recall that:

Definition 6.6 [17] A *wall* on a nonempty set Z is a partition of Z into two subsets. Z is a *space with walls* if Z is endowed with a collection of walls such that any two points of Z are separated by finitely many walls.

Remark 6.7 The two subsets that define a wall are known as half-spaces in the literature. Note that we have already used this terminology for separating lines. Separating lines in $\widetilde{X}$ do define walls, as we will show below. We will refer to a half-space associated to a wall as a half-space of the space with walls.

We quickly recall some terminology of spaces with walls before going to the proof of Lemma 6.4. We refer the reader to [26] for further details.

Definition 6.8 Let Z be a space with walls. An *ultrafilter* on Z is a nonempty collection ω of half-spaces of Z that satisfy the following conditions:

(1) $A \in \omega$ and $A \subset B$ imply that $B \in \omega$, and

(2) exactly one of A and A^c is contained in ω.

Lemma 6.9 *Let ω be an ultrafilter on Z and $A, B \in \omega$. Then A and B are not disjoint.* □

For a $z \in Z$, the *principal ultrafilter* σ_z is defined to be the set of half-spaces of Z that contain z. An ultrafilter ω of Z is said to be *almost principal* if for some (and therefore for any) $z \in Z$, the symmetric difference between ω and σ_z is finite.

Proof of Lemma 6.4 Let $Z_L = \widetilde{X} \setminus \bigcup_{g \in G} gL$. Then each half-space Y of gL defines a wall $\{Y \cap Z_L, Y^c \cap Z_L\}$, which we will denote as $\{Y, Y^c\}$. It is easy to see that Z_L is a space with walls. By [26, Theorem 4.1], there exists a connected graph T_L whose vertices are the principal and almost principal ultrafilters of Z_L. Two vertices are adjacent if the cardinality of their symmetric difference is two. T_L is then the 1–skeleton of a unique CAT(0) cube complex (see [32, Section 3], for instance).

We claim that T_L is a tree. If not, then it is the 1–skeleton of a cube complex of dimension at least 2. Thus there exists a cycle (w_1, w_2, w_3, w_4) of length 4 in T_L. Since ω_1 and ω_2 are adjacent, there exists a wall $\{Y, Y^c\}$ of Z_L such that $Y \in \omega_1$ and $Y^c \in \omega_2$. Similarly, there exists a wall $\{Y', Y'^c\}$ such that $Y' \in \omega_1$ and $Y'^c \in \omega_4$. Note that $Y' \in \omega_2$, $Y \in \omega_4$ and $Y^c, Y'^c \in \omega_3$, by definition. We will show below that this is not possible. Assume that Y and Y' are half-spaces of the lines gL and $g'L$. By Lemma 6.9, Y and Y' are not disjoint. This implies that either $gL \subset Y'$ or $g'L \subset Y$, by Lemma 6.5. Assume the former. Either $Y \subset Y'$ or not. If $Y \subset Y'$, then no ultrafilter can contain both Y and Y'^c and hence ω_4 cannot exist. On the other hand, if $Y \not\subseteq Y'$, then $g'L$ meets Y in its interior and hence $Y'^c \subset Y$. This then implies that no ultrafilter can contain both Y^c and Y'^c and hence ω_3 cannot exist. This proves the claim.

There exists a natural action of G on T_L. An element $g \in G$ sends an ultrafilter ω to an ultrafilter $g\omega$ where $g\omega$ is the set of half-spaces gY of Z_L, where $Y \in \omega$.

We claim that there exists an ultrafilter whose stabiliser is the stabiliser of L. Let $Y_1, \dots, Y_n$ be the set of half-spaces of L, and let ω_L be the set of half-spaces of Z_L consisting of $Y_1^c, \dots, Y_n^c$ and all half-spaces (of proper translates of L) which contain L.

Note that ω_L is an ultrafilter. (If $n=2$, then $\omega_L \setminus \{Y_1\}$ and $\omega_L \setminus \{Y_2\}$ are ultrafilters and hence ω_L is a vertex of T_L after subdivision.) We claim that ω_L is almost principal. Indeed, choose $y_1 \in Y_1 \cap Z_L$. Then $\sigma_{y_1} = \{Y_1, Y_2^c, \dots, Y_n^c\} \cup \{Y \mid y_1 \in Y\}$. There exist at most finitely many lines $g_1 L, \dots, g_k L$ that separate y_1 from L in $\widetilde{X}$. Except for the half-spaces of these lines, a half-space contains y_1 if and only if it contains L. Hence $\sigma_{y_1} \triangle \omega_L$ is finite. It is straightforward to check that $\mathrm{stab}(L) = \mathrm{stab}(\omega_L)$.

Lemma 4.10 implies that T_L is unbounded because a vertex at maximal distance from ω_L will contain a half-space that does not contain a translate of L in its interior.

There is no proper G–invariant subtree of T_L. Now T_L is spanned as a tree by the principal ultrafilters of Z_L, by [26, Proposition 4.8]. It thus suffices to prove that no subtree spanned by a proper subset of the set of principal ultrafilters is G–invariant. Choose $y_i \in Y_i \cap Z_L$ such that there exists a path α from y_i to L with $\mathring{\alpha} \subset Z_L$. Then observe that any principal ultrafilter σ_y is a translate of σ_{y_i}, for some i. Thus, if a proper subtree is G–invariant, then it has to miss at least one σ_{y_i}, say σ_{y_1}. But this is not possible as the interior of Y_1 contains at least one translate of L, by Lemma 4.10. □

Proposition 6.10 *Let H be a cyclic subgroup of G and L an axis of H in $\widetilde{X}$. Suppose that L is a separating line that does not cross any of its translates and that H is equal to the stabiliser of a proper subset of the set of half-spaces of L. Then G splits over H.*

Proof Let T_L be the dual tree of L obtained from Lemma 6.4. Let $Y_1, \dots, Y_n$ be the half-spaces of L and ω_i be the vertex adjacent to ω_L such that $\omega_i \triangle \omega_L = \{Y_i, Y_i^c\}$. Let T be the quotient simplicial graph of T_L obtained by first identifying for each $h \in H$ and $i \in \{1, \dots, n\}$, vertices ω_i and $h\omega_i$, and then extending equivariantly. It is easy to check that T is a tree that satisfies the conditions of Proposition 6.2 for H. □

A cyclic subgroup H of G that satisfies the hypothesis of Proposition 6.10 is a *geometric splitting subgroup*.

Proposition 6.11 *Let H be a cyclic subgroup of G over which G splits. If H has a vertical axis in $\widetilde{X}$ then there exists a geometric splitting subgroup H' commensurable with H.*

Proof Let L be an axis of H satisfying Lemma 6.3. Observe that H is contained in the stabiliser of L, which is a cyclic subgroup. Choose a half-space Y of L and let H' be the largest subgroup of the stabiliser of L that preserves Y. Then, by Proposition 6.10, H' is as required. □

7 Vertical cycles and cyclic splittings

In this section, we will examine splittings induced by vertical lines in $\widetilde{X}$. Recall that a cycle (Definition 2.12) is an immersion of graphs $\phi\colon C \to \Gamma$, where C is a subdivided circle. From now on, throughout the text, unless mentioned otherwise, Γ will be a vertex graph X_s of X and so C is a vertical cycle.

Remark 7.1 The map ϕ is π_1–injective. Indeed, $\pi_1(C)$ injects into $\pi_1(X_s)$ [36] and $\pi_1(X_s)$ injects into the fundamental group of X in the graph of groups setup [35].

Fact 7.2 A line in $\widetilde{X}$ is periodic and vertical if and only if it is a lift of a vertical cycle in X.

By abuse of notation, we will often refer to the lift $\widetilde{\phi}\colon \widetilde{C} \to \widetilde{X}$ as the line $\widetilde{C}$. Since the projection of the regular neighbourhood of a path onto the path is a deformation retraction, we have:

Lemma 7.3
$$\partial N(C) \cong \partial N(\widetilde{C})/\pi_1(C).$$
□

Definition 7.4 A *cyclic path* is an immersed combinatorial path $\rho\colon P \to X_s$ such that the initial and terminal vertices of P have the same image while the initial and terminal edges of P have distinct images.

A cyclic path P induces a quotient cycle $\phi_P\colon C_P \to X_s$, where C_P is the quotient of P obtained by gluing the initial and terminal vertices and defining $\phi_P([x]) := \rho(x)$.

Definition 7.5 (fundamental domain of a cycle) Let $\phi\colon C \to X_s$ be a cycle. A cyclic path $\rho_C\colon P_C \to X_s$ with induced quotient cycle C_{P_C} is said to be a *fundamental domain of* C if the following diagram commutes:

$$\begin{array}{ccc} C_{P_C} & \xrightarrow{\cong} & C \\ & {\scriptstyle \phi_{P_C}} \searrow \quad \swarrow {\scriptstyle \phi} & \\ & X_s & \end{array}$$

Remark 7.6 It is easy to see that for the action of $\pi_1(C)$ on $\widetilde{C}$, a lift $\widetilde{P}_C$ of P_C is a fundamental domain of $\widetilde{C}$ in the usual sense.

Definition 7.7 Let P_C be a fundamental domain of a cycle C. Let u and v be the initial and terminal vertices of P_C and a and b the initial and terminal edges. Let b_u be the vertex of $\partial N(u)$ that meets b and a_v the vertex of $\partial N(v)$ that meets a.

The *orthogonal sphere around* P_C, denoted by $\partial_{\mathrm{orth}} N(P_C)$, is defined as the closure of $\partial N(P_C) \setminus (\{b_u\}^{+2} \cup \{a_v\}^{+2})$, where $\{b_u\}^{+2}$ (respectively $\{a_v\}^{+2}$) denotes the second cubical neighbourhood in $\partial N(P_C)$ of b_u (respectively a_v).

Let $\widetilde{C}$ be a lift of C and $\widetilde{P}_C \subset \widetilde{C}$ of P_C. Then note that:

Fact 7.8 The natural map given by $P_C \cong \widetilde{P}_C \hookrightarrow \widetilde{C}$ induces an embedding of graphs $\partial_{\mathrm{orth}} N(P_C) \hookrightarrow \partial N(\widetilde{P}_C) \setminus \widetilde{C} \subset \partial N(\widetilde{C})$ as a deformation retract. Further, $\partial_{\mathrm{orth}} N(P_C)$ is connected if and only if $\partial N(\widetilde{P}_C) \setminus \widetilde{C}$ is connected.

It thus follows from Lemma 7.3 that:

Lemma 7.9 *The regular sphere around a cycle C is isomorphic to the quotient of the orthogonal sphere around a fundamental domain P_C of C with the natural gluing induced by $\pi_1(C)$.*

Let $\rho_C \colon P_C \to X_s$, and let e' be an edge in P_C. By Corollary 4.8, we have:

Lemma 7.10 *Let K be a component of $\partial_{\mathrm{orth}} N(P_C)$. Then there exists a square s in $N(P_C)$ that meets e' and $\mathsf{s} \cap \partial_{\mathrm{orth}} N(P_C) \subset K$.* □

Definition 7.11 A cycle C is a *UC-separating* (*universal cover separating*) *cycle* if $\widetilde{C}$ is a separating line. C is *strongly UC-separating* if $\partial N(C)$ is not connected.

By Lemma 4.4, we have:

Lemma 7.12 *If C is a UC-separating cycle, then $\partial_{\mathrm{orth}} N(P_C)$ is not connected.*

Lemma 7.13 *C is strongly UC-separating if and only if the following two conditions are satisfied*:

(1) *C is a UC-separating cycle*, *and*

(2) *$\pi_1(C)$ does not act transitively on the set of half-spaces of $\widetilde{C}$.*

Proof Recall that $\partial N(C) \cong \partial N(\widetilde{C})/\pi_1(C)$ (Lemma 7.3). A component of $\partial N(C)$ lifts to a component of $\partial N(\widetilde{C})$. So, if $\partial N(C)$ is connected, then either $\partial N(\widetilde{C})$ is itself connected or every component of $\partial N(\widetilde{C})$ projects onto $\partial N(C)$. So $\pi_1(C)$ acts transitively on the components of $\partial N(\widetilde{C})$ and therefore on the components of $\widetilde{X} \setminus \widetilde{C}$ (Lemma 4.7). The converse is clear. □

Definition 7.14 A cycle $\phi'\colon C' \to X_s$ is an n^{th} *power* of the cycle $\phi\colon C \to X_s$ if there exists an n–fold covering map $\psi\colon C' \to C$ such that the following diagram commutes:

$$\begin{array}{ccc} C' & & \\ \big\downarrow \psi & \overset{\phi'}{\searrow} & \\ & \underset{\phi}{\nearrow} & X_s \\ C & & \end{array}$$

Lemma 7.15 *Let N be such that the thickness of any edge of X is at most N. Given a UC-separating cycle C, there exists $n \leq N$ such that the regular sphere around an n^{th} power of C is not connected.*

Proof By Corollary 4.8, the number of half-spaces of $\widetilde{C}$ is at most N. Thus there exists a subgroup H of index at most N of $\pi_1(C)$ that does not act transitively on the set of half-spaces of $\widetilde{C}$. The required cycle C' is the quotient of $\widetilde{C}$ by H. □

Definition 7.16 Let $\rho_C\colon P_C \to X_s$ be a fundamental domain of a cycle C. A *subcycle* of C is the quotient cycle of a cyclic path $\rho_C|_P\colon P \to X_s$ with $P \subset P_C$.

Observe that if C' is an n^{th} power of C, then C is a subcycle of C'. We will often use this fact.

Definition 7.17 A UC-separating cycle C has a *self-crossing* if $\widetilde{C}$ and a translate cross.

Definition 7.18 A cycle C is a *splitting cycle* if the following conditions are satisfied:

(1) C is strongly UC-separating,

(2) $\pi_1(C)$ equals the stabiliser of a proper subset of the set of half-spaces of $\widetilde{C}$, and

(3) C has no self-crossings.

Remark 7.19 By Proposition 6.10, G splits over $\pi_1(C)$ whenever C is a splitting cycle.

We will now examine when C can have self-crossings. We start with the following.

Lemma 7.20 *Let L_1 and L_2 be vertical lines of $\widetilde{X}$ stabilised by H_1 and H_2 respectively. Let n_i be the translation length of a generator of H_i. If the length of $P = L_1 \cap L_2$ is at least $\mathrm{LCM}(n_1, n_2)$, then $P = L_1$.*

Recall that L_i is an axis of H_i. Hence a generator of H_i translates every point of L_i by n_i (see [6, Theorem II.6.8(i)]). Since any element of H_i takes vertices to vertices, n_i is indeed an integer.

Proof Suppose that P contains a segment of length $\mathrm{LCM}(n_1, n_2) = k$. Let v be a terminal point of P. Choose generators $h_1 \in H_1$ and $h_2 \in H_2$ such that $h_i(v) \in P$. Since the length of P is at least k, $h_i^{k/n_i}(v) \in P$ and hence $h_1^{k/n_1}(v) = h_2^{k/n_2}(v)$. As G acts freely on $\widetilde{X}$, $h_1^{k/n_1} = h_2^{k/n_2}$ and hence $L_1 = L_2$. □

Corollary 7.21 *If L_2 is a translate of a periodic vertical line L_1, then either $L_2 = L_1$ or P embeds in L_1/H_1, where H_1 is the stabiliser of L_1.*

In particular, for a cycle C, if $g \cdot \widetilde{C} \neq \widetilde{C}$, then $P = g \cdot \widetilde{C} \cap \widetilde{C}$ embeds in C.

Definition 7.22 A segment $P \subsetneq C$ is said to be a *component of self-intersection* of C if there exists a translate $g\widetilde{C} \neq \widetilde{C}$ such that the projection to C of $\widetilde{C} \cap g\widetilde{C}$ is equal to P. We say that there is a *self-crossing of C at P* if there exists a $g \in G$ such that $\widetilde{C} \cap g\widetilde{C} = P$ and $\widetilde{C}$ and $g\widetilde{C}$ cross.

Fact 7.23 Let $P \subset C$ be a segment so that a lift of P in $\widetilde{C}$ is isomorphic to P (and hence also denoted by P). Then $\partial N(P) \cap \partial N(C) \simeq \partial N(P) \setminus \widetilde{C}$. In other words, there is a self-crossing at P only if $g\widetilde{C}$ meets different components of $\partial N(P) \cap \partial N(C)$, by Proposition 5.2.

Lemma 7.24 *A splitting cycle is strongly UC-separating and has no self-crossing at any component of self-intersection.* □

Also, splitting cycles capture all "vertical" splittings up to commensurability:

Lemma 7.25 *Let H be a cyclic subgroup over which G splits with a vertical axis. Then there exists a splitting cycle C such that $\pi_1(C)$ is commensurable with a conjugate of H.*

Proof By Proposition 6.11, H is commensurable with a geometric splitting subgroup H'. Let L be a vertical axis for H. The required splitting cycle is obtained by taking the quotient of L by H'. □

As an immediate consequence, we have:

Lemma 7.26 *Given a UC-separating cycle C with no self-crossings, there exists a splitting cycle C' such that $\pi_1(C)$ and $\pi_1(C')$ are commensurable.* □

8 Universally elliptic splittings

Recall that a subgroup H of G is *elliptic* in a G–tree T if H fixes a point in T.

Lemma 8.1 (elliptic splittings) *Let H_1 and H_2 be cyclic subgroups over which G splits. Let L_i be an axis of H_i in $\widetilde{X}$. H_1 is elliptic in the Bass–Serre tree of the basic splitting over H_2 if and only if L_1 and gL_2 don't cross for any translate gL_2 of L_2.*

Proof Note that if L_1 and gL_2 don't cross for any g, then L_1 is contained in a half-space of gL_2 for each g. Thus, for $x \in L_1 \setminus L_2$, the stabiliser of σ_x in the dual tree T_{L_2} (Lemma 6.4) of L_2 contains H_1 and hence H_1 is elliptic in T_{L_2}. The Bass–Serre tree T_2 of the basic splitting over H_2 is obtained from T_{L_2} by a sequence of G–equivariant gluings of edges of T_{L_2}. Thus H_1 remains elliptic in T_2. Conversely, if there exists g such that L_1 and gL_2 cross, then $g^{-1}H_1g$ is hyperbolic in the dual tree T_{L_2} and hence in T_2. □

Remark 8.2 Since G is one-ended, H_1 is elliptic in the Bass–Serre tree of the basic splitting over H_2 if and only if H_2 is elliptic in the Bass–Serre tree of the basic splitting over H_1; see [29, Theorem 2.1].

Definition 8.3 [16] A cyclic splitting of G over the subgroup H is *universally elliptic* if H is elliptic in the Bass–Serre tree of any cyclic splitting of G. We then say that H is a universally elliptic subgroup. Analogously, a splitting cycle C is universally elliptic if $\pi_1(C)$ is universally elliptic.

A splitting induced by a transversal line can never be universally elliptic:

Lemma 8.4 *Let H be a cyclic subgroup over which G splits. Suppose that an axis of H is transversal in $\widetilde{X}$. Then H is not universally elliptic.*

Proof Let L be a transversal axis of H. By definition, there exists a vertical hyperplane h such that $L \cap \mathsf{h}$ is a singleton. Since h separates $\widetilde{X}$ and is either equal to or disjoint from its translates, it induces a splitting of G. Let T be the Bass–Serre tree of the splitting. Let e be the edge stabilised by the stabiliser of h. Note that the image of e under H then spans a line of T. Hence, H is not elliptic in T. □

Splittings induced by vertical lines need more careful treatment. They may or may not cross other vertical or transversal lines which induce splittings. We present below one sufficient condition for a splitting induced by a vertical line (cycle) to be universally elliptic.

Proposition 8.5 *Let L be a line that separates $\widetilde{X}$ into at least three half-spaces. Then a subgroup of the stabiliser of L is universally elliptic.*

Proof Let L' be a separating line such that L and L' meet at a compact segment P. Since $\partial N(P) \setminus L$ has at least three components, by Corollary 5.8, $L' \cap \partial N(P)$ lies in a component of $\partial N(P) \setminus L$. Hence, L and L' don't cross. In particular, L does not cross any of its translates. Let H be a maximal subgroup of the stabiliser of L that preserves a half-space of L. Then, by Proposition 6.10, G splits over H and H is universally elliptic. □

9 Repetitive cycles

Definition 9.1 Let $\rho\colon P \to X$ be a combinatorial path. Let e be an edge in X and e' an edge in $\rho^{-1}(e)$. Denote also by e' the image of e' in $N(P)$. Given a square s in X containing e, denote by s' (Figure 5) the union of all squares meeting e' in $N(P)$ whose image in X is contained in s. Then the *preimage of* s *around* e' in $\partial N(P)$ is defined as the segment $\mathsf{s}' \cap \partial N(P)$.

Recall that the orthogonal sphere of any fundamental domain P_C of a UC-separating cycle C contains at least two components (Lemma 7.12). By Lemma 7.10, for each component K of $\partial_{\mathrm{orth}} N(P_C)$ and each edge e' in P_C with image e in X, there exists a square s containing e such that the preimage of s around e' lies in K.

Definition 9.2 (repetitive cycles) Let C be a UC-separating vertical cycle. C is a *k–repetitive cycle* if there exists a vertical edge e in X and a fundamental domain P_C of C such that

(1) at least k distinct edges $e_1, \dots, e_k$ of P_C are mapped to e, and

(2) for each square s containing e, there exists a component K of $\partial_{\mathrm{orth}} N(P_C)$ such that for each $i \in \{1, \dots, k\}$, the preimage of s around e_i in $\partial N(P_C)$ lies in K.

Intuitively, if C is k–repetitive, then the squares at e do not "mix" in the components of $\partial_{\mathrm{orth}} N(P_C)$. In other words, the notion of repetitiveness not only requires the cycle "repeat" itself along some edges (condition (1)) but also ensures that the partitions induced by the cycle on the set of squares containing e_i coincide. Note that the definition depends on the choice of a fundamental domain, as illustrated in Figure 6.

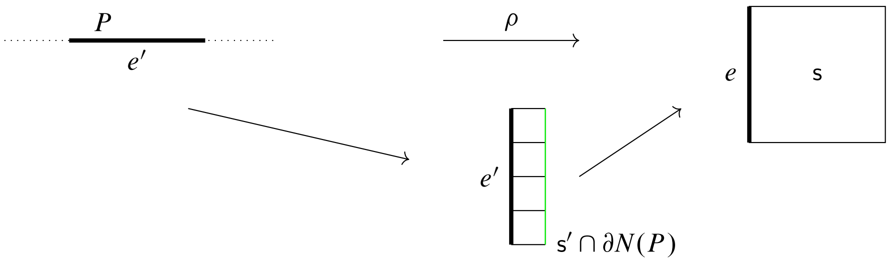

Figure 5: A preimage of a square in the regular sphere.

Fact 9.3 A k–repetitive cycle is k'–repetitive for $1 \le k' \le k$.

The following property of lifts of repetitive cycles will be crucial for the rest of the article. In fact, this is the only property of repetitive cycles that we will use.

Lemma 9.4 *Let C be a k–repetitive cycle. Then there exists an edge $\tilde{e}$ in $\widetilde{X}$ and distinct elements $g_1, \dots, g_k \in G$ such that*

(1) *for each $i \in \{1, \dots, k\}$, the translate $g_i \widetilde{C}$ contains $\tilde{e}$,*

(2) *for each $i \in \{1, \dots, k\}$, the translation length of g_i is strictly less than the length of C, and*

(3) *any two squares $\tilde{\mathsf{s}}$ and $\tilde{\mathsf{s}}'$ that contain $\tilde{e}$ are either separated by all translates $g_i \widetilde{C}$ for $i \in \{1, \dots, k\}$ or by none of them.*

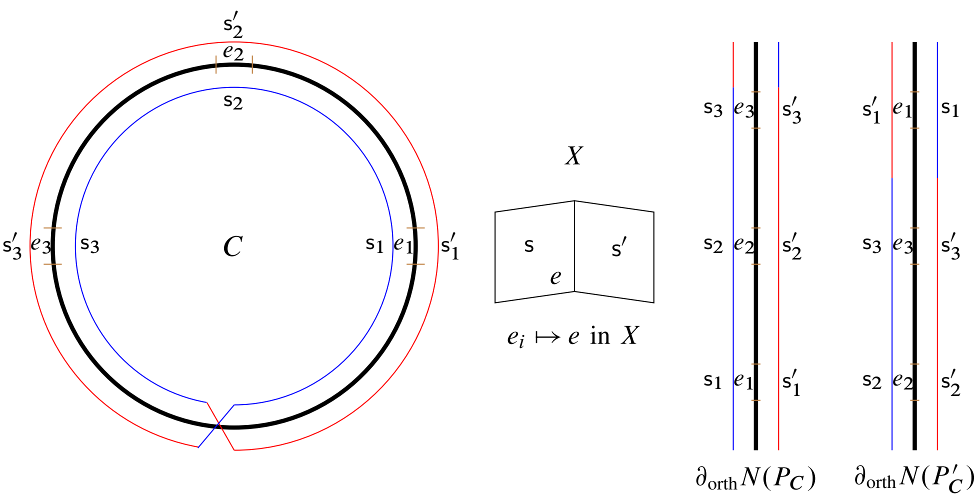

Figure 6: C is 3–repetitive with the fundamental domain P_C but not with the fundamental domain P'_C. The red and blue curves represent $\partial N(C)$.

Proof Let C be k–repetitive with fundamental domain P_C so that there exist edges $e_1, \dots, e_k$ in P_C that satisfy the conditions of Definition 9.2. Let $\widetilde{P}_C \subset \widetilde{C}$ denote a lift of P_C in $\widetilde{X}$. Denote $\tilde{e}_1$ in $\widetilde{P}_C$ by $\tilde{e}$. Since the edges $\tilde{e}_i$ in $\widetilde{P}_C$ all have the same image e in X, there exist $1 = g_1, \dots, g_k$ such that $g_i \tilde{e}_i = \tilde{e}$. Then, clearly, for each $i \in \{1, \dots, k\}$, the translation length of g_i is strictly less than the length of C and $g_i \widetilde{C}$ contains $\tilde{e}$.

Let $\tilde{\mathsf{s}}$ and $\tilde{\mathsf{s}}'$ be two squares that contain $\tilde{e}$. Then the squares $g_i^{-1}(\tilde{\mathsf{s}})$ and $g_i^{-1}(\tilde{\mathsf{s}}')$ in $\widetilde{P}_C$ are lifts of squares s_i and s_i' in P_C. Let D and D' be components of $\partial_{\mathrm{orth}} N(P_C)$ such that the preimage s_1 of s around e_1 meets D and the preimage s_1' of s' around e_1 meets D'. By definition, the corresponding preimage s_i meets D and s_i' meets D' for all i. Now $\tilde{\mathsf{s}} = \tilde{\mathsf{s}}_1$ and $\tilde{\mathsf{s}}' = \tilde{\mathsf{s}}_1'$ lie in different half-spaces of $g_1 \widetilde{C} = \widetilde{C}$ if and only if $D, D' \subset \partial N(\widetilde{C})$ (since $\partial_{\mathrm{orth}} N(P_C) \hookrightarrow \partial N(\widetilde{C})$, by Fact 7.8) meet different half-spaces of $\widetilde{C}$. Also, $\tilde{\mathsf{s}}$ and $\tilde{\mathsf{s}}'$ lie in different components of $g_i \widetilde{C}$ if and only if $\tilde{\mathsf{s}}_i = g_i^{-1} \tilde{\mathsf{s}}$ and $\tilde{\mathsf{s}}_i' = g_i^{-1} \tilde{\mathsf{s}}'$ lie in different half-spaces of $\widetilde{C}$ if and only if D and D' induce different half-spaces of $\widetilde{C}$. □

As a consequence, we have the following useful result when at least two of the translates in the above lemma are not equal. Let $k \geq 2$ and assume that C is a k–repetitive cycle. Let $g_1, \dots, g_k \in G$ be as in Lemma 9.4.

Lemma 9.5 *Suppose that at least two translates $g_i \widetilde{C}$ and $g_j \widetilde{C}$ are distinct. Then $\widetilde{C}$ separates $\widetilde{X}$ into exactly two half-spaces.*

We will need the following result on graphs with no cut points to prove the lemma.

Lemma 9.6 *Let Γ be a graph with no cut points. Let $\{a, b\}$ and $\{a', b'\}$ be cut pairs. Suppose there exist points $h_1, h_2, h_3 \in \Gamma \setminus \{a, b, a', b'\}$ that are pairwise separated by $\{a, b\}$ and also pairwise separated by $\{a', b'\}$. Then $\{a, b\} = \{a', b'\}$.*

Compare with [4, Lemma 3.8].

Proof By Corollary 5.8, $\{a, b\}$ lies in a half-space Y' of $\{a', b'\}$ and Lemma 5.7 implies that $\{a', b'\}$ lies in a half-space Y of $\{a, b\}$. Let $Y', Y_1', \dots, Y_n'$ be the half-spaces of $\{a', b'\}$. If $\{a, b\} \neq \{a', b'\}$, then $Y_1' \cup \dots \cup Y_n'$ lies in the half-space Y of $\{a, b\}$ that contains $\{a', b'\}$. This is a contradiction as at most one h_i lies in Y'. □

Proof of Lemma 9.5 After a reordering if necessary, we assume $\widetilde{C} = g_1 \widetilde{C}$ and $g_2 \widetilde{C}$ are distinct. By Lemma 4.7, it suffices to show that $\partial N(C)$ has exactly two components.

By Corollary 7.21, the segment $S = \widetilde{C} \cap g_2\widetilde{C}$ is compact. Suppose that $\widetilde{C}$ has at least three half-spaces. Then both $\partial N(S) \setminus \widetilde{C}$ and $\partial N(S) \setminus g_2\widetilde{C}$ have at least three components, by Lemma 3.15. This means that there exist three squares $\tilde{s}$, $\tilde{s}'$ and $\tilde{s}''$ containing $\tilde{e}$ that meet different components of both $\partial N(S) \setminus \widetilde{C}$ and $\partial N(S) \setminus g_2\widetilde{C}$, by Corollary 4.8 (as any pair of squares at $\tilde{e}$ are either separated by both lines or by none, by Lemma 9.4). Lemma 9.6 then implies that $\partial N(S) \cap \widetilde{C} = \partial N(S) \cap g_2\widetilde{C}$, which is a contradiction. □

Definition 9.7 A vertical cycle $\phi\colon C \to X_s$ is *maximal* if it is not a nontrivial power of any cycle.

Note that any cycle which induces a maximal cyclic subgroup has to be maximal, justifying the name.

Corollary 9.8 *A lift in $\widetilde{X}$ of a maximal k–repetitive cycle separates $\widetilde{X}$ into exactly two half-spaces whenever $k \geq 2$.*

Proof When C is maximal, the element g_2 that moves $\tilde{e}_2$ to $\tilde{e}_1$ in $\widetilde{P}_C$ does not preserve $\widetilde{C}$. This is because the translation length of g_2 is strictly less than the translation length of the generator of $\pi_1(C)$. If $g_2 \in \operatorname{stab}(\widetilde{C})$, then $\langle g_2 \rangle$ and $\pi_1(C)$ are contained in a common cyclic subgroup and hence $\pi_1(C) \lneq \operatorname{stab}(\widetilde{C})$, contradicting the fact that C is maximal. Hence $g_2\widetilde{C} \neq \widetilde{C}$. Lemma 9.5 then gives the result. □

We will end the section with a crucial result that bounds the length of nonrepetitive UC-separating cycles. Let E denote the number of vertical edges of X and F denote the number of squares of X.

Proposition 9.9 (long cycles are repetitive) *Let C be a vertical UC-separating cycle with length at least $E(k-1)2^{F(F+1)/2}+1$. Then C is k–repetitive.*

Proof The key ingredient in the proof is the pigeonhole principle. We apply it twice to show that C satisfies both the conditions of Definition 9.2. We give the details below.

Fix a fundamental domain P_C of C. Since there are E vertical edges in X, by the pigeonhole principle, there exists an edge e in X such that n oriented edges $e_1, \dots, e_n$ of C are mapped to e, with $n \geq (k-1)2^{F(F+1)/2}+1$.

Let $\lambda \leq F$ be the thickness of e. Note that the number of components μ of $\partial_{\text{orth}} N(P_C)$ is at most λ, by Lemma 7.10. We would like to show that there exist k edges out of $e_1, \dots, e_n$ for which the conditions of Definition 9.2 are satisfied. Denote by $A(\lambda, \mu)$ the number of ways in which the squares $\mathsf{s}_1, \dots, \mathsf{s}_\lambda$ containing e can be partitioned into exactly μ nonempty subsets. If $n > (k-1)A(\lambda, \mu)$, then, by the pigeonhole principle, k edges which satisfy the conditions of Definition 9.2 exist. Since $n \geq (k-1)2^{F(F+1)/2} + 1$ and $\lambda \leq F$, it is enough to show that $A(\lambda, \mu) \leq 2^{\lambda(\lambda+1)/2}$. Note that if $\mu = 1$, then $A(\lambda, \mu) = 1$ for any λ. Also, since no subset of the partition of the squares can be empty, any subset can contain at most $\lambda - \mu + 1$ squares. Hence

$$A(\lambda, \mu) = \sum_{r=1}^{\lambda-\mu+1} \binom{\lambda}{r} A(\lambda - r, \mu - 1).$$

Observe $A(\lambda_2, \mu) \geq A(\lambda_1, \mu)$ whenever $\lambda_2 \geq \lambda_1 \geq \mu$. Thus $A(\lambda, \mu) \leq 2^{\lambda(\lambda+1)/2}$. □

10 3–repetitive cycles and crossings

The main result of the section is the following.

Proposition 10.1 *Let C be a maximal UC-separating vertical cycle that is 3–repetitive. Then there exists a periodic separating line L' in $\widetilde{X}$ such that L' and $\widetilde{C}$ cross.*

If C has self-crossings, then, by definition, $\widetilde{C}$ and a translate cross, and there is nothing to show. The nontrivial part is to show that one such line exists even when C has no self-crossings. Henceforth, till the end of this section, C refers to a maximal 3–repetitive cycle C with no self-crossings.

The key idea behind the proof is the following. By Corollary 9.8, $\widetilde{C}$ separates $\widetilde{X}$ into exactly two half-spaces. Further, by Lemma 9.4, since C is 3–repetitive, there exists an edge in $\widetilde{X}$ along which three translates of $\widetilde{C}$ meet. We will show that one of these translates separates the other two. The periodic line L' will then be constructed by ensuring that it meets both the separated translates outside the central translate. This implies that L' crosses the central translate of $\widetilde{C}$. We give details below.

10.1 Lifts of C

Fix a fundamental domain P_C of C and edges e_1, e_2 and e_3 with image e in X_s satisfying the conditions of Definition 9.2. Let d be the thickness of e. Denote also by e a lift of the edge e in a vertex graph $\widetilde{X}_s$ of $\widetilde{X}$ such that for $1 \leq i \leq 3$, there exist $L_i = g_i \widetilde{C}$ that satisfy the conclusions of Lemma 9.4. Thus L_1, L_2 and L_3 contain e and if $\mathscr{S} = \{\mathsf{s}_1, \dots, \mathsf{s}_d\}$ is the set of squares containing e, then:

Lemma 10.2 *There exists a partition $A \sqcup B$ of $\mathscr{S}$ such that two squares s and s' in $\mathscr{S}$ lie in A (or B) if and only if for each $i \in \{1,2,3\}$, they lie in a single half-space of L_i.* □

By Corollary 9.8, we have:

Lemma 10.3 *Each L_i separates $\widetilde{X}$ into exactly two half-spaces.* □

Fix a square $\mathsf{s} \in A \subset \mathscr{S}$. For $1 \le i \le 3$, let Y_i be the half-space of L_i that contains s. Let $\overline{Y}_i$ be its complementary half-space. A set-theoretic consequence of Lemma 10.2 is that the half-spaces of L_1, L_2 and L_3 are nested. Namely:

Lemma 10.4 *For $i, j \in \{1,2,3\}$, either $Y_i \subset Y_j$ or $Y_j \subset Y_i$.*

Proof We will prove the lemma for Y_1 and Y_2. Observe that neither $\overline{Y}_1 \subset Y_2$ nor $Y_2 \subset \overline{Y}_1$, as otherwise a square in B lies in Y_2 or a square in A lies in $\overline{Y}_1$. Thus, by Lemma 5.5, either $Y_1 \subset Y_2$ or $Y_2 \subset Y_1$. □

After a reordering if necessary, assume that $Y_1 \subset Y_2 \subset Y_3$. We then have:

Lemma 10.5 *L_1 and L_3 lie in complementary half-spaces of L_2.*

Proof First, $L_1 \subset Y_2$ as $Y_1 \subset Y_2$. Similarly, $\overline{Y}_3 \subset \overline{Y}_2$ (as $Y_2 \subset Y_3$) implies $L_3 \subset \overline{Y}_2$. □

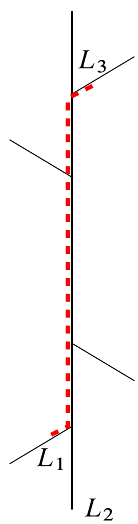

Figure 7: L' crosses L_2 if it meets L_1 and L_3 outside L_2.

10.2 The main result

Fix orientations on L_1, L_2 and L_3 such that they agree on e. As $L_1 \cap L_2$ is bounded, we can choose an element h_1 in the stabiliser of L_1 such that

(1) $h_1(Y_1) = Y_1$ (and thus $h_1(\overline{Y}_1) = \overline{Y}_1$), and

(2) $h_1(e)$ lies in $L_1 \setminus L_2$ before e in the orientation of L_1.

Recall that h_1 acts by translation on L_1. Similarly, choose an element h_3 in the stabiliser of L_3 such that

(1) $h_3(Y_3) = Y_3$, and

(2) $h_3(e)$ lies in $L_3 \setminus L_1$ after e in the orientation of L_3.

Let L' be the axis of $h' = h_3 \cdot h_1^{-1}$ in the vertical tree $\widetilde{X}_s$ that contains e. By definition, L' is periodic. Observe that:

Lemma 10.6 *L' contains $h_1(e)$ and, therefore, $h_3(e) = h'(h_1(e))$.*

Proof Suppose that $h_1(e)$ does not belong to L'. Let $h_1(u)$ be the initial vertex of $h_1(e)$ and $h_1(v)$ the final vertex. Let α be the geodesic from $h_1(u)$ to L'. Then $d(h_1(u), h' \cdot h_1(u)) = 2\ell(\alpha) + |h'|$, where $\ell(\alpha)$ is the length of α and $|h'|$ is the translation length of h' (see [35, Proposition 24 in Section I.6.4]). We also have $d(h_1(v), h' \cdot h_1(v)) = d(h_1(u), h' \cdot h_1(u)) - 2$. But since $h' \cdot h_1(u) = h_3(u)$ and $h' \cdot h_1(v) = h_1(v)$ are adjacent, $d(h_1(v), h_3(v)) = d(h_1(u), h_3(u))$, which is a contradiction. Hence $h_1(e)$ lies in L'. □

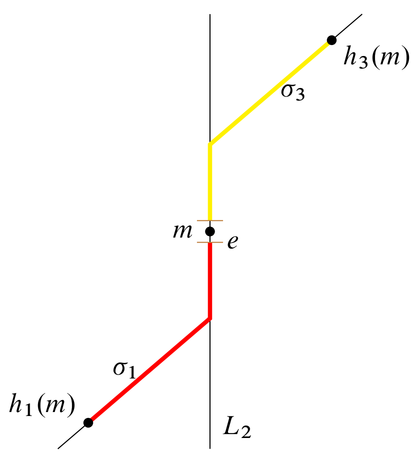

Figure 8: The segments σ_1 and σ_3.

Lemma 10.7 *L' is a separating line.*

Proof Let C' be the cycle obtained by taking the quotient of L' by the action of $\langle h' \rangle$. We will show that C' is strongly UC-separating. This will prove that L' is separating (see Lemma 7.13).

Let m be the midpoint of e. Subdivide $\widetilde{X}$ so that m, $h_1(m)$ and $h_3(m)$ are vertices of L'. Let σ be the geodesic segment from $h_1(m)$ to $h_3(m)$. Since $h' \cdot h_1(m) = h_3(m)$ and h' sends every element in the interior of σ outside σ, the segment σ is a fundamental domain for h' acting on L' and hence a fundamental domain of C'. We will first show that L' separates $\partial N(\sigma)$. Note that:

(1) $\sigma = \sigma_1 \cdot e \cdot \sigma_3$, where σ_1 is the segment (see Figure 8) in L_1 from $h_1(m)$ to the initial vertex of e and σ_3 is the segment in L_3 from the final vertex of e to $h_3(m)$.

(2) By Lemma 3.9, $\partial N(\sigma) \cong \partial N(\sigma_1) \oplus \partial N(\sigma_2)$, with labelling induced by the squares containing e.

(3) $\partial N(\sigma_i) \setminus L_i = \partial N(\sigma_i) \setminus L'$ as both L_i and L' meet $\partial N(\sigma_i)$ at e and $h_i(e)$.

Recall that $\partial N(\sigma_i) \setminus L_i$ is not connected (Lemma 4.4) and L_i has exactly two half-spaces (Lemma 10.3). Thus L_i induces a partition $A_i \sqcup B_i$ on the set of components of $\partial N(\sigma_i) \setminus L_i$ such that the components in A_i meet one half-space of L_i and the components in B_i meet the other half-space of L_i. Further, by Lemma 10.2, for each square $\mathsf{s} \in \mathscr{S}$, $\mathsf{s} \cap \partial N(\sigma_1)$ meets A_1 if and only if $\mathsf{s} \cap \partial N(\sigma_3)$ meets A_3. Thus, there exists no path between a point in A_1 and a point in B_3 in the spliced graph $\partial N(\sigma) \setminus L'$. Hence $\partial N(\sigma)$ is separated by L'. Thus $\partial_{\mathrm{orth}} N(\sigma)$ is not connected (Fact 7.8).

Because h_1^{-1} preserves half-spaces of L_1, it sends a square containing $h_1(e)$ in A_1 (respectively B_1) to a square containing e in A (respectively B). Similarly, h_3 sends a square in A (respectively B) to A_3 (respectively B_3). In other words, there is no path between a point in A_1 and a point in B_3 in the quotient of $\partial_{\mathrm{orth}} N(\sigma)$ by the action of h'. By Lemma 7.9, $\partial N(C')$ is not connected. □

Proof of Proposition 10.1 By Lemma 10.7, L' is a separating line. By Lemma 10.6, L' crosses L_2. □

11 An algorithm of double exponential time

The main result of this section is the following theorem.

Theorem 11.1 *There exists an algorithm of double exponential time complexity that takes a Brady–Meier tubular graph of graphs X with hyperbolic fundamental group G as input and returns a finite list of splitting cycles that contains all universally elliptic cycles up to commensurability.*

For the rest of the article, we will also assume that G is δ–hyperbolic. Denote by $\partial\widetilde{X}$ the *Gromov boundary* of $\widetilde{X}$. We refer the reader to [6] for background on hyperbolic groups and the Gromov boundary.

As G is hyperbolic and one-ended, [4, Lemma 5.21] and Proposition 10.1 imply that:

Proposition 11.2 *If C is a maximal 3–repetitive UC-separating cycle, then $\pi_1(C)$ is not universally elliptic.*

Proof of Theorem 11.1 Let G and X be as in the statement. By Lemma 8.4, every universally elliptic subgroup H has a vertical axis in $\widetilde{X}$. This implies that there exists a splitting cycle C in X such that $\pi_1(C)$ and H are commensurable (Lemma 7.25).

Let F be the number of squares of X and E the number of edges. By Proposition 9.9, any UC-separating cycle of length greater than $M = 2E(2^{F(F+1)/2}) \leq 8F(2^{F(F+1)/2})$ is 3–repetitive. By Proposition 11.2, a universally elliptic cycle C is either of length at most M, or it is a power of a maximal UC-separating cycle C' of length at most M. If it is the latter, then C' is not universally elliptic because it is not strongly UC-separating. By Lemma 7.15, C is an n^{th} power of C', where n is bounded by the maximal thickness of an edge of C'. Thus the length of C is at most $F \cdot M$.

There exist finitely many cycles of length at most $F \cdot M$ in X. Thus our algorithm takes each cycle from this finite list as input and returns whether this cycle is strongly UC-separating with no self-crossings or not. By Lemma 7.26, we thus have a list of all universally elliptic cycles up to commensurability.

The time taken by this algorithm is calculated as follows:

(1) The number of cycles of length at most $F \cdot M$ is bounded by a number which is exponential in $F \cdot M$ (see [1] for instance). This is of the order of a double exponential in F as M is itself exponential in F.

(2) The regular sphere around a cycle C of length k is a spliced graph of the regular spheres around its k vertices (Lemma 3.9), and the number of vertices and edges in this regular sphere is bounded by a constant times the number F of squares of X. Finding whether this sphere is connected is linear in F, by [19].

(3) A cycle C has a self-crossing if there exists a self-crossing at a component of self-intersection $P \subset C$ (Definition 7.22). There is a self-crossing at P only if a subpath of C meets $\partial N(P) \cap \partial N(C)$ in different components (Fact 7.23). This information is available when the regular sphere around C is computed and does not cost any additional time.

The algorithm thus takes double exponential time in the number of squares of X. □

12 Constructing a JSJ complex

The goal of this section is to construct from X a tubular graph of graphs X_{jsj} whose graph of groups structure gives the JSJ decomposition of G.

12.1 Splitting cycles as hyperplanes

Let $\phi\colon C \to X_s \subset X$ be a splitting vertical cycle. We will show how to modify X to a tubular graph of graphs X_C such that $\pi_1(X) \cong \pi_1(X_C)$ and $\pi_1(C)$ is commensurable with the cyclic group generated by a vertical hyperplane of X_C.

We perform this construction at the level of universal covers using the machinery of spaces with walls [17] (utilised earlier in Section 6). We refer the reader to [26] and [10] for details on constructing CAT(0) cube complexes from spaces with walls.

Recall that $\widetilde{X}$ is the cube complex dual to the space with walls $(\widetilde{X}^0, \mathcal{H})$ [31], where $\widetilde{X}^0$ denotes the 0–skeleton of $\widetilde{X}$ and $\mathcal{H}$ the set of hyperplanes of $\widetilde{X}$. For our purposes, we slightly modify the space with walls as follows. First we attach a strip $S_L = \mathbb{R} \times [0,1]$ isomorphically along $\mathbb{R} \times \{0\}$ to each translate L of $\widetilde{C}$. Note that there is a natural square structure on S_L so that every horizontal hyperplane of $\widetilde{X}$ that meets L naturally extends to S_L. Let Z be the set of open horizontal half-edges of the union of $\widetilde{X}$ and the attached strips. Then the vertical and horizontal hyperplanes of $\widetilde{X}$ induce a space with walls (Z, W).

Note that we do not add the vertical hyperplanes through the strips S_L to the collection W. Thus the dual cube complex of (Z, W) is nothing but $\widetilde{X}$ [31, Theorem 10.3]. We now enrich W to W_C. The walls in W_C are determined by

(i) the horizontal hyperplanes of $\widetilde{X}$,

(ii) the vertical hyperplanes of $\widetilde{X}$, and

(iii) the G–translates of $\widetilde{C}$.

Note that the elements of type (i) and (ii) induce W, where each half-space Y in $\widetilde{X}$ of an element of type (i) or (ii) defines a wall $\{Y \cap Z, Y^c \cap Z\}$. Given a translate L of $\widetilde{C}$ in $\widetilde{X}$, each half-space Y in $\widetilde{X}$ of L defines a wall $\{Y \cap Z, Y^c \cap Z\}$ of type (iii). Thus L induces exactly K walls in Z if it has K half-spaces in $\widetilde{X}$.

Lemma 12.1 *(Z, W_C) is a space with walls.* □

Denote by $\widetilde{X}_C$ the CAT(0) cube complex dual to Z.

Lemma 12.2 *$\widetilde{X}_C$ is a $\mathcal{VH}$–complex.*

The proof uses the following observation, which is a consequence of Lemma 5.5.

Lemma 12.3 *Let L and L' be two noncrossing lines of $\widetilde{X}$. Then, given half-spaces Y of L and Y' of L', at least one of the following four intersections is empty: $\mathring{Y} \cap \mathring{Y}'$, $Y^c \cap \mathring{Y}'$, $\mathring{Y} \cap Y'^c$ and $Y^c \cap Y'^c$.*

Two walls $\{Y, Y^c\}$ and $\{Y', Y'^c\}$ in a space with walls *cross* [10] if all four intersections $Y \cap Y'$, $Y^c \cap Y'$, $Y \cap Y'^c$ and $Y^c \cap Y'^c$ are nonempty.

Proof of Lemma 12.2 Two walls of type (i) don't cross as two horizontal hyperplanes of $\widetilde{X}$ are either equal or disjoint. Similarly, two walls of type (ii) don't cross. By Lemma 12.3, two walls of type (iii) don't cross either. Further, a wall of type (ii) and a wall of type (iii) don't cross since a vertical line is disjoint from any vertical hyperplane. By [26, Proposition 4.6], there exists a bijective correspondence between the hyperplanes of $\widetilde{X}_C$ and the walls of (Z, W_C). Further, two hyperplanes in $\widetilde{X}_C$ intersect if and only if the corresponding walls cross. Declare an edge e of $\widetilde{X}_C$ to be vertical if and only if the hyperplane through e corresponds to a wall of type (i). Otherwise, declare the edge to be horizontal. No square contains two adjacent edges of the same type as otherwise two hyperplanes of the same type or two hyperplanes of type (ii) and (iii) intersect. □

Observe that there exists a natural G–equivariant map $\hat{\eta}_C \colon \widetilde{X}_C \to \widetilde{X}$ such that the following diagram commutes:

$$\begin{array}{ccc} (Z, W_C) & \xrightarrow{\ \mathrm{id}\ } & (Z, W) \\ \downarrow & & \downarrow \\ \widetilde{X}_C & \xrightarrow{\ \hat{\eta}_C\ } & \widetilde{X} \end{array}$$

Since $W \subset W_C$, the map $\hat{\eta}_C$ takes any vertex (ultrafilter) σ' of $\widetilde{X}_C$ to a vertex $\sigma' \cap W$ of $\widetilde{X}$. But every vertex of $\widetilde{X}$ is a principal ultrafilter, and hence $\hat{\eta}_C(\sigma') = \sigma_z$ for some z. By the way the set of walls W_C was defined on Z, we have the following result:

Lemma 12.4 *The map $\widehat{\eta}_C$ has the following properties:*

(1) *Let c be a cell of $\widetilde{X}$ that does not meet any translate of $\widetilde{C}$. Then $\widehat{\eta}_C$ restricted to $\widehat{\eta}_C^{-1}(\mathsf{c})$ is injective.*

(2) *It sends vertical edges to vertical edges and horizontal edges to either horizontal edges or vertices.*

(3) *A horizontal edge is mapped to a vertex if and only if the vertical hyperplane through this edge is induced by a wall of type* (iii). □

Lemma 12.5 *For any $z \in Z$, the preimage $\widehat{\eta}_C^{-1}(\sigma_z)$ is a finite horizontal tree. Further, the edges in the preimage of σ_z are dual to vertical hyperplanes induced by translates of $\widetilde{C}$ that meet σ_z in $\widetilde{X}$.*

Proof Let σ'_1 and σ'_2 be two vertices of $\widehat{\eta}_C^{-1}(\sigma_z)$. Let $\{Y, Y^c\}$ be a wall such that $Y \in \sigma'_1$ and $Y^c \in \sigma'_2$. Then, clearly, $\{Y, Y^c\}$ is a wall of type (iii). Let L be the line that defines $\{Y, Y^c\}$. We claim that L passes through the vertex σ_z in $\widetilde{X}$. If not, then let h be a hyperplane of $\widetilde{X}$ that separates L from σ_z. Let Y_{h} be a half-space of h that contains the vertex σ_z. Then $Y_{\mathsf{h}} \in \sigma_z$, the ultrafilter. Clearly, this implies that $Y_{\mathsf{h}} \in \sigma'_1$ and $Y_{\mathsf{h}} \in \sigma'_2$. Since L and h are disjoint, either $Y_{\mathsf{h}} \subset Y$ or $Y_{\mathsf{h}} \subset Y^c$. Thus either $Y \in \sigma'_1$ and $Y \in \sigma'_2$ or $Y^c \in \sigma'_1$ and $Y^c \in \sigma'_2$, a contradiction. So L has to pass through σ_z. There are only finitely many translates of $\widetilde{C}$ that meet at any given point of $\widetilde{X}$. This proves the result. □

Since $\widehat{\eta}_C$ is a finite-to-one G–equivariant map, we conclude that:

Lemma 12.6 *G acts geometrically on $\widetilde{X}_C$.* □

Lemma 12.7 *Every vertical hyperplane of $\widetilde{X}_C$ is a line.*

Proof The stabiliser of a vertical hyperplane is the stabiliser of a wall of either type (ii) or type (iii), and hence is a cyclic subgroup. Thus every vertical hyperplane is a line. □

The complex $\widetilde{X}_C$ consists of two types of subcomplexes:

- Denote by $\widetilde{Z}_C$ a connected component of the subcomplex of $\widetilde{X}_C$ consisting of the union of the first cubical neighbourhood of all hyperplanes corresponding to walls of type (iii). In other words, $\widetilde{Z}_C$ is a connected component of the closed strips in $\widetilde{X}_C$ induced by half-spaces of translates of $\widetilde{C}$.
- The second type of subcomplex, denoted by $\widetilde{Y}_C$, is the closure of the complement in $\widetilde{X}_C$ of the G–translates of $\widetilde{Z}_C$.

Lemma 12.8 *The subcomplex $\widetilde{Z}_C$ is a tree of finite trees whose underlying tree $\widehat{\eta}_C(\widetilde{Z}_C)$ is a copy of $\widetilde{\phi(C)}$.*

Proof By Lemma 12.5, $\widehat{\eta}_C^{-1}(\sigma_z)$ is a finite tree for every vertex of $\widetilde{X}$, and thus after subdivision, for the midpoint of every edge of $\widetilde{X}$. The horizontal edges of $\widetilde{Z}_C$ are all dual to hyperplanes of type (iii). Note that $\widehat{\eta}_C$ sends such horizontal edges to vertices and vertical edges to vertical edges (Lemma 12.4). Thus $\widehat{\eta}_C(\widetilde{Z}_C)$ is a tree and $\widetilde{Z}_C$ is a tree of finite trees.

We now claim that $\widehat{\eta}_C(\widetilde{Z}_C)$ is a copy of the universal cover of $\phi(C)$. Note that $\widetilde{\phi(C)}$ is a connected union of lines which are translates of $\widetilde{C}$. Since $\widehat{\eta}_C(\widetilde{Z}_C)$ is also a union of translates of $\widetilde{C}$ with image $\phi(C)$ in X, $\widehat{\eta}_C(\widetilde{Z}_C) \subset \widetilde{\phi(C)}$. Conversely, if a vertex v of a translate L of $\widetilde{C}$ is contained in $\widehat{\eta}_C(\widetilde{Z}_C)$, then $\widehat{\eta}_C^{-1}(v)$ meets the strips induced by half-spaces of L and thus these strips are contained in $\widetilde{Z}_C$. The image of any such strip under $\widehat{\eta}_C$ is L and thus $L \subset \widehat{\eta}_C(\widetilde{Z}_C)$. □

Define $X_C := \widetilde{X}_C/G$. By Lemma 12.7, X_C is a tubular graph of graphs. The space X_C is called the *opened-up space of X along C*. The G–equivariant map $\widehat{\eta}_C\colon \widetilde{X}_C \to \widetilde{X}$ induces a map $\eta_C\colon X_C \to X$. Let Y_C and Z_C denote the respective images of $\widetilde{Y}_C$ and $\widetilde{Z}_C$ in X_C. We have proved that:

Lemma 12.9 *Z_C is a graph of finite trees with underlying graph $\phi(C)$ and with the following property: If $u \in \phi(C)$ is a vertex (or a midpoint of an edge), then the vertex (edge) tree $T(u)$ is the tree dual to Z with the walls induced by translates of $\widetilde{C}$ passing through a lift $\tilde{u}$ of u in $\widetilde{X}$.*

We conclude with the following observation:

Lemma 12.10 *The opened-up space X_C is a union of the subcomplexes Y_C and Z_C with $Y_C \cap Z_C$ consisting of those cells of Y_C that are mapped by η_C to $\phi(C)$.* □

12.2 Algorithmic construction of X_C

The main result of this subsection is the following:

Theorem 12.11 *There exists an algorithm of exponential time complexity that takes a Brady–Meier tubular graph of graphs X and a splitting cycle $\phi\colon C \to X$ as input and returns the opened-up space X_C along C as output.*

Define a complex Y'_C as the square complex obtained from $X \setminus \phi(C)$ by "completing the missing cells" as follows: for each vertex or edge x of $\phi(C)$, take as many copies of x as the number of squares of X that contain x and add them to the semi-open squares of $X \setminus \phi(C)$ to obtain closed squares. Observe that:

Lemma 12.12 *Y'_C is isomorphic to Y_C.* □

Therefore, in order to construct X_C algorithmically, it only remains to construct Z_C. The first result we will need is the following. Fix a lift $\widetilde{C}$ of the splitting cycle $\phi\colon C \to X_s \subset X$. Let K be the number of half-spaces of $\widetilde{C}$ in $\widetilde{X}$.

Proposition 12.13 *There exists $D' \in \mathbb{N}$ such that for any vertex or (midpoint of an edge) $v \in \widetilde{C}$ and any $D \geq D'$, the D^{th} cubical neighbourhood $\{v\}^{+D}$ of v has the following properties:*

(1) *Every translate $g\widetilde{C}$ that satisfies $v \in g\widetilde{C}$ separates $\{v\}^{+D}$ into exactly K components.*

(2) *$g\widetilde{C} \cap \{v\}^{+D} \neq g'\widetilde{C} \cap \{v\}^{+D}$ for any $g, g' \in G$ that satisfy $g\widetilde{C} \neq g'\widetilde{C}$ and $v \in g\widetilde{C} \cap g'\widetilde{C}$.*

The main ingredient for proving Proposition 12.13 is the following result. Let N be such that the thickness of any edge of X is at most N.

Lemma 12.14 *Let C_N be a $2^{N^{\text{th}}}$ power of C and P_N a fundamental domain of C_N. Then there exists a natural bijection between the set of half-spaces of $\widetilde{C}$ and the set of components of $\partial_{\text{orth}} N(P_N)$.*

We first prove a preliminary result on the number of connected components of graphs. Let Γ be a graph with no cut points and $\{a, b\}$ a cut pair. Assume that the valence n of a is equal to the valence of b. We will construct a spliced graph (Definition 3.7) of finitely many copies of Γ. Let $\phi_a\colon \{1, \dots, n\} \to \operatorname{adj}(a)$ and $\phi_b\colon \{1, \dots, n\} \to \operatorname{adj}(b)$ denote labellings of vertices adjacent to a and b. For each $i \in \mathbb{N}$, let Γ_i be a copy of Γ with the corresponding cut pair $\{a_i, b_i\}$. We will denote the labellings on the adjacent vertices by ϕ_{a_i} and ϕ_{b_i}. Let $\Gamma'_i := \Gamma_1 \,{}_{(b_1,\phi_{b_1})}\!\oplus_{(a_2,\phi_{a_2})} \Gamma_2 \,{}_{(b_2,\phi_{b_2})}\!\oplus \cdots \oplus_{(a_i,\phi_{a_i})} \Gamma_i$.

Lemma 12.15 *Suppose that the number of components of $\Gamma \setminus \{a, b\}$ is equal to the number of components of $\Gamma'_2 \setminus \{a_1, b_2\}$. Then, for each i, the number of components of $\Gamma'_i \setminus \{a_1, b_i\}$ is equal to the number of components of $\Gamma'_2 \setminus \{a_1, b_2\}$.*

Proof Let k be the number of components of $\Gamma \setminus \{a,b\}$. $\Gamma'_2 \setminus \{a_1,b_2\}$ has the same number of components as $\Gamma \setminus \{a,b\}$ if and only if there is a partition into k subsets of $\{1,\dots,n\}$ such that the corresponding partition induced by $\phi_a\colon \{1,\dots,n\} \to \mathrm{adj}(a)$ and $\phi_b\colon \{1,\dots,n\} \to \mathrm{adj}(b)$ on the vertices adjacent to a and b coincides with the partition induced by the k components of $\Gamma \setminus \{a,b\}$. Continuing iteratively, we obtain the result. □

Proof of Lemma 12.14 Denote by P_k a fundamental domain of a $2^{k^{\text{th}}}$ power C_k of C. Assume that for each k, P_k has been chosen such that P_k is a concatenation of two copies of P_{k-1}. The result then follows from Lemmas 12.15 and 3.9. □

Let L be a line in $\widetilde{X}$ and $v \in L$ a vertex. Let $D \in \mathbb{N}$. Note that $\{v\}^{+D}$ is a CAT(0) cube (sub)complex [18]. We will assume that $\partial N(L) \subset \widetilde{X}$. By Lemma 4.7, we have:

Lemma 12.16 *There exists a bijection between the half-spaces of $L \cap \{v\}^{+D}$ in $\{v\}^{+D}$ and the components of $\partial N(L) \cap \{v\}^{+D}$.* □

Proof of Proposition 12.13 Let $D = \ell(C)2^N$, where $\ell(C)$ denotes the length of C. Then $\{v\}^{+D}$ contains a lift of P_N, a fundamental domain of a $2^{N^{\text{th}}}$ power of C. By Lemma 12.14, $\widetilde{C}$ separates $\partial N(P_N) \subset \{v\}^{+D}$ into exactly K components. Lemma 12.16 then implies (1). Conclusion (2) follows from Lemma 7.20. □

12.2.1 Construction of X_C Choose a basepoint $v \in \phi(C)$ with lift $\tilde{v}$ in $\widetilde{C}$. Let $\mathcal{B} := \{\tilde{v}\}^{+D}$ in $\widetilde{X}$. Note that:

Lemma 12.17 *There exists an algorithm that takes X, C and v as input and returns $\mathcal{B}$ in exponential time.* □

For each translate L meeting $\mathcal{B}$, we attach a finite strip $S_L \cong L \cap \mathcal{B} \times [0,1]$. Let $Z = Z_{\mathcal{B}}$ be the set of open horizontal half-edges in the union of $\mathcal{B}$ and the collection of finite strips. For each vertex u (respectively edge e) in $\phi(C)$, we define a set of walls W_u (respectively W_e) on Z as follows. First fix a fundamental domain P of C in $\widetilde{X}$ such that $\tilde{v} \in P$. Choose a lift of u (respectively e) in P. Let $\widetilde{C} = L_1,\dots,L_n$ be the translates of $\widetilde{C}$ that pass through $\tilde{u}$ (respectively $\tilde{e}$). By Proposition 12.13, each line L_i separates $\mathcal{B}$ into exactly K components. Thus each line L_i induces K walls of W_u (respectively W_e) on Z, where each half-space Y in $\widetilde{X}$ of L_i defines a wall $\{Y \cap Z, Y^c \cap Z\}$.

Since C has no self-crossing, no two walls of W_u (respectively W_e) cross (Lemma 12.3). Thus the dual cube complex of (Z, W_u) (respectively (Z, W_e)) is a tree, denoted by $T(u)$ (respectively $T(e)$). Note that the definition of W_u (and W_e) is independent of the choice of $\tilde{v}$ and of the choice of $\tilde{u}$ (respectively $\tilde{e}$) in $P \subset \mathcal{B}$.

Suppose that an edge e is incident to a vertex u in $\phi(C)$. Let $\tilde{e}$ and incident vertex $\tilde{u}$ be the corresponding lifts in $P \subset B$. Since every translate of $\widetilde{C}$ that passes through $\tilde{e}$ also passes through $\tilde{u}$, there exists a natural inclusion $W_e \subset W_u$. Further, given translates L_1 and L_2 that contain $\tilde{e}$ with half-spaces Y_1 and Y_2 such that $Y_1 \subset Y_2$, suppose there exists a translate L' that meets $\tilde{u}$ with half-space Y' such that $Y_1 \subset Y' \subset Y_2$. Then it is easy to see that L' contains $\tilde{e}$. Thus we have:

Lemma 12.18 *Given an edge e in $\phi(C)$ incident to a vertex u, there exists a natural inclusion $T(e) \hookrightarrow T(u)$.*

Let Z'_C denote the geometric realisation of the graph of trees $(\phi(C), \{T(u)\}, \{T(e)\})$.

Proposition 12.19 *There exists a natural isomorphism between the square complexes Z_C and Z'_C.*

Proof By Lemma 12.9, Z_C is a graph of finite trees with underlying graph $\phi(C)$. So is Z'_C. Further, the wall structures that define vertex and edge trees of Z_C and Z'_C are isomorphic: Indeed, the walls that define $T(u)$ for $u \in \phi(C)$ in Z_C are induced by half-spaces in $\widetilde{X}$ of translates of $\widetilde{C}$ that pass through a lift $\tilde{u}$ of u. In Z'_C, the tree is defined by walls induced by half-spaces of translates of $\widetilde{C}$ in a finite ball $\mathcal{B}$ of $\widetilde{X}$ containing $\tilde{u}$. Since there exists a bijection between the half-spaces of $\widetilde{C}$ in $\mathcal{B}$ and the half-spaces of $\widetilde{C}$ in $\widetilde{X}$ (Lemma 12.16), Z_C is isomorphic to Z'_C. □

Proof of Theorem 12.11 The compact space $\mathcal{B}$ can be constructed in exponential time from X (Lemma 12.17). It costs exponential time to calculate the number of half-spaces of any translate of $\widetilde{C}$ (Lemma 12.14).

Since the number of translates of $\widetilde{C}$ meeting at any point of $\widetilde{X}$ is bounded by the length of C (by Lemma 7.20), the dual trees $T(u)$ (respectively $T(e)$) of all vertices u (respectively edges e) in $\phi(C)$ can be constructed in polynomial time. Thus Z'_C is constructed in exponential time. Y'_C is constructed in linear time in X and X_C is obtained in linear time from $Y_C \cong Y'_C$ and $Z_C \cong Z'_C$. Hence the result. □

12.2.2 Structure of X_C

Lemma 12.20 *The tree $T(u)$ (or $T(e)$) is a bipartite tree with black vertices having valence exactly K.*

Proof We will first show that there exist vertices of valence K in $T(u)$ (respectively $T(e)$) and then show that the tree is bipartite. Let L be a translate of $\widetilde{C}$ passing through $\tilde{u}$ (respectively $\tilde{e}$) in $\mathcal{B}$. Let $Y_1, \dots, Y_K$ be the half-spaces of L. Let z be an open horizontal half-edge in the strip S_L. Denote by σ_L the ultrafilter σ_z in $T(u)$ (respectively $T(e)$). Thus σ_L contains $\{Y_1^c, \dots, Y_K^c\}$ and exactly those half-spaces of translates of $\widetilde{C}$ passing through $\tilde{u}$ that contain L. Observe that the valence of σ_L is at least K: switching each half-space Y_i of L gives an edge incident to σ_L. It is easy to see that it is exactly K as no edge of the form $\{Y', Y'^c\}$ can be incident to σ_L when $Y' \neq Y_i$.

Let σ' be a vertex at distance two from σ_L. Let $\sigma' \triangle \sigma_L = \{Y_1, Y_1^c, Y_1', Y_1'^c\}$, where Y_1' is a half-space of a translate L' of $\widetilde{C}$. We will show that $\sigma' = \sigma_{L'}$. Assume that $Y_1' \in \sigma_L$. This implies that for each half-space Y_i' of L' with $i \neq 1$, $Y_i'^c$ is in σ_L and hence in σ'. Further, $Y_1'^c \in \sigma'$ by assumption. If $\sigma' \neq \sigma_{L'}$, then any path from σ' to $\sigma_{L'}$ involves a change of half-spaces of the type $\{Y_i', Y_i'^c\}$. Hence we conclude that $\sigma' = \sigma_{L'}$. The case when $Y_1' \notin \sigma_L$ is similar and we leave it as an exercise. This proves that a vertex in $T(u)$ (respectively $T(e)$) is of the form $\sigma_{L'}$ if and only if it is at even distance from σ_L. Thus the tree is bipartite. □

Let u (respectively e) be in $\phi(C)$. Let v_1 and v_2 be black vertices in $T(u)$ (respectively $T(e)$).

Lemma 12.21 *Given an edge e_1 incident to v_1 in $T(u)$ (respectively $T(e)$), there exists an edge e_2 incident to v_2 in $T(u)$ (respectively $T(e)$) such that the hyperplane in Z_C' dual to e_1 is equal to the hyperplane dual to e_2.*

Proof Let v_i be the ultrafilter σ_{L_i}, where L_1 and L_2 are translates of $\widetilde{C}$ passing through $\tilde{u}$ (respectively $\tilde{e}$) in $\mathcal{B}$. Let e_1 correspond to the wall $\{Y_1, Y_1^c\}$, where Y_1 is a half-space of L_1. Since L_2 is a translate of L_1, there exists a fundamental domain P' of C' in L_1 containing $\tilde{u}$ (respectively $\tilde{e}$) such that there exists $g \in G$ and $\tilde{u}'$ (respectively $\tilde{e}'$ in P') with $g\tilde{u}' = \tilde{u}$ (respectively $g\tilde{e}' = \tilde{e}$) and $gP' \subset L_2$. The segment from $\tilde{u}$ to $\tilde{u}'$ (respectively $\tilde{e}$ to $\tilde{e}'$) projects to $\phi(C)$ as a subgraph. Since L_1 passes through every vertex and edge in this segment, there is an edge corresponding to $\{Y_1, Y_1^c\}$ in the dual tree of the image in $\phi(C)$ of each vertex and edge of this segment. By the way Z_C' was defined, this defines a unique hyperplane in Z_C'. Since $g\tilde{u}' = \tilde{u}$ (respectively $g\tilde{e}' = \tilde{e}$) and $gL_1 = L_2$, the required edge incident to σ_{L_2} is $\{gY_1, gY_1^c\}$. □

Let $\phi'\colon C' \to X_s$ be a maximal cycle such that C is a power of C'.

Lemma 12.22 *There exists a natural embedding of C' in X_C such that the vertical graph that contains C' is isomorphic to C'. Further, for a vertex u (respectively edge e) in $\phi(C)$, the embedded copy of C' meets every black vertex of $T(u)$ (respectively $T(e)$) exactly once.* □

12.3 The tubular graph of graphs X'

Let $\mathscr{C} = \{C_1, \dots, C_n\}$ be the set of splitting cycles of X furnished by Theorem 11.1.

Remark 12.23 It is easy to see that a vertical cycle induced by the attaching map of a tube is a splitting cycle that is not 3–repetitive. Hence each such cycle is included in $\mathscr{C}$.

Procedure 12.24 (construction of X') The tubular graph of graphs X' is constructed from X using the cycles in $\mathscr{C}$ as follows:

- Start with $X = X_0$.
- For $1 \leq i \leq n$, check whether $\phi_i\colon C_i \to X$ factors through a vertical cycle $\psi_i\colon C_i \to X_{i-1}$. If it doesn't, then declare $X_i = X_{i-1}$. Otherwise, define X_i to be the opened-up space of X_{i-1} along the cycle $\psi_i\colon C_i \to X_{i-1}$.
- Declare $X' = X_n$.

Lemma 12.25 *The cycle C_i in $\mathscr{C}$ factors through a vertical cycle in X_{i-1} if and only if for $1 \leq j \leq i$, lifts of C_j and C_i don't cross in $\widetilde{X}$.* □

By Theorem 1.7, we will assume that X' is a Brady–Meier tubular graph of graphs.

Theorem 12.26 *There exists an algorithm of double exponential time complexity that takes a Brady–Meier tubular graph of graphs X with hyperbolic fundamental group G as input and returns a homotopy equivalent Brady–Meier tubular graph of graphs whose vertical hyperplanes generate all universally elliptic subgroups of G up to commensurability.*

Proof The algorithm of Theorem 11.1 takes X as input and returns a set of cycles $\mathcal{C}$ that contains all universally elliptic cycles up to commensurability. Given $\mathcal{C}$, X' is constructed using Procedure 12.24. This procedure consists of applying the algorithm of Theorem 12.11 repeatedly. The algorithm of Theorem 11.1 takes double exponential time to return a finite set of cycles. The number of cycles in this set is bounded by a number of double exponential magnitude. Given this data, the algorithm

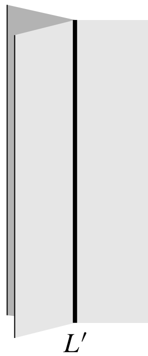

Figure 9: L' in $\widetilde{X}'$ when L has three half-spaces.

of Theorem 12.11 operates by taking exponential time for each cycle, and hence obtaining X' costs double exponential time in the input data. □

12.3.1 Structure of X' Note that the maps $\hat{\eta}_i\colon \widetilde{X}_i \to \widetilde{X}_{i-1}$ induce maps $\hat{\eta}\colon \widetilde{X}' \to \widetilde{X}$ and $\eta\colon X' \to X$. Denote by Γ' the underlying graph of the graph of spaces X' and by T' the underlying tree of the tree of spaces $\widetilde{X}'$. Let L be a lift of an element C_i of $\mathscr{C}$ such that C_i factors through a vertical cycle in X_{i-1}. From Lemma 12.22, it follows (see Figure 9) that:

Lemma 12.27 *There exists a vertical tree in $\widetilde{X}'$ whose stabiliser is equal to the stabiliser of L.*

Let h be a vertical hyperplane in $\widetilde{X}'$. Let L_1 and L_2 be the two *boundary lines* of h, that is, the two vertical lines on either side of h at distance $\frac{1}{2}$ from h, and parallel to h. As G is hyperbolic, it follows from the flat plane theorem (see [6, Theorem Γ.3.1]) that:

Lemma 12.28 *The stabiliser of h is equal to either* $\mathrm{stab}(L_1)$ *or* $\mathrm{stab}(L_2)$*.* □

Observe that $\widetilde{X}'$ is the cube complex dual to a space with walls, where walls are defined on the set Z of open horizontal half-edges of $\widetilde{X}$ along with open horizontal half-edges of strips attached to translates of $\widetilde{C}_i$ whenever C_i factors through a vertical cycle in X_{i-1}. The set of walls W' in Z are thus of three types:

- *Walls of type* (i), induced by horizontal hyperplanes of $\widetilde{X}$.
- *Walls of type* (ii), induced by vertical hyperplanes of $\widetilde{X}$.
- *Walls of type* (iii), induced by translates of $\widetilde{C}_i$, where $C_i \in \mathscr{C}$ factors through a vertical cycle in X_{i-1}.

A *vertical half-space* of $\widetilde{X}$ (or $\widetilde{X}'$) is a half-space of a wall of type (ii) or (iii). Let h be a vertical hyperplane in $\widetilde{X}'$ and L_1 and L_2 its boundary lines. Then:

Lemma 12.29 *Either the vertical tree containing L_1 or the vertical tree containing L_2 is a line.*

Proof Let $\{Y, Y^c\}$ be the wall of type (ii) or (iii) in Z corresponding to h. Let L be the (boundary) line in $\widetilde{X}$ (of the vertical hyperplane) that defines $\{Y, Y^c\}$. By Lemma 12.27, there exists a linear vertical tree L' in $\widetilde{X}'$ such that $\operatorname{stab}(L) = \operatorname{stab}(L')$. Note that any vertical half-space of $\widetilde{X}$ contained in a vertex (ultrafilter) of L' contains either Y or Y^c, by Lemma 6.5. Thus any vertical half-space of $\widetilde{X}'$ (except perhaps a half-space of L') that contains L' contains h. So no vertical tree separates L' and h. Hence the result. □

Definition 12.30 A vertex of a G–tree is a *cyclic vertex* if its stabiliser is a cyclic subgroup of G.

Thus, in the underlying tree T' of $\widetilde{X}'$, at least one of the two vertices of any edge is a cyclic vertex.

Lemma 12.31 *Let L' be a line in $\widetilde{X}'$. Suppose that $\partial N(L')$ contains at least three components. Then the vertical tree containing L' is equal to L'.*

We need two observations to prove the lemma. Let L be a line in $\widetilde{X}$ that defines a wall of type (iii). Suppose that the number of half-spaces of L is K. Let L' be a vertical tree in $\widetilde{X}'$ such that $\hat{\eta}(L') = L$ (Lemma 12.27).

Lemma 12.32 *Exactly K vertical strips are attached to L' in $\widetilde{X}'$. Furthermore, if $\hat{\eta}(L'') = L$ for any vertical line L'', then L'' is contained in one of these K strips.*

Proof The fact that exactly K strips are attached to L' follows from Lemma 12.22. Further, each of the K strips above is contained in $\hat{\eta}^{-1}(L)$, by Lemma 12.5.

Let L'' be a vertical line such that $\hat{\eta}(L'') = L$. Denote by $\sigma'_{L''}$ the set of vertical half-spaces contained in any vertex (ultrafilter) of L''. Note that $\hat{\eta}(L'') = L$ implies that the vertical half-spaces of type (ii) in $\sigma'_{L''}$ consists of the half-spaces of vertical hyperplanes in $\widetilde{X}$ that contain L.

If $\sigma'_{L''} = \sigma'_{L'}$, we have nothing to prove as the vertical tree that contains L' is equal to L'. Let σ'_i be the set of vertical half-spaces such that $\sigma'_i \triangle \sigma'_{L'} = \{Y_i, Y_i^c\}$, where Y_i is a half-space of L. Either $\sigma'_{L''} = \sigma'_i$ for some i or $\sigma'_{L''} \neq \sigma'_i$ for any i.

First assume the latter. Then there exists a half-space Y_0 of a line L_0 such that $\{Y_0, Y_0^c, Y_i, Y_i^c\} \subset \sigma'_{L''} \triangle \sigma'_{L'}$, with $Y_i \subset \sigma'_{L''}$. Let σ be a vertex in $L \setminus L_0$. Then there exists $\sigma'' \in L''$ such that $\widehat{\eta}(\sigma'') = \sigma$. This implies that the vertex σ contains Y_i, which is not possible. Assume now that $\sigma'_{L''} = \sigma'_i$, for some i. Let $L'_i \neq L'$ be a boundary line of the strip that separates L'' from L'. The result follows from the following observation. Let γ denote a geodesic between L'' and L'_i. Since γ consists of vertical edges, $\widehat{\eta}(\gamma)$ has the same length as γ (by Lemma 12.4) and is a geodesic between $\widehat{\eta}(L'')$ and $\widehat{\eta}(L'_i)$. □

Let $L'_1, \dots, L'_K$ be the boundary lines of strips attached to L' such that $L'_i \neq L'$. Then:

Lemma 12.33 *$\partial N(L'_i)$ has exactly two components.*

Proof Note that L'_i is a separating line (Fact 4.9). Let Y'_i be a half-space of L'_i that does not contain L'. Then $\widehat{\eta}(\mathring{Y}'_i)$ does not contain L (Lemma 12.32) and is connected. Thus $\widehat{\eta}(Y'_i)$ lies in the half-space Y_i of L. Further, if there exist two half-spaces of L'_i that do not contain L', then $\widehat{\eta}^{-1}(Y_i \setminus L)$ contains these half-spaces. By Lemma 12.5, one of these half-spaces is at finite distance from L'_i, contradicting Lemma 4.10. □

Proof of Lemma 12.31 If $\partial N(L')$ contains three or more components, then a subgroup H of $\operatorname{stab}(L')$ is universally elliptic, by Proposition 8.5. Let $L = \widehat{\eta}(L')$. Then H stabilises L as $\widehat{\eta}$ is G–equivariant, and L defines a wall of type (iii). Lemmas 12.32 and 12.33 then give the result. □

12.4 Modification of X'

The next step in the modification of X to X_{jsj} is the construction of an intermediate tubular graph of graphs X'' from X'.

12.4.1 Construction of X'' Remove an open tube of X' if both the vertex graphs bounding the tube are circles, and then identify the vertex graphs. This is possible as, by Lemma 12.28, one of the attaching maps of such a tube is an isomorphism of graphs. Successively remove all such tubes of X'. Let X'' be the tubular graph of graphs obtained after removing all such tubes.

12.4.2 Structure of X'' Let T'' be the underlying tree of $\widetilde{X}''$ and let T_{jsj} denote the Bass–Serre tree of the canonical JSJ decomposition of G.

Lemma 12.34 *For each cyclic vertex u of T_{jsj} there exists a cyclic vertex v of T'' such that* $\operatorname{stab}(u) = \operatorname{stab}(v)$.

Proof Fix an axis L in $\widetilde{X}$ of $\mathrm{stab}(u)$. Note that $\mathrm{stab}(u) = \mathrm{stab}(L)$ as $\mathrm{stab}(u)$ is maximal cyclic. We can assume L is vertical since $\mathrm{stab}(u)$ is commensurable with a universally elliptic subgroup (Lemma 8.4). By Lemma 7.25, there exists a splitting cycle C in X such that $\pi_1(C)$ is commensurable with a conjugate of $\mathrm{stab}(u)$. Hence $C \in \mathscr{C}$ and there exists a vertical tree in $\widetilde{X}'$ whose stabiliser is $\mathrm{stab}(\widetilde{C})$. Hence the result. □

Let v be a cyclic vertex of T'' and L'' the corresponding vertical tree (line) in $\widetilde{X}''$. Denote by ΛH the limit set in ∂G of a subgroup H of G.

Lemma 12.35 *The number of components of $\partial G \setminus \Lambda\,\mathrm{stab}(v)$ is equal to the number of edges incident to v.*

Proof The number of edges incident to v is equal to the number of strips attached to L'', which is equal to the number of components of $\partial N(L'')$. The number of components of $\partial G \setminus \Lambda\,\mathrm{stab}(v)$ is equal to the supremum of the number of components of $\widetilde{X}'' \setminus L''^{+k}$, where $k \in \mathbb{N}$. Let K be the number of strips attached to L''. Let L''_i be the boundary line of the i^{th} strip such that $L'' \neq L''_i$. Note that the vertical tree containing L''_i is not equal to L''_i as otherwise the corresponding strip would have been removed to obtain $\widetilde{X}''$. By construction, L''_i has exactly two half-spaces. Let Y''_i be the half-space of L''_i that does not contain L''. Note that no strip of Y''_i contains L''_i. By Lemma 4.15, $Y''_i \setminus L''^{+k}_i$ and hence $Y''_i \setminus L''^{+(k+1)}$ is connected, for every $k \in \mathbb{N}$. □

Proposition 12.36 *For each edge stabiliser H of T_{jsj}, there exists an edge of T'' whose stabiliser is H.*

Proof Since each edge of T_{jsj} is incident to a cyclic vertex, let $\mathrm{stab}(u)$ be the cyclic vertex group of T_{jsj} that contains H. Then H is the stabiliser of a component of $\partial G \setminus \Lambda\,\mathrm{stab}(u)$. Let v be a vertex of T'' such that $\mathrm{stab}(v) = \mathrm{stab}(u)$ (Lemma 12.34). Then the number of edges incident to v is equal to the number of components of $\partial G \setminus \Lambda\,\mathrm{stab}(u)$, by Lemma 12.35. Further, H is the stabiliser of an edge incident to v as each edge incident to v induces a unique component of $\partial G \setminus \Lambda\,\mathrm{stab}(u)$. □

Definition 12.37 [16] A G–tree $\widehat{T}$ is a *refinement* of a G–tree T if there exists a G–equivariant map $p\colon \widehat{T} \to T$ such that p sends any segment $[x, y]$ in $\widehat{T}$ onto the segment $[p(x), p(y)]$. In other words, $\widehat{T}$ is obtained by blowing up vertices of T.

By Proposition 12.36 and the properties of the JSJ decomposition (Definition 1.1), we have:

Corollary 12.38 *T'' is a refinement of T_{jsj}.* □

Lemma 12.39 *The stabiliser H of an edge e of T'' is not an edge stabiliser of T_{jsj} if and only if the cyclic vertex u incident to e in T'' is of valence two and both the noncyclic vertices adjacent to u are stabilised by hanging surface groups.*

Proof No edge of T_{jsj} is such that the cyclic vertex incident to this edge is adjacent to exactly two hanging surface group vertices, by definition. The converse follows from Lemma 12.35 and Corollary 12.38. □

So we can modify X'' to X_{jsj} by removing tubes which connect hanging surface groups. This requires an identification of such groups, which is done in the next subsection.

12.5 Surface graphs

Definition 12.40 A vertex graph of a tubular graph of graphs is a *surface graph* if the graph is not a circle and the fundamental group of the graph is a surface group whose peripheral subgroups are precisely the subgroups induced by the incident edge graphs.

Thus a vertex graph is a surface graph if its fundamental group is a hanging surface group. Recall that:

Definition 12.41 The *double* of a graph Γ with a finite family of immersed cycles $\{C_1, \dots, C_n\}$ is a tubular graph of graphs whose underlying graph consists of two vertices with n edges between them, each vertex space is a copy of Γ and the i^{th} tube attaches as C_i on both sides.

Lemma 12.42 *A vertex graph of a Brady–Meier tubular graph of graphs is a surface graph if and only if every edge of its double is of thickness two.*

Proof Let D_s be the double of the vertex graph X_s with incident edge cycles. It is a standard fact that D_s is homeomorphic to a surface if and only if X_s is a surface graph. Note that D_s is Brady–Meier as every vertex of X_s satisfies the Brady–Meier conditions. Thus every edge of D_s is of thickness at least two. If each edge is of thickness two, then the fact that every vertex link is connected implies that every vertex link is a circle. This implies that D_s is homeomorphic to a closed surface and we are done.

Conversely, suppose that there exists an edge e of thickness at least three in D_s. Let $\tilde{e}$ be a lift of e in $\widetilde{D}_S$ and h the horizontal hyperplane through $\tilde{e}$. Note that h is a tree. Let L be a line in h passing through the midpoint m of $\tilde{e}$. Note that L does not separate $\partial N(m)$ as $\tilde{e}$ is of thickness at least three. By Lemma 4.4, L does not separate $\widetilde{D}_s$. But this implies that D_s is not homeomorphic to a closed surface. □

12.6 Construction of X_{jsj}

We are now ready to construct X_{jsj}. Remove from X'' every cyclic vertex graph and the (open) tubes attached to it whenever exactly two tubes are attached to the vertex graph and both the tubes are attached to surface graphs on the other side. Call the resulting complex X'''. Denote by T''' the underlying tree of $\widetilde{X}'''$. By Lemma 12.39, we have:

Proposition 12.43 *T''' is isomorphic to T_{jsj} as G–trees.* □

The proposition proves that X''' is the required X_{jsj}. We now have the main result:

Theorem 12.44 *There exists an algorithm of double exponential time complexity that takes a Brady–Meier tubular graph of graphs X with hyperbolic fundamental group G as input and returns a Brady–Meier tubular graph of graphs whose underlying graph of groups structure is the JSJ decomposition of G.*

Proof Using Theorem 12.26, we obtain the tubular graph of graphs X' in double exponential time. Constructing X'' from X' involves identifying which tubes are attached to only cyclic vertex graphs and takes at most polynomial time in the number of squares of X. The construction of X_{jsj} from X'' involves removing pairs of tubes adjacent to surface graphs. Detecting surface graphs involves constructing doubles of vertex graphs (Lemma 12.42) and also takes at most polynomial time in the number of squares of X. □

13 Relative JSJ decompositions

Let F be a finite-rank free group and $\mathcal{H}$ be a finite family of maximal cyclic subgroups in F. Recall that F *splits relative to* $\mathcal{H}$ if there exists a nontrivial splitting of F in which each element of $\mathcal{H}$ is elliptic. Similarly F is *freely indecomposable relative to* $\mathcal{H}$ if F does not split freely relative to $\mathcal{H}$.

Definition 13.1 A *relative JSJ decomposition* of $(F,\mathcal{H})$ is a graph of groups splitting of F relative to $\mathcal{H}$ which satisfies the conditions of a JSJ decomposition (Definition 1.1), with the constraints that vertex groups of type (3) are rigid relative to $\mathcal{H}$ and elements of $\mathcal{H}$ can intersect vertex groups of type (2) only in their peripheral subgroups.

Recall that a subgroup F' of F is *rigid relative to* $\mathcal{H}$ if F' is elliptic in every splitting of F relative to $\mathcal{H}$.

Theorem 13.2 [7, Theorem 4.25] *Given a finite-rank free group F and a finite family $\mathcal{H}$ of maximal cyclic subgroups of F such that F is freely indecomposable relative to $\mathcal{H}$, a relative JSJ decomposition of $(F,\mathcal{H})$ exists and is unique.*

In [9], the authors implement an algorithm that returns the JSJ decomposition of F relative to $\mathcal{H}$, though they do not give an estimate of its time complexity. The main result of this section is the following.

Theorem 13.3 *There exists an algorithm of double exponential time complexity that takes a finite-rank free group F and a finite family of maximal cyclic subgroups $\mathcal{H}$ such that F is freely indecomposable relative to $\mathcal{H}$ as input and returns the relative JSJ decomposition of F relative to $\mathcal{H}$.*

We will construct a suitable tubular graph of graphs $X_{F,\mathcal{H}}$ to prove Theorem 13.3. There exists a *central vertex graph* X_{s_c} in $X_{F,\mathcal{H}}$ such that $\pi_1(X_{s_c}) = F$. If $\mathcal{H} = \{H_1,\dots,H_n\}$, then for each H_i there exists an immersed cycle $\phi_i\colon C_i \to X_{s_c}$ such that C_i induces a conjugate of the group H_i in $\pi_1(X_{s_c}) = F$. Note that the word generated by C_i is cyclically reduced in F as ϕ_i is an immersion of graphs. There exist exactly n tubes in $X_{F,\mathcal{H}}$ that are attached to X_{s_c} in the following way. The edge graph of the i^{th} tube is isomorphic to C_i and the attaching map is given by ϕ_i. We subdivide X_{s_c} and the n edge graphs sufficiently to make all graphs simplicial. The other end of the i^{th} tube is attached by an isomorphism to a circular vertex graph X_i. There are exactly two other tubes attached to X_i, with both attaching maps being isomorphisms. Each of these two tubes connects X_i to a copy of a surface graph whose fundamental group is the fundamental group of the oriented surface of genus two with exactly one boundary component. Thus the underlying graph of $X_{F,\mathcal{H}}$ is a tree with one "central" vertex s_c of valence n, n cyclic vertices adjacent to s_c, of valence three each, and $2n$ surface vertices of valence one each.

Let G be the fundamental group of $X_{F,\mathcal{H}}$. Since each vertex group is freely indecomposable relative to its incident edge groups, G is one-ended, by [39, Theorem 18]. Hence, we can assume that $X_{F,\mathcal{H}}$ is Brady–Meier, by Theorem 1.7.

As a consequence of the Bestvina–Feighn combination theorem [3], we have:

Lemma 13.4 *G is δ–hyperbolic.* □

Let G_s be a vertex group of the graph of groups structure of G induced by $X_{F,\mathcal{H}}$.

Lemma 13.5 *Either $G_s = G_{s_c}$ or G_s is a conjugate of either a cyclic vertex group of the JSJ decomposition of G or a maximal hanging surface group.*

Proof Let G_s be a cyclic vertex group adjacent to G_{s_c}, with X_s the corresponding vertex graph. Note that $\widetilde{X}_s$ is a line and $\partial N(\widetilde{X}_s)$ contains three components as there are three tubes attached to X_s. By Proposition 8.5, a subgroup of G_s is universally elliptic. The result then follows as G_s is maximal cyclic.

Now suppose that G_s is a hanging surface group corresponding to the surface graph X_s. Note that no separating transversal line can meet $\widetilde{X}_s$ as such a line will have to first cross a vertical tree (line) adjacent to $\widetilde{X}_s$, which is not possible, as seen above. Hence any line which crosses a line contained in $\widetilde{X}_s$ is itself contained in $\widetilde{X}_s$. Thus G_s is maximal hanging. □

Corollary 13.6 *If T_{jsj} is the JSJ tree of G and T the underlying tree of $\widetilde{X}_{F,\mathcal{H}}$, then T_{jsj} is a refinement of T obtained by blowing up lifts of the central vertex s_c.*

□

Proof of Theorem 13.3 Given $(F,\mathcal{H})$, the tubular graph of graphs $X_{F,\mathcal{H}}$ can be constructed algorithmically in polynomial time in the rank of F and the lengths of $\mathcal{H}$. Let X_{jsj} be the tubular graph of graphs obtained from $X_{F,\mathcal{H}}$ in double exponential time by Theorem 12.44. Let Γ_{jsj} and Γ be the underlying graphs of X_{jsj} and $X_{F,\mathcal{H}}$ respectively. Note that since vertex graphs other than X_{s_c} induce vertex groups of the JSJ, Γ_{jsj} is obtained from Γ by a blow-up of the vertex s_c. Let Y be the subgraph of groups of X_{jsj} with underlying graph $f^{-1}(s_c)$. Then it is straightforward to check that Y is the relative JSJ of $(F,\mathcal{H})$. □

14 The case of graphs of free groups with cyclic edge groups

We can now extend our result to general hyperbolic one-ended graphs of free groups with cyclic edge groups:

Theorem 14.1 *There exists an algorithm of double exponential time complexity that takes a graph of free groups with cyclic edge groups with hyperbolic one-ended fundamental group as input and returns its JSJ decomposition.*

We recall that a graph of free groups with cyclic edge groups is the group-theoretic analogue of a graph of spaces (see Definition 2.7) whose vertex and edge spaces are finite connected one-dimensional CW complexes, with the additional restriction that all the edge spaces are circles. We will further assume that the fundamental group of the geometric realisation is one-ended and hyperbolic. Again, we will use the same notation X to denote both the graph of groups under consideration and its geometric realisation. We caution the reader that in this section X need not be a tubular graph of graphs.

Remark 14.2 Each vertex space with its incident edge spaces gives rise to a free group with a "marked" family of cyclic subgroups at the level of fundamental groups. This marked family may contain nonmaximal cyclic subgroups, but this can be rectified by a normalisation process as done in [7, Section 2.4]. The normalisation process introduces a new vertex group H corresponding to each nonmaximal cyclic subgroup H' in the marked family such that H is the maximal cyclic subgroup containing H. So, without loss of generality, we will assume that we have a graph of free groups with cyclic edge groups such that whenever a vertex group is not cyclic, then the incident edge groups are all maximal cyclic. Since every cyclic subgroup of a (torsion-free) hyperbolic group is contained in a unique maximal cyclic subgroup (see [11, chapitre 10, proposition 7.1]), we note that any edge group injects into at most one of its vertex groups in a nonmaximal cyclic subgroup. Thus, the normalisation procedure does not produce two new adjacent cyclic vertices.

Proof of Theorem 14.1 Since the fundamental group G of the input graph of groups X is one-ended, each noncyclic vertex group with its incident edge groups gives a free group F with a finite family of maximal cyclic subgroups $\mathcal{H}$ such that F is freely indecomposable relative to $\mathcal{H}$. We now apply the algorithm of Theorem 13.3 to obtain the relative JSJ decomposition of $(F, \mathcal{H})$. Replace the vertex group F in X with the graph of groups corresponding to the relative JSJ decomposition of $(F, \mathcal{H})$. Repeat the procedure at each noncyclic vertex group of X to obtain a new graph of groups decomposition X', where each vertex group has been replaced by its relative JSJ decomposition. Observe that X' is still a graph of free groups with cyclic edge groups. We now modify X' to an intermediate graph of groups X'' by removing edges when the corresponding edge groups are attached to cyclic vertex groups on both sides, analogous to the procedure in Section 12.4. What remains is to identify surface graphs in X'' using Lemma 12.42 and gluing them together to obtain X_{jsj}, as done in Section 12.6. The resulting graph of groups is the JSJ decomposition of X. □

15 The isomorphism problem

An important consequence of Theorem 14.1 is that the isomorphism problem for hyperbolic fundamental groups of graphs of free groups with cyclic edge groups is reduced to the Whitehead problem [38] and can be solved in double exponential time.

Theorem 15.1 *There exists an algorithm of double exponential time complexity that takes two graphs of free groups with cyclic edge groups and hyperbolic fundamental group as input and decides whether they are isomorphic.*

We refer the reader to [13, Section 4] for the relevant definitions. The main result that we invoke from that paper is Proposition 4.4, which states that given two graphs of groups X_1 and X_2 on the same underlying graph, and a collection of group isomorphisms between the corresponding vertex groups, X_1 is isomorphic to X_2 as graphs of groups if and only if there exists an "extension adjustment" [13, Definition 4.3].

Proof of Theorem 15.1 Let X_1 and X_2 be the input graphs of free groups with cyclic edge groups. We will denote their (hyperbolic) fundamental groups by G_1 and G_2 respectively.

By [39, Theorem 18], G_i is one-ended if and only if each vertex group of X_i is freely indecomposable relative to its incident edge groups. Corollary E of [37] gives a polynomial time algorithm to detect whether a given vertex group is freely indecomposable relative to its incident edge groups. One then obtains the Grushko decomposition of G_i. Since the Grushko decomposition of a group is unique, G_1 and G_2 are isomorphic if and only if there is a one-to-one correspondence between the factors of their Grushko decompositions such that the corresponding factors are isomorphic as groups.

So assume that G_1 and G_2 are one-ended. Using Theorem 14.1, we can construct the JSJ decompositions of G_1 and G_2. Then, by the uniqueness of the JSJ decomposition and [13, Proposition 4.4], G_1 and G_2 are isomorphic if and only if

(1) their JSJ decompositions have a common underlying graph Γ,

(2) for each vertex v of Γ, there exists an isomorphism ϕ_v between the vertex group $G_{v,1}$ of X_1 and the vertex group $G_{v,2}$ of X_2, and

(3) there exists an extension adjustment.

In our case, (3) boils down, for each vertex v, to ϕ_v taking the set of incident edge subgroups of $G_{v,1}$ to the set of incident edge subgroups of $G_{v,2}$. We refer the reader to [13, Definition 4.3 and commutative diagram (8)] for a precise formulation.

Thus, the isomorphism problem is reduced to solving the Whitehead problem for each vertex group. There exist algorithms that provide solutions to the Whitehead problem in at most exponential time (see [23; 22; 30]). Combining the above gives us the required algorithm. □

References

[1] **N Alon**, **R Yuster**, **U Zwick**, *Finding and counting given length cycles*, Algorithmica 17 (1997) 209–223 MR

[2] **B Barrett**, *Computing JSJ decompositions of hyperbolic groups*, J. Topol. 11 (2018) 527–558 MR

[3] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, J. Differential Geom. 35 (1992) 85–101 MR

[4] **B H Bowditch**, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. 180 (1998) 145–186 MR

[5] **N Brady**, **J Meier**, *Connectivity at infinity for right angled Artin groups*, Trans. Amer. Math. Soc. 353 (2001) 117–132 MR

[6] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR

[7] **C H Cashen**, *Splitting line patterns in free groups*, Algebr. Geom. Topol. 16 (2016) 621–673 MR

[8] **C H Cashen**, **N Macura**, *Line patterns in free groups*, Geom. Topol. 15 (2011) 1419–1475 MR

[9] **C H Cashen**, **J F Manning**, *Virtual geometricity is rare*, LMS J. Comput. Math. 18 (2015) 444–455 MR

[10] **I Chatterji**, **G Niblo**, *From wall spaces to* CAT(0) *cube complexes*, Internat. J. Algebra Comput. 15 (2005) 875–885 MR

[11] **M Coornaert**, **T Delzant**, **A Papadopoulos**, *Géométrie et théorie des groupes: les groupes hyperboliques de Gromov*, Lecture Notes in Mathematics 1441, Springer (1990) MR

[12] **F Dahmani**, **V Guirardel**, *The isomorphism problem for all hyperbolic groups*, Geom. Funct. Anal. 21 (2011) 223–300 MR

[13] **F Dahmani**, **N Touikan**, *Deciding isomorphy using Dehn fillings, the splitting case*, Invent. Math. 215 (2019) 81–169 MR

[14] **M Dehn**, *Über unendliche diskontinuierliche Gruppen*, Math. Ann. 71 (1911) 116–144 MR

[15] **M Gromov**, *Hyperbolic groups*, from "Essays in group theory" (S M Gersten, editor), Math. Sci. Res. Inst. Publ. 8, Springer (1987) 75–263 MR

[16] **V Guirardel**, **G Levitt**, *JSJ decompositions of groups*, Astérisque 395, Soc. Math. France (2017) MR

[17] **F Haglund**, **F Paulin**, *Simplicité de groupes d'automorphismes d'espaces à courbure négative*, from "The Epstein birthday schrift" (I Rivin, C Rourke, C Series, editors), Geom. Topol. Monogr. 1, Geom. Topol. Publ., Coventry (1998) 181–248 MR

[18] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR

[19] **J Hopcroft**, **R Tarjan**, *Algorithm 447: efficient algorithms for graph manipulation*, Commun. ACM 16 (1973) 372–378

[20] **W Jaco**, **P B Shalen**, *A new decomposition theorem for irreducible sufficiently-large 3–manifolds*, from "Algebraic and geometric topology, II" (R J Milgram, editor), Proc. Sympos. Pure Math. 32, Amer. Math. Soc., Providence, RI (1978) 71–84 MR

[21] **K Johannson**, *Homotopy equivalences of 3–manifolds with boundaries*, Lecture Notes in Mathematics 761, Springer (1979) MR

[22] **I Kapovich**, **P Schupp**, **V Shpilrain**, *Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups*, Pacific J. Math. 223 (2006) 113–140 MR

[23] **D Lee**, *A tighter bound for the number of words of minimum length in an automorphic orbit*, J. Algebra 305 (2006) 1093–1101 MR

[24] **J F Manning**, *Virtually geometric words and Whitehead's algorithm*, Math. Res. Lett. 17 (2010) 917–925 MR

[25] **J P McCammond**, **D T Wise**, *Coherence, local quasiconvexity, and the perimeter of 2–complexes*, Geom. Funct. Anal. 15 (2005) 859–927 MR

[26] **B Nica**, *Cubulating spaces with walls*, Algebr. Geom. Topol. 4 (2004) 297–309 MR

[27] **P Papasoglu**, *Quasi-isometry invariance of group splittings*, Ann. of Math. 161 (2005) 759–830 MR

[28] **P Papasoglu**, *Splittings and the asymptotic topology of the lamplighter group*, Trans. Amer. Math. Soc. 364 (2012) 3861–3873 MR

[29] **E Rips**, **Z Sela**, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. 146 (1997) 53–109 MR

[30] **A Roig**, **E Ventura**, **P Weil**, *On the complexity of the Whitehead minimization problem*, Internat. J. Algebra Comput. 17 (2007) 1611–1634 MR

[31] **M Roller**, *Poc sets, median algebras and group actions*, preprint (2016) arXiv

[32] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 MR

[33] **P Scott**, **T Wall**, *Topological methods in group theory*, from "Homological group theory" (C T C Wall, editor), London Math. Soc. Lecture Note Ser. 36, Cambridge Univ. Press (1979) 137–203 MR

[34] **Z Sela**, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups, II*, Geom. Funct. Anal. 7 (1997) 561–593 MR

[35] **J-P Serre**, *Trees*, Springer (1980) MR

[36] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565 MR

[37] **Suraj Krishna, M S**, *Vertex links and the Grushko decomposition*, preprint (2018) arXiv

[38] **J H C Whitehead**, *On equivalent sets of elements in a free group*, Ann. of Math. 37 (1936) 782–800 MR

[39] **H Wilton**, *One-ended subgroups of graphs of free groups with cyclic edge groups*, Geom. Topol. 16 (2012) 665–683 MR

[40] **D T Wise**, *Non-positively curved squared complexes: aperiodic tilings and non-residually finite groups*, PhD thesis, Princeton University (1996) MR Available at `https://search.proquest.com/docview/304259249`

School of Mathematics, Tata Institute of Fundamental Research
Mumbai, India

suraj@math.tifr.res.in

Received: 23 December 2018 Revised: 14 September 2019