



HAL
open science

Constructing Grushko and JSJ decompositions : a combinatorial approach

Suraj Krishna Meda Satish

► **To cite this version:**

Suraj Krishna Meda Satish. Constructing Grushko and JSJ decompositions : a combinatorial approach. Group Theory [math.GR]. Université Paris-Saclay, 2018. English. NNT : 2018SACLS227 . tel-01924753

HAL Id: tel-01924753

<https://theses.hal.science/tel-01924753>

Submitted on 16 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT

de

L'UNIVERSITÉ PARIS-SACLAY

École doctorale de mathématiques Hadamard (EDMH, ED 574)

Établissement d'inscription : Université Paris-Sud

Laboratoire d'accueil : Laboratoire de mathématiques d'Orsay, UMR 8628 CNRS

Spécialité de doctorat : Mathématiques fondamentales

Suraj Krishna MEDA SATISH

Constructing Grushko and JSJ decompositions: a
combinatorial approach

Date de soutenance : 12 septembre 2018

Après avis des rapporteurs : MAHAN MJ (Tata Institute of Fundamental Research)
HENRY WILTON (University of Cambridge)

Jury de soutenance :
INDIRA CHATTERJI (Université de Nice) Présidente du jury
THOMAS DELZANT (Université de Strasbourg) Codirecteur de thèse
FRÉDÉRIC HAGLUND (Université Paris-Sud) Directeur de thèse
FRÉDÉRIC PAULIN (Université Paris-Sud) Examineur
HENRY WILTON (University of Cambridge) Rapporteur



Titre : Construction de scindements de Grushko et JSJ : une approche combinatoire

Mots Clefs : Théorie géométrique des groupes, groupes hyperboliques, complexes cubiques $CAT(0)$, scindements de Grushko, scindements JSJ

Resume : La classe des graphes de groupes libres à groupes d'arêtes cycliques constitue une source importante d'exemples en théorie géométrique des groupes, en particulier dans le cadre des groupes hyperboliques. Un résultat récent de Wilton montre qu'un tel groupe à un bout et hyperbolique contient un sous-groupe de surface, répondant à une question attribuée à Gromov. Cette thèse est consacrée à l'étude de ces groupes lorsqu'ils se présentent comme des groupes fondamentaux de certains complexes carrés à courbure négative ou nulle. Les complexes carrés en question, appelés graphes tubulaires de graphes, sont obtenus en attachant des tubes (un tube est un produit cartésien d'un cercle avec l'intervalle unitaire) à une collection finie de graphes finis. Le but principal de cette thèse est de construire deux décompositions de base pour les groupes fondamentaux de graphes tubulaires de graphes: leur décomposition de Grushko et leur décomposition JSJ.

Dans la première partie de la thèse, nous développons un algorithme en temps polynomial, dont l'entrée est un graphe tubulaire de graphes, et qui produit le scindement de Grushko de son groupe fondamental. Comme application, nous obtenons une version alternative d'un algorithme de Stallings, qui prend un ensemble fini de mots W dans un groupe libre F de rang fini, et décide s'il existe ou non un scindement libre de F relatif à W . Dans la deuxième partie de la thèse, nous développons un algorithme en temps doublement exponentiel, dont l'entrée est un graphe tubulaire de graphes avec un groupe fondamental hyperbolique à un bout, et qui produit le scindement JSJ du groupe fondamental. Nous remarquons qu'il s'agit du premier algorithme sur les scindements JSJ de groupes avec une borne effective sur la complexité de temps. La principale raison de l'efficacité de cet algorithme est que certaines propriétés asymptotiques du groupe, qui déterminent si le groupe se scinde au-dessus un sous-groupe cyclique, admettent des caractérisations locales en raison de la structure cubique $CAT(0)$. Comme application de ce résultat, nous obtenons un algorithme en temps doublement exponentiel, dont l'entrée est un groupe libre F de rang fini muni d'un ensemble fini de sous-groupes cycliques W tels que F est librement indécomposable relativement à W , et qui produit le scindement JSJ de F relativement à W .

Une conséquence des résultats ci-dessus est que le problème d'isomorphisme pour les groupes considérés se réduit à l'algorithme de Whitehead.

Title: Constructing Grushko and JSJ decompositions: a combinatorial approach

Keys words: Geometric group theory, hyperbolic groups, CAT(0) cube complexes, Grushko decompositions, JSJ decompositions

Abstract: The class of graphs of free groups with cyclic edge groups constitutes an important source of examples in geometric group theory, particularly of hyperbolic groups. A recent result of Wilton shows that any such group which is one-ended and hyperbolic contains a surface subgroup, answering a question attributed to Gromov. This thesis is devoted to the study of these groups when they arise as fundamental groups of certain nonpositively curved square complexes. The square complexes in question, called tubular graphs of graphs, are obtained by attaching tubes (a tube is a Cartesian product of a circle with the unit interval) to a finite collection of finite graphs. The main goal of this thesis is to construct two fundamental decompositions, the Grushko decomposition and the JSJ decomposition, of the fundamental groups of tubular graphs of graphs.

In the first part of the thesis we develop an algorithm of polynomial time-complexity that takes a tubular graph of graphs as input and returns the Grushko decomposition of its fundamental group. As an application, we obtain an alternative version of an algorithm of Stallings, which takes a finite set of words W in a finite rank free group F as input, and decides whether or not there exists a free splitting of F relative to W .

In the second part of the thesis we develop an algorithm of double exponential time-complexity that takes a tubular graph of graphs with one-ended hyperbolic fundamental group as input and returns the JSJ decomposition of the fundamental group. We remark that this is the first algorithm on JSJ decompositions of groups with an effective bound on the time-complexity. The main reason for the efficiency of this algorithm is that certain asymptotic properties of the group, which determine whether the group splits over a cyclic subgroup, admit local characterisations due to the CAT(0) cubical structure of these groups. As an application of this result, we obtain an algorithm of double exponential time-complexity that takes a finite rank free group F and a finite set of maximal cyclic subgroups W such that F is freely indecomposable relative to W as input and returns the JSJ decomposition of F relative to W .

A consequence of the above results is that the isomorphism problem for the groups under consideration is reduced to the Whitehead algorithm.

Dedicated to the memory of Mythri Srinagesh.

Acknowledgments

I am grateful to my advisors Frédéric Haglund and Thomas Delzant for their valuable guidance, help, support and availability during my PhD. I will always cherish this experience of learning from two brilliant mathematicians, with unique working styles. I cannot thank them enough, especially for their patience and kindness.

Many thanks to Mahan Mj and Henry Wilton for taking the time to read the thesis and for having written favorable reports. I also thank Indira Chatterji, Frédéric Paulin and Henry Wilton for being part of the defence jury.

I take this opportunity to express my gratitude:

To Markus, for being an informal mentor and a close friend. His help, both professional and personal, has been invaluable.

To IRMA, Strasbourg, for hosting me for more than a year during my PhD, and my colleagues there- Audrey, Alex, Alix, Cuong and others.

To my colleagues at Orsay- Gabriel, Corentin, Davi, Claudio, Federico, Kajal and others, for the many engaging conversations and wonderful interactions.

To Felipe, for numerous stimulating discussions during our Masters, and for sharing the difficulties of adjusting to a new life in Paris.

To Sudarshan, for immense support and interesting times together.

To Shreyas, for the flute lessons despite me being an irregular student. To Shreyas, Ketan, Pooja, Prakhar and Jay, who made me forget that I was away from home. Thank you for the numerous tasty dinners, impromptu plans, weekend getaways and

most importantly, your friendship.

To Flavie, Yann and Noah; for their remarkable hospitality in Strasbourg. They provided me with a home when all I expected was a house.

To my many teachers, especially S. Kumaresan, for establishing the MTTS program which taught me how to do mathematics. To C S Aravinda, for believing in me and mentoring me at every step. To A P Jain, for his encouragement and backing. To Mukunda and Narayana Shevire, for helping me expand my intellectual horizons and to Ram Madhav, for encouraging me to apply abroad.

To Vivek and Janson, for our friendship, and for being practically the first people with whom I discussed mathematics.

To Kavya, among many other things, for helping me improve my writing (including this note!).

I cannot thank my family enough for their support, encouragement and freedom all through. To my father, who taught me multiplication, and my grandfather, who rewarded me with Cadbury's gems every time I solved problems as a kid. To amma and akka, for showering me with so much love and affection, and to bhava for his optimism and smiling demeanour. To atthe and maava, for inspiring me through their achievements.

Finally, to Punya, for being such an excellent travel partner, in Europe and in life, and for being a solid pillar of support. For teaching me how to \TeX , and for listening to my proofs even when they made no sense to her.

Contents

Acknowledgments	vii
List of Figures	xii
1 Introduction	1
1.1 Tubular graphs of graphs	2
1.2 Free splittings	3
1.3 Cyclic splittings	6
1.3.1 Coarse behaviour and Brady-Meier complexes	8
1.3.2 Repetitive cycles and JSJ splittings	9
1.3.3 Obtaining a JSJ complex	10
1.3.4 Identifying surfaces	11
1.3.5 Relative JSJ decompositions	12
1.4 The isomorphism problem	12
1.5 Research directions	13
1.6 Graphes tubulaires de graphes	16
1.7 Scindements libres	17
1.8 Scindements cycliques	21
1.8.1 Comportement asymptotique et complexes Brady-Meier	23
1.8.2 Cycles répétitifs et scindements JSJ	24
1.8.3 L'obtention d'un complexe JSJ	25
1.8.4 Identifier les surfaces	26
1.8.5 Scindements relatifs JSJ	26
1.9 Le problème d'isomorphisme	27

1.10	Directions de recherche	28
2	The setup	31
2.1	$\mathcal{V}\mathcal{H}$ -complexes	31
2.2	Graphs of spaces	32
2.3	Tubular graphs of graphs	33
I	Constructing the Grushko decomposition	37
3	Vertex links and ends	39
3.1	Ends	39
3.2	Not one-ended	39
3.3	The second Brady-Meier criterion	43
3.4	The Algorithm	46
3.5	Whitehead graphs and separability	50
3.5.1	Construction of the double	52
3.6	A counter-example	54
II	Constructing the JSJ decomposition	57
4	Cyclic splittings and Brady-Meier complexes	59
4.1	Regular neighbourhoods and regular spheres	61
4.1.1	The regular sphere around an edge	63
4.1.2	The regular sphere around a combinatorial path	65
4.1.3	Connected regular spheres	66
4.2	Separating and coarsely separating lines	68
4.2.1	Coarsely separating periodic lines	71
4.3	A crossing criterion for lines	76
4.3.1	Graphs with no cut points	79
4.3.2	The crossing criterion	80
4.3.3	Coarse crossings of lines	80
4.4	Cyclic splittings and separating lines	82

4.4.1	Trees dual to separating lines	83
4.5	Vertical cycles and cyclic splittings	89
4.6	Universally elliptic splittings	94
4.7	Repetitive cycles	96
4.7.1	Long cycles are repetitive	100
4.8	3-repetitive cycles and crossings	102
4.8.1	Lifts of C	103
4.8.2	The main result	104
4.9	An algorithm of double exponential time	107
5	Constructing a JSJ complex	111
5.1	Splitting cycles as hyperplanes	111
5.2	Algorithmic construction of X_C	116
5.2.1	Construction of X'_C	118
5.2.2	Structure of X_C	121
5.3	The tubular graph of graphs X'	123
5.3.1	Structure of X'	124
5.4	Modification of X'	128
5.4.1	Structure of X''	128
5.5	Surface graphs	131
5.6	Construction of X_{jsj}	132
6	Relative JSJ decompositions	133
6.1	A tubular model for (F, \mathcal{H})	133
	Bibliography	137

List of Figures

2.1	An example of a hanging tree in X_s .	34
2.2	Removing rudimentary edges.	35
3.1	Removing squares containing thickness-one edges.	41
3.2	$\text{link}(u)$	44
3.3	Opening $\text{star}(u)$ to T_u	45
3.4	A highlighted path of each colour indicates a part of the image of an attaching map	46
3.5	A Whitehead graph	51
3.6	$X_{s_1}(= X_{s_2})$	54
3.7	$\text{link}(v)$	54
3.8	Step 1 of the algorithm	55
3.9	Step 2 of the algorithm	55
3.10	Step 3 of the algorithm	55
3.11	Step 4 of the algorithm	55
4.1	Cubical subdivision of a square	61
4.2	Two disjoint subpaths of P are mapped to the yellow square	62
4.3	P is a deformation retract of its regular neighbourhood	63
4.4	Regular spheres around two adjacent vertices. The star of e_a is highlighted in green.	64
4.5	The trichotomy when $v \in L^{+k} \setminus L^{+(k-1)}$	74
4.6	The edge e_i passes through the vertical hyperplane \mathbf{h}_i	74
4.7	L and L' meet at P	78
4.8	L_1 and L_2 coarsely cross	81

4.9	$L_1 \subset Y'_2, L_2 \subset Y'_1$	84
4.10	A pre-image of a square in the regular sphere	96
4.11	C is 3-repetitive with the fundamental domain P_C but not with the fundamental domain P'_C	97
4.12	L' crosses L_2 if it meets L_1 and L_3 outside L_2	104
4.13	The segments σ_1 and σ_3	106
5.1	L' in \tilde{X}' when L has three half-spaces	124

Chapter 1

Introduction

Given a group, a standard question to ask is if the group can be better understood by breaking it up into pieces of simpler groups in a meaningful way. This question has elicited a lot of interest and has been a subject of mathematical research for many decades. In geometric group theory, one often looks at free products with amalgamation and HNN extensions for this purpose.

A group is *freely indecomposable* if it does not split as a free product of two nontrivial groups. Analogous to the Kneser-Milnor prime decomposition for 3-manifolds [Mil62], the Grushko decomposition theorem [Gru40] states that a finitely generated group splits as a free product of a finite rank free group and finitely many freely indecomposable (non-free) groups, and this splitting is essentially unique.

The next step is to split a freely indecomposable group over its infinite cyclic subgroups. Zlil Sela [Sel97] showed that there exists a canonical decomposition of a freely indecomposable torsion-free hyperbolic group over its cyclic subgroups. Since this decomposition was motivated by the so-called JSJ decompositions of 3-manifolds (due to Jaco-Shalen [JS78] and Johannsen [Joh79]), Sela called it the *canonical JSJ decomposition*. We discuss JSJ decompositions in [Section 1.3](#) below.

The aim of this thesis is to explain how to construct the Grushko decomposition and the JSJ decomposition of a special class of groups, namely fundamental groups of tubular graphs of graphs, discussed in [Section 1.1](#). These groups have a CAT(0)

cubical structure of dimension two, and a $\mathcal{V}\mathcal{H}$ structure in the sense of Wise [Wis96]. Our approach to the above decomposition problems is geometric and strongly depends on the CAT(0) combinatorial structure of our groups. The algorithm for the construction of the Grushko decomposition runs in polynomial time while the algorithm for the construction of the JSJ decomposition runs in double exponential time.

As a result, one obtains a solution to the isomorphism problem for these groups (Section 1.4).

Application. Let F be a finite rank free group and \mathcal{H} a finite family of cyclic subgroups of F . Stallings obtained an algorithm in [Sta99] that detects whether or not the free group is freely indecomposable *relative to* \mathcal{H} , that is, whether or not there is a free splitting of F in which each element of \mathcal{H} conjugates into a vertex group. We give an alternate version of this algorithm in Section 3.5. Further, we also give a way to construct the relative JSJ decomposition of F relative to \mathcal{H} in Chapter 6.

We will first introduce the central objects of our study.

1.1 Tubular graphs of graphs

Finite graphs of free groups with cyclic edge groups (in the sense of Bass-Serre graphs of groups [Ser80]) are an important source of examples of hyperbolic groups. By the Bestvina-Feighn combination theorem [BF92], such a group is word-hyperbolic when it does not contain a non-trivial Baumslag-Solitar subgroup. Recently, Wilton [Wil17] showed that a graph of free groups with cyclic edge groups which is one-ended and hyperbolic contains a surface subgroup, answering a question attributed to Gromov.

When the underlying graph is a tree, and in certain other special cases, a graph of free groups with cyclic edge groups can be naturally associated to a compact square

complex, which we call a *tubular graph of graphs*. Tubular graphs of graphs are in fact \mathcal{VH} -complexes in which vertical hyperplanes are homeomorphic to circles. A \mathcal{VH} -complex, introduced by Wise in his PhD thesis [Wis96], is a square complex in which each square has edges which are alternately labelled as vertical and horizontal. We give precise definitions in [Chapter 2](#).

The universal covers of tubular graphs of graphs are CAT(0) square complexes and our methods strongly depend on this fact. A CAT(0) square complex is a simply connected square complex with the property that there is no triangle in any vertex link.

Fundamental groups of tubular graphs of graphs encompass a fairly large class of groups. For instance, all surface groups and doubles of free groups can be realised as such groups. They also have various interesting properties. For example, they are biautomatic [NR98a], a-T-menable [NR98b] and satisfy the Tits alternative [SW05].

A typical example of the groups that we deal with is a group of the form $\langle a_1, \dots, a_n \rangle \ast_{\langle w \rangle} \langle w' \rangle \ast_{\langle b_1, \dots, b_m \rangle}$, a free product with amalgamation over cyclic groups of finitely generated free groups. The main issue we address is to describe all free and cyclic splittings of such groups.

Our strategy follows 3-manifold theory, where one first cuts a 3-manifold along incompressible spheres to obtain its prime decomposition and then along incompressible tori to obtain its JSJ decomposition. Similarly, we first show how to construct the Grushko decomposition and then the JSJ decomposition by cutting along subspaces which induce free and cyclic splittings respectively.

1.2 Free splittings

The first part is devoted to questions of free splittings of fundamental groups of tubular graphs of graphs. We first note that these groups are torsion-free.

By a celebrated theorem of Stallings, a torsion-free finitely generated group is one-ended if and only if it is freely indecomposable [Sta68].

We will use a key result by N. Brady and J. Meier which imposes local conditions on a CAT(0) cube complex that imply that the complex is one-ended.

Theorem 1.2.1. [BM01] *Let X be a finite connected locally-CAT(0) cube complex. Suppose that*

(BM1) *for each vertex $v \in X$, the link of v is connected and*

(BM2) *for each vertex $v \in X$ and each simplex $\sigma \subset \text{link}(v)$, the complement of σ is (non-empty and) connected.*

Then \tilde{X} is one-ended.

We say that a square complex is *Brady-Meier* if it satisfies the conditions (BM1) and (BM2) above.

The main result of the first part of the thesis gives a geometric/combinatorial procedure that modifies a given tubular graph of graphs to a homotopy equivalent tubular graph of graphs which is not Brady-Meier if and only if the fundamental group is not one-ended:

Theorem 1.2.2 (Theorem A). *There is an algorithm of polynomial time-complexity which takes a tubular graph of graphs as input and returns a homotopy equivalent tubular graph of graphs which is either a Brady-Meier complex or contains a locally disconnecting edge or vertex which splits the fundamental group as a free product.*

The key step in the construction of our algorithm involves a simplification of the input tubular graph of graphs by ‘opening-up’ at a vertex which does not satisfy (BM2). This opening-up keeps the number of squares in the complex constant, while simplifying its vertex links. We call such an opening procedure as an *SL-move* (simplified link).

We immediately obtain a partial converse to Theorem 1.2.1:

Corollary 1.2.3. *A tubular graph of graphs has a one-ended universal cover if and only if it can be modified in finitely many SL-moves to a Brady-Meier tubular graph of graphs with an isomorphic fundamental group.*

As a consequence of Theorem A, we have:

Corollary 1.2.4 (Corollary B). *There is an algorithm of polynomial time-complexity which takes as input a tubular graph of graphs and decides whether or not its fundamental group is one-ended. In addition, it returns the Grushko decomposition of the fundamental group.*

We point out that our proof does not use Stallings' theorem nor assume the existence of a Grushko decomposition. In fact, our procedure yields a new proof of Stallings' theorem for fundamental groups of tubular graphs of graphs as well as the existence of a Grushko decomposition for these groups.

In fact, we obtain the following analogue of a result in 3-manifold theory due to Jaco [Jac69] which states that if the fundamental group of a compact 3-manifold is a free product, then each free factor is itself the fundamental group of a 3-manifold.

Corollary 1.2.5 (Corollary C). *Let X be a tubular graph of graphs with fundamental group G . If $G = A * B$, then there exist tubular graphs of graphs X_1 and X_2 such that A and B are fundamental groups of X_1 and X_2 respectively. Moreover, X_1 and X_2 can be so chosen such that the total number of squares in X_1 and X_2 is bounded by the number of squares in X .*

The Grushko decomposition may be found algorithmically in other situations. Jaco, Letscher and Rubinstein [JLR02] gave an algorithm of polynomial time-complexity to compute the prime decomposition of a 3-manifold from a triangulation. Gerasimov [Ger99] showed that the Grushko decomposition can be computed for hyperbolic groups, but his algorithm is a Turing machine without a solution to the halting problem. Dahmani and Groves [DG08] extended Gerasimov's ideas to groups which are hyperbolic relative to abelian subgroups. Diao and Feighn [DF05] gave an algorithm for graphs of free groups using cocompact actions of the groups on products of trees as studied by Fujiwara-Papasoglu [FP06].

We conclude the first part with another application of our algorithm. As defined by Stallings in [Sta99], a finite set of words W of a finite rank free group F is *separable* if there exists a nontrivial free splitting of F such that each word of W conjugates into a free factor. In other words, W is separable if and only if F admits a free

splitting relative to the cyclic subgroups generated by elements of W .

Stallings obtained an algorithm to detect separability in [Sta99]. He constructs a Whitehead graph for the given set of words in a chosen basis. He then uses a Whitehead automorphism to modify the basis whenever there is a cut vertex in the Whitehead graph to reduce the total length of the given set of words. We give an alternate version of this algorithm using [Theorem A](#). In fact, our method is strongly related to Stallings' method.

Corollary 1.2.6 (Stallings, [Corollary D](#)). *There exists an algorithm of polynomial time-complexity that takes a finite set of words in a finite rank free group as input and decides whether it is separable.*

For our algorithm, we first construct the tubular graph of graphs associated to a 'double' of the free group with the given set of words. In [Lemma 3.5.9](#), we show that that the vertex link of a special vertex in the double that we construct is isomorphic to a Whitehead graph associated to the free group and the given set of words. So far, we are in a similar situation as Stallings. However, we apply the algorithm of [Theorem A](#) at this stage, and hence we perform an SL-move if there is a vertex link with a cut vertex, whereas Stallings chooses a Whitehead automorphism.

We then use Wilton's characterization [Wil12] of free splittings of graphs of free groups with cyclic edge groups to conclude that the set of words is separable if the fundamental group of the double is not one-ended.

1.3 Cyclic splittings

The second part of the thesis is devoted to the explicit geometric/combinatorial construction of JSJ decompositions of one-ended fundamental groups of tubular graphs of graphs in the hyperbolic case. In order to define a JSJ decomposition, we will need a few definitions.

We adopt Sela's terminology [Sel97]. 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 non-cyclic 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 [Bow98]) if it is not properly contained as a finite index subgroup in any subgroup of G .

We are now ready to define JSJ decompositions in the sense of Sela ([Sel97]), modified by Bowditch [Bow98].

Definition 1.3.1 (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 1.3.2 ([Sel97]). *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 1.3.3 (Theorem 5.6.2). *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 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 [DG11], Dahmani and Guirardel give an algorithm to compute JSJ decompositions of one-ended hyperbolic groups over maximal virtually cyclic subgroups with infinite centre. In [DT13], Dahmani and Touikan give an algorithm to compute JSJ decompositions of torsion-free hyperbolic groups over its cyclic subgroups. In [Bar16], Barrett gives an algorithm to compute JSJ decompositions of one-ended hyperbolic groups over virtually 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.

1.3.1 Coarse behaviour and Brady-Meier complexes

Let X be a tubular graph of graphs (Section 1.1) endowed with its \mathcal{VH} structure. Each vertical hyperplane of X is a circle (Proposition 2.3.2). If \tilde{X} denotes the CAT(0) universal cover of X , then the vertical hyperplanes of \tilde{X} are lines. Let G denote the one-ended fundamental group of X . By Corollary 1.2.3, we can assume that \tilde{X} is Brady-Meier. Adopting the terminology of Scott and Wall [SW79], X has a structure of a graph of spaces (see Section 2.2 for details), where each vertex space is itself a graph. Similarly, \tilde{X} has a structure of a tree of spaces, where each vertex space is a (vertical) tree.

Let L be a subset of \tilde{X} . We will denote by $N_R(L)$ the set of all points in \tilde{X} at distance at most R from a point of L . We say L *separates* \tilde{X} if $\tilde{X} \setminus L$ is not connected. It *coarsely separates* \tilde{X} [Pap12] if there exists $R > 0$ such that $\tilde{X} \setminus N_R(L)$ contains at least two components which are not contained in $N_{R'}(L)$ for any $R' > 0$.

An *axis* in \tilde{X} of an element $g \in G$ is a geodesic line in \tilde{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 \tilde{X} [BH99]. If G splits over $\langle g \rangle$, then L coarsely separates \tilde{X} (as a consequence of [Pap05, Lemma 1.8]). When L is contained in a vertical tree, the fact that \tilde{X} is a Brady-Meier complex implies that L in fact separates \tilde{X} (Lemma 4.2.15). In the simply connected space \tilde{X} , L separates \tilde{X} if and only if it separates $N_{\frac{1}{4}}(L)$

(Lemma 4.2.8).

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 $\langle g \rangle$, implies the following result:

Lemma 1.3.4 (Lemma 4.5.25). *C separates its regular neighbourhood.*

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 [GL16]. The edge groups of the JSJ decomposition are universally elliptic.

Two coarsely separating lines L_1 and L_2 *coarsely cross* if there exists $R > 0$ such that L_1 meets different components of $\tilde{X} \setminus N_R(L_2)$ and L_2 meets different components of $\tilde{X} \setminus N_R(L_1)$ (Definition 4.3.9). 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 coarsely cross (Lemma 4.6.1). This property has a local characterization in the Brady-Meier complex \tilde{X} :

Proposition 1.3.5 (Proposition 4.3.2). *Two separating lines L_1 and L_2 coarsely cross if and only if*

1. $L_1 \cap L_2$ is non-empty and compact, and
2. L_2 meets two components of $N_{\frac{1}{4}}(L_1 \cap L_2) \setminus L_1$.

1.3.2 Repetitive cycles and JSJ splittings

In Section 4.7, we introduce an important notion, namely repetitivity, that bounds the length of a cycle that induces a universally elliptic splitting. Let \tilde{C} denote a lift of a cycle C in \tilde{X} .

Definition 1.3.6 (Definition 4.7.2, Lemma 4.7.4). A cycle C is *k-repetitive* if \tilde{C} is a separating line and there exists an edge e in \tilde{X} and distinct elements $g_1, \dots, g_k \in G$

such that

1. each translate $g_i\tilde{C}$ contains e ,
2. the distance between e and g_ie is strictly less than the length of C , and
3. any two squares \mathfrak{s} and \mathfrak{s}' that contain e are either separated by all translates $g_i\tilde{C}$ or by none of them.

There are two important reasons for defining repetitive cycles. The first reason is that any cycle that is longer than a certain bound is k -repetitive ([Proposition 4.7.9](#)). Here, the bound depends only on k and the number of squares of X .

The second reason is the following:

Proposition 1.3.7 ([Proposition 4.8.1](#)). *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 periodic separating line L in \tilde{X} such that L and \tilde{C} coarsely cross.*

This implies that $\pi_1(C)$ conjugates into a hanging surface subgroup of the JSJ splitting of G , by [Proposition 5.30](#) of [Bow98]. Hence, $\pi_1(C)$ is not universally elliptic.

Therefore, the length of a cycle which induces a universally elliptic cyclic subgroup is bounded. This leads to the following result:

Theorem 1.3.8 ([Theorem 4.9.1](#)). *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 upto commensurability.*

1.3.3 Obtaining a JSJ complex

In [Chapter 5](#), we modify the given tubular graph of graphs X to a tubular graph of graphs X_{jsj} such that the graph of groups structure of X_{jsj} is the JSJ decomposition of G .

The first step involves modifying the initial tubular graph of graphs X to a tubular

graph of graphs X' by cutting along the finite list of cycles supplied by [Theorem 1.3.8](#). We do this cutting procedure using the machinery of *spaces with walls* [HP98]. The vertex set of \tilde{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 1.3.8](#). The square complex dual to this new space with walls is \tilde{X}' (see [Section 5.3](#)).

In the second step, we perform a simplification on X' by removing tubes which are attached to cyclic vertex graphs on both sides. Call the new tubular graph of graphs as X'' . In [Proposition 5.4.3](#), 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 X'' .

Thus, an edge stabiliser of the underlying tree of \tilde{X}'' 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. So what remains is to identify the maximal surface subgroups that appear as vertex groups in the JSJ decomposition.

1.3.4 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 1.3.9 ([Lemma 5.5.2](#)). *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.*

Armed with this result, we remove tubes in X'' which are attached to surface graphs on both sides. The resulting tubular graph of graphs will then have the JSJ decomposition as its underlying graph of groups, proving the main result ([Theorem 1.3.3](#)).

1.3.5 Relative JSJ decompositions

Let \mathcal{H} be a family of subgroups of a group G . A *splitting of G relative to \mathcal{H}* is a graph of groups decomposition of G in which each element of \mathcal{H} is elliptic.

A *relative JSJ splitting* of a finite rank free group F relative to a family \mathcal{H} of maximal cyclic subgroups is a graph of groups splitting of F relative to \mathcal{H} that satisfies the conditions of a JSJ decomposition ([Definition 1.3.1](#)).

In [[Cas16](#), Theorem 4.25], Cashen shows that a relative JSJ splitting exists and is unique.

As an application of our algorithm to construct the JSJ decomposition ([Theorem 1.3.3](#)), we obtain an algorithm to construct the relative JSJ decomposition:

Theorem 1.3.10 ([Theorem 6.0.3](#)). *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} .*

Our approach is to construct a Brady-Meier tubular graph of graphs X such that F is the fundamental group of a vertex graph of X and the incident edge groups are precisely members of \mathcal{H} (see [Chapter 6](#) for details). The construction ensures that $\pi_1(X)$ is hyperbolic. The algorithm of our main result ([Theorem 1.3.3](#)) then constructs the JSJ decomposition of $\pi_1(X)$, and thus the relative JSJ of (F, \mathcal{H}) .

1.4 The isomorphism problem

The isomorphism problem is the algorithmic problem of deciding whether two finite presentations of groups present isomorphic groups [[Deh11](#)].

An important consequence of the work done in this thesis is that the isomorphism problem for hyperbolic fundamental groups of tubular graphs of graphs is reduced to the Whitehead algorithm ([[Whi36](#)]):

Let G_1 and G_2 be hyperbolic fundamental groups of two tubular graphs of graphs.

Using the algorithm of [Corollary 1.2.4](#), we can construct tubular graphs of graphs which give Grushko decompositions of G_1 and G_2 . By the uniqueness of the Grushko decomposition, G_1 and G_2 are isomorphic if and only if there is a one-to-one correspondence between the vertex groups of their Grushko decompositions with an isomorphism between the corresponding vertex groups. The free groups in the Grushko decompositions have the same rank if and only if the tubular graphs of graphs which represent these free groups can be modified to graphs with the same number of edges outside a maximal tree. This is straightforward. It remains to check if the one-ended factors of the Grushko decomposition are isomorphic.

So assume that G_1 and G_2 are one-ended. Using [Theorem 1.3.3](#), we can construct the JSJ decomposition of G_1 and G_2 . Then G_1 and G_2 are isomorphic if and only if there is a one-to-one correspondence between the vertex groups of their JSJ decompositions such that the corresponding vertex groups are isomorphic with the isomorphism respecting the incident edge groups. But all vertex groups are free groups and their incident edge groups are cyclic subgroups. The Whitehead algorithm [Whi36] precisely decides if there exists an isomorphism between free groups that takes one finite set of cyclic subgroups to another.

1.5 Research directions

We list a few directions in which the current work can be taken forward.

1. One question is to construct a JSJ decomposition of the fundamental group G of a tubular graph of graphs when G is not hyperbolic. In the hyperbolic case, each edge group is maximal cyclic in at least one of the vertex groups at the endpoints of the edge. This is not true when G is not hyperbolic. Another difficulty is to decide when a cyclic subgroup conjugates into a surface subgroup. We believe that the methods used in this thesis can be extended to the non-hyperbolic case as well.
2. A second question is whether it is possible to give an algorithm of bounded time-

complexity to construct the JSJ decomposition of the hyperbolic fundamental group of a general finite graph of free groups with cyclic edge groups, in the absence of a tubular graph of graphs structure.

3. Most of our results only use the fact that \tilde{X} is a Brady-Meier CAT(0) cube complex. It is interesting to ask which results hold for general Brady-Meier CAT(0) cube complexes with or without a \mathcal{VH} structure. In particular, it is an open question as to which classes of one-ended groups arise as fundamental groups of Brady-Meier cube complexes, as in [Corollary 1.2.3](#). We used Brady-Meier complexes to analyse free and cyclic splittings in this thesis. It would be interesting to see if the Brady-Meier structure will be useful to study other kinds of splittings.

Introduction (en français)

Étant donné un groupe, une question standard à poser est de savoir si le groupe peut être mieux compris en le scindant en sous-groupes plus simples. Cette question a suscité beaucoup d'intérêt et fait l'objet de recherches en mathématiques depuis plusieurs décennies. Pour cela, en théorie géométrique des groupes, on se penche souvent sur les produits libres amalgamés et sur les extensions HNN.

Un groupe est *librement indécomposable* s'il ne se scinde pas en un produit libre de deux groupes non triviaux. Comme la décomposition primaire de Kneser-Milnor pour les 3-variétés [Mil62], le théorème de décomposition de Grushko [Gru40] indique qu'un groupe de type fini se scinde en un produit libre d'un groupe libre de rang fini et d'un nombre fini de groupes librement indécomposables. De plus, ce scindement est essentiellement unique.

L'étape suivante consiste en l'étude des scindements d'un groupe librement indécomposable au-dessus de ses sous-groupes cycliques infinis. Sela [Sel97] a montré qu'il existe un scindement canonique d'un groupe hyperbolique sans torsion librement indécomposable au-dessus de ses sous-groupes cycliques. Puisque ce scindement était motivé par les décompositions JSJ des 3-variétés (d'après Jaco-Shalen [JS78] et Johannsen [Joh79]), Sela l'appelait le *scindement JSJ canonique*. Nous discuterons les scindements JSJ en Section 1.8 ci-dessous.

Le but de cette thèse est d'expliquer comment construire le scindement de Grushko et le scindement JSJ d'une classe particulier de groupes, à savoir des groupes fondamentaux de graphes tubulaires de graphes, discutés en Section 1.6. Ces groupes ont une structure $CAT(0)$ cubique de dimension deux, et possèdent une structure \mathcal{VH}

au sens de Wise [Wis96]. Notre approche pour résoudre des problèmes de scindement ci-dessus est géométrique, et dépend fortement de la structure combinatoire $CAT(0)$ de nos groupes. L'algorithme pour la construction du scindement de Grushko s'exécute en temps polynomial tandis que l'algorithme pour la construction du scindement JSJ s'exécute en temps doublement exponentiel.

Par conséquent, on obtient une solution au problème d'isomorphisme pour ces groupes (Section 1.9).

Application. Soient F un groupe libre de rang fini et \mathcal{H} une famille finie de sous-groupes cycliques de F . Stallings a obtenu un algorithme [Sta99] qui détecte si le groupe libre est librement indécomposable *relativement à \mathcal{H}* , c'est-à-dire s'il y a ou non un scindement libre de F dans lequel chaque élément de \mathcal{H} se conjugue dans un groupe de sommet. Nous donnons une version alternative de cet algorithme en Section 3.5. De plus, nous donnons un moyen de construire le scindement JSJ de F relatif à \mathcal{H} en Chapitre 6.

Nous allons d'abord présenter les objets centraux de notre étude.

1.6 Graphes tubulaires de graphes

Les graphes finis de groupes libres à groupes d'arêtes cycliques (au sens de Bass-Serre [Ser80]) sont une source importante d'exemples de groupes hyperboliques. Par le théorème de combinaison de Bestvina-Feighn [BF92], un tel groupe est hyperbolique s'il ne contient pas un sous-groupe Baumslag-Solitar non-trivial. Récemment, Wilton [Wil17] a montré qu'un graphe de groupes libres aux groupes d'arêtes cycliques ayant un bout et hyperbolique, contient un sous-groupe de surface, répondant à une question attribuée à Gromov.

Lorsque le graphe sous-jacent est un arbre, et dans certains autres cas particuliers, un graphe de groupes libres aux groupes d'arêtes cycliques peut être naturellement associé à un complexe carré compact, que nous appelons un *graphe tubulaire de*

graphes. Les graphes tubulaires des graphes sont en fait des complexes \mathcal{VH} dans lesquels les hyperplans verticaux sont homéomorphes à des cercles. Un complexe \mathcal{VH} , introduit par Wise dans sa thèse [Wis96], est un complexe carré dans lequel chaque carré a des arêtes qui sont alternativement étiquetées comme verticales et horizontales. Nous donnons des définitions précises en Chapitre 2.

Les revêtements universels des graphes tubulaires de graphes sont des complexes carrés $\text{CAT}(0)$ et nos méthodes dépendent fortement de ce fait. Un complexe carré $\text{CAT}(0)$ est un complexe carré simplement connexe avec la propriété qu'il n'y a pas de triangles dans les *links* de sommets.

Les groupes fondamentaux de graphes tubulaires de graphes incluent une classe assez grande de groupes. Par exemple, tous les groupes de surface et les doubles de groupes libres (le long de sous-groupes cycliques) peuvent être réalisés en tant que tels groupes. Ils possèdent des propriétés intéressantes. Par exemple, ils sont bi-automatiques [NR98a] et satisfont l'alternative de Tits [SW05].

Un exemple typique de groupes que nous traitons dans cette thèse est un groupe de la forme $\langle a_1, \dots, a_n \rangle *_{\langle w \rangle = \langle w' \rangle} \langle b_1, \dots, b_m \rangle$. Le problème principal que nous abordons est de décrire tous les scindements libres et cycliques de ces groupes.

Notre stratégie suit la théorie des 3-variétés, où l'on coupe d'abord une 3-variété le long de sphères incompressibles pour obtenir sa décomposition primaire et ensuite le long de tores incompressibles pour obtenir sa décomposition JSJ. De la même façon, nous montrons d'abord comment construire le scindement de Grushko puis le scindement JSJ en coupant le long de sous-espaces qui induisent respectivement des scindements libres et cycliques.

1.7 Scindements libres

La première partie est consacrée à des questions de scindements libres de groupes fondamentaux de graphes tubulaires de graphes. Notons d'abord que ces groupes sont sans torsion.

Par un théorème célèbre de Stallings, un groupe fini sans torsion a un bout si et seulement s'il est librement indécomposable [Sta68].

Nous allons utiliser un résultat-clé de Brady et Meier. Ces auteurs donnent des conditions locales sur un complexe cubique $\text{CAT}(0)$ qui impliquent que le complexe a un bout.

Théorème 1.7.1 ([BM01]). *Soit X un complexe cubique $\text{CAT}(0)$ localement fini, tel que*

(BM1) *pour chaque sommet $v \in X$, le link de v ($\text{link}(v)$) est connexe, et*

(BM2) *pour chaque $v \in X$ et chaque simplexe $\sigma \subset \text{link}(v)$, le complément de σ est (non-vide et) connexe.*

Alors \tilde{X} a un bout.

Nous dirons qu'un complexe carré est *Brady-Meier* s'il satisfait les conditions (BM1) et (BM2) ci-dessus.

Le résultat principal de la première partie de thèse donne une procédure géométrique/combinatoire pour modifier un graphe tubulaire de graphes en un graphe tubulaire de graphes homotopiquement équivalent qui n'est pas Brady-Meier si, et seulement si, le groupe fondamental a plus qu'un bout :

Théorème 1.7.2 (Théorème A). *Il existe un algorithme en temps polynomial qui prend en entrée un graphe tubulaire de graphes et qui renvoie un graphe tubulaire de graphes homotopiquement équivalent à celui d'entrée, et de plus le complexe produit ou bien est Brady-Meier, ou bien contient une arête ou un sommet localement déconnectant qui scinde le groupe fondamental en un produit libre.*

L'étape clé dans la construction de notre algorithme implique une simplification du graphe tubulaire de graphes par «ouverture» en un sommet qui ne satisfait pas (BM2). Cette ouverture ne change pas le nombre de carrés dans le complexe, tout en simplifiant ses links de sommets. Nous appelons une telle procédure d'ouverture un *mouvement-SL* (link simplifié).

Nous obtenons immédiatement une réciproque partielle du Théorème 1.7.1:

Corollaire 1.7.3. *Un graphe tubulaire de graphes a un revêtement universel à un bout si et seulement s'il peut être modifié par un nombre fini de mouvements-SL en un graphe tubulaire de graphes Brady-Meier avec le même groupe fondamental.*

Comme conséquence du Théorème 1.7.2, nous obtenons :

Corollaire 1.7.4 (Corollaire B). *Il existe un algorithme en temps polynomial qui prend en entrée un graphe tubulaire de graphes et décide si son groupe fondamental a un bout. De plus, il renvoie la décomposition de Grushko du groupe fondamental.*

Nous rappelons que notre démonstration n'utilise pas le théorème de Stallings et ne présuppose pas l'existence d'une décomposition de Grushko. En fait, notre procédure fournit une nouvelle preuve du théorème de Stallings pour les groupes fondamentaux de graphes tubulaires de graphes ainsi que l'existence d'une décomposition de Grushko pour ces groupes.

En fait, nous obtenons l'analogie suivant d'un résultat dans la théorie des 3-variétés d'après Jaco [Jac69] qui stipule que si le groupe fondamental d'une 3-variété compact est un produit libre, alors chaque facteur libre est lui-même le groupe fondamental d'une 3-variété.

Corollaire 1.7.5 (Corollaire C). *Soit X un graphe tubulaire de graphes, de groupe fondamental G . Si $G = A * B$, alors il existe des graphes tubulaires de graphes X_1 et X_2 tels que A et B sont des groupes fondamentaux de X_1 et X_2 respectivement. De plus, X_1 et X_2 peuvent être choisis de telle sorte que le nombre total de carrés dans X_1 et X_2 est borné par le nombre de carrés dans X .*

La décomposition de Grushko peut être trouvée de manière algorithmique dans d'autres situations. Jaco, Letscher et Rubinstein [JLR02] ont donné un algorithme en temps polynomial pour calculer la décomposition primaire d'une 3-variété à partir d'une triangulation. Gerasimov [Ger99] a montré que la décomposition de Grushko peut être calculée pour des groupes hyperboliques, mais son algorithme est une machine de Turing sans solution au problème d'arrêt. Dahmani et Groves [DG08] ont étendu les idées de Gerasimov aux groupes hyperboliques relatifs aux sous-groupes abéliens. Diao et Feighn [DF05] ont donné un algorithme pour les

graphes de groupes libres en utilisant des actions co-compactes de ces groupes sur les produits d'arbres étudiés par Fujiwara-Papasoglu [FP06].

Nous concluons la première partie avec une autre application de notre algorithme. Suivant la terminologie de Stallings dans [Sta99], un ensemble fini W de mots d'un groupe libre F de rang fini est *séparable* s'il existe un scindement libre non trivial de F tel que chaque mot de W se conjugue dans un facteur libre. Autrement dit, W est séparable si et seulement si F admet un scindement libre relatif aux sous-groupes cycliques engendrés par les éléments de W .

Stallings a obtenu un algorithme pour détecter la séparabilité dans [Sta99]. Il construit un graphe de Whitehead pour l'ensemble de mots donné dans une base choisie. Il utilise ensuite un automorphisme de Whitehead pour modifier la base chaque fois qu'il y a un sommet déconnectant dans le graphe de Whitehead pour réduire la longueur totale de l'ensemble de mots donné. Nous donnons une version alternative de cet algorithme en utilisant Théorème 1.7.2. En fait, notre méthode est fortement liée à la méthode de Stallings.

Corollaire 1.7.6 (Stallings, Théorème D). *Il existe un algorithme en temps polynomial qui prend en entrée un ensemble fini de mots dans un groupe libre de rang fini et décide s'il est séparable.*

Pour notre algorithme, nous construisons d'abord le graphe tubulaire des graphes associé à un «double» du groupe libre avec l'ensemble de mots donné. Dans le Lemme 3.5.9, nous montrons que le link de sommet d'un sommet spécial dans le double que nous construisons est isomorphe à un graphe de Whitehead associé au groupe libre et à l'ensemble de mots donné. Jusqu'à présent, nous sommes dans une situation similaire à celle de Stallings. Cependant, nous appliquons l'algorithme donné par Théorème 1.7.2 à ce stade, et donc nous effectuons un mouvement-SL s'il y a un link de sommet avec un sommet déconnectant, alors que Stallings choisit un automorphisme de Whitehead.

Nous utilisons ensuite la caractérisation de Wilton [Wil12] des scindements libres de graphes de groupes libres aux groupes d'arêtes cycliques pour conclure que l'ensemble

des mots est séparable si le groupe fondamental du double n'a pas seulement un bout.

1.8 Scindements cycliques

La deuxième partie de cette thèse est consacrée à la construction géométrique/combinatoire explicite de décompositions JSJ de groupes fondamentaux à un bout de graphes tubulaires de graphes dans le cas hyperbolique. Afin de définir une décomposition JSJ, nous aurons besoin de quelques définitions.

Nous adoptons la terminologie de Sela [Sel97]. Soit G un groupe hyperbolique sans torsion. Un *sous-groupe de surface suspendu* G' de G est un sous-groupe isomorphe au groupe fondamental d'une surface à bord tel qu'il existe une décomposition de G en graphe de groupes dans lequel G' est un groupe de sommet dont les groupes d'arêtes incidents sont précisément les sous-groupes périphériques de G' . Un *sous-groupe de surface suspendu maximal* est un sous-groupe de surface suspendu qui n'est pas un sous-groupe propre d'un sous-groupe de surface suspendu. Un sous-groupe non-cyclique G' de G est *rigide* s'il est elliptique dans chaque scindement cyclique de G .

Un sous-groupe est *plein* (au sens de Bowditch [Bow98]) s'il n'est pas un sous-groupe propre d'indice fini dans un sous-groupe de G .

Nous sommes maintenant prêts à définir la décomposition JSJ au sens de Sela ([Sel97]), modifié par Bowditch [Bow98].

Définition 1.8.1 (Décomposition JSJ). *Soit G un groupe hyperbolique sans torsion. Un scindement JSJ de G est un scindement fini de G en graphe de groupes où chaque groupe d'arête est cyclique et chaque groupe de sommet est plein et de l'un des trois types suivants:*

1. *un sous-groupe cyclique,*
2. *un sous-groupe de surface suspendu maximal, ou*

3. un sous-groupe rigide.

Si un sommet v de type (1) a degré un, alors le groupe d'arête incident est un sous-groupe propre du groupe de sommet G_v . De plus, exactement une extrémité de chaque arête est de type (1) et les groupes d'arêtes qui se connectent à un groupe de sommet de type (2) sont précisément les sous-groupes périphériques de ce groupe.

Théorème 1.8.2 ([Sel97]). *Soit G un groupe hyperbolique à un bout sans torsion, qui n'est pas de groupe fondamental d'une surface fermée. Alors une décomposition JSJ de G existe et elle est unique.*

Nous sommes maintenant prêts à énoncer notre résultat principal.

Théorème 1.8.3 (Théorème 5.6.2). *Il existe un algorithme en temps doublement exponentiel qui prend en entrée un graphe tubulaire de graphes Brady-Meier à groupe fondamental hyperbolique G et qui renvoie un graphe tubulaire de graphes dont la structure de graphe de groupes est celle de la décomposition JSJ de G .*

D'autres auteurs ont obtenu des algorithmes pour calculer des décompositions JSJ de groupes dans des conditions différentes. Dans [DG11], Dahmani et Guirardel donnent un algorithme pour calculer les décompositions JSJ de groupes hyperboliques à un bout au-dessus des sous-groupes maximaux virtuellement cycliques à centre infini. Dans [DT13], Dahmani et Touikan donnent un algorithme pour calculer les décompositions JSJ de groupes hyperboliques sans torsion au-dessus de ses sous-groupes cycliques. Dans [Bar16], Barrett donne un algorithme pour calculer les décompositions JSJ de groupes hyperboliques à un bout au-dessus des sous-groupes virtuellement cycliques. Nous remarquons que la complexité en temps de ces algorithmes n'est pas connue.

Notre approche est combinatoire/géométrique. Nous allons maintenant brièvement décrire cette approche.

1.8.1 Comportement asymptotique et complexes

Brady-Meier

Soit X un graphe tubulaire de graphes (Section 1.6) doté de sa structure \mathcal{VH} . Chaque hyperplan vertical de X est un cercle (Proposition 2.3.2). Si \tilde{X} dénote le revêtement universel $\text{CAT}(0)$ de X , alors les hyperplans verticaux de \tilde{X} sont des droites. Soit G le groupe fondamental à un bout de X . Par Théorème 1.7.3, on peut supposer que \tilde{X} est Brady-Meier. En adoptant la terminologie de Scott et Wall [SW79], X a une structure d'un graphe d'espaces (voir Section 2.2 pour les détails), où chaque espace de sommet est lui-même un graphe. De même, \tilde{X} a la structure d'un arbre d'espaces, où chaque espace de sommet est un arbre (vertical).

Soit L un sous-ensemble de \tilde{X} . Notons $N_R(L)$ l'ensemble de tous les points de \tilde{X} à distance au plus R d'un point de L . On dit L *sépare* \tilde{X} si $\tilde{X} \setminus L$ n'est pas connexe. Il *sépare* \tilde{X} *grossièrement* [Pap12] s'il existe $R > 0$ tel que $\tilde{X} \setminus N_R(L)$ contient au moins deux composants connexes qui ne sont pas contenu dans $N_{R'}(L)$ pour tout $R' > 0$.

Un *axe* dans \tilde{X} d'un élément $g \in G$ est une droite géodésique dans \tilde{X} qui est invariante sous l'action du sous-groupe cyclique $\langle g \rangle$. Étant donné $g \in G$, un axe L de g existe toujours dans \tilde{X} [BH99]. Si G se scinde au-dessus de $\langle g \rangle$, alors L sépare \tilde{X} grossièrement (en conséquence de [Pap05, Lemme 1.8]). Quand L est contenu dans un arbre vertical, le fait que \tilde{X} soit un complexe Brady-Meier implique que L en fait sépare \tilde{X} (Lemme 4.2.15). Dans l'espace \tilde{X} simplement connexe, L sépare \tilde{X} si et seulement s'il sépare $N_{\frac{1}{4}}(L)$ (Lemme 4.2.8).

Le quotient de L par $\langle g \rangle$ est un cercle C immergé dans X , que nous appelons un *cycle*. Le *voisinage régulier* de C est le quotient de $N_{\frac{1}{4}}(L)$ par l'action de $\langle g \rangle$. Le fait que L sépare $N_{\frac{1}{4}}(L)$, avec une condition satisfaite puisque G se scinde au-dessus de $\langle g \rangle$, implique le résultat suivant:

Lemme 1.8.4 (Lemma 4.5.25). *C sépare son voisinage régulier.*

Nous avons besoin d'une autre propriété pour construire la décomposition JSJ. Un sous-groupe cyclique au-dessus duquel G se scinde est *universellement elliptique* s'il

est elliptique dans l'arbre de Bass-Serre de tout scindement cyclique de G [GL16]. Les groupes d'arêtes de la décomposition JSJ sont universellement elliptiques.

Deux droites L_1 et L_2 qui séparent \tilde{X} se croisent grossièrement s'il existe $R > 0$ tel que L_1 rencontre différentes composantes connexes de $\tilde{X} \setminus N_R(L_2)$ et L_2 rencontre différents composants de $\tilde{X} \setminus N_R(L_1)$ (Définition 4.3.9). Soit L_1 (respectivement L_2) un axe de g_1 (respectivement de g_2) tel que G se scinde au-dessus de $\langle g_1 \rangle$ et $\langle g_2 \rangle$. Alors $\langle g_1 \rangle$ est elliptique dans l'arbre de Bass-Serre du scindement au-dessus de $\langle g_2 \rangle$ seulement si L_1 et toute translation de L_2 ne se croisent pas grossièrement (Lemme 4.6.1). Cette propriété a une caractérisation locale dans le complexe Brady-Meier \tilde{X} :

Proposition 1.8.5 (Proposition 4.3.2). *Deux droites séparantes L_1 et L_2 se croisent grossièrement si et seulement si*

1. $L_1 \cap L_2$ est non-vide et compact, et
2. L_2 rencontre deux composants connexes distinctes de $N_{\frac{1}{4}}(L_1 \cap L_2) \setminus L_1$.

1.8.2 Cycles répétitifs et scindements JSJ

En Section 4.7, nous introduisons une notion importante, à savoir la répétitivité, qui borne la longueur d'un cycle qui induit un scindement universellement elliptique. Soit \tilde{C} une élévation d'un cycle C dans \tilde{X} .

Définition 1.8.6 (Définition 4.7.2, Lemme 4.7.4). Un cycle C est k -répétitif si \tilde{C} est une droite séparante et s'il existe un arête e dans \tilde{X} et des éléments distincts $g_1, \dots, g_k \in G$ tels que

1. chaque translation $g_i \tilde{C}$ contient e ,
2. la distance entre e et $g_i e$ est strictement inférieure à la longueur de C , et
3. deux carrés quelconques \mathbf{s} et \mathbf{s}' qui contiennent e sont séparés soit par tous les translats $g_i \tilde{C}$ soit par aucun.

Il y a deux raisons importantes de définir des cycles répétitifs. La première est que

tout cycle plus long qu'une certaine borne est k -répétitif (Proposition 4.7.9). Ici, la borne ne dépend que de k et du nombre de carrés de X .

La deuxième est donnée par la proposition suivante:

Proposition 1.8.7 (Proposition 4.8.1). *Soit C un cycle k -répétitif avec $k \geq 3$. Supposons que $\pi_1(C)$ est un sous-groupe cyclique maximal de G . Alors il existe une droite séparante périodique L dans \tilde{X} telle que L et \tilde{C} se croisent grossièrement.*

Ceci implique que $\pi_1(C)$ se conjugue dans un sous-groupe de surface suspendu du scindement JSJ de G , par la Proposition 5.30 de [Bow98]. Par conséquent, $\pi_1(C)$ n'est pas universellement elliptique.

Donc la longueur d'un cycle qui induit un sous-groupe cyclique universellement elliptique est bornée. Cela mène au résultat suivant:

Théorème 1.8.8 (Théorème 4.9.1). *Il existe un algorithme en temps doublement exponentiel qui prend en entrée un graphe tubulaire de graphes Brady-Meier dont le groupe fondamental hyperbolique est G et qui renvoie une liste finie de cycles de scindements qui contient tous les sous-groupes universellement elliptiques de G à commensurabilité près.*

1.8.3 L'obtention d'un complexe JSJ

En Section 5, nous modifions le graphe tubulaire de graphes X en un graphe tubulaire de graphes X_{jsj} tel que la structure du graphe de groupes de X_{jsj} donne le scindement JSJ de G .

La première étape consiste à modifier X en un graphe tubulaire de graphes X' en coupant le long de la liste finie de cycles fournie par Théorème 1.8.8. Nous effectuons ce découpage en utilisant la machinerie d'espaces à murs [HP98]. L'ensemble des sommets de \tilde{X} est un espace à murs, avec des murs définis par ses hyperplans verticaux et horizontaux. Nous enrichissons l'ensemble de murs en ajoutant les relevés des cycles fournies par Théorème 1.8.8 dans \tilde{X} . Le complexe carré dual à ce nouvel espace à murs est \tilde{X}' (voir Section 5.3).

Dans la deuxième étape, nous effectuons une simplification sur X' en enlevant les tubes qui sont attachés aux graphes de sommets cycliques des deux côtés. Appelons le nouveau graphe tubulaire de graphes X'' . En Proposition 5.4.3, nous montrons que chaque groupe d'arête du scindement JSJ de G se conjugue dans un groupe d'arête du graphe de groupes sous-jacent de X'' .

Ainsi, un stabilisateur d'arête de l'arbre sous-jacent de \tilde{X}'' est soit un stabilisateur d'arête de l'arbre de JSJ, soit un sous-groupe cyclique qui se conjugue dans un sous-groupe maximal du scindement JSJ. Il reste donc à identifier les sous-groupes de surfaces suspendus maximaux qui apparaissent comme des groupes de sommets dans le scindement JSJ.

1.8.4 Identifier les surfaces

Nous donnons un critère pour identifier les surfaces dans la configuration de Brady-Meier. Un graphe de sommet d'un graphe tubulaire de graphes est un *graphe de surface* si le groupe fondamental du graphe est un groupe de surface dont les sous-groupes périphériques sont précisément les sous-groupes d'arêtes incidents. Alors

Lemme 1.8.9 (Lemme 5.5.2). *Un graphe de sommet d'un graphe tubulaire de graphes Brady-Meier est un graphe de surface si et seulement si chaque arête de son double est contenue dans exactement deux carrés.*

Armé de ce résultat, nous enlevons les tubes en X'' qui sont attachés aux graphes de surface des deux côtés. Le graphe tubulaire de graphes ainsi obtenu aura alors la décomposition JSJ comme graphe sous-jacent des groupes, prouvant le résultat principal (Théorème 1.8.3).

1.8.5 Scindements relatifs JSJ

Soit \mathcal{H} une famille de sous-groupes d'un groupe G . Un *scindement de G relatif à \mathcal{H}* est une décomposition de G en graphe de groupes dans lequel chaque élément de \mathcal{H}

est elliptique.

Un *scindement JSJ* d'un groupe libre de rang fini F relativement à une famille \mathcal{H} de sous-groupes cycliques maximaux est une décomposition de F en graphe de groupes relatif à \mathcal{H} qui satisfait les conditions d'un scindement JSJ (Définition 1.8.1).

Dans [Cas16, Theorem 4.25], Cashen montre qu'un scindement JSJ de F relativement à \mathcal{H} existe et il est unique.

Comme application de notre algorithme pour construire le scindement JSJ (Théorème 1.8.3), nous obtenons un algorithme pour construire le scindement JSJ relatif:

Théorème 1.8.10 (Théorème 6.0.3). *Il existe un algorithme en temps doublement exponentiel qui prend en entrée un groupe libre de rang fini F et une famille finie de sous-groupes cycliques maximaux \mathcal{H} tels que F est librement indécomposable relativement à \mathcal{H} et qui renvoie le scindement JSJ de F relativement à \mathcal{H} .*

Notre approche consiste à construire un graphe tubulaire de graphes Brady-Meier X tel que F est le groupe fondamental d'un graphe de sommet de X et les groupes d'arêtes incidents sont précisément les membres de \mathcal{H} (voir Chapitre 6 pour plus de détails). La construction assure que $\pi_1(X)$ est hyperbolique. L'algorithme de notre résultat principal (Théorème 1.8.3) construit alors le scindement JSJ de $\pi_1(X)$, et donc le JSJ relatif de (F, \mathcal{H}) .

1.9 Le problème d'isomorphisme

Le problème d'isomorphisme est le problème algorithmique de décider si deux présentations finies de groupes fournissent des groupes isomorphes [Deh11].

Une conséquence importante du travail effectué dans cette thèse est que le problème d'isomorphisme pour les groupes fondamentaux hyperboliques de graphes tubulaires de graphes est réduit à l'algorithme de Whitehead ([Whi36]):

Soit G_1 et G_2 des groupes fondamentaux hyperboliques de deux graphes tubulaires de graphes. En utilisant l'algorithme de Corollaire 1.7.4, nous pouvons construire des

graphes tubulaires de graphes qui donnent des scindements de Grushko de G_1 et G_2 . Par l'unicité du scindement de Grushko, G_1 et G_2 sont isomorphes si et seulement s'il y a une correspondance bijective entre les groupes de sommets de leurs scindements de Grushko avec un isomorphisme entre les groupes de sommets correspondants. Les groupes libres dans les scindements de Grushko ont le même rang si et seulement si les graphes tubulaires de graphes qui représentent ces groupes libres peuvent être modifiés en des graphes ayant le même nombre d'arêtes à l'extérieur d'un arbre maximal. Il reste à vérifier si les facteurs à un bout du scindement de Grushko sont isomorphes.

Supposons donc que G_1 et G_2 sont à un bout. En utilisant Théorème 1.8.3, nous pouvons construire les scindements JSJ de G_1 et G_2 . Alors G_1 et G_2 sont isomorphes si et seulement s'il y a une correspondance bijective entre les groupes de sommets de leurs scindements JSJ de telle sorte que les groupes de sommets correspondants sont isomorphes avec l'isomorphisme respectant les groupes d'arêtes incidents. Mais tous les groupes de sommets sont des groupes libres et leurs groupes d'arêtes incidents sont des sous-groupes cycliques. L'algorithme de Whitehead [Whi36] détermine précisément s'il existe un isomorphisme entre deux groupes libres qui envoie un ensemble fini de sous-groupes cycliques sur un autre.

1.10 Directions de recherche

Nous énumérons quelques directions dans lesquelles le présent travail peut être poursuivi.

1. Une question est de construire un scindement JSJ du groupe fondamental G d'un graphe tubulaire de graphes lorsque G n'est pas hyperbolique. Dans le cas hyperbolique, chaque groupe d'arête est cyclique maximal dans au moins un des groupes de sommets aux extrémités de l'arête. Ce n'est pas vrai quand G n'est pas hyperbolique. Une autre difficulté consiste à décider quand un sous-groupe cyclique se conjugue dans un sous-groupe de surface. Nous pensons que les méthodes utilisées dans cette thèse peuvent être également

- étendues au cas non hyperbolique.
2. Une deuxième question est de savoir s'il est possible de développer un algorithme en temps borné pour construire le scindement JSJ du groupe fondamental hyperbolique d'un graphe fini de groupes libres à groupes d'arêtes cycliques, en l'absence de la structure de graphe tubulaire de graphes.
 3. La plupart de nos résultats utilisent uniquement le fait que \tilde{X} est un complexe cubique CAT(0) Brady-Meier. Il est intéressant de se demander quels résultats subsistent pour les complexes cubiques CAT(0) Brady-Meier généraux avec ou sans la structure \mathcal{VH} . En particulier, la question de savoir quelles classes de groupes à un bout apparaissent en tant que groupes fondamentaux de complexes cubiques Brady-Meier, comme en Théorème 1.7.3, est une question ouverte. Nous avons utilisé des complexes Brady-Meier pour analyser les scindements libres et cycliques dans cette thèse. Il serait intéressant de voir si la structure de Brady-Meier sera utile pour étudier d'autres types de scindements.

Chapter 2

The setup

We will briefly describe our objects of study in this chapter.

2.1 $\mathcal{V}\mathcal{H}$ -complexes

The main objects of our study, tubular graphs of graphs, form a subclass of square complexes known as $\mathcal{V}\mathcal{H}$ -complexes.

The notion of $\mathcal{V}\mathcal{H}$ -complexes was first introduced in [Wis96]. Details can also be found in [Wis06].

Definition 2.1.1. 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 2.1.2 (Vertex links). Let $v \in X$ be a vertex of a square complex. The link of v is a graph whose vertex set is the set $\{e \mid e \text{ is a half-edge incident at } v\}$. The number of edges between two vertices e, f is the number of squares of X in which e, f are adjacent half-edges.

Definition 2.1.3 (CAT(0) square complexes). A square complex is *nonpositively curved* if the length of a closed path in the link of any of its vertices is at least 4. A

nonpositively curved square complex is said to be CAT(0) if it is simply connected.

Definition 2.1.4 ([Wis96]). A $\mathcal{V}\mathcal{H}$ -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 $\mathcal{V}\mathcal{H}$ -complex as horizontal and vertical induces a labelling of vertices in the link of any vertex as horizontal and vertical, thus making the link a bipartite graph.

Remark 2.1.5. Since the link of any vertex of a $\mathcal{V}\mathcal{H}$ -complex is bipartite, the length of a closed path is even. Thus a $\mathcal{V}\mathcal{H}$ -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, first explained by Serre [Ser80].

It was studied from a topological perspective in [SW79] by looking at graphs of spaces instead of graphs of groups. We will adopt this point of view.

Definition 2.2.1. By a *graph of spaces*, we mean the following data: Γ is a connected graph, called the underlying graph. For each vertex s (edge a) of Γ , X_s (X_a) is a topological space. Further, whenever a is incident to s , $\partial_{a,s} : X_a \rightarrow X_s$ is a π_1 -injective continuous map. The *geometric realisation* of the above graph of spaces is the space $X = (\bigsqcup_{s \in \Gamma^{(0)}} X_s \sqcup \bigsqcup_{a \in \Gamma^{(1)}} X_a \times [0, 1]) / \sim$, 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 a tree. Moreover, the underlying tree is the Bass-Serre tree of the associated graph of groups structure of X [SW79].

2.3 Tubular graphs of graphs

Definition 2.3.1. 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, we have

Proposition 2.3.2 ([Wis06]). *The geometric realisation of a tubular graph of graphs is a compact, connected nonpositively curved \mathcal{VH} -complex whose vertical hyperplanes are circles.*

Proof. Indeed, the geometric realisation X is a square complex whose cell structure we describe below.

Every vertex graph X_s inherits the cell structure induced by the graph: its vertices are X_s^0 and its edges are X_s^1 . All edges X_s^1 are *vertical*.

Given an edge graph X_a (homeomorphic to a circle), $X_a \times [0, 1]$ has the product cell structure: its vertices are $X_a^0 \times \{0, 1\}$, its *vertical* edges are $X_a^1 \times \{0, 1\}$ and its *horizontal* edges are $X_a^0 \times [0, 1]$.

Since the underlying graph along with all vertex graphs are assumed to be connected, X is connected. X is compact as the underlying graph and all vertex and edge graphs are finite.

Thus X is a compact and connected \mathcal{VH} -complex. A vertical hyperplane in X is a hyperplane which is dual to a horizontal edge of X . Since the edge spaces are circles, vertical hyperplanes of X are circles.

Finally, X is nonpositively curved. Indeed, since the attaching maps of edge graphs are immersions, two adjacent edges of an edge graph do not have the same image. This removes the possibility of the existence of cycles of length 2 in vertex links. \square

A *tube* in a tubular graph of graphs X is the image of the Cartesian product of an edge graph and the unit interval in the geometric realisation of X .

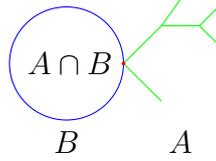


Figure 2.1: An example of a hanging tree in X_s .

Convention. Throughout this text, we will use the same notation for a graph of graphs and the \mathcal{VH} -complex which is its geometric realisation. X will denote a tubular graph of graphs with underlying graph Γ_X . Let $s \in \Gamma_X$ be a vertex. Then X_s will denote the vertex graph at s and if a is an edge of Γ_X , we will denote the edge graph at a by X_a .

Definition 2.3.3 (Thickness). For an edge e in X , the *thickness* of e is the number of squares of X which contain e .

We will define below notions of “hanging trees” and “rudimentary edges” in tubular graphs of graphs. One can always simplify a given tubular graph of graphs by removing hanging trees and rudimentary edges.

Definition 2.3.4. Let X_s be a vertex graph of a tubular graph of graphs X . We say that X_s (and hence X) has a *hanging tree* if X_s is a wedge of two subgraphs A and B such that one of them, say A , is a tree. Here, A is called a hanging (sub)tree of X_s (see [Figure 2.1](#)).

Remark 2.3.5. Since the attaching maps of edge graphs are immersions, we observe that no edge of a hanging tree is in the image of any attaching map. In other words, an edge in a hanging tree of X has thickness 0.

Observe that

Lemma 2.3.6. *A tubular graph of graphs is homotopy equivalent to a tubular graph of graphs with no hanging trees.*

Definition 2.3.7. Let e be an edge in X_s . We call e a *rudimentary edge* if X_s is homeomorphic to S^1 and e has thickness 1.

Lemma 2.3.8. *A tubular graph of graphs is homotopically equivalent to a tubular graph of graphs with no rudimentary edges.*

Proof. Let X be a tubular graph of graphs. If X has no rudimentary edges, we have

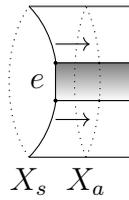


Figure 2.2: Removing rudimentary edges.

nothing to show.

If an edge e in a vertex graph X_s of X is a rudimentary edge, then all edges in X_s are rudimentary edges: First note that X_s is a circle, by definition. Further, the fact that the attaching maps of edge graphs to X_s are graph immersions implies that the thickness of every edge of X_s is one.

Thus there exists exactly one edge a incident to s in the underlying graph Γ_X and the attaching map from X_a to X_s is a graph isomorphism (see Figure 2.2). Then X is homotopic to X' obtained by removing the tube containing X_a . $\Gamma_{X'}$ is the graph obtained from Γ_X by collapsing $a = (s, s')$ to s' . Repeating this procedure at each rudimentary edge gives the result. \square

The following definition, introduced by Brady and Meier (see Theorem 1.2.1 above) is crucial throughout this work.

Definition 2.3.9 ([BM01]). A square complex is said to be *Brady-Meier* if

1. the link of each vertex is connected and
2. for each vertex, the complement of any simplex in the vertex link is connected.

Observe that the Brady-Meier property is local. Thus,

Fact 2.3.10. Given any covering map $X \rightarrow Y$ of square complexes, X is Brady-Meier if and only if Y is Brady-Meier.

Part I

Constructing the Grushko decomposition

Chapter 3

Vertex links and ends

3.1 Ends

The theory of ends of a topological space was first studied by Hans Freudenthal in his thesis [Fre31]. The notion we require for this work is that of “connectivity at infinity”, or “one-endedness”. We will use the following definition due to Specker (see [Spe49] or [Ray60]).

Definition 3.1.1. A locally finite CW complex X is *one-ended* if for every compact set K , $X \setminus K$ has exactly one unbounded component.

It is a well-known fact that being one-ended is a quasi-isometry invariant (see Proposition I.8.29 of [BH99], for instance). Then by an application of the Švarc-Milnor Lemma (see Proposition I.8.19 of [BH99]), for instance), we have the following definition of one-endedness of a finitely presented group.

Proposition 3.1.2. *Let G be a finitely presented group and X be a finite connected CW complex such that $G \cong \pi_1(X)$. G is one-ended if and only if \tilde{X} is one-ended.*

3.2 Not one-ended

In this section, we will collect a few results that lead to a group being not one-ended.

We first recall a useful definition.

Definition 3.2.1. Let Z be a CW complex and $v \in Z$ be a vertex. The set $\text{star}(v)$ is defined as the set of all cells σ such that $v \in \sigma$.

Remark 3.2.2. We note that the open star of v , denoted by $\overset{\circ}{\text{star}}(v)$, is a contractible open neighbourhood of v which is evenly covered in the universal cover of Z .

We first recall a classical result due to Hopf [Hop44]:

Lemma 3.2.3. *Let G be a torsion-free finitely generated group and $G = H * K$ be a nontrivial free splitting of G . Then G is not one-ended.*

Corollary 3.2.4. *Let $G = A *_1$ be an HNN extension of a finitely generated group A over its trivial subgroup. Then G is not one-ended.*

Proof. By definition, an HNN extension of A over the trivial subgroup adds a free generator to A . This implies that $G = A * \mathbb{Z}$. The result follows as \mathbb{Z} is not one-ended. \square

As a consequence, we obtain the following two standard results ([Lemma 3.2.5](#) and [Lemma 3.2.6](#)).

Lemma 3.2.5. *Let Z be a connected CW complex. Let $c \in Z$ be a vertex or the midpoint of an edge. Suppose that $Z \setminus \{c\}$ is not connected, i.e., $Z = Z_1 \vee_{\{c\}} Z_2$, and that neither Z_1 nor Z_2 is simply connected. Then \tilde{Z} is not one-ended.*

Lemma 3.2.6. *Let Z be a connected CW complex. Let $c \in Z$ be a vertex or the midpoint of an edge. Let $G = \pi_1(Z, c)$. Suppose that $\text{star}(c) \setminus \{c\}$ is not connected, but $Z \setminus \{c\}$ is connected. Then \tilde{Z} is not one-ended.*

Proposition 3.2.7. *Let X be a tubular graph of graphs with no hanging trees. Suppose there exists an edge of thickness 0. Then \tilde{X} is not one-ended.*

Proof. Since all horizontal edges have thickness 2, an edge of thickness 0 has to be vertical. Let e in X_s be such an edge. Let c be the midpoint of e .

If $X \setminus \{c\}$ is connected, then by [Lemma 3.2.6](#), \tilde{X} is not one-ended.

If $X \setminus \{c\}$ is not connected, then X is a wedge of two subcomplexes (after subdivision). Let $X = X_1 \vee_c X_2$ and let $X_s = A \vee_c B$ be the induced decomposition of X_s . Since e

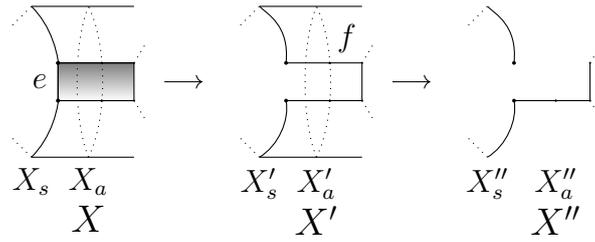


Figure 3.1: Removing squares containing thickness-one edges.

is not in a hanging tree, neither A nor B is a tree. Thus, X_1 and X_2 are not simply connected (because $\pi_1(A) \hookrightarrow \pi_1(X_1)$ and $\pi_1(B) \hookrightarrow \pi_1(X_2)$ in the graph of groups setup [Ser80]). By [Lemma 3.2.5](#), \tilde{X} is not one-ended. \square

Proposition 3.2.8. *Let X be a tubular graph of graphs with no hanging trees and no rudimentary edges. Suppose there exists an edge of thickness 1. Then \tilde{X} is not one-ended.*

Proof. By [Proposition 3.2.7](#), we can first assume that there is no edge of thickness 0. Let e in X_s be an edge of thickness 1. Since e is not a rudimentary edge, first observe that X_s is not a circle. Since there are no edges of thickness 0, there are at least two tubes attached to X_s .

Let Q be the lone square containing e . Then X is homotopic to X' , the complex obtained by removing the open square Q and the open edge e . Let f be an edge adjacent to e in Q . In X' , since the horizontal edge f ([Figure 3.1](#)) is of thickness 1, we can similarly remove the square containing f . We repeat this process consecutively for all squares intersecting vertices of X'_a until we are left with just one vertex, to obtain X'' . X'' has no hanging trees since X had neither hanging trees nor rudimentary edges, but now has an edge of thickness 0. Observe that X is of the same homotopy type as X'' . By [Proposition 3.2.7](#), \tilde{X} is not one-ended. \square

We wish to prove now that if all edges of X have thickness at least two, then \tilde{X} is one-ended whenever every vertex of X has a connected vertex link ([\(BM1\)](#)). In order to do so, we will need the following result.

Lemma 3.2.9. *Let Z be a compact, connected nonpositively curved square complex which has at least one edge. If each edge of Z is contained in at least two squares,*

then $\pi_1(Z)$ is infinite.

We will first recall the definition of hyperplanes. We adopt below the terminology of elementary parallel and parallelism from [Hag07].

Definition 3.2.10 ([Sag95]). Let X be a square complex. Say two edges e and f are *elementary parallel* if they are opposite edges of some square of X . We denote by *parallelism* the equivalence relation on the set of edges of X generated by elementary parallelism.

A *mid-edge* of a square \mathfrak{s} is an edge (after subdivision of \mathfrak{s}) running through the center of \mathfrak{s} and parallel to two of the edges of \mathfrak{s} . Given an equivalence class $[e]$ of parallel edges, the *hyperplane dual to e* , denoted by \mathfrak{h}_e is the collection of mid-edges which intersect edges in $[e]$.

Proof of Lemma 3.2.9. Let e be an edge of Z . The hyperplane \mathfrak{h}_e dual to e is a finite connected graph in which every vertex is of valence at least 2, by the hypothesis on Z . This implies that $\pi_1(\mathfrak{h}_e)$ is a free group of positive rank. To see this, first start from a vertex, say $v_1 \in \mathfrak{h}_e$ (after subdivision). Let $e_1 = (v_1, v_2)$ be an edge at v_1 . Since v_2 is of valence at least 2, let $e_2 = (v_2, v_3)$ be another edge at v_2 . Continuing in this manner, since there are only finitely many vertices, by the pigeon-hole principle, there exists $e_n = (v_n, v)$ where $v = v_i$, for some $i < n$, giving a cycle in \mathfrak{h}_e .

It is a standard result that any lift of \mathfrak{h}_e embeds as a hyperplane in \tilde{Z} , since hyperplanes of CAT(0) square complexes are convex subcomplexes (see Appendix B of [HW08], for example). This implies that $\pi_1(\mathfrak{h}_e) \hookrightarrow \pi_1(Z)$. Hence the result. \square

We are now ready for the case when all edges of X have thickness at least 2.

Proposition 3.2.11. *Let X be a tubular graph of graphs. Let each edge have thickness at least 2. If \tilde{X} is one-ended, then the link of every vertex of X is connected.*

Recall that X satisfies the first Brady-Meier criterion (BM1) if every vertex link is connected.

Proof. Let $u \in X_s$ be a vertex whose link is not connected. This implies that

$\text{star}(u) \setminus \{u\}$ is not connected. If $X \setminus \{u\}$ is connected, then by [Lemma 3.2.6](#), \tilde{X} is not one-ended.

Suppose $X \setminus \{u\}$ is not connected. Then $X = X_1 \vee_u X_2$. By [Lemma 3.2.9](#), both X_1 and X_2 have nontrivial fundamental groups. The result follows from [Lemma 3.2.5](#). \square

Lemma 3.2.12. *If X is not equal to a vertex and satisfies the Brady-Meier criteria ((BM1) and (BM2)), then each edge of X has thickness at least two.*

Proof. Indeed, there cannot be hanging trees as any vertex v of a hanging tree with valence at least two is such that $\text{link}(v)$ is not connected. There cannot be a rudimentary edge because of the following: If $e = (u, v)$ in X_s is a rudimentary edge, then $\eta_e \in \text{link}(u)$, the vertex induced by e , has valence one and the complement in $\text{link}(u)$ of η_x is disconnected, where η_x is the sole adjacent vertex of η_e . Finally, there can be no other edge of valence zero by [Proposition 3.2.7](#) and none of thickness one by [Proposition 3.2.8](#). \square

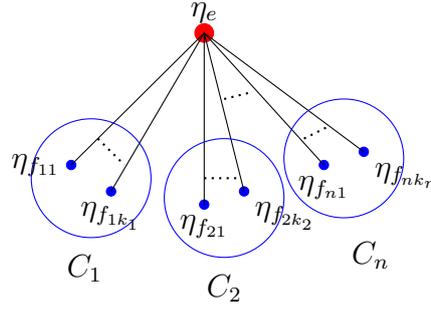
3.3 The second Brady-Meier criterion

In this section, we assume that each edge of X has thickness at least two and every vertex link is connected, but X does not satisfy the second Brady-Meier criterion (BM2). We will explain how to simplify X in this case by defining an opening of the complex at a vertex whose link does not satisfy (BM2).

Let $u \in X_s \subset X$ be a vertex such that for a simplex (vertex or edge) $\sigma \subset \text{link}(u)$, $\text{link}(u) \setminus \sigma$ is not connected.

A vertex of $\text{link}(u)$ is *vertical* (*horizontal*) if it is a vertical (horizontal) half-edge incident to u in X . Observe that the horizontal vertices in the link of any vertex are of valence exactly two. Also, $\text{link}(u)$ is not a segment of length two as each edge of X is contained in at least two squares. This leads us to a lemma:

Lemma 3.3.1. *X does not satisfy (BM2) at u if and only if a vertical vertex of $\text{link}(u)$ disconnects $\text{link}(u)$.*

Figure 3.2: $\text{link}(u)$

Proof. If a vertical vertex of $\text{link}(u)$ disconnects $\text{link}(u)$, then clearly, X does not satisfy (BM2).

For the converse, there are two cases:

Case 1. Suppose a vertex v of $\text{link}(u)$ disconnects $\text{link}(u)$. Then either this vertex is vertical or it is horizontal. If it is vertical, we have nothing to prove. If it is horizontal, let v_1 and v_2 be the two vertical vertices adjacent to v . Then any path between v_1 and v_2 meets v . Let C_1, C_2 be the two components of $\text{link}(u) \setminus \{v\}$ with $v_i \in C_i$. Since $\text{link}(u)$ is not a segment of length 2, one of the components, say C_1 , is not a singleton. Let $x \neq v_1 \in C_1$ be a vertex. Then any path in $\text{link}(u)$ from x to v_2 meets v , and hence meets v_1 . Thus v_1 disconnects $\text{link}(u)$.

Case 2. Suppose an edge e of $\text{link}(u)$ disconnects $\text{link}(u)$. Let v be the unique horizontal vertex incident to e . Then v disconnects $\text{link}(u)$. The proof then follows from the first case.

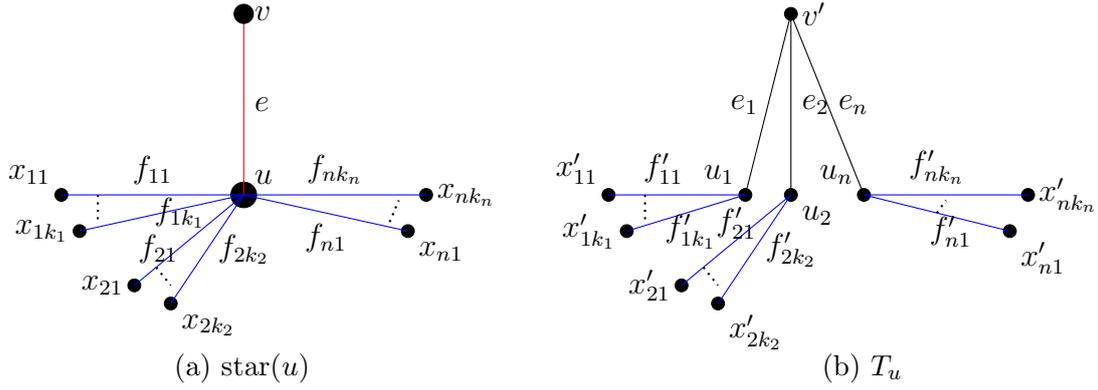
□

Let e be a vertical edge incident to u and let η_e denote the half-edge of e in $\text{link}(u)$. Suppose that η_e disconnects $\text{link}(u)$. Let C_1, \dots, C_n denote maximal connected subgraphs of $\text{link}(u) \setminus \eta_e$ (Figure 3.2), where maximality is by inclusion.

Since X_s is a simplicial graph, the star of u in X_s is a tree (Figure 3.3).

We now explain how to open $\text{star}(u)$:

Definition 3.3.2. We define a tree T_u (Figure 3.3) as follows: There is one ‘primary’ vertex v' out of which emit n edges e_1, \dots, e_n (corresponding to the components


 Figure 3.3: Opening $\text{star}(u)$ to T_u

C_1, \dots, C_n). For each i , we label the other endpoint of e_i as u_i . From each u_i , we have k_i branches to the vertices $x'_{i1}, \dots, x'_{ik_i}$, where k_i is the number of vertices adjacent to η_e in C_i (compare with $\text{star}(u)$). We label these edges as $f'_{i1}, \dots, f'_{ik_i}$.

We also define a new graph X'_s by replacing $\mathring{\text{star}}(u)$ by T_u , with the obvious identifications.

Definition 3.3.3. The graph X'_s is called an *opened-up graph* of X_s .

Clearly, X'_s is connected, $X_s \setminus \mathring{\text{star}}(u) \hookrightarrow X'_s$ and $T_u \hookrightarrow X'_s$. There is a natural surjective map from X'_s to X_s which sends each e_i in T_u to e . Further,

Lemma 3.3.4. *The graphs X'_s and X_s are homotopy equivalent.*

Construction. We now construct a new tubular graph of graphs X' with the same underlying graph Γ_X as X and the only change is that X'_s replaces X_s . An attaching map of an edge graph is unchanged if u is not in the image, as $X_s \setminus \mathring{\text{star}}(u)$ embeds in X'_s . If u is in the image, we do the obvious modification (see Figure 3.4 for an illustration).

Definition 3.3.5. The tubular graph of graphs X' is called an *SL-complex* (simplified link) of X .

There exists a natural map from X' to X . Further,

Proposition 3.3.6. *The tubular graphs of graphs X and X' are homotopy equivalent.*

Lemma 3.3.7. *X' is not isomorphic to X as square complexes.*

Proof. Since the number of edges of X'_s is strictly greater than the number of edges

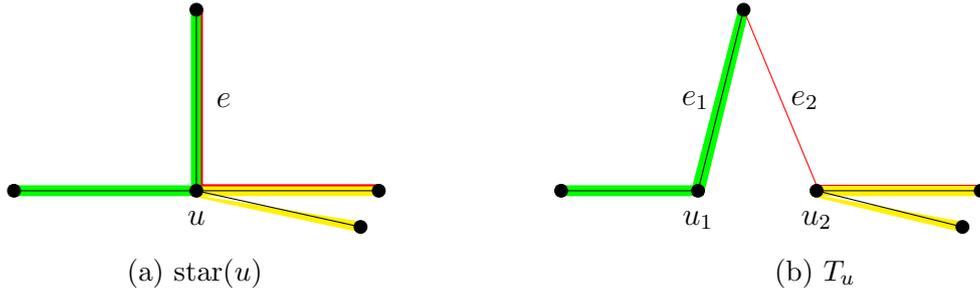


Figure 3.4: A highlighted path of each colour indicates a part of the image of an attaching map

of X_s , we have $X' \not\cong X$. \square

Lemma 3.3.8. *Every edge of X' belongs to at least one square and the number of squares in X' is the same as the number of squares in X .*

Proof. Every horizontal edge of X' belongs to two squares by definition. If a vertical edge belongs to X_φ , with $\varphi \neq s$, then since the attaching maps are unchanged from X , and each edge of X belonged to at least two squares, this vertical edge also belongs to at least two squares. The same argument works for edges in X'_s not in T_u . Since f_{ij} belongs to at least two squares, so does f'_{ij} , by the way the attaching maps are defined. The edge e_i belongs to a square if and only if a pair of adjacent edges $y_{\lambda_1}, y_{\lambda_2}$ in some edge graph is mapped to the pair e, f_{ij} for some j . But since there is an edge between η_e and $\eta_{f_{ij}}$ in $\text{link}(u)$, such a pair exists.

The number of squares in X' is equal to the total number of edges in the cyclic edge graphs, which is equal to the number of squares of X . \square

Remark 3.3.9. It is possible to have edges of thickness one in X' even though there aren't any such edges in X . The only edges that can be of thickness one are the ones that are incident to some u_i . However, these cannot be rudimentary edges as no edge of X_s was of thickness one to begin with.

3.4 The Algorithm

Definition 3.4.1. A tubular graph of graphs is *wedge-like* if either

1. it has no hanging trees or rudimentary edges, but has an edge of thickness zero or one, or
2. every edge is of thickness two but there exists a vertex whose link is not connected.

Remark 3.4.2. By [Proposition 3.2.7](#), [Proposition 3.2.8](#) and [Proposition 3.2.11](#), the fundamental group of a wedge-like tubular graph of graphs is not one-ended.

Theorem A (Main Theorem). *There is an algorithm of polynomial time-complexity that takes a tubular graph of graphs as input and returns a homotopy equivalent tubular graph of graphs which is either Brady-Meier, or is a point, or is wedge-like.*

Proof. We will prove the theorem by constructing the algorithm. Start with a tubular graph of graphs $X = X_0$. Let $k \in \mathbb{N} \cup \{0\}$.

- Step 1 Check if X_k has hanging subtrees. If yes, collapse each hanging subtree to a point and call the new complex as X_k . Go to the next step.
- Step 2 Check if X_k has a rudimentary edge. If yes, remove tubes attached to rudimentary edges as in [Lemma 2.3.8](#) and call the resulting complex also as X_k . Go to the next step.
- Step 3 Check if X_k has at least one square. If yes, go to the next step. Otherwise, X_k is either a point or wedge-like. Stop.
- Step 4 Check if X_k has an edge of thickness zero. If yes, X_k is wedge-like. Stop. If not, go to the next step.
- Step 5 Check if X_k has an edge of thickness one. If yes, stop. If not, go to the next step.
- Step 6 Check if the link of a vertex of X_k is not connected. If yes, stop. If not, go to the next step.
- Step 7 Check if X_k satisfies both [\(BM1\)](#) and [\(BM2\)](#). If yes, stop. If not, go to the next step.
- Step 8 Replace X_k by $X_{k+1} = X'_k$, an SL-complex of X_k , and go to Step 5.

We observe that

- (i) X_k and X_{k+1} are not isomorphic cube complexes (Lemma 3.3.7). Further, for $k \neq k'$, $X_k \not\cong X_{k'}$, as each opening increases the number of edges.
 - (ii) X_k and X_{k+1} have the same number of squares (Lemma 3.3.8).
 - (iii) There is no edge of thickness zero in any X_k for $k \geq 1$ (Lemma 3.3.8).
- (i) implies that the algorithm does not return a tubular graph of graphs from an earlier step. Since there are only a finite number of connected square complexes with a fixed number of squares (ii) and no thickness zero edges (iii), the algorithm cannot proceed indefinitely.

Checking if a graph has hanging trees can be performed in linear time in the number of vertices and edges of X [HT73]. Similarly, checking for edges of thickness zero or one or for rudimentary edges takes linear time in the number of edges and squares of X . Thus steps 1 through 4 run in linear time in the number of vertices, edges and squares of X .

From step 5 onwards, the number of vertices and edges of X_k is bounded by the number of squares of X_k , as each edge is contained in a square. Step 5 runs in linear time. Steps 6 and 7 run in polynomial time in the number of squares: indeed, the size of a vertex link in X_k is bounded by the number of squares of X_k and checking for connectedness and disconnecting vertices is linear (see [HT73] for details) in the number of vertices and edges of the graph.

If n is the number of squares in X , we claim that the number of times the algorithm goes back to step 5 is at most n .

Indeed, the algorithm performs the k^{th} opening-up only if every square of X_{k-1} is of thickness at least two. When each edge is of thickness at least two, the number of vertical edges (as well as horizontal edges) of a tubular graph of graphs can be at most equal to the number of squares.

Observe that the opening procedure in Step 8 increases the number of vertical edges of X_k by at least one, while decreasing the thickness of certain vertical edges. Thus,

the algorithm continues at most until each vertical edge is contained in exactly two squares. In other words, if $N \leq n$ is the number of vertical edges of $X = X_0$, the algorithm stops at most when the number of vertical edges is equal to n . Hence the result. \square

As an immediate consequence, we have:

Corollary B. *There exists an algorithm of polynomial time-complexity that takes a tubular graph of graphs as input and decides whether its fundamental group is one-ended or not. In addition, the algorithm returns the Grushko decomposition of the fundamental group with each free factor itself being the fundamental group of a tubular graph of graphs.*

Proof. Let X be the input tubular graph of graphs with fundamental group G . Apply the algorithm of [Theorem A](#) to X . Let X_N be the output. If X_N is a point, then G is trivial. If X_N is Brady-Meier, G is one-ended and has trivial Grushko decomposition.

If X_N is wedge-like, G is not one-ended ([Remark 3.4.2](#)). In the first step, we remove open edges of thickness one and the open squares that contain them (see [Figure 3.1](#)). We hence assume that no edge of X_N is of thickness one.

Cut X_N along an edge of thickness zero or a locally disconnecting vertex. Either we get a connected tubular graph of graphs X'_1 or we get a disconnected space with components X'_1, \dots, X'_n , where each X'_i is a tubular graph of graphs. In the first case, $G = G_1 * \mathbb{Z}$. In the latter case, $G = G_1 * \dots * G_n$. Apply the algorithm again to each X'_i . If each G_i is one-ended, we are done. Otherwise, cut again at an X'_i with a many-ended G_i and repeat.

This procedure terminates in polynomial time. Indeed, at each step we get tubular graphs of graphs whose total number of squares is bounded by the number of squares of X . \square

Remark 3.4.3. We point out that we do not use Stallings' theorem for our proof. In fact, our procedure yields an alternate proof of Stallings' theorem about ends for

fundamental groups of tubular graphs of graphs. Similarly, we do not assume the existence of the Grushko decomposition either. Our algorithm proves its existence for the groups under consideration.

It is immediate from [Corollary B](#) that

Corollary C. *Let X be a tubular graph of graphs with fundamental group G . Suppose that G admits a free splitting as $G = A * B$. Then there exist tubular graphs of graphs X_1 and X_2 such that A and B are fundamental groups of X_1 and X_2 respectively. Moreover, X_1 and X_2 can be so chosen such that the total number of squares in X_1 and X_2 is bounded by the number of squares in X .*

3.5 Whitehead graphs and separability

In this section we will relate our work to works of Whitehead and Stallings for free groups.

Let F_n be a free group of rank $n \geq 2$ and let W be a finite set of non-trivial elements of F_n . Let B be a basis of F_n .

We will first define Whitehead graphs ([Whi36]). Let H_n denote the orientable 3 dimensional handlebody of genus n . Fix an identification of F_n with the fundamental group of H_n . The basis B corresponds to a system of embedded disks $D = \{d_1, \dots, d_n\}$ such that for an element $b_i \in B$, b_i is represented by a closed path in H_n which starts from the chosen basepoint, intersects d_i transversely and returns to the basepoint without touching any other d_j . Cutting open H_n along these disks results in a 3-ball with $2n$ disks d_i^\pm (such that the chosen representative b_i enters along d_i^+ and leaves along d_i^-).

W is represented by a set of curves in H_n . After cutting, the set of curves is now a set of arcs between these discs.

Definition 3.5.1 ([Whi36]). The *Whitehead graph* $\Gamma_{F_n, B}(W)$ is a graph with $2n$ vertices labelled $\{b_1^\pm, \dots, b_n^\pm\}$. There is an edge between two vertices b_i^+ (respectively, b_i^-) and b_j^+ (b_i^-) for every arc between the corresponding discs d_i^+ (d_i^-) and d_j^+ (d_j^-).

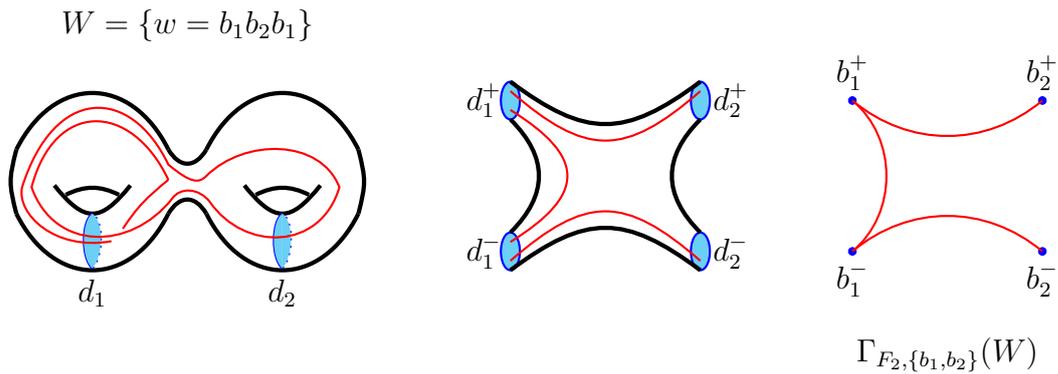


Figure 3.5: A Whitehead graph

in the cut up handlebody.

Figure 3.5 illustrates an example when $n = 2$, $B = \{b_1, b_2\}$ and $W = \{b_1 b_2 b_1\}$.

Definition 3.5.2 (Stallings). W is *separable* if there exists a non-trivial free splitting of $F_n = H * K$ such that each element of W is either a conjugate of an element of H or a conjugate of an element of K .

We recall another definition.

Definition 3.5.3. Let Y be a topological space. A *cut point* $y \in Y$ is a point such that $Y \setminus \{y\}$ is not connected.

Remark 3.5.4. A disconnected graph with at least three vertices necessarily has a cut vertex.

There is a well-known result about the separability of W .

Theorem 3.5.5 ([Whi36]). *If W is separable, then $\Gamma_{F_n, B}(W)$ has a cut vertex for any basis B .*

More details can be found in [Sta99]. The goal of this section is to construct an algorithm to detect separability (Corollary D). Stallings constructs one such algorithm by choosing a Whitehead automorphism whenever there is a cut vertex in a Whitehead graph. Our strategy is to use the machinery of Theorem A when a Whitehead graph contains a cut vertex. We will do so using a construction called ‘double’ in the literature. We will first re-prove Theorem 3.5.5 above using this construction.

3.5.1 Construction of the double

Let R_n denote an oriented rose with petals $\{a_1, \dots, a_n\}$. Fix an identification of F_n with the fundamental group of R_n such that each petal of R_n in the positive direction represents a distinct element of the basis $B = \{b_1, \dots, b_n\}$.

The double $X_{F,W}$ is a tubular graph of graphs constructed in the following way. The underlying graph Γ_X of the double $X_{F,W}$ is a (multi-)graph with two vertices s_1, s_2 and k edges between them, where k is the cardinality of $W = \{w_1, \dots, w_k\}$. The vertex graph X_{s_i} is a subdivided copy of R_n , subdivided as many times as necessary to make all attaching maps simplicial, for $i = 1, 2$. The attaching map of the edge graph on both sides at the j^{th} edge is a simplicial immersion which induces the word w_j at the level of fundamental groups.

Definition 3.5.6. A vertex $v_i \in X_{s_i} \subset X_{F,W}$ is said to be *special* if v_i is the descendant of the unique vertex of R_n .

Remark 3.5.7. We will henceforth drop the subscript i for vertex graphs (and special vertices) as our arguments hold true for both vertex graphs (and special vertices) by the symmetry in $X_{F,W}$.

Lemma 3.5.8. *Let $u \in X_s \subset X_{F,W}$ be a vertex and let $v \in X_s$ be the special vertex. Assume that $\text{link}(v)$ is connected with no cut vertex. Then $\text{link}(u)$ is also connected with no cut vertex.*

Proof. If $u = v$, we have nothing to prove.

If $u \neq v$, then u is in the interior of a petal a_i of R_n . Hence, there are exactly two vertical vertices η_{e_1}, η_{e_2} in $\text{link}(u)$. There is one horizontal vertex for each vertex of an incident edge graph whose image is u . Further, there is a path of length 2 between η_{e_1}, η_{e_2} for every horizontal vertex in $\text{link}(u)$. Thus, the number of reduced paths between η_{e_1}, η_{e_2} is equal to the number of horizontal vertices in $\text{link}(u)$.

Suppose that there is no horizontal vertex in $\text{link}(u)$. This means that $\text{link}(u)$ is disconnected. Then $\text{link}(v)$ is also disconnected as the valence of the two vertices of $\text{link}(v)$ corresponding to the petal a_i is zero, which is a contradiction.

Assume now that there is at least one horizontal vertex in $\text{link}(u)$. This implies that η_{e_i} is not a cut vertex. By [Lemma 3.3.1](#), there is a cut vertex in $\text{link}(u)$ only if there is just one horizontal vertex in $\text{link}(u)$ and so η_{e_i} is of valence one. This means that in $\text{link}(v)$, both the vertical vertices corresponding to a_i are also of valence one. Any vertex adjacent to one of these vertices is then a cut vertex, a contradiction. \square

Lemma 3.5.9. *The link of the special vertex v is isomorphic as graphs to the first subdivision of the Whitehead graph $\Gamma_{F_n, B}(W)$.*

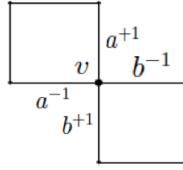
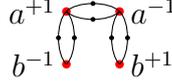
Proof. There are two vertical vertices in $\text{link}(v)$ for each petal a_i of R_n and hence there are $2n$ vertical vertices $\{a_1^\pm, \dots, a_n^\pm\}$, where the signs agree with the fixed orientation of the petal a_i . There is a segment of length 2 between a_i^+ (respectively, a_i^-) and a_j^+ (respectively, a_j^-) in $\text{link}(v)$ exactly when two consecutive edges of an incident edge graph X_a are mapped to the edges induced by a_i^+ (respectively, a_i^-) and a_j^+ (respectively, a_j^-). This corresponds to an occurrence of $a_i^{-1}.a_j$ or $a_j^{-1}.a_i$ (similar strings respectively) in the cyclic word $w_i \in W$. This gives an edge between b_i^+ (respectively, b_i^-) and b_j^+ (respectively, b_j^-) in $\Gamma_{F_n, B}(W)$. The isomorphism is then clear. \square

Proof of [Theorem 3.5.5](#). Let $X_{F, W}$ be the double of (F, W) , with fundamental group G . G is not one-ended as W is separable and hence a vertex group of $X_{F, W}$ splits freely relative to its incident edge groups, giving a free splitting of G . This implies that $\tilde{X}_{F, W}$ is not one-ended. Thus, the link of a vertex of $\tilde{X}_{F, W}$ and therefore of $X_{F, W}$ is either disconnected or has a cut vertex. By [Lemma 3.5.8](#), this implies that $\text{link}(v)$ has a cut vertex. [Lemma 3.5.9](#) now gives the result. \square

Corollary D (Stallings). *There exists an algorithm of polynomial time-complexity that detects the separability of a finite set of words in a finite rank free group.*

We will need a result by Wilton [[Wil12](#), Theorem 18]:

Theorem 3.5.10. *The fundamental group of a graph of free groups with cyclic edge groups is one-ended if and only if every vertex group is freely indecomposable relative to the incident edge groups.*

Figure 3.6: $X_{s_1} (= X_{s_2})$ Figure 3.7: $\text{link}(v)$

Note that a vertex group is freely indecomposable relative to the incident edge groups if and only if the set of words induced by the generators of these edge groups is not separable.

Proof of Corollary D. Let W be the given set of words of the free group F . Let $X_{F,W}$ be the double and G its fundamental group. Apply the algorithm of Corollary B to detect whether G is one-ended. By Theorem 3.5.10 above, G is one-ended if and only if W is not separable. \square

3.6 A counter-example

We now construct a tubular graph of graphs X with no hanging trees or rudimentary edges and such that it satisfies only (BM1) and not (BM2). We will apply our algorithm to X and show that $\pi_1(X)$ is one-ended, resulting in a counter-example to the converse of Theorem 1.2.1.

Let $F_2 = \langle a, b \rangle$ and $W = \{aba^3b\}$. Let X be the double of this data, as constructed in the previous section. We thus have $X_{s_1} = X_{s_2} = R_2''$, the second barycentric subdivision of the bouquet of two circles (Figure 3.6). Then the link of the special vertex $v \in X_s$ is given in Figure 3.7.

We will illustrate the steps of the algorithm of Theorem A by pictures. We obtain a Brady-Meier tubular graph of graphs in the fourth step (Figure 3.11).

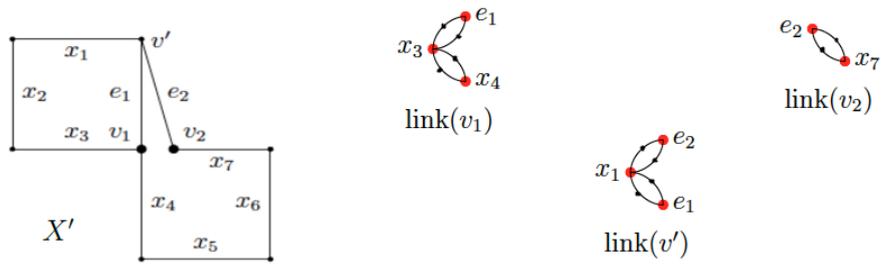


Figure 3.8: Step 1 of the algorithm

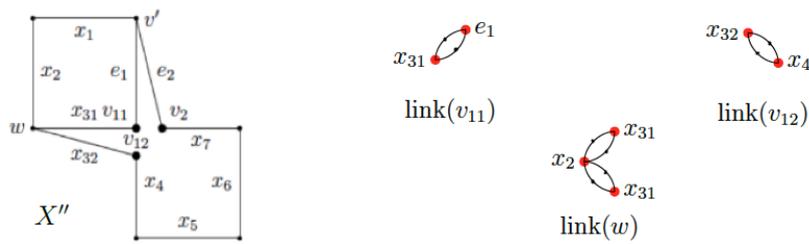


Figure 3.9: Step 2 of the algorithm

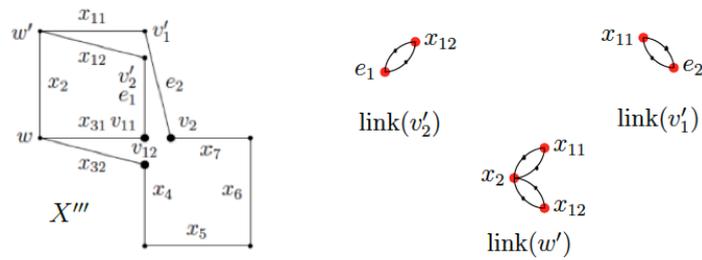


Figure 3.10: Step 3 of the algorithm

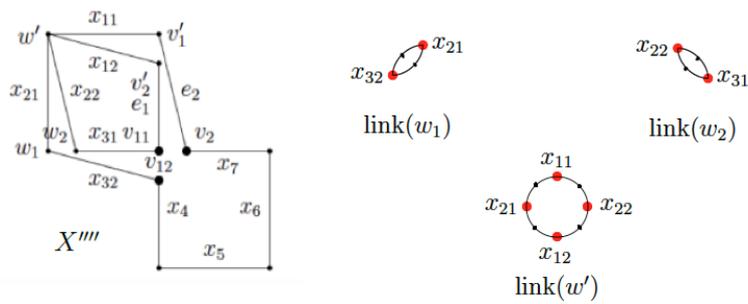


Figure 3.11: Step 4 of the algorithm

Part II

Constructing the JSJ decomposition

Chapter 4

Cyclic splittings and Brady-Meier complexes

The main goal in this chapter is to construct the JSJ decomposition of a one-ended hyperbolic fundamental group of a tubular graph of graphs.

By [Corollary 1.2.3](#), we that the tubular graph of graphs under consideration is Brady-Meier whenever its fundamental group is one-ended.

Convention. Henceforth, X will denote a Brady-Meier tubular graph of graphs, \tilde{X} its CAT(0) universal cover and G its fundamental group. X_s will denote a vertex graph (a component of the vertical 1-skeleton) in X . Unless mentioned otherwise, we work with the CAT(0) metric in \tilde{X} .

Definition 4.0.1 (Paths, lines). Recall that a path in a space Z is a continuous map from a closed interval to Z .

A *combinatorial path* ([MW05]) is a map of graphs $\rho : P \rightarrow \Gamma$, 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 $\rho : P \rightarrow \Gamma$ as the path P .

A combinatorial path is called a *segment* if it is an embedding. Note that any

compact graph homeomorphic to an interval is the image of a segment. We will often call such graphs as segments.

Unless mentioned to the contrary, a path between two vertices of X or \tilde{X} is a combinatorial path, though we will often not mention it explicitly.

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

A *line* is an isometric embedding $\mathbb{R} \hookrightarrow \tilde{X}$ (with the CAT(0) metric), while a ray is an isometric embedding of $[0, \infty)$.

A *combinatorial line* is an isometric embedding of graphs $R \rightarrow \tilde{X}^1$, where R is the real line subdivided at integer intervals. We will only consider combinatorial lines that are also lines. In other words, $R \rightarrow \tilde{X}^1$ is also an isometric embedding in \tilde{X} with the CAT(0) metric.

Since horizontal edges of \tilde{X} are of valence two, vertical hyperplanes of \tilde{X} are lines. Further,

Fact 4.0.2. The first cubical neighbourhood of a vertical hyperplane \mathfrak{h} , or the set of all closed squares of \tilde{X} that meet \mathfrak{h} , is convex ([Sag95]) and hence isometric to a Euclidean strip $[0, 1] \times \mathbb{R}$ with $\mathfrak{h} \cong \{\frac{1}{2}\} \times \mathbb{R}$.

Thus maximal geodesics in such a strip are of the form either $\{t_0\} \times \mathbb{R}$ or segments from $(0, x)$ to $(1, y)$.

We next divide the set of lines into the following three types.

Definition 4.0.3. 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 certain lines can be both vertical and tubular.

We note that a tubular line that is not vertical is disjoint from the vertical 1-skeleton (including vertices) and hits horizontal edges at most at one point, while a transversal line hits at least one vertical hyperplane (in exactly one point).

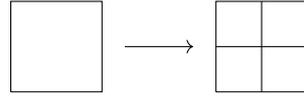


Figure 4.1: Cubical subdivision of a square

4.1 Regular neighbourhoods and regular spheres

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

Definition 4.1.1 (Cubical neighbourhoods). The *first cubical neighbourhood* Y^{+1} of a subspace 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 first cubical neighbourhood of a subcomplex Y is also known as the *star of Y* . We will often use the *open star of Y* , denoted by $\overset{\circ}{\text{star}}(Y)$, to denote the interior of $\text{star}(Y)$.

The n^{th} cubical neighbourhood Y^{+n} is defined inductively as $(Y^{+(n-1)})^{+1}$.

Definition 4.1.2 (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 (see [Figure 4.1](#)) by taking the center of the square as a new vertex and taking four new edges between the center of the square and 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 nonpositively curved 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 .

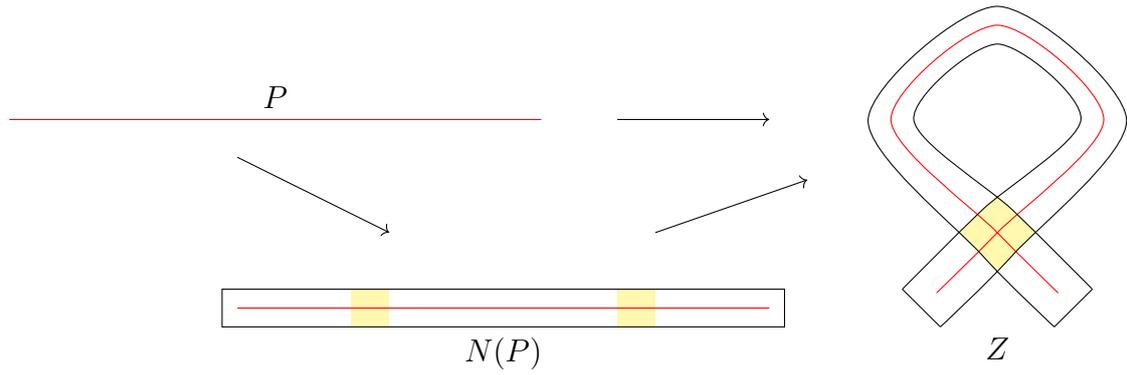


Figure 4.2: Two disjoint subpaths of P are mapped to the yellow square

Definition 4.1.3. 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 $\rho^{-1}(c)$ (see Figure 4.2). The adjacency of cells is given by the adjacency of arcs of P , where each arc is a component of the pre-image of a cell of $Z^{(2)}$.

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

Fact 4.1.4. There is a natural embedding of P in $N(P)$ such that ρ factors through this embedding.

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

Thus, if ρ is an embedding, then $N(P)$ embeds in Z , since $\rho^{-1}(c)$ 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 that we would like P and $N(P)$ to have the same topology, as illustrated in Figure 4.3.

Definition 4.1.5. 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 4.1.6. The regular sphere around a vertex is homeomorphic to the link of the vertex. The regular sphere is in fact isomorphic as graphs to the first barycentric

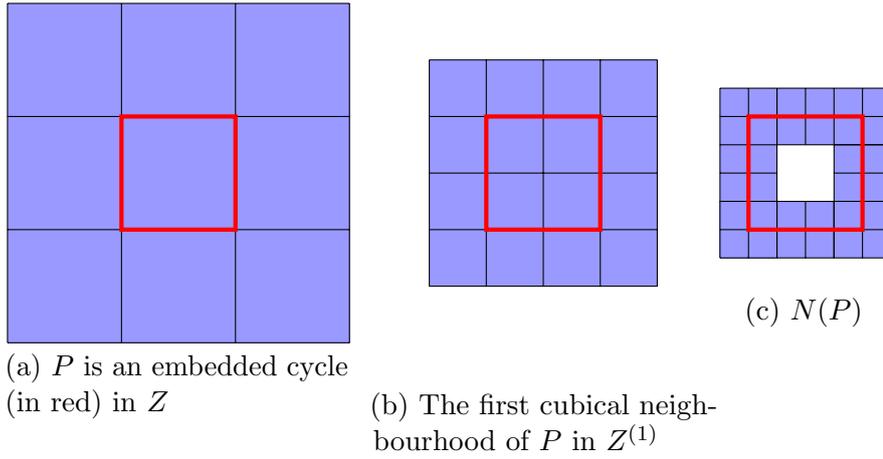


Figure 4.3: P is a deformation retract of its regular neighbourhood

subdivision of the vertex link.

4.1.1 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.

In order to state the result precisely, we will first define a notion of connected sum of graphs.

Definition 4.1.7. 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 $\phi_i : \{1, \dots, k\} \rightarrow \text{adj}(v_i)$ be a labelling of the vertices adjacent to v_i , $i = 1, 2$. Then the *connected sum* $\Gamma_1 \bigoplus_{(v_1, \phi_1), (v_2, \phi_2)} \Gamma_2$ is defined as a quotient of $\Gamma_1 \setminus \mathring{\text{star}}(v_1) \sqcup \Gamma_2 \setminus \mathring{\text{star}}(v_2)$, where $\phi_1(j)$ is glued to $\phi_2(j)$, for $1 \leq j \leq k$. Recall that $\mathring{\text{star}}(v_1)$ refers to the open star of v_1 . If $v_1 \neq v_2$ are vertices in Γ_1 as above, then we define the *self-connected sum* $\bigoplus_{(v_1, \phi_1), (v_2, \phi_2)} \Gamma_1$ as a quotient of $\Gamma_1 \setminus (\mathring{\text{star}}(v_1) \cup \mathring{\text{star}}(v_2))$, where $\phi_1(j)$ is glued to $\phi_2(j)$, for $1 \leq j \leq k$.

We also recall the definition of a dipole graph.

Definition 4.1.8. The *dipole graph of order d* is a multigraph consisting of two vertices and d edges joining them.

Let $e = (u, v)$ be an (oriented) edge of a nonpositively curved square complex Z .

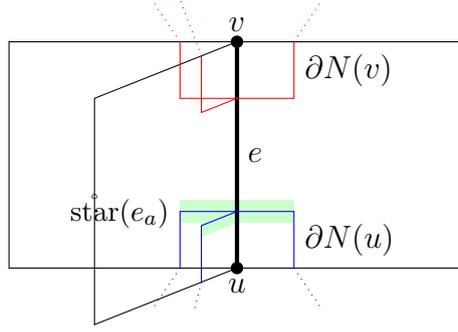


Figure 4.4: Regular spheres around two adjacent vertices. The star of e_a is highlighted in green.

We can now state the main result of this subsection.

Lemma 4.1.9. *The regular sphere around e is homeomorphic to a connected sum of the regular spheres around u and v , with 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, by abuse of notation, as e_a and e_b respectively. 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 (e_b) meets $\partial N(u)$ ($\partial N(v)$) at a vertex of valence d , see Figure 4.4. So $\partial N(u) \setminus \mathring{\text{star}}(e_a)$ ($\partial N(v) \setminus \mathring{\text{star}}(e_b)$) is a graph with d ‘hanging’ edges: edges with one of their endpoints having valence 1.

We thus see that

$$\partial N(e) \cong \partial N(u) \setminus \mathring{\text{star}}(e_a) \sqcup \partial N(m) \setminus (\mathring{\text{star}}(e_a) \cup \mathring{\text{star}}(e_b)) \sqcup \partial N(v) \setminus \mathring{\text{star}}(e_b) / \sim$$

where the gluing is defined on vertices of valence 1. A vertex of valence 1 in $\partial N(u) \setminus \mathring{\text{star}}(e_a)$ is glued to a vertex of valence 1 in $\partial N(m) \setminus \mathring{\text{star}}(e_a)$ if and only if they are contained in a common square. Similarly, a vertex of valence 1 in $\partial N(m) \setminus \mathring{\text{star}}(e_b)$ is glued to a vertex of valence 1 in $\partial N(v) \setminus \mathring{\text{star}}(e_b)$ if and only if

they are contained in a common square.

Observe that the squares containing e induce a labelling ϕ_u of the vertices adjacent to e_a in $\partial N(u)$ and a labelling ϕ_v of vertices adjacent to e_b in $\partial N(v)$. Since $\partial N(m \setminus (\mathring{\text{star}}(e_a) \cup \mathring{\text{star}}(e_b)))$ is a union of segments that join the regular sphere around u and v , $\partial N(e)$ is homeomorphic to $\partial N(u) \underset{(e_a, \phi_u)}{\oplus} \underset{(e_b, \phi_v)}{\oplus} \partial N(v)$. \square

4.1.2 The regular sphere around a combinatorial path

Henceforth, till the end of [Section 4.1](#), Z is either X or \tilde{X} .

Assume that P is not a vertex. Suppose first that P is not a cycle. Let e be an edge in P . Note that P is then a concatenation of paths P_1 , e and P_2 , where P_1 and P_2 are the components of the complement of the open edge \mathring{e} . Similarly, if P is a cycle, we denote the connected complement of \mathring{e} by just P_1 . Let m be the midpoint of e . Then note that

Remark 4.1.10. $N(P) = N(P_1) \cup N(m) \cup N(P_2)$.

Lemma 4.1.11. *The regular sphere around P is homeomorphic to a*

1. *connected sum 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-connected sum of the regular sphere around P_1 (with labelling induced by the squares containing e) if P is a cycle.*

Proof. The proof is analogous to the proof of [Lemma 4.1.9](#). We prove the lemma for the non-cyclic case. The cyclic case is similar.

The initial half-edge e_a of e hits $\partial N(P_1)$ at a vertex (of valence d) whose adjacent vertices have a labelling ϕ_1 induced on them by the squares containing e . Similarly, e_b hits $\partial N(P_2)$ at a vertex (of valence d) with a labelling ϕ_2 induced on its adjacent vertices by the squares containing e . Since $\partial N(m) \setminus (\mathring{\text{star}}(e_a) \cup \mathring{\text{star}}(e_b))$ is a disjoint union of d segments, $\partial N(P) \cong \partial N(P_1) \underset{(e_a, \phi_1)}{\oplus} \underset{(e_b, \phi_2)}{\oplus} \partial N(P_2)$. \square

4.1.3 Connected regular spheres

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

Fact 4.1.12. A graph has no cut points if and only if it has no cut vertices and no open edge separates the graph. In particular, every vertex of a connected graph with no cut points has valence at least two.

Lemma 4.1.13. *The regular sphere around a vertex or midpoint of an edge of Z has no cut points if and only if Z is Brady-Meier.* \square

We now state the main result of the section.

Proposition 4.1.14. *If P is compact, then the regular sphere around P has no cut points.*

The proof requires the following lemma.

Lemma 4.1.15. *Let Γ_1 and Γ_2 be connected graphs with no cut points. Suppose that Γ is the connected sum $\Gamma_1 \bigoplus_{(v_1, \phi_1)(v_2, \phi_2)} \Gamma_2$. Then Γ has no cut points.*

Proof. First observe that $\Gamma_i \setminus \mathring{\text{star}}(v_i)$ is connected by assumption. Recall that $\mathring{\text{star}}(v_i)$ refers to the open star of v_i . Let $v \in \Gamma$. We will show that v is not a cut point. Suppose that $v \in \Gamma_1$.

Let x and y be two points in Γ .

- If $x, y \in \Gamma_2$, then there exists a path in $\Gamma_2 \setminus \mathring{\text{star}}(v_2)$ between them and hence in $\Gamma \setminus \{v\}$.
- Suppose $x \in \Gamma_1$ and $y \in \Gamma_2$. There exists a path in $\Gamma_1 \setminus \{v\}$ from x to a vertex u in Γ_1 adjacent to v_1 . Now u is glued to a vertex in Γ_2 and hence there exists a path from u to y disjoint from v .
- Suppose both x and y belong to Γ_1 . There exists a path in $\Gamma \setminus \{v\}$ between x and y . If this path does not meet v_1 , we are done. Otherwise, there exist two vertices u_1 and u_2 adjacent to v_1 such that the paths from x to u_1 and u_2 to

y are disjoint from both v and v_1 . Now u_1 and u_2 are glued to vertices in Γ_2 and hence there exists a path between them disjoint from v .

□

Proof of Proposition 4.1.14. 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 , with the length of P_i being strictly less than the length of P . By induction, the regular sphere around P_i has no cut points. Lemma 4.1.11 and Lemma 4.1.15 then give the result for P .

□

Lemma 4.1.16. *If P is not a line, then the regular sphere around P is connected.*

Proof. Let p_0 denote the initial point of P . Then P meets the regular sphere around p_0 at a unique point, also denoted by P .

By Lemma 4.1.13, $\partial N(p_0) \setminus \mathring{\text{star}}(P)$ is connected. We will show that given any v of $\partial N(P)$, there exists a path in $\partial N(P)$ from v to $\partial N(p_0) \setminus \mathring{\text{star}}(P)$.

There exists $p \in P$ such that $v \in \partial N(p)$. P is thus a concatenation of two paths P_1 and P_2 , where $P_1 = [p_0, p]$. Note that $v \in \partial N(P_1)$ and further, if we denote the point at which P_2 meets $\partial N(P_1)$ by P_2 , then $v \neq P_2$. By Proposition 4.1.14, $\partial N(P_1) \setminus \mathring{\text{star}}(P_2)$ is connected. This implies that there exists a path from v to $\partial N(p_0) \setminus \mathring{\text{star}}(P)$. But $\partial N(P_1) \setminus \mathring{\text{star}}(P_2)$ embeds in $\partial N(P)$. Hence the result. □

As an immediate corollary, we have

Corollary 4.1.17 (Rays don't separate). *Let γ be a ray in \tilde{X} . Then $\tilde{X} \setminus \gamma$ is connected.*

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

Lemma 4.1.18 (Path-abundance lemma). *Let P be a combinatorial geodesic in \tilde{X} and $x \in \tilde{X} \setminus P$. 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 \tilde{X} by Fact 4.1.4.

Let γ be a path from x to p . Let γ' be the maximal initial subpath of γ such that $\overset{\circ}{\gamma}' \cap P$ 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 4.1.14](#), 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 $\partial N([p, p']) \setminus P_2$. 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 the path in $N(p)$ from h to p . \square

4.2 Separating and coarsely separating lines

In this section, we will define the notion of separating lines and coarsely separating lines in \tilde{X} , examine the relation between separating combinatorial lines and their regular spheres and show that combinatorial lines separate if and only if they coarsely separate.

Definition 4.2.1 (Separation). A subspace Y of a topological space Z is said to *separate* two points z_1 and z_2 in Z if z_1 and z_2 lie in different components of $Z \setminus Y$. Y is said to separate $Y' \subset Z$ if Y separates two points of Y' .

Recall that a line L is an isometric embedding of \mathbb{R} in \tilde{X} .

Definition 4.2.2 (Separating lines). A *separating line* in \tilde{X} is a line that separates \tilde{X} .

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

Definition 4.2.3 (Coarsely separating lines, [Pap12]). We say that L *coarsely separates* \tilde{X} if there exists $R > 0$ such that

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

2. there exist at least two components Y_1 and Y_2 of $\tilde{X} \setminus N_R(L)$ such that Y_i contains points at arbitrarily large distances from L .

Since a line L is an isometric embedding, note that whenever L is combinatorial, the regular sphere around L embeds in \tilde{X} , by [Fact 4.1.4](#).

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

Definition 4.2.4. The regular sphere around a non-vertical tubular line L at distance at most $\frac{1}{2}$ from a vertical hyperplane h in \tilde{X} is defined to be the regular sphere around h in the first cubical subdivision of \tilde{X} .

Henceforth, till the end of this section, a line L will be assumed to be either combinatorial or tubular.

Lemma 4.2.5. *Let L be a combinatorial separating line in \tilde{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 $x, y \in \tilde{X} \setminus L$. Fix $p \in P$. First assume that both x and y lie outside $N(P)$. By [Lemma 4.1.18](#), 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 $x, y \in \tilde{X}$, a contradiction. \square

Definition 4.2.6 (Half-spaces of a line). Let L be a line in \tilde{X} . A *half-space* of L is the closure in \tilde{X} of a component of $\tilde{X} \setminus L$.

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

One important application of the Brady-Meier property of \tilde{X} is the following.

Lemma 4.2.7. *Let Y be a half-space of a line L . Then $L \subset Y$.*

Proof. Fix $y \in Y \setminus L$. Then given any $l \in L$, there exists a path α_l from y to l such that $\alpha_l \cap L = \{l\}$, by [Lemma 4.1.18](#). Hence $l \in Y$. \square

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

Lemma 4.2.8. *There exists a natural map from the set of half-spaces of a line L to the set of components of the regular sphere around L . Further, this map is bijective.*

Proof. The required map is the one that sends a half-space Y of L to $Y \cap \partial N(L)$.

Observe that each component of $\partial N(L)$ lies in a half-space of L . Thus the number of half-spaces of L is at most the number of components of $\partial N(L)$.

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 \tilde{X} is simply connected, and $D \cap \partial N(L)$ gives a path between h_1 and h_2 in $\partial N(L)$. \square

Corollary 4.2.9. *Given an edge e in L , for each component K of $\partial N(L)$, there exists a square \mathfrak{s} containing e such that $\mathfrak{s} \cap \partial N(L) \subset K$.*

Proof. Let Y be the half-space of L corresponding to K , by [Lemma 4.2.8](#).

By [Lemma 4.2.7](#), Y meets e . Let m be the midpoint of e . By [Lemma 4.1.18](#), 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 \mathfrak{s} that contains e and is as required. \square

Fact 4.2.10. It is easy to see that a tubular line L is a separating line. 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 .

We also observe that

Lemma 4.2.11. *A separating line L of \tilde{X} coarsely separates \tilde{X} .*

The proof is immediate from [Lemma 4.2.12](#).

Lemma 4.2.12. *Let Y be a half-space of a separating line L . Then Y contains points at arbitrarily large distances from L .*

Proof. A hyperplane of a CAT(0) cube complex is, after subdivision, a CAT(0) subcomplex [Sag95]. Thus every hyperplane of \tilde{X} is a tree. But since each edge of

\tilde{X} is of thickness at least 2 (Lemma 3.2.12), 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 .

Case 1. L is vertical or tubular. A horizontal hyperplane through S meets L at a single point.

Case 2. L is transversal. The vertical hyperplane through S meets L at a single point.

□

4.2.1 Coarsely separating periodic lines

Definition 4.2.13. An *axis* in \tilde{X} of an element $g \in G$ is a line L in \tilde{X} such that $gL \subset L$ and g moves an element of L by its translation length. A line L in \tilde{X} is said to be *periodic* if it is an axis of some element of G .

Note that given an element $g \in G$, there exists an axis in \tilde{X} of g . We refer the reader to Theorem II.6.8 of [BH99] for details. We first observe that

Lemma 4.2.14. *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

Lemma 4.2.15. *A periodic coarsely separating combinatorial line L of \tilde{X} separates \tilde{X} .*

Recall that \mathfrak{h}^{+1} is the first cubical neighbourhood of \mathfrak{h} . The main ingredient in the proof is the following:

Lemma 4.2.16. *Let Y be a half-space of a periodic combinatorial line L such that for any vertical hyperplane \mathfrak{h} in Y , L is not contained in \mathfrak{h}^{+1} . Then for each $k \in \mathbb{N}$, $Y \setminus L^{+k}$ is connected.*

Proof of Lemma 4.2.15. First assume that L is not contained in \mathfrak{h}^{+1} for any vertical hyperplane \mathfrak{h} . Suppose that L does not separate. Let $Y = \tilde{X}$ be the unique half-space of L . By Lemma 4.2.16, $Y \setminus L^{+k}$ is connected for all k , implying that L does not coarsely separate.

Now suppose that there exists a vertical hyperplane \mathfrak{h} such that $L \subset \mathfrak{h}^{+1}$. By Fact 4.0.2, L is tubular. By Fact 4.2.10, every tubular line is a separating line and there is nothing to prove. \square

The proof of Lemma 4.2.16 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 \mathfrak{h}^{+1} for any vertical hyperplane in Y .

Definition 4.2.17. Let Z be a convex subcomplex of \tilde{X} . A hyperplane \mathfrak{h} is said to be *tangent* to Z if Z is disjoint from \mathfrak{h} but meets \mathfrak{h}^{+1} .

Remark 4.2.18. By Lemma 13.15 of [HW08], L^{+k} is convex for any k .

Fact 4.2.19. As L^{+K} is convex, any element of $L^{+(k+1)}$ is contained either in L^{+k} or in \mathfrak{h}^{+1} for some hyperplane \mathfrak{h} tangent to L^{+k} .

Lemma 4.2.20. *Given a vertical hyperplane \mathfrak{h} in Y and $k \in \mathbb{N}$, $\mathfrak{h} \cap L^{+k}$ is compact.*

Proof. Let T be the underlying tree of the tree of spaces structure of \tilde{X} . By Lemma 4.2.14, the image of L in T is either a point or a line.

Suppose that there exists a vertical hyperplane \mathfrak{h} such that $\mathfrak{h} \cap L^{+k}$ is not compact. As L^{+k} and \mathfrak{h} are both convex, this means that a ray of \mathfrak{h} is in L^{+k} and so a ray of \mathfrak{h} is at finite Hausdorff distance from a ray of L . Since the image of \mathfrak{h} is a point in T , a ray of \mathfrak{h} can be at bounded distance from a ray of L only if the image of L in T is a point. Thus L is vertical. Let α be the path in T between the image of \mathfrak{h} and

the image of L . Let \mathfrak{h}' be the unique vertical hyperplane tangent to L such that its image in T lies in α .

Then $\mathfrak{h}'^{+1} \cap L$ is not compact: This is because the geodesic from any point of \mathfrak{h} to any point of L has to meet \mathfrak{h}' , since α meets the image in T of \mathfrak{h}' . Thus, if $\mathfrak{h}'^{+1} \cap L$ is compact, then the Hausdorff distance between a ray of L and a ray of \mathfrak{h} cannot be finite.

If a ray of L lies in \mathfrak{h}'^{+1} , then $L \subset \mathfrak{h}'^{+1}$: Let h be an element of the stabiliser of \mathfrak{h}' . Then h preserves a ray of L . So if v is a vertex of L at distance $\frac{1}{2}$ from \mathfrak{h}' and g an element of the stabiliser of L , there exist m and n such that $g^m v = h^n v$. Since G acts freely on \tilde{X} , we conclude that $g^m = h^n$. Let $v' \in \mathfrak{h}'$ be such that the distance between v and v' is $\frac{1}{2}$. Then for any integer k , $(h^n)^k v' \in \mathfrak{h}'$ is at distance $\frac{1}{2}$ from $(h^n)^k v$, as G acts by isometries. But $h^n = g^m$ and therefore $(h^n)^k v \in L$. Hence, \mathfrak{h}' and L have Hausdorff distance $\frac{1}{2}$. But this is not possible, and hence $\mathfrak{h} \cap L^{+k}$ is compact. \square

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

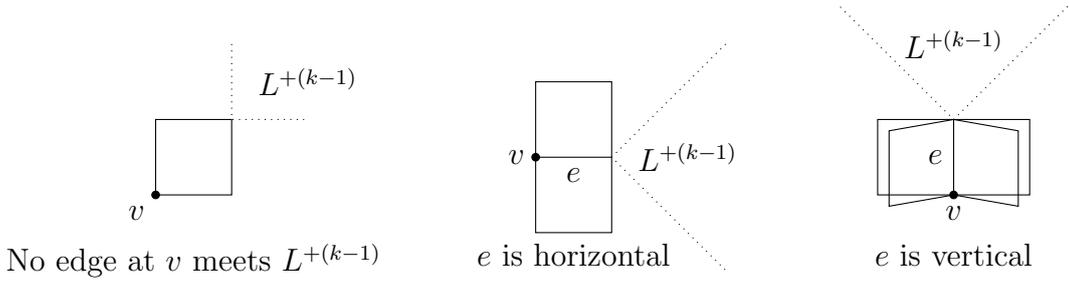
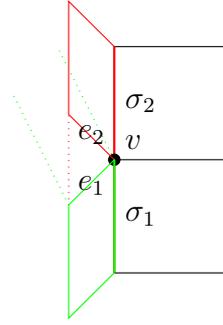
Lemma 4.2.21. *Let v be a vertex in ∂L^{+k} . Then exactly one of the following holds.*

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 $L^{+(k-1)}$. If no such edge exists, then a unique square containing v lies in L^{+k} and (1) holds (Figure 4.5).

Otherwise let e be the unique edge incident to v that meets $L^{+(k-1)}$. Let \mathfrak{s} be a square in L^{+k} that contains v . Then \mathfrak{s} contains e as $L^{+(k-1)}$ is convex.

Suppose that e is horizontal. Then exactly two squares in L^{+k} contain e (Figure 4.5) and (2) holds.

Figure 4.5: The trichotomy when $v \in L^{+k} \setminus L^{+(k-1)}$ Figure 4.6: The edge e_i passes through the vertical hyperplane h_i

If e is vertical, then m squares in L^{+k} contain e , where m is the thickness of e (Figure 4.5). Thus (3) holds. \square

Before proving Lemma 4.2.16, we will have to prove

Lemma 4.2.22. *Let h_1 and h_2 be hyperplanes in Y tangent to L^{+k} . Suppose that $h_1^{+1} \cap L^{+k}$ and $h_2^{+1} \cap L^{+k}$ intersect. Then $h_1^{+1} \cap \partial L^{+(k+1)}$ and $h_2^{+1} \cap \partial L^{+(k+1)}$ lie in a component of $\partial L^{+(k+1)}$.*

We will denote $h_i^{+1} \cap L^{+k}$ by σ_i . Note that σ_i is compact (Lemma 4.2.20) whenever it is vertical (Figure 4.6).

Lemma 4.2.23. *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 $\sigma_1 \cap \sigma_2$. Then either (1) or (2) or (3) of Lemma 4.2.21 holds at $\{v\}$. If (1) or (2) holds, then 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} . Thus $h \neq h_2$. Hence the result. \square

Thus, $\sigma_1 \cap \sigma_2$ is always compact.

Lemma 4.2.24. *Let v be a terminal vertex in $\sigma_1 \cap \sigma_2$. Then either $\sigma_1 \cap \sigma_2 = \{v\}$ or v is also a terminal vertex of σ_1 or σ_2 .*

Proof. If σ_1 is vertical and σ_2 horizontal, then clearly, $\sigma_1 \cap \sigma_2 = \{v\}$.

If σ_1 and σ_2 are horizontal, then by [Lemma 4.2.23](#), $\sigma_1 \cap \sigma_2 = \{v\}$.

Now suppose that σ_1 and σ_2 are vertical. Suppose that their intersection contains an edge. Then either (1) or (2) of [Lemma 4.2.21](#) holds. If (1) holds, then v is terminal in both σ_1 and σ_2 . If (2) holds (see [Figure 4.6](#)), then either no edge or one edge or both edges incident to v lie in σ_1 . If it is the first two, then v is terminal in σ_1 . If it is the last case, then note that both the edges cannot lie in σ_2 as well since v is terminal in $\sigma_1 \cap \sigma_2$. This implies that v is terminal in σ_2 . \square

Proof of [Lemma 4.2.22](#). 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 4.2.21](#).

Case 1. Only one vertical edge f incident to v lies in $L^{+k} \setminus L^{+(k-1)}$. Since \tilde{X} is Brady-Meier, there exists a path β in $\text{link}(v) \setminus \{f\}$ 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 $L^{+k} \setminus L^{+(k-1)}$. Without loss of generality, we assume that either σ_1 is horizontal, or v is terminal in σ_1 , by [Lemma 4.2.24](#). Thus one of the edges, say f_2 does not lie in σ_1 . Let β be a path in $\text{link}(v) \setminus \{f_1\}$ between e_1 and e_2 . Since f_2 does not lie in σ_1 , β and its projection to $\partial\{v\}^{+1}$ is 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 and so on, until the path no longer meets σ_2 . Since σ_2 is compact, the procedure stops in a finite number of steps. Hence the result.

Case 3. Only horizontal edges incident to v lie in $L^{+k} \setminus L^{+(k-1)}$. Let e be the vertical edge incident to v and contained in L^{+k} . Let β be a path between e_1 and e_2 in

$\text{link}(v) \setminus \{e\}$. Then β is disjoint from L^{+k} and so is its projection to $\partial\{v\}^{+1}$.

□

We are now ready to prove [Lemma 4.2.16](#).

Proof of Lemma 4.2.16. 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 4.2.8](#). 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 4.2.19](#). Thus given two vertices u and u' in $Y \cap \partial L^{+(k+1)}$, there exist hyperplanes \mathfrak{h} and \mathfrak{h}' tangent to L^{+k} such that $u \in \mathfrak{h}^{+1}$ and $u' \in \mathfrak{h}'^{+1}$. Let $\sigma = \mathfrak{h}^{+1} \cap L^{+k}$ and $\sigma' = \mathfrak{h}'^{+1} \cap L^{+k}$. 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 $\mathfrak{h} = \mathfrak{h}_1, \dots, \mathfrak{h}_n = \mathfrak{h}'$ such that if $\sigma_i = \mathfrak{h}_i^{+1} \cap L^{+k}$, then $\sigma_i \cap \sigma_{i+1}$ is nonempty. [Lemma 4.2.22](#) then implies that u and u' lie in a component of ∂L^{k+1} . Hence the result. □

4.3 A crossing criterion for lines

Definition 4.3.1 (Crossing of lines). Let L and L' be two separating lines of \tilde{X} . We say that L *crosses* L' if for every half-space Y' of L' , $L \not\subseteq Y'$. 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 a criterion for the crossing of two lines. The criterion will show that while it is necessary for two crossing lines to intersect, it is not sufficient.

Proposition 4.3.2 (Crossing criterion). *Let L and L' be two separating lines in \tilde{X} . L crosses L' if and only if*

1. $L \cap L' = P$ is non-empty and compact and
2. $L' \cap \partial N(P)$ separates $L \cap \partial N(P)$.

Throughout this section, L and L' are two separating lines and P denotes their intersection.

Lemma 4.3.3. *If P is either empty or non-compact, then L and L' don't cross.*

Proof. If P is empty or equal to L , then the result is obvious. So suppose that P is a ray. Then P is contained in all half-spaces of both L and L' , by [Lemma 4.2.7](#). The complementary sub-ray of P in L is contained in a single half-space of L' (as it is connected and disjoint from $L' \setminus P$) and vice-versa. Hence the result. \square

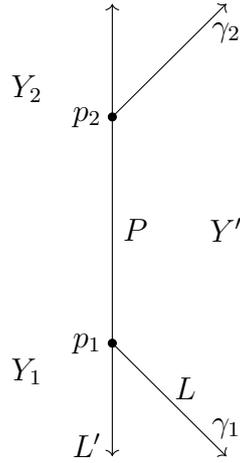
Lemma 4.3.4 (Crossing is symmetric). *L is contained in a half-space of L' if and only if L' is contained in a half-space of L .*

Proof. If P is empty or non-compact, then by [Lemma 4.3.3](#), the result is always true. This leaves us with the case when P is compact.

Assume that L is contained in a half-space Y' of L' ([Figure 4.7](#)) while $L' \subset Y_1 \cup Y_2$, where Y_1 and Y_2 are half-spaces of L . Assume that $Y_1 \neq Y_2$. We will then show that $Y' = \tilde{X}$, a contradiction to the fact that L' separates \tilde{X} .

Let Y be a half-space of L that is not equal to Y_1 or Y_2 . Then $Y \subset Y'$. Indeed, [Lemma 4.1.18](#) gives a path from any point in $Y \setminus L$ to any point in $L \setminus L'$ that is disjoint from L' as the path is contained in the interior of Y . Since $L \subset Y'$, we have that $Y \subset Y'$. We will now deal with Y_1 and Y_2 .

Let p_1 and p_2 be the endpoints of $P = L \cap L'$. Denote by γ_i the closure of each component of $L \setminus P$ so that the initial point of γ_i is p_i .

Figure 4.7: L and L' meet at P

Let $x \in Y_1 \setminus L$. Suppose $x \notin Y'$. Then every path from x to L hits L' . Let α_1 and α_2 be paths given by the path-abundance lemma (Lemma 4.1.18) from x to p_1 and p_2 respectively such that $\alpha_i \cap L' = \{p_i\}$. This implies that $\alpha_i \cap L = \{p_i\}$.

Observe that every path in $\partial N(p_i) \setminus P$ between α_i and γ_i hits L' . Indeed, if not, then there exists a path from x to L disjoint from L' .

Now either $L' \cap \partial N(p_1)$ or $L' \cap \partial N(p_2)$ is in Y_2 and hence $x \in Y_2$. This means that $Y_1 = Y_2$, which is not possible. Hence, $x \in Y'$.

For the converse, argue as above after reversing the roles of L and L' . □

Lemma 4.3.5. *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 $Y' \subset Y$.*

Proof. We start with the easy direction. Suppose there exist half-spaces Y and Y' such that $Y \subset Y'$, say. Then $L \subset Y'$ and thus by Lemma 4.3.4, L and L' don't cross.

Conversely, suppose L and L' don't cross. Let Y_0 and Y'_0 be such that $L \subset Y'_0$ and $L' \subset Y_0$.

Let $Y \neq Y_0$ be a half-space of L . Then $\overset{\circ}{Y}$ is disjoint from L' and hence contained in

a half-space of L' . Further, since $\partial Y = L \subset Y'_0$, $Y \subset Y'_0$.

Similarly, if $Y' \neq Y'_0$ is a half-space of L' , then $Y' \subset Y_0$. Hence the result. \square

Before we go to the proof of [Proposition 4.3.2](#), 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 4.1.14](#)).

4.3.1 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 \tilde{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 4.2.7](#).

Lemma 4.3.6. *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 . \square

The second result is analogous to [Lemma 4.3.4](#).

Lemma 4.3.7. *Let $\{a, b\}$ and $\{a', b'\}$ be cut pairs in Γ . Then $\{a', b'\}$ separates $\{a, b\}$ if and only if $\{a, b\}$ separates $\{a', b'\}$.*

Proof. If $\{a, b\}$ and $\{a', b'\}$ are not disjoint, then neither pair separates the other pair. So assume that they are disjoint. Suppose that $\{a, b\}$ is contained in a half-space Y' of $\{a', b'\}$ while a' lies in a half-space Y_1 of $\{a, b\}$ and b' lies in a half-space Y_2 . If $Y_1 \neq Y_2$, we will show that $\{a', b'\}$ is not a cut pair.

Let $x \in \Gamma \setminus \{a', b'\}$. If $x \notin Y_1, Y_2$, then there exists a path from x to a disjoint from a' and b' and hence $x \in Y'$. Let $x \in Y_1$. Since Γ has no cut points, there exists a

path from x to a in $\Gamma \setminus \{a'\}$. If this path hits b' , then it first hits b as $x \in Y_1$ and $b' \in Y_2$. Hence $x \in Y'$.

Similarly, if $x \in Y_2$, then there exists a path from x to a or b disjoint from a' and b' and hence $x \in Y'$. This implies that $Y' = \Gamma$.

The converse follows by interchanging $\{a, b\}$ and $\{a', b'\}$. \square

Corollary 4.3.8. *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 $\{a, b\}$.*

Proof. Let Y be a half-space of $\{a, b\}$ that contains neither a' nor b' . Then $(\{a, b\} \subset Y)$ is contained in a half-space of $\{a', b'\}$. The result then follows from [Lemma 4.3.7](#). \square

4.3.2 The crossing criterion

We are now ready for the proof of [Proposition 4.3.2](#).

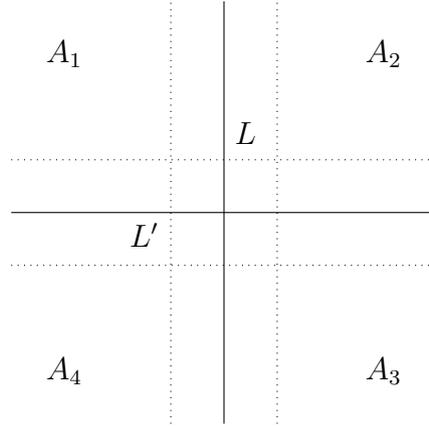
Proof of Proposition 4.3.2. By [Lemma 4.2.5](#), both L and L' separate $\partial N(P)$.

We will start with the ‘if’ direction. Suppose that P is compact and $L' \cap \partial N(P)$ separates $L \cap \partial N(P)$. Then by [Corollary 4.3.8](#), 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 4.2.7](#). 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 4.3.3](#). 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' . \square

4.3.3 Coarse crossings of lines

Definition 4.3.9 (Coarse crossing). Let Z be a CW complex and L, L' two subspaces which coarsely separate Z . L and L' *coarsely cross* if for every large enough number

Figure 4.8: L_1 and L_2 coarsely cross

$R > 0$, there exist connected subspaces A_1, A_2, A_3, A_4 (Figure 4.8) such that

1. for any $R' > 0$ and $i \in \{1, \dots, 4\}$, $A_i \not\subseteq N_{R'}(L)$ and $A_i \not\subseteq N_{R'}(L')$,
2. $N_R(L)$ separates A_1 from A_2 but not A_1 from A_4 ,
3. $N_R(L)$ separates A_3 from A_4 but not A_3 from A_2 , while
4. $N_R(L')$ separates A_1 from A_4 but not A_1 from A_2 and finally,
5. $N_R(L')$ separates A_2 from A_3 but not A_3 from A_1 .

Lemma 4.3.10. *Two separating lines L and L' of \tilde{X} cross if and only if they coarsely cross.*

Proof. Suppose L and L' cross. Let Y_1 and Y_2 be the two half-spaces of L which contain L' and Y'_1 and Y'_2 the half-spaces of L' which contain L . Then L separates $Y_1 \cap Y'_i$ from $Y_2 \cap Y'_i$ while L' separates $Y_i \cap Y'_1$ from $Y_i \cap Y'_2$. Each of the intersections contains points at arbitrarily large distances from both L_1 and L_2 since there exist hyperplanes which intersect L_1 or L_2 at a single point in these subspaces.

Conversely, suppose L and L' cross coarsely. Let R be large enough so that there exist A, B, C and D as in the definition. Since A and D are separated by L' and not L , A and D lie in a half-space of L , say Y_1 . Similarly, B and C lie in a half-space of L , say Y_2 . Note that $L' \cap Y_i$ is non-compact for $i = 1, 2$. If not, then suppose $L' \cap Y_1$ is compact. This implies that $L' \cap Y_1$ is contained in L as L and L' are geodesics. Then L' does not separate A from D as there exists a path between a point on A

and a point on D disjoint from L .

Thus L' hits Y_1 and Y_2 non-compactly. Further, $Y_1 \neq Y_2$ as L separates A from B and $A \subset Y_1, B \subset Y_2$. Thus L and L' cross. \square

4.4 Cyclic splittings and separating lines

Definition 4.4.1. Let H be a subgroup of G . Recall that G *splits* over H if G decomposes as a nontrivial free product with amalgamation over H or as an HNN extension over H .

Recall that a group G acts *minimally* on a tree T if there exists no proper G -invariant subtree of T .

Proposition 4.4.2. 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 G acts minimally on T .*

A decomposition of G either as a free product with amalgamation or as an HNN extension is known as an *elementary splitting* of G .

Proof. One direction is clear. The Bass-Serre tree of the amalgamated product or HNN extension over H satisfies the hypothesis.

For the converse, a quotient graph of groups of the G -tree T has an edge e whose edge group is H . Collapse the (two) component(s) of the complement of the open edge e in the underlying graph to the endpoint(s) of e . The resulting graph of groups is either a free product with amalgamation or an HNN over H . If it is the latter, we are done. If it is the former, say $A *_H B$, then we claim that $H \leq A$ and $H \leq B$.

Suppose not, say $A = H$. Note that $G = B$ in this case. Then in the G -tree T , let \tilde{e} be a lift of e which is stabilised by H . If $\tilde{e} = (v, w)$, then one of them, say v , is stabilised by only H . Let T'_v be the component of $T \setminus \tilde{e}$ containing v . Note that H contains the stabiliser of each vertex and edge of T'_v , as otherwise an element of $B = G$ not in H fixes v . But this implies that G leaves the proper

subtree T'_w invariant, where T'_w is the component of $T \setminus \hat{e}$ containing w . This is a contradiction. \square

4.4.1 Trees dual to separating lines

Let L be a periodic separating line of \tilde{X} such that L does not cross any of its translates. The goal of this subsection is to show that there exists a G -tree in which the stabiliser of a certain vertex is equal to the stabiliser of L . This G -tree will be useful in determining splittings of G by subgroup(s) of the stabiliser of L .

Lemma 4.4.3. *Let L be a periodic line that separates \tilde{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, G acts minimally on T_L .*

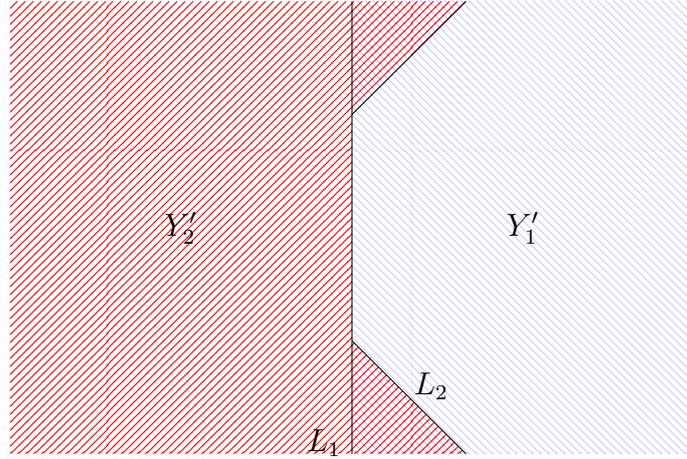
The construction of such a *dual tree* when L has exactly 2 half-spaces and is disjoint from all its translates is standard. In that case, a bipartite graph is constructed as follows: each component of $\tilde{X} \setminus \sqcup_{g \in G} gL$ defines a black vertex while each translate of L defines 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. One can then check that the bipartite graph is in fact a tree.

In our case, L may not be disjoint from its translates. We only ask that L does not cross any of its translates. In addition, L 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 4.4.4. *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 non-empty, 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 $L_2 \subset Y'_1$, see [Figure 4.9](#). We claim that either $Y'_1 = Y_1$ or

Figure 4.9: $L_1 \subset Y'_2$, $L_2 \subset Y'_1$

$Y'_2 = Y_2$. Suppose not.

Since $L_1 \subset Y'_2$, L_1 is disjoint from $\tilde{X} \setminus Y'_2 \supset Y_2 \setminus L_2$. But if L_1 is disjoint from the connected subspace $Y_2 \setminus L_2$, then $Y_2 \setminus L_2$ is contained in a half-space Y''_1 of L_1 . But the fact that the boundary L_2 of $Y_2 \setminus L_2$ is contained in Y'_1 implies that $Y''_1 = Y'_1$ and hence $Y_2 \setminus L_2$ is disjoint from $Y_1 \setminus L_1$, a contradiction. \square

The required tree T_L will be the CAT(0) cube complex dual to a space with walls [HP98]. Recall that

Definition 4.4.5. 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 4.4.6. 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 \tilde{X} do define walls, as we will show below. To avoid confusion, 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 4.4.3](#). We refer the reader to [Nic04] for further details.

Definition 4.4.7. 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 .

Observe that

Lemma 4.4.8. *If ω is an ultrafilter of Z and $A, B \in \omega$, then A and B are not disjoint.* \square

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 4.4.3. Let $Z_L = \tilde{X} \setminus \cup_{g \in G} gL$. Then each half-space Y of gL defines a wall $\{Y \cap Z_L, Y^c \cap Z_L\}$ of Z_L . By abuse of notation, we will denote the wall defined by Y as $\{Y, Y^c\}$. It is easy to see that Z_L is a space with walls. Note that if L has n half-spaces, then there are n distinct walls defined by these half-spaces when $n > 2$.

By theorem 4.1 of [Nic04], 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 section 3 of [Sag95], for instance).

Claim. T_L is a tree.

Proof of claim. If T_L is not a tree, then it is the 1-skeleton of a cube complex of dimension at least 2. This implies that 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 half-space Y of Z_L such that $Y \in \omega_1$ and $Y^c \in \omega_2$. Similarly, there exists a half-space Y' of Z_L such that $Y' \in \omega_1$ and $Y'^c \in \omega_4$. Note that $Y' \in \omega_2$ and $Y \in \omega_4$ as otherwise they would not be adjacent to ω_1 . Further, $Y^c, Y'^c \in \omega_3$ as both ω_2 and ω_4 are adjacent to ω_3 . We will show below that this is not possible.

Now either Y or Y^c is a half-space of a translate of L . Assume without loss of generality that Y and Y' are half-spaces of the lines gL and $g'L$.

By [Lemma 4.4.8](#), Y and Y' are not disjoint. This implies that either $gL \subset Y'$ or $g'L \subset Y$, by [Lemma 4.4.4](#). 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 $Y \not\subset Y'$, 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. \square

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 . Indeed, if Y_1, \dots, Y_n is the set of half-spaces of L , then 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. Indeed, if $Y \in \omega_L$ and $Y \subset Y'$, then $Y' \in \omega_L$ as clearly, $L \subset Y'$. Further, if Y is not a half-space of L , then exactly one of Y and Y^c contains L .

We also claim that ω_L is almost principal. Indeed, choose $y_1 \in Y_1 \cap Z_L$. Then $\sigma_{y_1} = \{Y_1, Y_2^c, \dots, Y_n^c\} \cup \{Y | y_1 \in Y\}$. There exist at most finitely many lines $g_1L \cdots g_kL$ that separate y_1 from L in \tilde{X} . Except for the half-spaces of these lines, a half-space contains y_1 if and only if it contains L . Hence $\sigma_{y_1} \Delta \omega_L$ is finite.

Note that any element in the stabiliser of L permutes $\{Y_1^c, \dots, Y_n^c\}$ and sends a half-space containing L to a half-space containing L . Further, if $g \notin \text{stab}(L)$, then gY_i is a half-space of $gL \neq L$. At least one gY_i contains L , implying that L is not contained in gY_i^c . This implies that $g\omega_L \neq \omega_L$. Hence, $\text{stab}(L) = \text{stab}(\omega_L)$.

Also, T_L is unbounded. Indeed, if not, let ω be a vertex at maximal distance from the vertex ω_L . Then there exists a half-space Y of gL in ω such that Y does not contain any translate of L in its interior. Since G acts geometrically on \tilde{X} , there exists $m > 0$ such that the orbit of $N_m(L)$ covers \tilde{X} . Hence $Y \subset N_m(gL)$, a contradiction to [Lemma 4.2.12](#).

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 Proposition 4.8 of [Nic04]. It thus suffices to prove that no subtree spanned by a proper subset of the set of principal ultrafilters is G -invariant.

As above, let Y_i be a half-space of L . 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 any principal ultrafilter σ_y is a translate of σ_{y_i} , for some i . Indeed, take a path from y to a translate of L such that the interior of the path lies in Z_L . Then the path can be translated to end in L so that σ_y is a translate of σ_{y_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 if even one translate gL of L is contained in the interior of Y_1 as that would mean that σ_{gy_i} for each i is cut off from the subtree. Now, if Y_1 contains no translate of L , then Y_1 is contained in a bounded neighbourhood of L , since G acts geometrically on \tilde{X} . This is a contradiction to [Lemma 4.2.12](#). \square

Proposition 4.4.9. *Let H be a cyclic subgroup of G and L an axis of H in \tilde{X} . Suppose that*

1. L separates \tilde{X} ,
2. L does not cross any of its translates and
3. 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 L satisfy the hypothesis of the Proposition. Using [Lemma 4.4.3](#), we construct a dual tree T_L that has vertices which are stabilised by the stabilisers of translates of L . We will now do a G -equivariant gluing of edges to obtain a tree T that satisfies the conditions of [Proposition 4.4.2](#) for H . This will complete the proof.

Let Y_1, \dots, Y_n be the list of half-spaces of L . Let ω_i be a vertex adjacent to ω_L such that $\omega_i \Delta \omega_L = \{Y_i, Y_i^c\}$. T is a 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. In other words, for each $g \in G$, we identify $g\omega_i$ with $gh\omega_i$.

T_L is connected and hence T is connected. T is a tree since any reduced cycle in T has a reduced cycle as preimage in T_L .

We will now show that T is unbounded. It is enough to show that each vertex of T has valence at least two. Indeed, for a vertex $\bar{\omega}$ such that ω is not a translate of ω_L , $\bar{\omega}$ has valence at least 2 as no pair of vertices adjacent to ω in T_L is identified in T . If ω is a translate of ω_L , then since H does not act transitively on the half-spaces of L , all vertices adjacent to ω are not identified in T .

Assume without loss of generality that the proper subset of half-spaces of condition (3) contains Y_1 . Consider the edge $e = (\omega_L, \omega_1)$ of T_L . Then the stabiliser of \bar{e} in T is clearly equal to H .

There is no proper G -invariant subtree of T . Indeed, since there exists no proper G -invariant subtree of T_L , there exists no proper subtree of T that is G -invariant. \square

Let H be a cyclic subgroup of G over which G splits as in [Definition 4.4.1](#). We call H an *algebraic splitting subgroup*. If H is such that it satisfies the hypothesis of [Proposition 4.4.9](#), we say that H is a *geometric splitting subgroup*. We will now show that every cyclic algebraic splitting subgroup of G is commensurable with a cyclic geometric splitting subgroup.

Lemma 4.4.10. *Let H be a cyclic subgroup over which G splits. Suppose that a vertical line L is an axis of H . Then*

1. L separates \tilde{X} ,
2. L does not cross any of its translates.

Proof. The fact that L coarsely separates \tilde{X} is a consequence of Lemma 1.8 of [Pap05]. We give a proof here for the sake of completeness.

Either G splits as a free product with amalgamation $A *_H B$ or an HNN extension $A *_H$. Then G is isomorphic to the fundamental group of a graph (edge) of spaces Z where the vertex space(s) are presentation complex(es) for A (and B). The unique edge space is a circle. Then the 1-skeleton of \tilde{Z} is a Cayley graph of G . The underlying tree T of \tilde{Z} is the Bass-Serre tree of $A *_H (B)$. Denote by e the edge of

T which is stabilised by H . Note that H coarsely separates \tilde{Z}^1 as e separates T (coarsely).

Let gH be a translate of H . Then H and gH don't cross coarsely. If not, let R be large enough such that there exist subspaces A, B, C and D in \tilde{Z}^1 which satisfy the conditions of [Definition 4.3.9](#). Since each pair A, D and B, C lies in a component of $\tilde{Z}^1 \setminus N_R(H)$, the projection of each pair to T lies on one side of e . Similarly, the projection of each pair A, B and C, D lies on one side of $g \cdot e$. But also, A and B lie on different sides of e . This is not possible.

Fix $x_0 \in L$. Then the map from \tilde{Z}^1 to \tilde{X} which sends g to gx_0 is a quasi-isometry, by the Švarc-Milnor Lemma (Proposition I.8.19 of [BH99]). Note that H is then mapped into L . Thus L coarsely separates \tilde{X} (and hence separates \tilde{X} , by [Lemma 4.2.15](#)). To see this, note that if g_1x_0 and g_2x_0 are two elements chosen in such a way that g_1 and g_2 are in different components of a large neighbourhood of H , then a path α between them in \tilde{X} has a coarse pre-image in Γ and hence hits a bounded neighbourhood of H . This implies that α itself lies in a bounded neighbourhood of L . Similarly, if L and gL coarsely cross, then the pre-images of A, B, C, D have points at arbitrarily large distances from H and gH and appropriate pairs lie in the same coarse components of H and gH . Hence, L and gL don't cross. \square

Proposition 4.4.11. *Let H be a cyclic algebraic splitting subgroup of G with a vertical axis in \tilde{X} . Then there exists a geometric splitting subgroup H' commensurable with H .*

Proof. Let L be an axis of H satisfying [Lemma 4.4.10](#). 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 4.4.9](#), H' is as required. \square

4.5 Vertical cycles and cyclic splittings

In this section, we will examine splittings induced by vertical lines in \tilde{X} .

Recall that a *cycle* (Definition 4.0.1) is an immersion of graphs $\phi : C \rightarrow \Gamma$, 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 4.5.1. The map ϕ is π_1 -injective. Indeed, $\pi_1(C)$ injects into $\pi_1(X_s)$ [Sta83] and $\pi_1(X_s)$ injects into the fundamental group of X in the graph of groups setup [Ser80].

Note that the quotient of a periodic vertical line in \tilde{X} by the action of a nontrivial subgroup of its stabiliser is a cycle in X . The converse is true as well:

Fact 4.5.2. Let $\phi : C \rightarrow X_s$ be a cycle and \tilde{C} denote the universal cover of C . Then \tilde{C} embeds as a vertical line in \tilde{X} .

By abuse of notation, we will often call the lift $\tilde{\phi} : \tilde{C} \rightarrow \tilde{X}$ as the line \tilde{C} .

Lemma 4.5.3. $\partial N(C) \cong \partial N(\tilde{C})/\pi_1(C)$.

Proof. The projection onto C (\tilde{C}) of the regular neighbourhood of C (\tilde{C}) is a deformation retraction. Hence the result. \square

Definition 4.5.4. A *cyclic path* is an immersed combinatorial path $\rho : P \rightarrow X_s$ 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 $\phi_P : C_P \rightarrow X_s$, where C_P is the quotient of P obtained by gluing the initial and terminal vertices and defining $\phi_P([x]) := \rho(x)$.

Definition 4.5.5 (Fundamental domain of a cycle). Let $\phi : C \rightarrow X_s$ be a cycle. A cyclic path $\rho_C : P_C \rightarrow X_s$ with induced quotient cycle C_{P_C} is said to be a *fundamental domain of C* if the following diagram commutes.

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

Remark 4.5.6. It is easy to see that for the action of $\pi_1(C)$ on \tilde{C} , a lift \tilde{P}_C of P_C is a fundamental domain of \tilde{C} in the usual sense.

Definition 4.5.7. 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* is defined as

$$\partial_{\text{orth}}N(P_C) := \overline{\partial N(P_C) \setminus (\{b_u\}^{+2} \cup \{a_v\}^{+2})}$$

where $\overline{\partial N(P_C) \setminus (\{b_u\}^{+2} \cup \{a_v\}^{+2})}$ denotes the closure of $\partial N(P_C) \setminus (\{b_u\}^{+2} \cup \{a_v\}^{+2})$. Recall that $\{b_u\}^{+2}$ denotes the second cubical neighbourhood of $\{b_u\}$.

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

Fact 4.5.8. The natural map from $P_C \cong \tilde{P}_C \hookrightarrow \tilde{C}$ induces an embedding of graphs $\partial_{\text{orth}}N(P_C) \hookrightarrow \partial N(\tilde{P}_C) \setminus \tilde{C} \subset \partial N(\tilde{C})$. Further, $\partial_{\text{orth}}N(P_C)$ is connected if and only if $\partial N(\tilde{P}_C) \setminus \tilde{C}$ is connected as $\partial_{\text{orth}}N(P_C)$ embeds in $\partial N(\tilde{P}_C) \setminus \tilde{C}$ as a deformation retract.

Lemma 4.5.9. *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 $\pi_1(C)$.*

Proof. Recall that $\partial N(C) \cong \partial N(\tilde{C})/\pi_1(C)$ (Lemma 4.5.3). Fix an orientation on P_C . This induces an orientation on both C and \tilde{C} . Let c be a generator of $\pi_1(C)$ that moves an element of \tilde{C} in the positive direction. Then $\tilde{C} \cong \bigsqcup_{n \in \mathbb{Z}} c^n \tilde{P}_C / \sim$, where $\forall n \in \mathbb{Z}, c^n \tilde{u} \sim c^{n-1} \tilde{v}$. Hence

$$\partial N(\tilde{C}) \cong \bigsqcup_{n \in \mathbb{Z}} \partial N(c^n \tilde{P}_C) \setminus \{c^{n+1} \tilde{e}\}^2 \cup \{c^{n-1} \tilde{f}\}^2 / \sim$$

with the following obvious gluing: A vertex (edge) of $\partial N(c^n \tilde{u}) \setminus \{c^{n-1} \tilde{f}\}^2$ is glued to a vertex (edge) of $\partial N(c^{n-1} \tilde{v}) \setminus \{c^n \tilde{e}\}^2$ if and only if they have the same image in \tilde{X} .

Since $c^n \tilde{u}$ is glued to $c^{n-1} \tilde{v}$, $c^{n+1} \tilde{e}$ is identified to $c^n \tilde{e}_v$ and $c^{n-1} \tilde{f}$ is identified to $c^n \tilde{f}_u$. Further, $\partial N(c^n \tilde{P}_C) \setminus \{c^n \tilde{e}_v\}^2 \cup \{c^n \tilde{f}_u\}^2 \cong \partial_{\text{orth}}N(P_C)$. Hence the result. \square

Let e in X be an edge in the image of $\rho_C : P_C \rightarrow X_s$. Let e' be an edge in P_C such that $\rho_C(e') = e$. By Corollary 4.2.9, we have the following result:

Lemma 4.5.10. *Let K be a component of $\partial_{\text{orth}}N(P_C)$. Then there exists a square s in $N(P_C)$ that meets e' and $s \cap \partial_{\text{orth}}N(P_C) \subset K$.*

Definition 4.5.11. A cycle C is a *UC-separating cycle* if \tilde{C} is a separating line.

By [Lemma 4.2.5](#), we have

Lemma 4.5.12. *If C is a UC-separating cycle, then $\partial_{\text{orth}}N(P_C)$ is not connected.*

Definition 4.5.13. A cycle C is *strongly UC-separating* if $\partial N(C)$ is not connected.

Lemma 4.5.14. *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 \tilde{C} .

Proof. Recall that $\partial N(C) \cong \partial N(\tilde{C})/\pi_1(C)$ ([Lemma 4.5.3](#)). A component of $\partial N(C)$ lifts to a component of $\partial N(\tilde{C})$. So if $\partial N(C)$ is connected, then either $\partial N(\tilde{C})$ is itself connected or every component of $\partial N(\tilde{C})$ projects onto $\partial N(C)$. So $\pi_1(C)$ acts transitively on the components of $\partial N(\tilde{C})$ and therefore on the components of $\tilde{X} \setminus \tilde{C}$. The converse is clear. \square

Definition 4.5.15. A cycle $\phi' : C' \rightarrow X_s$ is said to be an n^{th} -power of the cycle $\phi : C \rightarrow X_s$ if there exists an n -fold covering map $\psi : C' \rightarrow C$ such that the following diagram commutes.

$$\begin{array}{ccc} C' & \xrightarrow{\phi'} & X_s \\ \downarrow \psi & \searrow \phi & \\ C & \xrightarrow{\phi} & X_s \end{array}$$

Lemma 4.5.16. *Let N be the maximal thickness of an edge of X . 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 4.2.9](#), the number of half-spaces of \tilde{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 \tilde{C} . The required cycle C' is the quotient of \tilde{C} by H . \square

Definition 4.5.17. Let $\rho_C : P_C \rightarrow X_s$ be a fundamental domain of a cycle C . A *subcycle* of C is the quotient cycle of a cyclic path $\rho_C|_P : P \rightarrow X_s$ with $P \subset P_C$.

Observe that if C' is an n^{th} -power of C , then C is a subcycle of C' .

Definition 4.5.18. A UC-separating cycle C has a *self-crossing* if there exists $g \in G$ such that \tilde{C} and $g \cdot \tilde{C}$ cross.

Definition 4.5.19. We say that C is a *splitting cycle* if the following conditions are satisfied:

1. C is a strongly UC-separating cycle,
2. $\pi_1(C)$ is equal to the stabiliser of a proper subset of the set of half-spaces of \tilde{C} , and
3. C has no self-crossings.

Remark 4.5.20. By [Proposition 4.4.9](#), 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 4.5.21. *Let L_1 and L_2 be vertical lines of \tilde{X} stabilised by the cyclic subgroups H_1 and H_2 respectively. Let n_i be the translation length of a generator of H_i . If $P = L_1 \cap L_2$ is compact, then the length of P is strictly less than $\text{LCM}(n_1, n_2)$.*

Proof. Suppose that P contains a segment of length $\text{LCM}(n_1, n_2) = k$. 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 $h_1^{k/n_1}(v) = h_2^{k/n_2}(v)$. Since G acts freely on \tilde{X} , $h_1^{k/n_1} = h_2^{k/n_2}$ and hence $L_1 = L_2$, a contradiction. \square

Corollary 4.5.22. *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 $g \cdot \tilde{C} \neq \tilde{C}$, then $P = g \cdot \tilde{C} \cap \tilde{C}$ embeds in C .

Definition 4.5.23. A segment $P \subsetneq C$ is said to be a *component of self-intersection* of C if there exists a translate $g\tilde{C} \neq \tilde{C}$ such that the projection to C of $\tilde{C} \cap g\tilde{C}$ 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 $\tilde{C} \cap g\tilde{C} = P$ and \tilde{C} and $g\tilde{C}$ cross.

Fact 4.5.24. Let $P \subset C$ be a segment so that the lift of P in \tilde{C} is isomorphic to P (and hence also denoted by P). Then $\partial N(P) \cap \partial N(C) \simeq \partial N(P) \setminus \tilde{C}$. In other words, there is a self-crossing at P only if $g\tilde{C}$ meets different components of $\partial N(P) \cap \partial N(C)$, by [Proposition 4.3.2](#).

We thus have

Lemma 4.5.25. *A splitting cycle is strongly UC-separating and has no self-crossing at any component of self-intersection.* \square

Also, splitting cycles capture all ‘vertical’ splittings upto commensurability:

Lemma 4.5.26. *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 4.4.11](#), 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' . \square

As an immediate consequence, we have

Lemma 4.5.27. *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.* \square

4.6 Universally elliptic splittings

Recall that a subgroup H of G is said to be elliptic in a G -tree T if H fixes a point in T .

Lemma 4.6.1 (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 \tilde{X} . H_1 is elliptic in the Bass-Serre tree of the elementary splitting over H_2 if and only if given any translate gL_2 of L_2 , L_1 and gL_2 don't cross.*

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 4.4.3) of L_2 contains H_1 and hence H_1 is elliptic in T_{L_2} . The Bass-Serre tree T_2 of the elementary splitting over H_2 is obtained from T_{L_2} by a sequence of G -equivariant gluings of edges of T_{L_2} . Thus elliptic elements remain elliptic, implying that H_1 is 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 . \square

Remark 4.6.2. Since G is one-ended, H_1 is elliptic in the Bass-Serre tree of the elementary splitting over H_2 if and only if H_2 is elliptic in the Bass-Serre tree of the elementary splitting over H_1 , see Theorem 2.1 of [RS97].

Definition 4.6.3 ([GL16]). 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 4.6.4. *Let H be a cyclic subgroup over which G splits. Suppose that an axis of H is transversal in \tilde{X} . Then H is not universally elliptic.*

Proof. Let L be a transversal axis of H . By definition, there exists a vertical hyperplane \mathfrak{h} such that $L \cap \mathfrak{h}$ is a singleton. Since \mathfrak{h} separates \tilde{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 \mathfrak{h} . Note that the image of e under H then spans a line of T . Hence, H is not elliptic in T . \square

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 4.6.5. *Let L be a line that separates \tilde{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

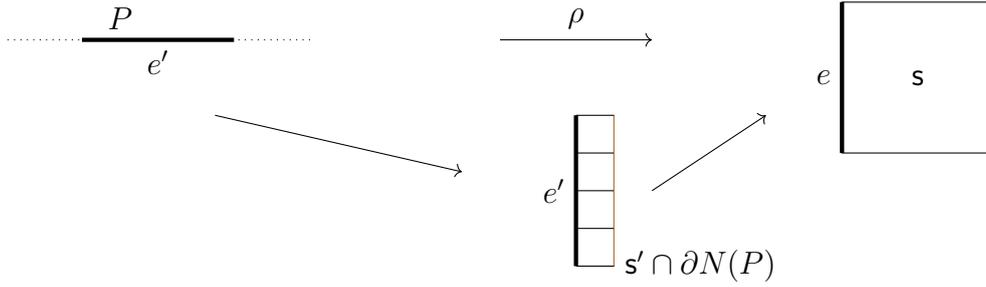


Figure 4.10: A pre-image of a square in the regular sphere

P . Since $\partial N(P) \setminus L$ has at least 3 components, by [Corollary 4.3.8](#), $L' \cap \partial N(P)$ 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 4.4.9](#), G splits over H and H is universally elliptic. \square

4.7 Repetitive cycles

Definition 4.7.1. Let $\rho : P \rightarrow X$ 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 4.10](#)) the union of all squares meeting e' in $N(P)$ whose image in X is contained in s . Then the *pre-image of s around e'* in $\partial N(P)$ is defined as the segment $s' \cap \partial N(P)$.

Recall that the orthogonal sphere of any fundamental domain P_C of a UC-separating cycle C contains at least two components ([Lemma 4.5.12](#)). By [Lemma 4.5.10](#), for each component K of $\partial_{orth}N(P_C)$ and each edge e' in P_C with image e in X , there exists a square s containing e such that the pre-image of s around e' lies in K .

Definition 4.7.2 (Repetitive cycles). Let C be a UC-separating vertical cycle. C is said to be a *k-repetitive cycle* if there exists an oriented vertical edge e in X and an oriented 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 in an orientation preserving way, and

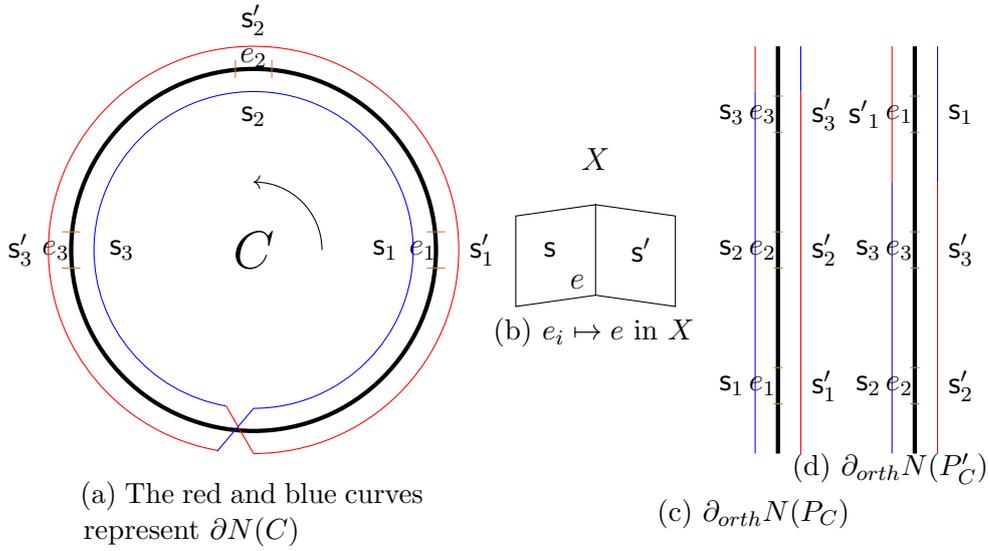


Figure 4.11: C is 3-repetitive with the fundamental domain P_C but not with the fundamental domain P'_C

2. for each square s containing e , there exists a component K of $\partial_{orth}N(P_C)$ such that for each $i \in \{1, \dots, k\}$, the pre-image 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 $\partial_{orth}N(P_C)$. In other words, the notion of repetitiveness requires the cycle to not only ‘repeat’ itself along some edges (Condition 1), but also to ensure that the partitions induced by the cycle on the set of squares containing e_i coincide.

Fact 4.7.3. A k -repetitive cycle is k' -repetitive for $1 \leq k' \leq k$.

Note that the definition depends on a choice of fundamental domain. As the example in Figure 4.11 shows, a cycle C maybe repetitive with respect to one fundamental domain but may not be repetitive with respect to another.

The following property of lifts of repetitive cycles will be crucial for the rest of the text. In fact, this is the only property of repetitive cycles that we will use in the proof of Proposition 4.8.1.

Lemma 4.7.4. *Let C be a k -repetitive cycle. Then there exists an edge \tilde{e} in \tilde{X} and distinct elements $g_1, \dots, g_k \in G$ such that*

1. for each $i \in \{1, \dots, k\}$, $g_i\tilde{C}$ 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{s} and \tilde{s}' that contain \tilde{e} are separated by one of the translates, say $g_1\tilde{C}$, if and only if they are separated by $g_i\tilde{C}$ for all $i \in \{1, \dots, k\}$.

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 4.7.2](#). Let $\tilde{P}_C \subset \tilde{C}$ denote a lift of P_C in \tilde{X} . Note that $\tilde{P}_C \cong P_C$. Denote \tilde{e}_1 in \tilde{P}_C by \tilde{e} .

Since the edges \tilde{e}_i in \tilde{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 $g_i\tilde{C}$ contains \tilde{e} .

Let \tilde{s} and \tilde{s}' be two squares that contain \tilde{e} . Then the squares $g_i^{-1}(\tilde{s})$ and $g_i^{-1}(\tilde{s}')$ in \tilde{P}_C are lifts of squares s_i and s'_i in P_C . Let D and D' be components of $\partial_{\text{orth}}N(P_C)$ such that the pre-image s_1 of s around e_1 meets D and the pre-image s'_1 of s' around e_1 meets D' . By definition, the corresponding pre-image s_i meets D and s'_i meets D' for all i .

Now $\tilde{s} = \tilde{s}_1$ and $\tilde{s}' = \tilde{s}'_1$ lie in different half-spaces of $g_1\tilde{C} = \tilde{C}$ if and only if $D, D' \subset \partial N(\tilde{C})$ (since $\partial_{\text{orth}}N(P_C) \leftrightarrow \partial N(\tilde{C})$, by [Fact 4.5.8](#)) meet different half-spaces of \tilde{C} . Also, \tilde{s} and \tilde{s}' lie in different components of $g_i\tilde{C}$ if and only if $\tilde{s}_i = g_i^{-1}\tilde{s}$ and $\tilde{s}'_i = g_i^{-1}\tilde{s}'$ lie in different half-spaces of \tilde{C} if and only if D and D' induce different half-spaces of \tilde{C} . \square

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 4.7.4](#).

Lemma 4.7.5. *Suppose that at least two translates $g_i\tilde{C}$ and $g_j\tilde{C}$ are distinct. Then \tilde{C} separates \tilde{X} into exactly two half-spaces.*

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

Lemma 4.7.6. *Let Γ be a graph with no cut points. Let $\{a, b\}$ and $\{a', b'\}$ be cut pairs. Suppose there exist points $h_1, h_2, h_3 \in \Gamma \setminus \{a, b, a', b'\}$ such that they are pairwise*

separated by $\{a, b\}$ and also pairwise separated by $\{a', b'\}$. Then $\{a, b\} = \{a', b'\}$.

Compare with Lemma 3.8 of [Bow98].

Proof. Observe that both the pairs $\{a, b\}$ and $\{a', b'\}$ separate Γ into at least 3 components. Thus by Corollary 4.3.8, $\{a, b\}$ lies in a half-space Y' of $\{a', b'\}$. Lemma 4.3.7 then implies that $\{a', b'\}$ lies in a half-space Y of $\{a, b\}$.

Let Y', Y'_1, \dots, Y'_n be the list of 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 $\{a', b'\}$. By assumption, at most one h_i lies in Y' . This implies that the other two lie in Y , a contradiction. \square

Proof of Lemma 4.7.5. After a re-ordering if necessary, we assume that $\tilde{C} = g_1\tilde{C}$ and $g_2\tilde{C}$ are distinct. We will show below that $\partial N(\tilde{C})$ has exactly two components. We then have by Lemma 4.2.8 that \tilde{C} separates \tilde{X} into exactly two half-spaces.

Since $\tilde{C} = g_1\tilde{C} \neq g_2\tilde{C}$, the segment $S = \tilde{C} \cap g_2\tilde{C}$ is compact (Corollary 4.5.22). Suppose that \tilde{C} has at least three half-spaces. Then both $\partial N(S) \setminus \tilde{C}$ and $\partial N(S) \setminus g_2\tilde{C}$ have at least three components, by Lemma 4.1.18. This means that there exist three squares \tilde{s} , \tilde{s}' and \tilde{s}'' containing \tilde{e} that meet different components of both $\partial N(S) \setminus \tilde{C}$ and $\partial N(S) \setminus g_2\tilde{C}$, by Corollary 4.2.9. Then Lemma 4.7.6 implies that $\partial N(S) \cap \tilde{C} = \partial N(S) \cap g_2\tilde{C}$, which is a contradiction. Hence the result. \square

Definition 4.7.7. A vertical cycle $\phi : C \rightarrow X_s$ is said to be *primitive* if it is not a non-trivial power of any cycle. In other words, if there exists a cycle $\phi' : C' \rightarrow X_s$ such that $\phi' = \phi \circ \psi$, then ψ is a homeomorphism of graphs.

Corollary 4.7.8. A lift in \tilde{X} of a primitive k -repetitive cycle separates \tilde{X} into exactly two half-spaces whenever $k \geq 2$.

Proof. When C is primitive, the element g_2 that moves \tilde{e}_2 to \tilde{e}_1 in \tilde{P}_C does not preserve \tilde{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 $g_2 \in \text{stab}(\tilde{C})$, then $\langle g_2 \rangle$ and $\pi_1(C)$ are contained in a common cyclic subgroup and hence $\pi_1(C) \leq \text{stab}(\tilde{C})$, contradicting the fact that C is primitive. Hence $g_2\tilde{C} \neq \tilde{C}$. Lemma 4.7.5 then gives the result. \square

4.7.1 Long cycles are repetitive

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

Proposition 4.7.9 (Long cycles are repetitive). *Let C be a vertical UC-separating cycle with length at least $2E(k-1)2^{F(F+1)/2} + 1$. Then C is k -repetitive.*

Proof. The key ingredient in the proof is the pigeonhole principle. We apply it twice, once to show that the first condition of [Definition 4.7.2](#) is satisfied and the second time to show that the second condition is satisfied. We give the details below.

Let P_C be a fundamental domain of C . Fix an orientation on C . Then each oriented edge of C is mapped to an oriented vertical edge of X . Since there are E vertical edges in X , there are $2E$ vertical oriented edges. Therefore, by the pigeonhole principle, there exists an oriented edge e in X such that n oriented edges e_1, \dots, e_n of C are mapped to e in an orientation preserving way, 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 4.5.10](#). We would like to show that there exist k edges out of e_1, \dots, e_n for which the conditions of [Definition 4.7.2](#) are satisfied. Denote by $A(\lambda, \mu)$ the number of ways in which the squares $\mathfrak{s}_1, \dots, \mathfrak{s}_\lambda$ 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 4.7.2](#) 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 we have

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

We provide a proof of the fact that $A(\lambda, \mu) \leq 2^{\lambda(\lambda+1)/2}$ in [Lemma 4.7.10](#) below. \square

Given natural numbers $\mu \leq \lambda$, let $A(\lambda, \mu) := \sum_{r=1}^{\lambda-\mu+1} \binom{\lambda}{r} A(\lambda - r, \mu - 1)$, whenever $\mu > 1$ and $A(\lambda, 1) = 1$. Then

Lemma 4.7.10. $A(\lambda, \mu) \leq 2^{\lambda(\lambda+1)/2}$.

Proof. We first establish a claim. Let $m_2 \geq m_1 \geq l$ be natural numbers. Then

Claim. $A(m_2, l) \geq A(m_1, l)$.

Proof. If $m_2 = m_1$, there is nothing to show. So assume $m_2 > m_1$. Let $z = m_2 - m_1 \geq 1$.

By definition,

$$A(m_2, l) = \sum_{r=1}^{m_2-l+1} \binom{m_2}{r} A(m_2 - r, l - 1)$$

Splitting the sum, and using the fact that $m_2 - z = m_1$, we have

$$A(m_2, l) = \sum_{r=1}^z \binom{m_2}{r} A(m_2 - r, l - 1) + \sum_{r=1}^{m_1-l+1} \binom{m_2}{z+r} A(m_1 - r, l - 1)$$

Since $\binom{m_1+z}{r+z} \geq \binom{m_1}{r}$, we have

$$\begin{aligned} A(m_2, l) &\geq \sum_{r=1}^z \binom{m_2}{r} A(m_2 - r, l - 1) + \sum_{r=1}^{m_1-l+1} \binom{m_1}{r} A(m_1 - r, l - 1) \\ &\geq \sum_{r=1}^z \binom{m_2}{r} A(m_2 - r, l - 1) + A(m_1, l) \end{aligned}$$

Hence the claim. \square

Thus, $A(\lambda - r, \mu - 1) \leq A(\lambda - 1, \mu - 1)$ for $1 \leq r \leq \lambda - \mu + 1$. Plugging this in the definition of $A(\lambda, \mu)$, we obtain

$$\begin{aligned}
A(\lambda, \mu) &\leq \sum_{r=1}^{\lambda-\mu+1} \binom{\lambda}{r} A(\lambda-1, \mu-1) \\
&= A(\lambda-1, \mu-1) \left(\sum_{r=1}^{\lambda-\mu+1} \binom{\lambda}{r} \right) \\
&\leq A(\lambda-1, \mu-1) \left(\sum_{r=1}^{\lambda} \binom{\lambda}{r} \right) \\
&= A(\lambda-1, \mu-1) 2^\lambda
\end{aligned}$$

Expanding $A(\lambda-1, \mu-1)$ and continuing until $\mu = 1$, we obtain $A(\lambda, \mu) \leq 2^{\lambda(\lambda+1)/2}$. \square

4.8 3-repetitive cycles and crossings

The main result of the section is the following.

Proposition 4.8.1. *Let C be a primitive UC-separating vertical cycle that is 3-repetitive. Then there exists a periodic separating line L' in \tilde{X} such that L' and \tilde{C} cross.*

If C has self-crossings, then by definition, \tilde{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 primitive 3-repetitive cycle C with no self-crossings.

The key idea behind the proof is the following. By [Corollary 4.7.8](#), \tilde{C} separates \tilde{X} into exactly two half-spaces. Further, by [Lemma 4.7.4](#), since C is 3-repetitive, there exists an edge in \tilde{X} along which three translates of \tilde{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 \tilde{C} . We give the details below.

4.8.1 Lifts of C

Fix a fundamental domain P_C of C , an orientation on P_C and edges e_1, e_2, e_3 with image e in X_s satisfying the conditions of [Definition 4.7.2](#). Let d be the thickness of e .

Denote also by e a lift of the edge e in a vertex graph \tilde{X}_s of \tilde{X} such that for $1 \leq i \leq 3$, there exist $L_i = g_i \tilde{C}$ that satisfy the conclusions of [Lemma 4.7.4](#). Recall that this means that L_1, L_2 and L_3 contain e and if $\mathcal{S} = \{s_1, \dots, s_d\}$ is the set of squares containing e , then by conclusion (3) of [Lemma 4.7.4](#),

Lemma 4.8.2. *There exists a partition $A \sqcup B$ of \mathcal{S} such that two squares s and s' in \mathcal{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 .*

Equivalently, s and s' lie in different half-spaces of one of L_1, L_2, L_3 if and only if they lie in different half-spaces of each L_i , if and only if, upto re-ordering, $s \in A$ and $s' \in B$.

By [Corollary 4.7.8](#), we have

Lemma 4.8.3. *Each L_i separates \tilde{X} into exactly two half-spaces.*

Fix a square $s \in A \subset \mathcal{S}$. For $1 \leq i \leq 3$, let Y_i be the half-space of L_i that contains s . Let \bar{Y}_i be its complementary half-space. A set-theoretic consequence of [Lemma 4.8.2](#) is that half-spaces of L_1, L_2 and L_3 are nested. Namely,

Lemma 4.8.4. *For $i, j \in \{1, 2, 3\}$, either $Y_i \subset Y_j$ or $Y_j \subset Y_i$.*

Proof. We will prove that either $Y_1 \subset Y_2$ or $Y_2 \subset Y_1$. The other cases are similar. Note that all the squares in A lie in Y_1 and all the squares in B lie in Y_1' . Similarly, all the squares in A lie in Y_2 and no square in B lies in Y_2 .

We first observe that neither $\bar{Y}_1 \subset Y_2$ nor $Y_2 \subset \bar{Y}_1$, as that would imply that a square in B lies in Y_2 or a square in A lies in \bar{Y}_1 . Thus, by [Lemma 4.3.5](#), either $Y_1 \subset Y_2$ or $Y_2 \subset Y_1$. □

Corollary 4.8.5. *The set $\{Y_1, Y_2, Y_3\}$ is totally ordered by inclusion.*

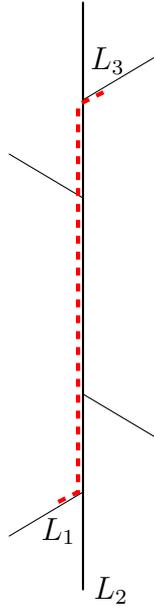


Figure 4.12: L' crosses L_2 if it meets L_1 and L_3 outside L_2

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

Lemma 4.8.6. 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, $\bar{Y}_3 \subset \bar{Y}_2$ (as $Y_2 \subset Y_3$) implies $L_3 \subset \bar{Y}_2$. \square

4.8.2 The main result

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 $h_1(Y'_1) = Y'_1$), 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. given a half-space Y_3 of L_3 , $h_3(Y_3) = Y_3$, 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 \tilde{X}_s that contains e . L' is periodic by definition. Observe that

Lemma 4.8.7. *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 Proposition 24 in §I.6.4 of [Ser80]). 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' . \square

Lemma 4.8.8. *L' is a separating line.*

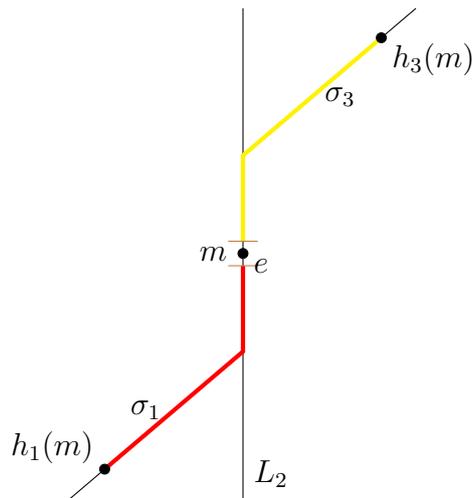
Proof. Let C' be the cycle obtained by taking the quotient of L' by the action of $\langle h' \rangle$. We will show that C' is strongly UC-separating. This will prove that L' is separating (see Lemma 4.5.14).

Let m be the midpoint of e . Subdivide \tilde{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 , 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 4.13) 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 4.1.11, $\partial N(\sigma) \cong \partial N(\sigma_1) \oplus \partial N(\sigma_3)$, where the labelling is 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 4.2.5) and L_i has exactly two half-spaces (Lemma 4.8.3). Thus L_i induces a partition $A_i \sqcup B_i$ on the set of components

Figure 4.13: The segments σ_1 and σ_3

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 4.8.2](#), for each square $s \in \mathcal{S}$, $s \cap \partial N(\sigma_1)$ meets A_1 if and only if $s \cap \partial N(\sigma_3)$ meets A_3 . Therefore, there exists no path between a point in A_1 and a point in B_3 in the connected sum $\partial N(\sigma) \setminus L'$. Hence $\partial N(\sigma)$ is separated by L' . Thus $\partial_{orth} N(\sigma)$ is not connected ([Fact 4.5.8](#)).

As h_1^{-1} preserves half-spaces of L_1 , h_1^{-1} sends a square containing $h_1(e)$ in A_1 (B_1) to a square containing e in A (B). Similarly, h_3 sends a square in A (B) to A_3 (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 $\partial_{orth} N(\sigma)$ by the action of h' . By [Lemma 4.5.9](#), $\partial N(C')$ is not connected. \square

Proof of Proposition 4.8.1. By [Lemma 4.8.8](#), the periodic line L' is a separating line. L' crosses L_2 ([Definition 4.3.1](#)) as L' is not contained in a half-space of L_2 , by [Lemma 4.8.7](#). Hence the result. \square

4.9 An algorithm of double exponential time

For the rest of the text, we will also assume that G is δ -hyperbolic. The main result of this section is the following theorem.

Theorem 4.9.1. *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 upto commensurability.*

Denote by $\partial\tilde{X}$ the Gromov boundary of \tilde{X} . Recall that the Gromov boundary of \tilde{X} is defined to be the set of equivalence classes of geodesic rays in \tilde{X} , where two rays are said to be equivalent if they are asymptotic. We refer the reader to [BH99] for more details.

By Theorem III.H.3.9 of [BH99], $\partial\tilde{X}$ is invariant under quasi-isometries and hence ∂G is well-defined and isomorphic to $\partial\tilde{X}$. Let H be a quasiconvex subgroup of G . By Proposition III.Γ.3.7 of [BH99], H is a hyperbolic group. Further, H is quasi-isometrically embedded in G (Corollary III.Γ.3.6 of [BH99]). Thus, $\partial H \hookrightarrow \partial G$. Note that a cyclic subgroup of G is always quasiconvex (Corollary III.Γ.3.10 of [BH99]). Let $H < G$ be quasiconvex. Then ∂H separates ∂G if and only if H coarsely separates G .

Proposition 4.9.2. *Let C be a primitive 3-repetitive UC-separating cycle in X . Then $\pi_1(C)$ is not universally elliptic.*

Proof. By (Corollary 4.7.8), \tilde{C} separates \tilde{X} into two half-spaces. Lemma 4.2.11 implies that $\pi_1(C)$ separates ∂G into at least two components. Proposition 5.30 of [Bow98] then implies that $\pi_1(C)$ conjugates into either a cyclic vertex group of the JSJ decomposition of G or a hanging surface vertex group.

By Proposition 4.8.1, there exists a periodic line L' such that L' and \tilde{C} cross. This implies that $\pi_1(C)$ conjugates into a hanging surface vertex group in the JSJ decomposition of G and further it is not peripheral in this vertex group (see Lemma 5.21 of [Bow98]). Since universally elliptic splittings are conjugates of exactly the

edge subgroups of the JSJ decomposition upto commensurability, $\pi_1(C)$ is not universally elliptic. \square

Let Γ be a finite simplicial graph with vertex set $V = \{v_1, \dots, v_m\}$. The *adjacency matrix* of Γ is a symmetric square matrix of order $|V|$ such that the $(i, j)^{th}$ entry is one if there is an edge between v_i and v_j , or zero otherwise. It is well known that (see [AYZ97] for instance):

Fact 4.9.3. The number of immersed cycles of length k in Γ is bounded by the trace of the matrix A^k .

Proof of Theorem 4.9.1. Let G and X be given as in the statement. First note that by Lemma 4.6.4, a cyclic subgroup that has a transversal axis in \tilde{X} is not universally elliptic. Thus every universally elliptic subgroup H has a vertical axis in \tilde{X} . This implies that there exists a splitting cycle C in X such that $\pi_1(C)$ and H are commensurable (Lemma 4.5.26).

Let F be the number of squares of X and E the number of edges. By Proposition 4.7.9, any UC-separating cycle of length greater than $M = 4E(2^{F(F+1)/2})$ is 3-repetitive. Since every edge is contained in a square, $E \leq 4F$ and thus M is at most $16F(2^{F(F+1)/2})$. By Proposition 4.9.2, any primitive 3-repetitive cycle is not universally elliptic. Thus a cycle C can be universally elliptic only if

1. its length is at most M , or
2. it is not primitive.

In other words, whenever the length of a universally elliptic cycle C is at least M , then it is a power of a primitive subcycle C' such that C' is not 3-repetitive and is UC-separating. Further, C' has no self-crossings as C has no self-crossings. But C' may or may not be strongly UC-separating (Lemma 4.5.14). By Lemma 4.5.16, an n^{th} power of C' is strongly UC-separating, where n is bounded by the maximal thickness of an edge of C' . This implies that $n \leq F$ and thus a power of C' whose length is at most $F.M$ is universally elliptic. Note that $\pi_1(C)$ is commensurable with $\pi_1(C')$ and thus with every infinite cyclic subgroup of $\pi_1(C')$. Hence any

universally elliptic subgroup is commensurable with the cyclic subgroup generated by a universally elliptic cycle of length at most $F.M$.

There exist finitely many cycles of length at most $F.M$ 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 4.5.27](#), we thus have a list of all universally elliptic cycles upto commensurability.

The time taken by this algorithm is calculated as follows:

1. The number of cycles of length at most $F.M$ is bounded by a number which is exponential in $F.M$, by [Fact 4.9.3](#). 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 connected sum of the regular spheres around its k vertices ([Lemma 4.1.11](#)). 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 [HT73].
3. A cycle C has a self-crossing if there exists self-crossing at a component of self-intersection $P \subset C$ ([Definition 4.5.23](#)). The information about the components of self-intersection of C is readily available with C and does not cost any additional time. 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 4.5.24](#)). This information is also 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 . \square

Chapter 5

Constructing a JSJ complex

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

5.1 Splitting cycles as hyperplanes

Let $\phi : C \rightarrow X_s \subset X$ 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 first perform this construction at the level of universal covers using the machinery of spaces with walls [HP98] (utilised earlier in [Section 4.4.1](#)). We refer the reader to [Nic04] and [CN05] for details on constructing CAT(0) cube complexes from spaces with walls.

First note that the 0-skeleton \tilde{X}^0 is a space with walls, where the walls are defined by the complementary half-spaces of vertical and horizontal hyperplanes. It is well known that the *dual* CAT(0) *cube complex* of \tilde{X}^0 with this wall structure is \tilde{X} (see Theorem 10.3 of [Rol16], for instance). 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 \tilde{C} . Note that there is a natural square structure on S_L so that every horizontal hyperplane of \tilde{X} that meets L naturally extends to S_L . Let Z be the set of open horizontal half-edges of the union of \tilde{X} and the attached strips. Then the vertical and horizontal hyperplanes of \tilde{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 of hyperplanes that define walls in W . Thus the dual cube complex of (Z, W) is nothing but \tilde{X} (Theorem 10.3 of [Rol16]). We now enrich W to W_C . The walls in W_C are determined by the following:

- (i) the horizontal hyperplanes of \tilde{X} ,
- (ii) the vertical hyperplanes of \tilde{X} , and
- (iii) the G -translates of \tilde{C} .

Note that the elements of type (i) and (ii) induce W , where each half-space Y in \tilde{X} of an element of type (i) or (ii) defines a wall $\{Y \cap Z, Y^c \cap Z\}$. Given a translate L of \tilde{C} in \tilde{X} , each half-space Y in \tilde{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 \tilde{X} .

Given $z_1 \neq z_2 \in Z$, observe that there are only finitely many walls in W_C between z_1 and z_2 . Thus

Lemma 5.1.1. (Z, W_C) is a space with walls. □

In fact, two elements z_1 and z_2 are not separated by a wall if and only if either $z_1 = z_2$ or the closures of z_1 and z_2 in \tilde{X} share a vertex and are not separated by any line of type (iii).

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

Lemma 5.1.2. \tilde{X}_C is a \mathcal{VH} -complex.

The proof uses the following observation.

Lemma 5.1.3. Let L and L' be two non-crossing lines of \tilde{X} . Then given half-spaces Y of L and Y' of L' , at least one of the following four intersections is empty: $\overset{\circ}{Y} \cap \overset{\circ}{Y}'$, $Y^c \cap \overset{\circ}{Y}'$, $\overset{\circ}{Y} \cap Y'^c$ and $Y^c \cap Y'^c$.

Proof. Since L and L' don't cross, by [Lemma 4.3.5](#), there exists a half-space Y'_0 of L' such that either $Y \subset Y'_0$ or $Y'_0 \subset Y$. Hence the result. \square

Two walls $\{Y, Y^c\}$ and $\{Y', Y'^c\}$ in a space with walls *cross* ([\[CN05\]](#)) if all four intersections $Y \cap Y'$, $Y^c \cap Y'$, $Y \cap Y'^c$ and $Y^c \cap Y'^c$ are non-empty.

Proof of [Lemma 5.1.2](#). Two walls of type (i) don't cross as two horizontal hyperplanes of \tilde{X} are either equal or disjoint. Similarly, two walls of type (ii) don't cross. By [Lemma 5.1.3](#), 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 Proposition 4.6 of [\[Nic04\]](#), there exists a bijective correspondence between the hyperplanes of \tilde{X}_C and the walls of (Z, W_C) . Further, two hyperplanes in \tilde{X}_C intersect if and only if the corresponding walls cross. Declare an edge e of \tilde{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. \square

Observe that there exists a natural G -equivariant map $\hat{\eta}_C : \tilde{X}_C \rightarrow \tilde{X}$ such that the

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

following diagram commutes:

Since $W \subset W_C$, $\hat{\eta}_C$ takes any vertex (ultrafilter) σ' of \tilde{X}_C to a vertex $\sigma' \cap W$ of \tilde{X} . But every vertex of \tilde{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 also have that

Lemma 5.1.4. *The map $\hat{\eta}_C$ has the following properties:*

1. *Let c be a cell of \tilde{X} that does not meet any translate of \tilde{C} . Then $\hat{\eta}_C$ restricted to $\hat{\eta}_C^{-1}(c)$ 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). \square

Lemma 5.1.5. *For any $z \in Z$, $\hat{\eta}_C^{-1}(\sigma_z)$ is a finite horizontal tree. Further, the edges in the pre-image of σ_z are dual to vertical hyperplanes induced by translates of \tilde{C} that meet σ_z in \tilde{X} .*

Proof. Let σ'_1 and σ'_2 be two vertices of $\hat{\eta}_C^{-1}(\sigma_z)$. 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 \tilde{X} . If not, then let \mathfrak{h} be a hyperplane of \tilde{X} that separates L from σ_z . Let $Y_{\mathfrak{h}}$ be a half-space of \mathfrak{h} that contains the vertex σ_z . Then $Y_{\mathfrak{h}} \in \sigma_z$, the ultrafilter. Clearly, this implies that $Y_{\mathfrak{h}} \in \sigma'_1$ and $Y_{\mathfrak{h}} \in \sigma'_2$. Since L and \mathfrak{h} are disjoint, either $Y_{\mathfrak{h}} \subset Y$ or $Y_{\mathfrak{h}} \subset Y^c$. 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 \tilde{C} that meet at any given point of \tilde{X} . This proves the result. \square

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

Lemma 5.1.6. *G acts geometrically on \tilde{X}_C .* \square

Lemma 5.1.7. *Every vertical hyperplane of \tilde{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. \square

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

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

Lemma 5.1.8. *The subcomplex \tilde{Z}_C is a tree of finite trees whose underlying tree $\hat{\eta}_C(\tilde{Z}_C)$ is a copy of $\widehat{\phi(C)}$.*

Proof. By Lemma 5.1.5, $\hat{\eta}_C^{-1}(\sigma_z)$ is a finite tree for every vertex of \tilde{X} , and thus after subdivision, for the midpoint of every edge of \tilde{X} . Hence \tilde{Z}_C is a tree of finite trees if $\hat{\eta}_C(\tilde{Z}_C)$ is a tree. Now $\hat{\eta}_C$ sends vertical edges to vertical edges and horizontal edges dual to hyperplanes of type (iii) to vertices (Lemma 5.1.4). The horizontal edges of \tilde{Z}_C are of exactly this type by definition. Thus $\hat{\eta}_C(\tilde{Z}_C)$ is a tree.

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

Define $X_C := \tilde{X}_C/G$. By Lemma 5.1.7, 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 $\hat{\eta}_C : \tilde{X}_C \rightarrow \tilde{X}$ induces a map $\eta_C : X_C \rightarrow X$.

Let Y_C and Z_C denote the respective images of \tilde{Y}_C and \tilde{Z}_C in X_C . We have proved that

Lemma 5.1.9. *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 \tilde{C} passing through a lift \tilde{u} of u in \tilde{X} .*

We conclude with the following observation:

Lemma 5.1.10. *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 $\phi(C)$. \square*

5.2 Algorithmic construction of X_C

The main result of this section is the following.

Theorem 5.2.1. *There exists an algorithm of exponential time-complexity that takes a Brady-Meier tubular graph of graphs X and a splitting cycle $\phi : C \rightarrow X$ as input and returns the opened-up space X_C along C as output.*

Our method is to first combinatorially construct two square complexes Y'_C and Z'_C from X and glue them to form X'_C . We will show that $Y'_C \cong Y_C$ and $Z'_C \cong Z_C$ and thus $X'_C \cong X_C$ can be constructed combinatorially.

The first result we will need is the following. Fix a lift \tilde{C} of the splitting cycle $\phi : C \rightarrow X_s \subset X$. Let K be the number of half-spaces of \tilde{C} in \tilde{X} .

Lemma 5.2.2. *There exists $D' \in \mathbb{N}$ such that for any vertex or (midpoint of an edge) $v \in \tilde{C}$ and $\forall D \geq D'$, the D^{th} cubical neighbourhood $\{v\}^{+D}$ of v has the following properties:*

1. *For each translate $g\tilde{C}$ such that $v \in g\tilde{C}$, $g\tilde{C}$ separates $\{v\}^{+D}$ into exactly K components.*
2. *For every $g, g' \in G$ such that $g\tilde{C} \neq g'\tilde{C}$ and $v \in g\tilde{C} \cap g'\tilde{C}$, $g\tilde{C} \cap \{v\}^{+D} \neq g'\tilde{C} \cap \{v\}^{+D}$.*

The main ingredient for proving [Lemma 5.2.2](#) is the following result. Let N be such that the thickness of any edge of X is at most N .

Lemma 5.2.3. *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 \tilde{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 connected sum ([Definition 4.1.7](#)) of finitely many copies of Γ . Let $\phi_a : \{1, \dots, n\} \rightarrow \text{adj}(a)$ and $\phi_b : \{1, \dots, n\} \rightarrow \text{adj}(b)$ 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}, ϕ_{b_i} . Let $\Gamma'_i := \Gamma_1 (b_1, \phi_{b_1}) \bigoplus_{(a_2, \phi_{a_2})} \Gamma_2 (b_2, \phi_{b_2}) \bigoplus \cdots \bigoplus_{(a_i, \phi_{a_i})} \Gamma_i$.

Lemma 5.2.4. *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 partition induced by $\phi_a : \{1, \dots, n\} \rightarrow \text{adj}(a)$ and $\phi_b : \{1, \dots, n\} \rightarrow \text{adj}(b)$ on the vertices adjacent to a and b coincides with the partition induced by the k components of $\Gamma \setminus \{a, b\}$.

Similarly, $\Gamma'_3 \setminus \{a_1, b_3\}$ has k components if and only if there is a partition into k subsets of $\{1, \dots, n\}$ such that the partition induced by $\phi_{a_3} : \{1, \dots, n\} \rightarrow \text{adj}(a_3)$ and $\phi_{b_2} : \{1, \dots, n\} \rightarrow \text{adj}(b_2)$ coincide with the partitions induced by the k components of $\Gamma'_2 \setminus \{a_1, b_2\}$ and $\Gamma_3 \setminus \{a_3, b_3\}$. Since $\Gamma_2 \setminus \{a_2, b_2\} \cong \Gamma \setminus \{a, b\}$, the partition induced on $\text{adj}(b_2)$ by the k components of $\Gamma'_2 \setminus \{a_1, b_2\}$ coincides with the partition induced by the k components of $\Gamma_2 \setminus \{a_2, b_2\}$. Hence $\Gamma'_3 \setminus \{a_1, b_3\}$ has k components. Continuing iteratively, we obtain the result. \square

Proof of Lemma 5.2.3. Let e be an edge in the image of C . Subdivide X so that the midpoint m of e is a vertex. Let e_1 and e_2 denote the new edges obtained from e . Let P be a fundamental domain of C . We will assume that $P \subset \tilde{C}$ and that the initial vertex of P is (a lift of) m . Let P_1 be the concatenation $P.cP$, where c is a generator of $\pi_1(C)$. We will assume that the regular spheres of P and P_1 are embedded in \tilde{X} as subsets. Note that $\partial N(P_1) \cong \partial N(P) \bigoplus_{(e_1, \phi_1)} \bigoplus_{(ce_2, \phi_2)} \partial N(cP)$, by Lemma 4.1.11. Suppose that the number of components of $\partial N(P) \setminus \tilde{C}$ is equal to the number of components of $\partial N(P_1) \setminus \tilde{C}$. By Lemma 5.2.4 we have that the number of half-spaces of \tilde{C} is equal to the number of components of $\partial N(P) \setminus \tilde{C}$. Suppose now that the number of components of $\partial N(P) \setminus \tilde{C}$ is strictly greater than the number of components of $\partial N(P_1) \setminus \tilde{C}$. By Lemma 4.5.10, the number of components of $\partial N(P) \setminus \tilde{C}$ has at most N components. Thus the number of components of

$\partial N(P_1) \setminus \tilde{C}$ is at most $N - 1$. Let $P_2 = P_1.c^2P_1$. Proceeding as before, either the number of half-spaces of \tilde{C} is equal to the number of components of $\partial N(P_1) \setminus \tilde{C}$ or the number of components of $\partial N(P_1) \setminus \tilde{C}$ is strictly greater than the number of components of $\partial N(P_2) \setminus \tilde{C}$ which is at most $N - 2$. Continuing this way, we get that the number of components of $\partial N(P_N) \setminus \tilde{C}$ is either equal to 1 or is equal to the number of half-spaces of \tilde{C} . Hence the result. \square

Let L be a line in \tilde{X} and $v \in L$ a vertex. Let $D \in \mathbb{N}$. Note that $\{v\}^{+D}$ is a CAT(0) cube (sub)complex [HW08]. We will assume that $\partial N(L) \subset \tilde{X}$. By Lemma 4.2.8, we have

Lemma 5.2.5. *There exists a bijection between the half-spaces of $L \cap \{v\}^{+D}$ in $\{v\}^{+D}$ and the components of $\partial N(L) \cap \{v\}^{+D}$.* \square

Proof of Lemma 5.2.2. Let $D = 2^N$. Then $\{v\}^{+D}$ contains a lift of P_N , a fundamental domain of a $2^{N^{\text{th}}}$ power of C . By Lemma 5.2.3, \tilde{C} separates $\partial N(P_N) \subset \{v\}^{+D}$ into exactly K components. This implies that $\partial N(\tilde{C}) \cap \{v\}^{+D}$ has exactly K components. Lemma 5.2.5 then implies conclusion 1. Conclusion 2 is obvious as either $g\tilde{C} = g'\tilde{C}$ or $g\tilde{C} \cap g'\tilde{C}$ has length at most the length of C , by Lemma 4.5.21. \square

5.2.1 Construction of X'_C

Let $D = \ell(C) + 2^N$, where $\ell(C)$ denotes the length of C . Choose a basepoint $v \in \phi(C)$ with lift \tilde{v} in \tilde{C} . Let $\mathcal{B} := \{\tilde{v}\}^{+D}$ in \tilde{X} . Note that

Lemma 5.2.6. *There exists an algorithm that takes X and v as input and returns \mathcal{B} in exponential time.* \square

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 (edge e) in $\phi(C)$, we define a set of walls W_u (W_e) on Z as follows. Choose a lift of u (e) in \mathcal{B} such that $\{\tilde{u}\}^{+2^N}$ (\tilde{e}^{+2^N}) is contained in \mathcal{B} .

Let $\tilde{C} = L_1, \dots, L_n$ be the translates of \tilde{C} that pass through \tilde{u} (\tilde{e}). By Lemma 5.2.2,

each line L_i separates \mathcal{B} into exactly K components. Thus each line L_i induces K walls of W_u (W_e) on Z , where each half-space Y in \tilde{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 (W_e) cross (Lemma 5.1.3). Thus the dual cube complex of (Z, W_u) (respectively (Z, W_e)) is a tree, denoted by $T(u)$ ($T(e)$).

Lemma 5.2.7. *Let $v \in \phi(C)$ be a vertex and \tilde{v}_1, \tilde{v}_2 be lifts of v in \tilde{X} . Then*

1. $\{\tilde{v}_1\}^{+D}$ and $\{\tilde{v}_2\}^{+D}$ are isomorphic as square complexes. Further,
2. if $u \in \phi(C)$ is a vertex or a midpoint of an edge and \tilde{u}_1 and \tilde{u}_2 are lifts of u in $\{\tilde{v}_1\}^{+D}$, then there exists a natural bijection between the translates of \tilde{C} that meet \tilde{u}_1 and the translates of \tilde{C} that meet \tilde{u}_2 . \square

Thus the definition of W_u (and W_e) is independent of the choice of \tilde{v} or of the choice of \tilde{u} (\tilde{e}) in \mathcal{B} .

Suppose that an edge e is incident to a vertex u in $\phi(C)$. Let \tilde{e} with incident vertex \tilde{u} be a lift of e incident to u . Since every translate of \tilde{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, L_2 that contain \tilde{e} with half-spaces Y_1, 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 5.2.8. *Given an edge e in $\phi(C)$ incident to a vertex u , there exists a natural inclusion $T(e) \hookrightarrow T(u)$.*

The space X'_C is constructed from two spaces Z'_C and Y'_C , defined below.

- Z'_C is the geometric realisation of the graph of trees $(\phi(C), \{T(u)\}, \{T(e)\})$.
- Y'_C is 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. Call the resulting space as Y'_C . The vertices and edges of Y'_C corresponding to $\phi(C)$ are the *boundary cells* of

Y'_C .

Lemma 5.2.9. *There exists a natural map from the boundary cells of Y'_C to Z'_C which defines a gluing of Z'_C to Y'_C .*

Proof. Consider a boundary cell u' (e') of Y'_C , which is a copy of the vertex u (edge e) of $\phi(C)$. Choose a horizontal open half-edge z' in Y'_C incident to u' (e') with corresponding edge z in X . Let \tilde{u} (\tilde{e}) be a lift of u (e) in \mathcal{B} and \tilde{z} be the half-edge incident to \tilde{u} (\tilde{e}) that projects to z in X . Then the required map is the one that sends z' to σ_z in $T(u) \subset Z'_C$ ($\sigma_z \subset T(e) \times e \subset Z'_C$). \square

Definition 5.2.10. The square complex X'_C is defined to be the complex $Z'_C \sqcup Y'_C / \sim$, where \sim is induced by the natural gluing.

Proposition 5.2.11. *There exists a natural isomorphism between the square complexes X_C and X'_C .*

Proof. By [Lemma 5.1.10](#), X_C is a union of Y_C and Z_C , while X'_C is a union of Y'_C and Z'_C by definition. Y_C is clearly isomorphic to Y'_C .

By [Lemma 5.1.9](#), 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 \tilde{X} of translates of \tilde{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 \tilde{C} in a finite ball \mathcal{B} of \tilde{X} containing \tilde{u} . Since there exists a bijection between the half-spaces of \tilde{C} in \mathcal{B} and the half-spaces of \tilde{C} in \tilde{X} ([Lemma 5.2.5](#)), Z_C is isomorphic to Z'_C .

Thus X_C is isomorphic to X'_C if the isomorphism from Y_C to Y'_C restricted to $Y_C \cap Z_C$ is an isomorphism to $Y'_C \cap Z'_C$. The cells of $Y_C \cap Z_C$ are precisely the cells of Y_C that are mapped to $\phi(C)$. These cells are mapped to the boundary cells of Y'_C by definition. Further, this mapping is bijective. Hence the result. \square

Proof of [Theorem 5.2.1](#). The compact space \mathcal{B} can be constructed in exponential time from X ([Lemma 5.2.6](#)). It costs exponential time to calculate the number of half-spaces of any translate of \tilde{C} ([Lemma 5.2.3](#)).

Since the number of translates of \tilde{C} meeting at any point of \tilde{X} is bounded by the length of C (by [Lemma 4.5.21](#)), the dual trees $T(u)$ ($T(e)$) of all vertices u (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 $X'_C \cong X_C$ is obtained in linear time from Y'_C and Z'_C . Hence the result. \square

5.2.2 Structure of X_C

Lemma 5.2.12. *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 $T(u)$ ($T(e)$) and then show that the tree is bipartite. Let L be a translate of \tilde{C} passing through \tilde{u} (\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 $T(u)$ ($T(e)$). Thus σ_L contains $\{Y_1^c, \dots, Y_K^c\}$ and exactly those half-spaces of translates of \tilde{C} passing through \tilde{u} that contain L .

We now claim that the valence of σ_L is exactly K . Indeed, it is at least K : switching each half-space Y_i of L gives an edge incident to σ_L . We now claim that there exists no ultrafilter σ' such that $\sigma' \Delta \sigma_L = \{Y', Y'^c\}$, with $Y' \neq Y_i$. Assume without loss of generality that $Y' \in \sigma_L$ (that is, $L \subset Y'$). Since Y'^c is disjoint from L , there exists Y_i of L such that $Y'^c \subset Y_i$. This implies that σ' contains Y_i . Thus $\sigma_L \Delta \sigma'$ contains $\{Y_i, Y_i^c, Y', Y'^c\}$.

Let σ' be a vertex at distance two from σ_L . Let $\sigma' \Delta \sigma_L = \{Y_1, Y_1^c, Y'_1, Y'^c_1\}$, where Y'_1 is a half-space of a translate L' of \tilde{C} . We will show that $\sigma' = \sigma_{L'}$.

Assume first that $Y'_1 \in \sigma_L$. This implies that for each half-space Y'_i of L' with $i \neq 1$, $Y'^c_i \in \sigma_L$ and hence in σ' . Further, $Y'^c_1 \in \sigma'$ by assumption. If $\sigma' \neq \sigma_{L'}$, then any path from σ' to $\sigma_{L'}$ involves a change of half-spaces of the type $\{Y'_i, Y'^c_i\}$. Hence we conclude that $\sigma' = \sigma_{L'}$.

Assume now that $Y'_1 \in \sigma'$. Thus $Y'^c_1 \in \sigma_L$. Since the edge which flips Y'_1 and Y'^c_1

is incident to $\sigma_{L'}$, $\sigma_{L'}$ lies in any path between σ' and σ_L . We will show that σ_L and $\sigma_{L'}$ are not adjacent. This will prove that σ' is not at distance two from σ_L whenever $Y'_1 \in \sigma'$. Indeed, there exists a half-space Y'_2 , say, of L' that contains L and hence $Y'_2 \in \sigma_L$. Thus there exists an edge which flips Y'_2 and $Y_2'^c$ in any path between σ_L and $\sigma_{L'}$. But such an edge is not adjacent to σ_L as it does not flip Y_i and Y_i^c .

This proves that a vertex in $T(u)$ ($T(e)$) is of the form $\sigma_{L'}$ if and only if it is at even distance from σ_L . Thus the tree is bipartite. \square

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

Lemma 5.2.13. *Given an edge e_1 incident to v_1 in $T(u)$ ($T(e)$), there exists an edge e_2 incident to v_2 in $T(u)$ ($T(e)$) 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 \tilde{C} passing through \tilde{u} (\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} (\tilde{e}) such that there exists $g \in G$ and \tilde{u}' (\tilde{e}' in P') such that $g\tilde{u}' = \tilde{u}$ ($g\tilde{e}' = \tilde{e}$) and $gP' \subset L_2$. The segment from \tilde{u} to \tilde{u}' (\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 $g\tilde{u}' = \tilde{u}$ ($g\tilde{e}' = \tilde{e}$) and $gL_1 = L_2$, the required edge incident to σ_{L_2} is $\{gY_1, gY_1^c\}$. \square

Let $\phi' : C' \rightarrow X_s$ be a primitive cycle such that C is a power of C' .

Lemma 5.2.14. *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 (edge e) in $\phi(C)$, the embedded copy of C' meets every black vertex of $T(u)$ ($T(e)$) exactly once.*

Proof. Choose a fundamental domain P' of C' in \tilde{C} in \mathcal{B} . For each vertex \tilde{u} and edge \tilde{e} in P' with images u and e respectively in $\phi(C)$, choose $\sigma_{\tilde{C}}$ in $T(u)$ and $T(e)$.

This induces an embedding of C' in Z'_C with the required properties. \square

5.3 The tubular graph of graphs X'

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be the set of splitting cycles of X furnished by [Theorem 4.9.1](#).

Remark 5.3.1. 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 \mathcal{C} .

Procedure 5.3.2 (Construction of X'). The tubular graph of graphs X' is constructed from X using the cycles in \mathcal{C} as follows:

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

Lemma 5.3.3. *The cycle C_i in \mathcal{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 \tilde{X} . \square*

By [Theorem A](#), we will assume that X' is a Brady-Meier tubular graph of graphs with fundamental group G .

Theorem 5.3.4. *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 upto commensurability.*

Proof. The algorithm of [Theorem 4.9.1](#) takes X as input and returns a set of cycles \mathcal{C} that contains all universally elliptic cycles upto commensurability. Given \mathcal{C} , X' is constructed using [Procedure 5.3.2](#). This procedure consists of applying the algorithm of [Theorem 5.2.1](#) repeatedly.

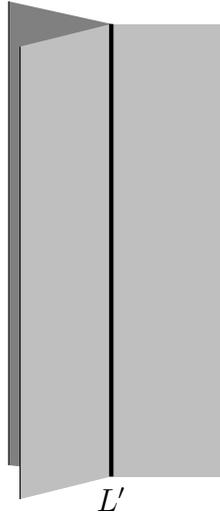


Figure 5.1: L' in \tilde{X}' when L has three half-spaces

The algorithm of [Theorem 4.9.1](#) 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 of [Theorem 5.2.1](#) operates by taking exponential time for each cycle, and hence obtaining X' costs double exponential time in the input data. \square

5.3.1 Structure of X'

Note that the maps $\hat{\eta}_i : \tilde{X}_i \rightarrow \tilde{X}_{i-1}$ induce maps $\hat{\eta} : \tilde{X}' \rightarrow \tilde{X}$ and $\eta : X' \rightarrow X$.

Denote by Γ' the underlying graph of the graph of spaces X' and by T' the underlying tree of the tree of spaces \tilde{X}' . Note that T' is the Bass-Serre tree of the cyclic splitting of G induced by X' . Let L be a lift of an element C_i of \mathcal{C} such that C_i factors through a vertical cycle in X_{i-1} .

Lemma 5.3.5. *There exists a vertical tree in \tilde{X}' whose stabiliser is equal to the stabiliser of L .*

Proof. By [Lemma 5.2.14](#), a primitive cycle C'_i such that C_i is a power of C'_i embeds in X_i as a vertical graph. Since the process of producing X' from X_i does not crush any horizontal edges, a lift L' of C' in \tilde{X}' is a vertical tree ([Figure 5.1](#)). \square

Let \mathfrak{h} be a vertical hyperplane in \widetilde{X}' . Let L_1 and L_2 be the two *boundary lines* of \mathfrak{h} , that is, the two vertical lines on either side of \mathfrak{h} at distance $\frac{1}{2}$ from \mathfrak{h} , and parallel to \mathfrak{h} .

Lemma 5.3.6. *The stabiliser of \mathfrak{h} is equal to either $\text{stab}(L_1)$ or $\text{stab}(L_2)$.*

Proof. Denote by H , H_1 and H_2 the stabilisers of \mathfrak{h} , L_1 and L_2 respectively. Suppose that H is not equal to either H_1 or H_2 . Fix $h_1 \in H_1 \setminus H$ and $h_2 \in H_2 \setminus H$. We will then show that there is a flat plane contained in \widetilde{X}' as a subcomplex.

Consider the hyperplane $h_1\mathfrak{h}$. Its boundary lines are $h_1L_1 = L_1$ and $h_1L_2 \neq L_2$. Thus the strips S and h_1S share exactly one boundary line, L_1 . Next, consider the hyperplane $h_1h_2\mathfrak{h}$. Its boundary lines are $h_1h_2L_1 \neq L_1$ and $h_1h_2L_2 = h_1L_2$. Thus the strip h_1h_2S shares exactly one boundary line, h_1L_2 with h_1S . Note that $h_1h_2S \neq S$ as $h_1h_2 \notin H$. Similarly, the strip $h_1h_2h_1S$ shares exactly one boundary line $h_1h_2L_1$ with h_1h_2S and so on. In this way, we get a closed half-plane (homeomorphic to $\mathfrak{h} \times [0, \infty)$) on one side of \mathfrak{h} . Similarly, we can construct another closed half-plane on the other side of \mathfrak{h} starting with the strip h_2S . The union of these two half-planes is a plane containing \mathfrak{h} . By the Flat Plane Theorem (see Theorem $\Gamma.3.1$ of [BH99]), G is not hyperbolic, which is a contradiction. \square

Recall that \widetilde{X}_C was defined as the cube complex dual to W_C , the set of walls induced by hyperplanes of \widetilde{X} and walls induced by translates of \widetilde{C} . Analogously, \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)* are induced by horizontal hyperplanes of \widetilde{X} .
- *Walls of type (ii)* are induced by vertical hyperplanes of \widetilde{X} .
- *Walls of type (iii)* are induced by translates of \widetilde{C}_i , where $C_i \in \mathcal{C}$ 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 \tilde{X}' and L_1, L_2 its boundary lines. Then

Lemma 5.3.7. *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 type (iii) in Z corresponding to h .

If $\{Y, Y^c\}$ is of type (iii), let L be the line in \tilde{X} that defines the wall. If not, then let L be a boundary line in \tilde{X} of the vertical hyperplane that defines $\{Y, Y^c\}$, so that L defines a wall of type (iii) (Remark 5.3.1).

By Lemma 5.3.5, there exists a linear vertical tree L' in \tilde{X}' such that $\text{stab}(L) = \text{stab}(L')$. Note that any vertical half-space of \tilde{X} contained in an ultrafilter of L' contains either Y or Y^c , by Lemma 4.4.4. Thus any vertical half-space of \tilde{X}' (except perhaps a half-space of L') that contains L' contains h . So no vertical tree separates L' and h . Hence the result. \square

Definition 5.3.8. 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 \tilde{X}' , at least one of the two vertices of any edge is a cyclic vertex.

Lemma 5.3.9. *Let L' be a line in \tilde{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 \tilde{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 \tilde{X}' such that $\hat{\eta}(L') = L$ (Lemma 5.3.5).

Lemma 5.3.10. *Exactly K vertical strips are attached to L' in \tilde{X}' . Further, 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 5.2.14. Further, each of the K strips above are contained in $\hat{\eta}^{-1}(L)$, by Lemma 5.1.5.

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 \tilde{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 $\sigma'_i \Delta \sigma'_{L'} = \{Y_i, Y_i^c\}$, where Y_i is a half-space of L . Since $\{Y_i, Y_i^c\}$ defines a vertical hyperplane \mathfrak{h}_i in \tilde{X}' , σ'_i is the set of all vertical half-spaces of some vertex of \tilde{X}' that is separated from L' by exactly \mathfrak{h}_i . So 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 $\{Y_0, Y_0^c, Y_i, Y_i^c\} \subset \sigma'_{L''} \Delta \sigma'_{L'}$, with $Y_i \subset \sigma'_{L''}$. Let σ be a vertex in $L \setminus L_0$. Then there exists $\sigma'' \in L''$ such that $\hat{\eta}(\sigma'') = \sigma$. 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 proof follows from the following observation. Let γ denote a geodesic between L'' and L'_i . Since γ consists of vertical edges, $\hat{\eta}(\gamma)$ has the same length as γ (by [Lemma 5.1.4](#)) and is a geodesic between $\hat{\eta}(L'')$ and $\hat{\eta}(L'_i)$. \square

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

Lemma 5.3.11. *$\partial N(L'_i)$ has exactly two components.*

Proof. Note that L'_i is a separating line since it is a tubular line ([Fact 4.2.10](#)).

Let Y'_i be a half-space of L'_i that does not contain L' . Then $\hat{\eta}(Y'_i)$ does not contain L ([Lemma 5.3.10](#)) and is connected. Thus $\hat{\eta}(Y'_i)$ 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 $\hat{\eta}^{-1}(Y_i \setminus L)$ contains these half-spaces. By [Lemma 5.1.5](#), one of these half-spaces is at finite distance from the other, and hence from L'_i . This is a contradiction to [Lemma 4.2.12](#). \square

Proof of [Lemma 5.3.9](#). If $\partial N(L')$ contains three or more components, then a subgroup H of $\text{stab}(L')$ is universally elliptic, by [Proposition 4.6.5](#).

Let $L = \hat{\eta}(L')$. Then H stabilises L as $\hat{\eta}$ is G -equivariant. Since the limit set of H

separates ∂G , L coarsely separates \tilde{X} and hence separates \tilde{X} . This implies that L defines a wall of type (iii). [Lemma 5.3.10](#) and [Lemma 5.3.11](#) then give the result. \square

5.4 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' by removing certain tubes of X' .

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 5.3.6](#), one of the attaching maps of such a tube is an isomorphism at the level of groups. Thus the attaching map is an isomorphism of graphs between the edge graph and the corresponding vertex graph. Successively remove all such tubes of X' . Call the new tubular graph of graphs as X'' .

5.4.1 Structure of X''

Let T'' be the underlying tree of \tilde{X}'' and let T_{jsj} denote the Bass-Serre tree of the canonical JSJ decomposition of G . We will show that every edge stabiliser of T'' is either an edge stabiliser of T_{jsj} or conjugates into a hanging surface subgroup of the JSJ decomposition. Thus T'' can be modified to T_{jsj} by removing the latter type of edges.

Lemma 5.4.1. *For each cyclic vertex u of T_{jsj} there exists a cyclic vertex v of T'' such that $\text{stab}(u) = \text{stab}(v)$.*

Proof. Fix an axis L in \tilde{X} of $\text{stab}(u)$. Note that $\text{stab}(u) = \text{stab}(L)$ as $\text{stab}(u)$ is a maximal cyclic subgroup of G . Let H be the edge stabiliser in T_{jsj} of an edge incident to u . Then H is a universally elliptic subgroup. We can thus assume that the line L is vertical ([Lemma 4.6.4](#)).

By [Lemma 4.5.26](#), there exists a splitting cycle C in X such that $\pi_1(C)$ is commensurable with a conjugate of H . Hence C is universally elliptic and $C \in \mathcal{C}$. Thus there exists a vertical tree in \tilde{X}' whose stabiliser is $\text{stab}(\tilde{C})$. Since \tilde{C} is a translate of L , there exists a vertical tree in \tilde{X}' whose stabiliser is $\text{stab}(L) = \text{stab}(u)$. Observe that the process of modifying X' to obtain X'' does not glue any linear tree of \tilde{X}' to a non-linear tree as tubes are removed only when both the bounding vertex graphs are circles. Hence the result. \square

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

Lemma 5.4.2. *The number of components of $\partial G \setminus \Lambda \text{stab}(v)$ is equal to the number of edges incident to v in T'' .*

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 $\partial G \setminus \Lambda \text{stab}(v)$ is equal to the supremum of the number of components of $\tilde{X}'' \setminus L''^{+k}$, 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 \tilde{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 as otherwise L''_i would have more than two half-spaces. By [Lemma 4.2.16](#), $Y''_i \setminus L''_i^{+k}$ and hence $Y''_i \setminus L''^{+(k+1)}$ is connected, for every $k \in \mathbb{N}$. Hence the result. \square

Proposition 5.4.3. *For each edge of T_{jsj} with stabiliser H , there exists an edge of T'' whose stabiliser is H .*

Proof. Since each edge of T_{jsj} is incident to a cyclic vertex, let $\text{stab}(u)$ be the cyclic vertex group of T_{jsj} that contains H . Then H is the stabiliser of a component of $\partial G \setminus \Lambda \text{stab}(u)$. Let v be a vertex of T'' such that $\text{stab}(v) = \text{stab}(u)$ ([Lemma 5.4.1](#)).

Then the number of edges incident to v is equal to the number of components of $\partial G \setminus \Lambda \text{stab}(u)$, by [Lemma 5.4.2](#). Further, H is the stabiliser of an edge incident to

v as each edge incident to v induces a unique component of $\partial G \setminus \Lambda \text{stab}(u)$. \square

Definition 5.4.4 ([GL16]). A G -tree \hat{T} is said to be a *refinement* of a G -tree T if there exists a G -equivariant map $p : \hat{T} \rightarrow T$ such that p sends any segment $[x, y]$ in \hat{T} onto the segment $[p(x), p(y)]$. In other words, \hat{T} is obtained by blowing up vertices of T .

Corollary 5.4.5. T'' is a refinement of T_{jsj} .

Proof. Indeed, every edge stabiliser of T_{jsj} is an edge stabiliser of T'' , by [Proposition 5.4.3](#). Further, any edge stabiliser of T'' that is not an edge stabiliser of T_{jsj} has to belong to a maximal hanging surface vertex group in T_{jsj} ([Definition 1.3.1](#)) and hence gives a relative splitting of the maximal hanging surface subgroup that it belongs to. Hence the result. \square

Recall that a surface subgroup G' of G is a *hanging surface group* ([Sel97]) if there exists a graph of groups decomposition of G such that G' is a vertex group and the peripheral subgroups of G' are precisely the incident edge subgroups.

Lemma 5.4.6. *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 non-cyclic vertices adjacent to u are stabilised by hanging surface groups.*

Proof. One direction is clear, since no edge of T_{jsj} is such that the cyclic vertex incident to this edge is adjacent to two hanging surface group vertices.

Conversely, if H is not an edge stabiliser of T_{jsj} , it is contained in a maximal hanging surface group of T_{jsj} . Thus u is of valence two, by [Lemma 5.4.2](#). Further, the stabilisers of both the vertices adjacent to the u are hanging surface groups as their image in T_{jsj} is contained in the maximal hanging surface group that contains H . \square

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

5.5 Surface graphs

Definition 5.5.1. A vertex graph of a tubular graph of graphs is said to be 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.

In other words, a vertex graph with its incident edge graphs is a surface graph if its fundamental group is a hanging surface group.

Lemma 5.5.2. *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.*

Recall that 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.

Proof. Let D_s be the double of the vertex graph X_s . 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 \tilde{D}_s and \mathfrak{h} the horizontal hyperplane through \tilde{e} . Note that \mathfrak{h} is a tree. Let L be a line in \mathfrak{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 4.2.5](#), L does not separate \tilde{D}_s . But this implies that D_s is not homeomorphic to a closed surface. \square

5.6 Construction of X_{jsj}

We are now ready to construct X_{jsj} . Let X''' be the tubular graph of graphs obtained from X'' by removing pairs of tubes of X'' whenever they are incident to the same cyclic vertex graph on one side and to surface graphs on the other. Surface graphs can be identified by [Lemma 5.5.2](#). Denote by T''' the underlying tree of \tilde{X}''' . By [Lemma 5.4.6](#), we have

Proposition 5.6.1. *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 of the article:

Theorem 5.6.2. *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 Brady-Meier tubular graph of graphs whose underlying graph of groups structure is the JSJ decomposition of G .*

Proof. Let X be the input tubular graph of graphs and G its fundamental group. Using [Theorem 5.3.4](#), 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. 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 5.5.2](#)) and also takes at most polynomial time. □

Chapter 6

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 6.0.1. A *relative JSJ decomposition* of (F, \mathcal{H}) is a graph of groups splitting of F which satisfies the conditions of a JSJ decomposition ([Definition 1.3.1](#)) with the additional property that each element of \mathcal{H} is elliptic.

Theorem 6.0.2 (Theorem 4.25, [Cas16]). *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.*

The main result of this section is the following.

Theorem 6.0.3. *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} .*

6.1 A tubular model for (F, \mathcal{H})

We will construct a suitable tubular graph of graphs $X_{F, \mathcal{H}}$ to prove [Theorem 6.0.3](#).

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 $\phi_i : C_i \rightarrow X_{s_c}$ 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 and each of valence three, 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 Theorem 18 of [Wil12]. Hence, we can assume that $X_{F,\mathcal{H}}$ is Brady-Meier, by [Theorem A](#).

Lemma 6.1.1. *G is δ -hyperbolic.*

Proof. Consider the graph of groups structure of the tubular graph of graphs $X_{F,\mathcal{H}}$. Each vertex group is either a free group of rank 1 or more and is hence hyperbolic. By construction, we have that each edge group is maximal cyclic in both its incident vertex groups. The result is then a consequence of the Bestvina-Feighn Combination Theorem (Corollary (torsion-free products over \mathbb{Z}), page 100 of [BF92]). \square

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

Lemma 6.1.2. *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} . Let X_s be the corresponding vertex graph in $X_{F, \mathcal{H}}$. Note that \tilde{X}_s is a line for any lift of X_s . Further, $\partial N(\tilde{X}_s)$ contains three components as there are three tubes attached to X_s by graph isomorphisms. By [Proposition 4.6.5](#), a subgroup of G_s is universally elliptic. This implies that G_s is commensurable with a cyclic vertex group of the JSJ splitting of G . The result then follows from the fact that G_s is maximal cyclic.

Now suppose that G_s is a hanging surface group induced by a surface graph X_s . Suppose that G_s is not maximal. Then there exists $g \in G \setminus G_s$ such that G splits over $\langle g \rangle$ and $\langle g \rangle$ is not elliptic in the Bass-Serre tree of a cyclic splitting induced by some element of G_s . Let L be an axis of g in $\tilde{X}_{F, \mathcal{H}}$. Then the above implies that L and a line $L' \subset \tilde{X}_s$ cross (by [Lemma 4.6.1](#)). If L is transversal and meets \tilde{X}_s , then L has to meet a vertex graph adjacent to \tilde{X}_s . But any vertex graph adjacent to \tilde{X}_s is a line whose stabiliser is a cyclic vertex group of the JSJ as seen above. Hence L cannot cross such a vertex graph. This implies that L is contained in \tilde{X}_s , contradicting the fact that $g \notin G_s$. \square

Corollary 6.1.3. *If T_{jsj} is the JSJ tree of G and T the underlying tree of $\tilde{X}_{F, \mathcal{H}}$, then T_{jsj} is a refinement of T obtained by blowing up lifts of the central vertex s_c .*

We are now ready to prove the main result of this section.

Proof of [Theorem 6.0.3](#). Given (F, \mathcal{H}) , the space $X_{F, \mathcal{H}}$ can be constructed algorithmically. Constructing X_{s_c} and the tubes attached to X_{s_c} takes at most polynomial time in \mathcal{H} . Constructing the tubular graph of graphs $X_{F, \mathcal{H}}$ then takes at most polynomial time.

Let X_{jsj} be the tubular graph of graphs obtained from $X_{F, \mathcal{H}}$ in double exponential time by [Theorem 5.6.2](#). 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 . In other words, there exists a map of graphs $f : \Gamma_{jsj} \rightarrow \Gamma$ such that for each vertex $s \neq s_c$, $f^{-1}(s)$ consists of a single vertex.

Let Y be the subgraph of groups of the JSJ decomposition of G induced by $f^{-1}(s_c)$. Then it is easy to see that Y is the relative JSJ of (F, \mathcal{H}) . Indeed, since each element of \mathcal{H} corresponds to an edge graph of $X_{F, \mathcal{H}}$, every element of \mathcal{H} is elliptic in Y . Y is bipartite since Γ_{jsj} is bipartite. The incident edge groups of a non-cyclic vertex group G_v of Y map onto maximal cyclic subgroups of G_v since $Y \subset \Gamma_{jsj}$ is the JSJ decomposition. In a similar fashion, all other conditions of the relative JSJ decomposition can be verified since $Y \subset \Gamma_{jsj}$. \square

Bibliography

- [AYZ97] N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.
- [Bar16] B. Barrett. Computing JSJ decompositions of hyperbolic groups. *ArXiv e-prints*, November 2016.
- [BF92] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.
- [BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [BM01] Noel Brady and John Meier. Connectivity at infinity for right angled Artin groups. *Trans. Amer. Math. Soc.*, 353(1):117–132, 2001.
- [Bow98] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [Cas16] Christopher H. Cashen. Splitting line patterns in free groups. *Algebr. Geom. Topol.*, 16(2):621–673, 2016.
- [CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.
- [Deh11] M. Dehn. über unendliche diskontinuierliche Gruppen. *Math. Ann.*, 71(1):116–144, 1911.

- [DF05] Guo-An Diao and Mark Feighn. The Grushko decomposition of a finite graph of finite rank free groups: an algorithm. *Geom. Topol.*, 9:1835–1880, 2005.
- [DG08] François Dahmani and Daniel Groves. Detecting free splittings in relatively hyperbolic groups. *Trans. Amer. Math. Soc.*, 360(12):6303–6318, 2008.
- [DG11] François Dahmani and Vincent Guirardel. The isomorphism problem for all hyperbolic groups. *Geom. Funct. Anal.*, 21(2):223–300, 2011.
- [DT13] F. Dahmani and N. Touikan. Deciding Isomorphy using Dehn fillings, the splitting case. *ArXiv e-prints*, November 2013.
- [FP06] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.
- [Fre31] Hans Freudenthal. Über die enden topologischer räume und gruppen. *Mathematische Zeitschrift*, 33(1):692–713, 1931.
- [Ger99] V Gerasimov. Detecting connectedness of the boundary of a hyperbolic group. Unpublished, 1999.
- [GL16] V. Guirardel and G. Levitt. JSJ decompositions of groups. *ArXiv e-prints*, February 2016.
- [Gru40] I. Gruschko. über die Basen eines freien Produktes von Gruppen. *Rec. Math. [Mat. Sbornik] N.S.*, 8 (50):169–182, 1940.
- [Hag07] F. Haglund. Isometries of CAT(0) cube complexes are semi-simple. *ArXiv e-prints*, May 2007.
- [Hop44] Heinz Hopf. Enden offener Räume und unendliche diskontinuierliche Gruppen. *Comment. Math. Helv.*, 16:81–100, 1944.
- [HP98] Frédéric Haglund and Frédéric Paulin. Simplicité de groupes d’automorphismes d’espaces à courbure négative. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 181–248. Geom. Topol. Publ., Coventry, 1998.

-
- [HT73] John Hopcroft and Robert Tarjan. Algorithm 447: Efficient algorithms for graph manipulation. *Commun. ACM*, 16(6):372–378, June 1973.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.
- [Jac69] William Jaco. Three-manifolds with fundamental group a free product. *Bull. Amer. Math. Soc.*, 75:972–977, 1969.
- [JLR02] William Jaco, David Letscher, and J. Hyam Rubinstein. Algorithms for essential surfaces in 3-manifolds. In *Topology and geometry: commemorating SISTAG*, volume 314 of *Contemp. Math.*, pages 107–124. Amer. Math. Soc., Providence, RI, 2002.
- [Joh79] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.
- [JS78] William Jaco and Peter B. Shalen. A new decomposition theorem for irreducible sufficiently-large 3-manifolds. In *Algebraic and geometric topology (Proc. Sympos. Pure Math., Stanford Univ., Stanford, Calif., 1976), Part 2*, Proc. Sympos. Pure Math., XXXII, pages 71–84. Amer. Math. Soc., Providence, R.I., 1978.
- [Mil62] J. Milnor. A unique decomposition theorem for 3-manifolds. *Amer. J. Math.*, 84:1–7, 1962.
- [MW05] J. P. McCammond and D. T. Wise. Coherence, local quasiconvexity, and the perimeter of 2-complexes. *Geom. Funct. Anal.*, 15(4):859–927, 2005.
- [Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.
- [NR98a] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.
- [NR98b] Graham A. Niblo and Martin A. Roller. Groups acting on cubes and Kazhdan’s property (T). *Proc. Amer. Math. Soc.*, 126(3):693–699, 1998.

- [Pap05] Panos Papasoglu. Quasi-isometry invariance of group splittings. *Ann. of Math. (2)*, 161(2):759–830, 2005.
- [Pap12] Panos Papasoglu. Splittings and the asymptotic topology of the lamplighter group. *Trans. Amer. Math. Soc.*, 364(7):3861–3873, 2012.
- [Ray60] Frank Raymond. The end point compactification of manifolds. *Pacific J. Math.*, 10:947–963, 1960.
- [Rol16] M. Roller. Poc Sets, Median Algebras and Group Actions. *ArXiv e-prints*, July 2016.
- [RS97] E. Rips and Z. Sela. Cyclic splittings of finitely presented groups and the canonical JSJ decomposition. *Ann. of Math. (2)*, 146(1):53–109, 1997.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [Sel97] Z. Sela. Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II. *Geom. Funct. Anal.*, 7(3):561–593, 1997.
- [Ser80] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin-New York, 1980. Translated from the French by John Stillwell.
- [Spe49] Ernst Specker. Die erste Cohomologiegruppe von Überlagerungen und Homotopie-Eigenschaften dreidimensionaler Mannigfaltigkeiten. *Comment. Math. Helv.*, 23:303–333, 1949.
- [Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.
- [Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.
- [Sta99] John R. Stallings. Whitehead graphs on handlebodies. In *Geometric group theory down under (Canberra, 1996)*, pages 317–330. de Gruyter, Berlin, 1999.

-
- [SW79] Peter Scott and Terry Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.
- [SW05] Michah Sageev and Daniel T. Wise. The Tits alternative for CAT(0) cubical complexes. *Bull. London Math. Soc.*, 37(5):706–710, 2005.
- [Whi36] J. H. C. Whitehead. On equivalent sets of elements in a free group. *Ann. of Math. (2)*, 37(4):782–800, 1936.
- [Wil12] Henry Wilton. One-ended subgroups of graphs of free groups with cyclic edge groups. *Geom. Topol.*, 16(2):665–683, 2012.
- [Wil17] H. Wilton. Essential surfaces in graph pairs. *ArXiv e-prints*, January 2017.
- [Wis96] Daniel T. Wise. *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)—Princeton University.
- [Wis06] Daniel T. Wise. Subgroup separability of the figure 8 knot group. *Topology*, 45(3):421–463, 2006.