


Detecting topological properties of boundaries of hyperbolic groups



Benjamin Barrett

Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Trinity College

November 2018

Detecting topological properties of boundaries of hyperbolic groups

Benjamin Barrett

In general, a finitely presented group can have very nasty properties, but many of these properties are avoided if the group is assumed to admit a nice action by isometries on a space with a negative curvature property, such as Gromov hyperbolicity. Such groups are surprisingly common: there is a sense in which a random group admits such an action, as do some groups of classical interest, such as fundamental groups of closed Riemannian manifolds with negative sectional curvature. If a group admits an action on a Gromov hyperbolic space then large scale properties of the space give useful invariants of the group. One particularly natural large scale property used in this way is the Gromov boundary.

The Gromov boundary of a hyperbolic group is a compact metric space that is, in a sense, approximated by spheres of large radius in the Cayley graph of the group. The technical results contained in this thesis are effective versions of this statement: we see that the presence of a particular topological feature in the boundary of a hyperbolic group is determined by the geometry of balls in the Cayley graph of radius bounded above by some known upper bound, and is therefore algorithmically detectable.

Using these technical results one can prove that certain properties of a group can be computed from its presentation. In particular, we show that there are algorithms that, when given a presentation for a one-ended hyperbolic group, compute Bowditch's canonical decomposition of that group and determine whether or not that group is virtually Fuchsian.

The final chapter of this thesis studies the problem of detecting Čech cohomological features in boundaries of hyperbolic groups. Epstein asked whether there is an algorithm that computes the Čech cohomology of the boundary of a given hyperbolic group. We answer Epstein's question in the affirmative for a restricted class of hyperbolic groups: those that are fundamental groups of graphs of free groups with cyclic edge groups. We also prove the computability of the Čech cohomology of a space with some similar properties to the boundary of a hyperbolic group: Otal's decomposition space associated to a line pattern in a free group.

Declaration

I hereby declare that this dissertation is the result of my own work and includes nothing that is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text

Benjamin Barrett
November 2018

Acknowledgements

First and foremost I would like to express my deep gratitude to my supervisor, Henry Wilton, whose guidance, encouragement, positivity and patience have made these four years a rewarding and (mostly) enjoyable experience. I could not have hoped for more dedicated supervisor. I also thank the many other people in the geometric group theory community, in Cambridge and more broadly, who have taken the time to discuss my work with me, including but certainly not limited to Martin Bridson, Brian Bowditch, Daniel Groves, Panos Papasoglu, Alan Logan, John Mackay and Alex Margolis. I am grateful to the organisers of Geometric Group Theory in the South East and of the Young Geometric Group Theory conference series: these events have broadened my understanding of geometric group theory and made me a better mathematician.

I thank the EPSRC for funding this PhD. I also thank Trinity College for being my home in Cambridge for seven years and Churchill College for welcoming me into its community.

I thank my friends from my time as an undergraduate and graduate student at Cambridge, including James Munro, Fiona Hamey, David Mestel, Jack Smith, Jo Evans, Tom Berrett, Lawrence Barrott, James Kilbane, Nigel Burke and Nicolas Dupré, for their conversation both mathematical and (especially) non-mathematical. Living with James Munro and Fiona Hamey has made writing this thesis a more pleasant experience, and I am particularly grateful to them.

I am also grateful to my friends outside the mathematical community, whose friendship has enriched my life and provided often-needed distraction from mathematics. I thank Cambridge Labour Party, and in particular George Owers, Dan Ratcliffe, Dan Swain and Ashley Walsh.

Finally I thank my family, and in particular my parents, to whom this thesis is dedicated. It is thanks to their early encouragement and support that I have reached this point.

Table of contents

1	Introduction	1
1.1	Detecting connectedness properties of the boundary	3
1.2	Computing splittings and JSJ decompositions	4
1.3	Decomposition spaces	8
2	Group actions on hyperbolic spaces	11
2.1	Word hyperbolicity	11
2.1.1	The word metric and the Švarc-Milnor Lemma	11
2.1.2	Hyperbolic geometry	13
2.1.3	Hyperbolic groups	16
2.2	The Gromov boundary	17
2.2.1	Definition and topology	18
2.2.2	The visual metric	21
2.2.3	Local connectedness	22
2.3	Relative hyperbolicity	22
2.3.1	A dynamical characterisation	23
2.3.2	The cusped space	24
2.3.3	The Bowditch boundary	26
3	Splittings and JSJ decompositions	27
3.1	Amalgams and HNNs	27
3.2	More general splittings	28
3.2.1	Graphs of groups	28
3.2.2	Graphs of spaces	29
3.2.3	Domination and refinement	31
3.3	Group actions on trees	32
3.4	Splittings over finite subgroups	34
3.5	JSJ decompositions	35

3.5.1	Guirardel and Levitt	36
3.5.2	Bowditch's canonical decomposition	37
3.5.3	Peripheral splittings	40
3.5.4	Cut pairs in the relative setting	41
3.6	Detecting splittings over finite groups	42
3.6.1	The double dagger condition	43
3.6.2	The computability of \ddagger	48
4	Cut points and cut pairs	49
4.1	Cylinders, cut points and cut pairs	50
4.1.1	Spaces and constants	50
4.1.2	Cylinders in the cusped space	51
4.2	Cylinders in the thin part of the cusped space	59
4.3	The geometry of finite balls in the cusped space	61
4.3.1	Cut points	62
4.3.2	Cut pairs	62
4.3.3	Circular boundaries	66
4.4	Generalisations	70
5	Algorithmic consequences	71
5.1	Detecting cut points, cut pairs and circular boundaries	72
5.1.1	Computing the constants	72
5.1.2	Algorithms	74
5.2	Computing splittings	75
5.2.1	Virtually cyclic subgroups	76
5.2.2	Enumerating the splittings	78
5.3	Splittings of groups with circular boundary	80
5.3.1	Groups with circular boundary	80
5.3.2	Orbifolds	82
5.3.3	Splittings of fundamental groups of orbifolds	85
5.3.4	Detecting whether or not an orbifold is small	86
5.4	Maximal splittings	87
5.4.1	Boundaries with cut pairs	88
5.4.2	Boundaries without cut points or pairs	90
5.4.3	Computing a maximal splitting	91
5.5	JSJ decompositions	92
5.5.1	Virtually cyclic edge groups and Bowditch's decomposition	93

Table of contents	xi
5.5.2 \mathcal{Z} edge groups	96
5.5.3 \mathcal{Z}_{\max} edge groups	98
5.6 Generalisations	99
6 The cohomology of decomposition spaces	101
6.1 Whitehead graphs and open covers	102
6.2 Computing $\check{H}^0(\mathcal{D})$	105
6.3 Computing $\check{H}^1(\mathcal{D})$	109
6.4 The cohomology of hyperbolic graphs of free groups	111
6.4.1 Boundaries of fundamental groups of graphs of groups	111
6.4.2 Computing the cohomology	113
Bibliography	119

Chapter 1

Introduction

Geometric group theory has its roots in the idea that one can learn about a group from the geometric properties of its Cayley graph, or more generally the geometry of other spaces on which the group admits a nice action. Fundamental to this idea is the Švarc-Milnor Lemma, which tells us that, given a group G acting nicely on a metric space X , the large scale geometry of X depends only on G , and so large scale properties of X can be powerful invariants of the group G . This raises the following question.

Question 1.0.1. *On how large a scale does one need to look at a space to see its large scale features? To what extent does the geometry of a ball of large but bounded radius in the Cayley graph of a group determine the large scale geometry of that group?*

In this thesis we answer this question for certain large scale properties of groups satisfying an important curvature condition, called Gromov hyperbolicity. This property was introduced by Gromov [Gro87] as a coarse notion of negative curvature. Groups that act on these spaces are well behaved in very many ways, and share many nice properties with free groups, but the definition is not narrow: many groups of classical interest—such as fundamental groups of hyperbolic manifolds—are hyperbolic, as are most finitely presented groups.

Here we are particularly interested in those geometric properties of a hyperbolic group that are derived from its Gromov boundary. See Section 2.2 for a definition of this space. This space can be thought as a limit of spheres in the Cayley graph of the group. Large spheres give good approximations to the boundary, and it is possible that, for a given question about the boundary, we might be able to say that some large sphere approximates the boundary sufficiently well to answer that question.

Question 1.0.2. *Given a property \mathcal{P} that the boundary of a hyperbolic group G might and might not possess, can we decide whether or not ∂G has \mathcal{P} by looking at a sphere in G of some predetermined size?*

Some important properties that the boundary of a hyperbolic group can possess are equivalent to abstract algebraic properties of the group, so that showing that such a property of the boundary is determined by a ball in the Cayley graph of predetermined size is equivalent to showing that such a finite ball determines an algebraic property of the group. The connectivity properties of the boundary studied in this thesis are closely connected to the structure of splittings of the group as amalgamated products and HNN extensions. Disconnectedness of the boundary is equivalent by a theorem of Stallings [Sta68, Sta71] to the existence of a splitting of the group over a finite subgroup. In the absence of splittings over finite subgroups, decompositions of the group over its virtually cyclic subgroups are governed by the structure of cut pairs in its boundary. (A *cut pair* in a connected topological space M is a pair of points p and q such that $M - \{p, q\}$ is disconnected.)

We shall also be interested in the property of having circular boundary. By the Convergence Group Theorem of Tukia, Gabai, Casson and Jungreis [Tuk88, Gab92, CJ94], a hyperbolic group has circular boundary if and only if it maps with finite kernel onto the fundamental group of a closed hyperbolic orbifold of dimension two.

Finally, by a theorem of Bestvina and Mess [BM91] the Čech cohomology of the boundary of a hyperbolic group is equal (modulo a shift in dimension) to the cohomology of that group with coefficients in the group ring. Čech cohomological properties are the final class of boundary properties studied in this thesis, now for a space closely related to the Gromov boundary of a hyperbolic group: Otal's decomposition space of a line pattern in a free group.

Answers to Questions 1.0.1 and 1.0.2 have implications for the computability theory of groups: if one can say in advance that a particular large scale geometric property would be visible at a certain scale then by looking at a ball in the Cayley graph of the group of that scale one can tell whether or not the group possesses that property. To do this one must be able to algorithmically construct arbitrarily large balls in the Cayley graphs of the groups under consideration. This is equivalent to requiring a solution to the word problem in the group. This requirement is no impediment to the results of this thesis, since the word problem is known to be solvable for hyperbolic groups.

This thesis is structured as follows. We begin by reviewing some preliminary material which will be required later in the thesis: in Chapters 2 and 3 we review some key properties

of hyperbolic groups and splittings of groups respectively. No part of either chapter is original, except for the proof of Proposition 3.6.9.

In Chapter 4 we answer Question 1.0.2 for some connectivity properties of the boundary of a hyperbolic group. This answer has applications to computability results: in Chapter 5 we show how to apply the results of Chapter 4 to prove the computability of certain canonical decompositions of hyperbolic groups. Most results of these two chapters appear in [Bar18].

In Chapter 6 we answer Question 1.0.1 for cohomological properties of a space closely related to the Gromov boundary of a hyperbolic group: the decomposition space associated to a line pattern in a free group. A line pattern is a collection of words in a free group; geometrically it can be thought of as a family of lines in the Cayley graph of the free group, which is a tree. The geometry of these lines is modelled locally by the Whitehead diagrams for the line pattern, and these diagrams take the role previously played by large balls in the Cayley graph. These results are taken from [Bar17]. As a corollary, we show that the Čech cohomology of the boundary of a hyperbolic graph of free groups with cyclic edge groups is computable.

1.1 Detecting connectedness properties of the boundary

In Chapter 4 we answer some questions about topological properties of the boundary of a hyperbolic group. In answering these questions Bestvina and Mess's [BM91] double dagger condition $\ddagger(n)$ is very useful. This is a geometric property of the Cayley graph of the group and is equivalent to a quantitative local connectedness property of the boundary. See Section 3.6.1 for a definition of the double dagger condition. The size of the sphere that one must look at to answer questions about connectivity properties of the boundary will depend on the value of the parameter n with which the double dagger condition holds.

The main geometric results of this thesis answer these types of questions. The following theorem will be proved in Chapter 4.

Theorem 1.1.1. *There are functions $N_1(\delta, n, k)$ and $N_2(\delta, n, k)$ such that, given a group G with a generating set of size k such that the Cayley graph with respect to that generating set is δ -hyperbolic and satisfies $\ddagger(n)$, we have the following equivalences.*

1. *There is a cut pair in ∂G if and only if the ball of radius N_1 in the Cayley graph of G with respect to the given generating set contains a particular feature, which we shall call a coarsely separating geodesic segment.*

2. Secondly, ∂G is homeomorphic to a circle if and only if the ball of radius N_2 in the Cayley graph of G with respect to the given generating set contains a second feature, which we shall call a coarsely non-separating geodesic segment.

1.2 Computing splittings and JSJ decompositions

As a corollary to Theorem 1.1.1, we obtain the following computability results.

Theorem 1.2.1. *There is an algorithm that takes as input a presentation for a hyperbolic group G and returns the answers to the following two questions.*

1. Does ∂G contain a cut pair?
2. Is ∂G homeomorphic to a circle?

After some precomputation to compute the constant δ with respect to which the Cayley graph is δ -hyperbolic and the constant n such that the Cayley graph satisfies $\ddagger(n)$, the algorithm of Theorem 1.2.1 only requires the computation of a ball in the Cayley graph of known size. Therefore (after precomputation of some constants associated to the group) we can give an explicit upper bound on the worst-case running time for these algorithms, although this upper bound is sufficiently high that these algorithms are infeasible in practice.

The topological properties of the boundary discussed in Theorem 1.1.1 have important implications for the algebraic properties of the group. First, recall that the Convergence Group Theorem of Tukia, Gabai, Casson and Jungreis [Tuk88, Gab92, CJ94] tells us that a hyperbolic group has circular boundary if and only if it maps with finite kernel onto the fundamental group of a closed hyperbolic orbifold of dimension two. Therefore, Theorem 1.2.1 shows that there is an algorithm that determines whether or not a given hyperbolic group is a two dimensional orbifold group. In other words, the class of virtually fuchsian groups is recursive in the class of all hyperbolic groups. Secondly, cut pairs in the boundary of a hyperbolic group are related by a theorem of Bowditch [Bow98b] to splittings of the group over its virtually cyclic subgroups.

When given a group, it is natural to try to cut it up into simpler pieces by splitting over its particularly easy to understand subgroups. By a *splitting* of a group G we mean a decomposition of G obtained by repeatedly taking amalgamated free products and HNN extensions. Such a splitting is recorded as a *graph of groups*: this is a graph with a group at each vertex and a group at each edge so that G is obtained by gluing the vertex groups together along the edge groups. We will usually want to place restrictions on the edge

groups: if \mathcal{A} is a class of subgroups of G then a splitting of G such that all edge groups are contained in \mathcal{A} is called a *splitting of G over \mathcal{A}* . See Section 3.2 for precise definitions.

In the first instance, one might try to split a given group over its finite subgroups. By a theorem of Dunwoody [Dun85] a finitely presented group admits a maximal splitting over finite subgroups, so that none of the vertex groups of the splitting admit further non-trivial splittings over finite subgroups. In a maximal splitting over finite subgroups one sees all possible splittings of that group over finite subgroups.

Furthermore, if the initial group is assumed to be hyperbolic then this maximal splitting can be computed from a presentation of the group. By Stallings's theorem a finitely generated group admits a non-trivial splitting if and only if it has more than one end. If the group is hyperbolic then it has more than one end if and only if its Gromov boundary is disconnected. Gerasimov [Ger] proved that there exists an algorithm that determines whether or not the boundary of a given hyperbolic group is connected. This algorithm is unpublished; see also [DG08]. Given a hyperbolic group with connected boundary, Gerasimov's algorithm detects this connectedness by computing a number n such that the Cayley graph of the group satisfies $\ddagger(n)$. Equipped with this algorithm and Stallings's theorem on ends of groups it is not difficult to compute a maximal decomposition of a given hyperbolic group over its finite subgroups. With Gerasimov's result in hand, we may restrict to the case of one-ended hyperbolic groups and consider splittings over virtually (infinite) cyclic groups.

When studying splittings over virtually cyclic groups one immediately encounters a difficulty that was not present for splittings over finite groups: splitting over some virtually cyclic subgroups can hide splittings over others, so a maximal splitting does not have to reveal all possible splittings of the group over virtually cyclic subgroups, as is the case for splittings over finite groups. A maximal splitting over virtually cyclic subgroups requires some choices to be made. This is unsatisfactory: in many applications a canonical decomposition associated to the group is vital.

The solution to this problem is inspired by the characteristic submanifold decomposition of Jaco, Shalen and Johannson [JS79, Joh79], in which the family of embedded tori along which a 3-manifold is cut is unique up to isotopy. Similar JSJ decompositions were introduced to group theory by Sela [Sel97] to answer questions about rigidity and the isomorphism problem for torsion-free hyperbolic groups. In [Bow98b] Bowditch developed a related type of decomposition for hyperbolic groups possibly with torsion. This decomposition is built from the structure of cut pairs in the boundary of the group and is therefore an automorphism invariant of the group; this property of Bowditch's decomposition was

used in Levitt's work [Lev05] on outer automorphism groups of one-ended hyperbolic groups. For other constructions of JSJ decompositions of groups see [RS97, DS99, FP06].

More recently, Guirardel and Levitt [GL17] have introduced a general framework for JSJ decompositions of groups. Their construction begins with a class \mathcal{A} of subgroups of the given group and defines a JSJ decomposition over that class of subgroups by formalising the idea that in a JSJ decomposition one should split along precisely those subgroups in \mathcal{A} that do not cross other splittings over elements in \mathcal{A} , so that the decomposition splits the group as much as possible without hiding any of the splittings of the group. Guirardel and Levitt show that such a JSJ decomposition exists in considerable generality. Those parts of a JSJ decomposition that admit further splittings are called *flexible*. Then a JSJ decomposition of the group over \mathcal{A} , together with an understanding of the splittings over \mathcal{A} of all flexible subgroups, gives a full description of the structure of the splittings of that group over \mathcal{A} .

An important consequence of Bowditch's construction is that if G is a one-ended hyperbolic group G such that ∂G is not homeomorphic to a circle, G admits a non-trivial splitting over a virtually cyclic subgroup if and only if its boundary contains a cut pair. Using this result and Theorem 1.2.1 it is possible to determine algorithmically whether or not a given one-ended hyperbolic group admits a non-trivial splitting over a virtually cyclic subgroup. Using this we prove the following theorem.

Theorem 1.2.2. *There is an algorithm that takes as input a presentation for a one-ended hyperbolic group G and returns the graph of groups associated to the three following JSJ decompositions:*

1. *A JSJ decomposition over virtually cyclic subgroups of G , which can be taken to be Bowditch's canonical decomposition for G .*
2. *A JSJ decomposition over virtually cyclic subgroups of G with infinite centre, which we shall call a \mathcal{Z} -JSJ.*
3. *A decomposition over maximal virtually cyclic subgroups of G with infinite centre that is universally elliptic over (not necessarily maximal) virtually cyclic subgroups of G and is maximal for domination in the class of such decompositions. We shall call this a \mathcal{Z}_{max} -JSJ.*

Many existing algorithms that compute related decompositions of groups make use of a variant due to Rips and Sela [RS95] of Makanin's algorithm [Mak82]. This is a procedure that determines whether or not a given system of equations in a hyperbolic group has a

solution. It is difficult to relate such an approach to the geometry of the group; in contrast to these methods, the proof of Theorem 1.2.2 is purely geometric.

In fact, the algorithm of Theorem 1.2.2 needs something a little stronger than Theorem 1.1.1. Given a splitting of a one-ended hyperbolic group over virtually cyclic subgroups, we need to be able to tell whether or not the group can be decomposed further. It is not enough to be able to tell whether or not each vertex group admits a splitting over a virtually cyclic subgroup: we need the splitting to be consistent with the splittings we have already taken, so we need to know whether a vertex group admits a splitting over a virtually cyclic subgroup that does not cross the incident edge groups. This difficulty is solved by replacing the Gromov boundary with a closely related object: the Bowditch boundary of a given group relative to a collection of virtually cyclic subgroups.

The main content of Theorem 1.2.2 is the computability of Bowditch's decomposition; it is shown in [DG11] to be closely related to the \mathcal{Z} -JSJ and \mathcal{Z}_{\max} -JSJ and can be converted into either algorithmically. The \mathcal{Z}_{\max} -JSJ is the decomposition shown to be computable by Dahmani and Guirardel [DG11].

The \mathcal{Z} -JSJ also plays a role in Dahmani and Guirardel's work, although they comment that their methods cannot compute this decomposition, since such a decomposition does not necessarily give rise to infinitely many distinct outer automorphisms of the group. For example, let G be a rigid hyperbolic group (such as the fundamental group of a closed hyperbolic 3-manifold) and let g be an element of G that is not a proper power. Let $k > 1$ and consider the group $G' = G *_{g=t^k} \langle t \rangle$ obtained by adjoining a k th root of g to G . In this case the \mathcal{Z}_{\max} decomposition computed by Dahmani and Guirardel is trivial while the \mathcal{Z}_{\max} -JSJ is not.

Like Gerasimov's algorithm, our approach uses the geometry of large balls in the Cayley graph. This is in contrast to existing algorithms computing JSJ decompositions over restricted families of virtually cyclic subgroups, most of which rely on Makanin's algorithm for solving equations in free groups.

In [DG11] Dahmani and Guirardel show that a canonical decomposition of a one-ended hyperbolic group over a particular family of virtually cyclic subgroups is computable; the family in question is the set of virtually cyclic subgroups with infinite centre that are maximal for inclusion among such subgroups. Crucial to this method is an algorithm that determines whether or not the outer automorphism group of such a group is infinite. If a group admits such a splitting then that splitting gives rise to an infinite set of distinct elements of the outer automorphism group that are analogous to Dehn twists in the mapping class group of a surface. The converse of this statement is a theorem of Paulin [Pau91] that is refined by Dahmani and Guirardel.

Dahmani and Guirardel comment that it is not known whether or not Bowditch's JSJ decomposition is computable. Their approach is not suitable to this problem: only central elements of the edge groups in a splitting contribute Dehn twists to the automorphism group, so it is quite possible for a group to admit a splitting over an infinite dihedral group, say, while having only a finite outer automorphism group; in this case the decomposition computed by Dahmani and Guirardel is trivial while Bowditch's JSJ decomposition is not. For examples of hyperbolic groups exhibiting this property see [MNS99].

In the absence of torsion, the JSJ decomposition of a hyperbolic group over its cyclic subgroups was shown to be computable by Dahmani and Touikan in [DT13]. Their result is based on Touikan's algorithm [Tou18], which determines whether or not a given one-ended hyperbolic group without 2-torsion splits acylindrically. Touikan's methods are based on application of the Rips machine.

It seems plausible that the techniques of this thesis might extend to the problem of detecting the presence of splittings of relatively hyperbolic groups with parabolic subgroups in some restricted class. In particular, it is natural to try to solve this problem for groups that are hyperbolic relative to finitely generated virtually nilpotent subgroups. Such groups arise as fundamental groups of complete finite volume Riemannian manifolds with pinched negative sectional curvature. However, it is of fundamental importance to the argument presented in this paper that the cusped space associated to the group satisfies Bestvina and Mess's double dagger condition, and we do not know under what circumstances this condition holds for virtually nilpotent parabolic subgroups. In [DG08] it is established that the double dagger condition holds if the parabolic subgroups are abelian, so it seems likely that the methods of this paper could be extended to the case of toral relatively hyperbolic groups. (A group is toral relatively hyperbolic if it is torsion free and hyperbolic relative to a finite collection of abelian subgroups.) However, the JSJ decomposition of a toral relatively hyperbolic group is shown to be computable in [DT13], so we do not introduce additional technical complexity by trying to give a new proof of this result here.

1.3 Decomposition spaces

In the final chapter, we turn our attention to a more complicated topological property of the boundary of a hyperbolic metric space: its Čech cohomology. By a theorem of Bestvina and Mess, for any ring R there is an isomorphism of G -modules $\check{H}_r^k(\partial G, R) \cong H^{k+1}(G, RG)$ between the reduced Čech cohomology of ∂G and the group cohomology of G with coefficients in the group ring RG . Epstein asked whether or not there is an algorithm that

computes these cohomology groups. This question is listed as Question 1.18 in Bestvina's problem list [Bes04].

In Chapter 6 we prove the computability of the Čech cohomology of a space related to the Gromov boundary: Otal's decomposition space. A line pattern in a free group F is the set of translates of a finite collection of bi-infinite lines in the Cayley graph of F . It is determined by a finite collection of cyclic subgroups; if these cyclic subgroups are assumed to be maximal and non-conjugate then Otal's decomposition space is equal to the Bowditch boundary of F relative to those cyclic subgroups.

Using the computability of the Čech cohomology of the decomposition space of a free group, we then show how to compute the Čech cohomology of the Gromov boundary for a class of hyperbolic groups that can be built out of free groups: fundamental groups of graphs of groups with free vertex groups and cyclic edge groups. Therefore we answer Epstein's question in the affirmative for this class of groups.

Theorem 1.3.1. *There is an algorithm that takes as input a presentation for a hyperbolic group G that is the fundamental group of a graph of free groups with cyclic edge groups and computes the Čech cohomology of the Gromov boundary of G as a G -module.*

Chapter 2

Group actions on hyperbolic spaces

In this chapter, we summarise some aspects of the theory of group actions on hyperbolic metric spaces. We begin by describing *hyperbolic groups*: these are the groups that admit the nicest possible action on a hyperbolic space, and are the main subject of this thesis. We shall then loosen our restriction on the types of action that we consider to those actions satisfying a geometric finiteness condition. The groups that admit such an action are called *relatively hyperbolic*.

Most of the ideas outlined in this chapter are originally due to Gromov [Gro87]. For additional details see [CDP90, GdlH90, BH99].

2.1 Word hyperbolicity

2.1.1 The word metric and the Švarc-Milnor Lemma

Begin by fixing a group G with a finite generating set S . Let $\Gamma(G, S)$ be the Cayley graph of G with respect to S ; this can be made into a metric space by isometrically identifying each edge of $\Gamma(G, S)$ with a copy of the interval $[0, 1]$ and endowing $\Gamma(G, S)$ with the path metric. The induced metric on the vertex set $G \subset \Gamma(G, S)$ is called the *word metric* and is denoted d_S . Then $d_S(g, h) = \min\{n \mid h = g s_1^{\pm 1} \dots s_n^{\pm 1}, s_i \in S\}$. Crucially, note that if S' is another finite generating set for G then the metrics d_S and $d_{S'}$ are closely related: they are Lipschitz equivalent with constant equal to the maximum length of any element of S or S' when written in the alphabet S' or S respectively.

We now extend this independence of generating set to the metric on the Cayley graph, and further to the geometry of a large class of metric spaces on which G acts, so that geometric properties of metric spaces admitting G -actions becomes an intrinsic property

G itself. To do this we loosen our idea of metric equivalence from Lipschitz equivalence to “Lipschitz equivalence up to an additive error”.

Definition 2.1.1. Let (X, d_X) and (Y, d_Y) be metric spaces, and let $f : X \rightarrow Y$ be a map. (Note that f is not assumed to be continuous.) Then f is a λ -*quasi-isometric embedding* for $\lambda \geq 1$ if

$$\frac{1}{\lambda}d_X(x_1, x_2) - \lambda \leq d_Y(f(x_1), f(x_2)) \leq \lambda d_X(x_1, x_2) + \lambda. \quad (2.1)$$

Furthermore, f is λ -*quasi-surjective* if for every $y \in Y$, $d_Y(y, f(X)) \leq \lambda$.

Finally, f is a *quasi-isometry* (and X and Y are *quasi-isometric*) if it is a λ -quasi-surjective λ -quasi-isometric embedding for some λ .

Lemma 2.1.2. *Quasi-isometry is an equivalence relation on metric spaces.*

Immediately we see that (G, d_S) is quasi-isometric (with $\lambda = 1/2$) to the Cayley graph $\Gamma(G, S)$; it follows that $\Gamma(G, S)$ and $\Gamma(G, S')$ are quasi-isometric for any finite generating sets S and S' for G . Indeed, we shall soon see that much more is true: any two sufficiently nice metric spaces on which G acts sufficiently nicely must be quasi-isometric.

Spaces that are quasi-isometric look similar on a sufficiently large scale, and we expect large scale metric properties of a space to be invariant under quasi-isometry.


Definition 2.1.3. Let (X, d_X) be a metric space. Then X is *proper* if all closed balls in X are compact, or, equivalently, if X is complete and locally compact.

We call X a *geodesic space* if for any pair of points x_1 and x_2 in X , there is an isometric embedding $\gamma : [0, d_X(x_1, x_2)] \rightarrow X$ with $\gamma(0) = x_1$ and $\gamma(d_X(x_1, x_2)) = x_2$. In this case we call γ a *geodesic segment*. An isometric embedding $\gamma : [0, \infty) \rightarrow X$ is called a *geodesic ray* and an isometric embedding $\gamma : (-\infty, \infty) \rightarrow X$ is called a *bi-infinite geodesic*.

Definition 2.1.4. An action of a group G on a topological space X is called *cocompact* if $G \backslash X$ is compact. An action is *properly discontinuous* if for any compact subset $K \subset X$, the set $\{g \in G \mid g \cdot K \cap K \neq \emptyset\}$ is finite. If X is a metric space, then an action of G on X by isometries is *geometric* if it is cocompact and properly discontinuous.

Then $\Gamma(G, S)$ is a proper geodesic space and the action of G on $\Gamma(G, S)$ by left multiplication is geometric.

Remark 2.1.5. Note that geodesics are not required to be unique when they exist. However, we shall often ignore this issue and denote by $[x_1, x_2]$ the image of some geodesic segment from x_1 to x_2 .

Figure 2.1: A δ -slim triangle.

Theorem 2.1.6 (The Švarc-Milnor Lemma). *Let a group G act geometrically on a proper geodesic space (X, d_X) . Then G has a finite generating set S and for any $v \in X$ the orbit map $g \mapsto g \cdot v$ from (G, d_S) to (X, d_X) is a quasi-isometry.*

It follows that any two proper geodesic spaces on which G acts geometrically are quasi-isometric, and the large scale geometric properties of any such space are properties intrinsic to G itself.

2.1.2 Hyperbolic geometry

We now define a large scale negative curvature property called (*Gromov*) *hyperbolicity*.

Definition 2.1.7. [Gro87] Let (X, d_X) be a metric space. A *geodesic triangle* in X is a triangle in X , the edges of which are geodesics. We shall sometimes denote by $\Delta(x_1, x_2, x_3)$ some geodesic triangle with vertices x_1 , x_2 and x_3 .


For $\delta \geq 0$, a geodesic triangle is said to be δ -*slim* if each of its three edges is contained in the union of its other two edges. See Figure 2.1.

The space X is δ -*hyperbolic* (or just *hyperbolic* if the constant δ is not important) if every geodesic triangle in X is δ -slim.

Example 2.1.8. 1. All metric trees are 0-hyperbolic: geodesic triangles in trees are tripods, and each of the edges of a triangle is contained in the union of the other two.

2. For any n , the hyperbolic space \mathbb{H}^n is $\log(1 + \sqrt{2})$ -hyperbolic.

This property of metric spaces is quasi-isometry invariant. This fact follows from the Morse property of geodesics in a hyperbolic metric space. See Theorem 2.1.16 below.

Figure 2.2: A δ -thin triangle

Lemma 2.1.9. *Let (X, d_X) and (Y, d_Y) be metric spaces and suppose that Y is δ -hyperbolic. Suppose also that $f : X \rightarrow Y$ is a λ -quasi-isometric embedding. Then X is δ' hyperbolic, where δ' depends only on δ and λ .*

Alternative formulations of hyperbolicity

There are very many equivalent definitions of hyperbolicity. In this section we recall two that will be used later on. The first is a condition on triangles similar to that of Definition 2.1.7.

Definition 2.1.10. Let X be a geodesic metric space and let x, y and z be points in X . We say that the triangle $\Delta(x, y, z)$ is δ -thin if there exist points i_x, i_y and i_z on edges $[y, z]$, $[z, x]$ and $[x, y]$ respectively such that $\{i_x, i_y, i_z\}$ has diameter at most δ . See Figure 2.2.

Since the three points i_x, i_y and i_z all lie close to all three edges of $\Delta(x, y, z)$, any one of them can be thought of as a “coarse incentre” for the triangle. Then a triangle is δ -thin if its “coarse inradius” is at most δ .

Lemma 2.1.11. *Let X be a geodesic metric space. If every geodesic triangle in X is δ -slim then every geodesic triangle is 4δ -thin. If every geodesic triangle is δ -thin then every geodesic triangle is δ -slim.*

This lemma shows that we could equivalently have defined a hyperbolic metric space to be a geodesic space in which all triangles are uniformly thin.

Secondly, we give a formulation of hyperbolicity in terms of a quantity called the Gromov product. This quantity will be used in Section 2.2.2 to define the visual metric on the boundary of a hyperbolic space.

Definition 2.1.12. Let X be a metric space containing a base point v and let x and y be points in X . Then the *Gromov product* of x and y with respect to v is

$$(x \cdot y)_v = \frac{1}{2}(d(v, x) + d(v, y) - d(x, y)) \quad (2.2)$$

For intuition on this quantity, note that if X is a tree then $(x \cdot y)_v$ is the distance from v to the (unique) geodesic $[x, y]$.

Definition 2.1.13. Let X be a metric space and let v be a base point in X . Then X is δ -*hyperbolic* if for any x, y and z in X the following inequality holds.

$$(x \cdot y)_v \geq \min\{(x \cdot z)_v, (z \cdot y)_v\} - \delta \quad (2.3)$$

Proposition 2.1.14. *A geodesic metric space is hyperbolic in the sense of Definition 2.1.7 if and only if it is hyperbolic in the sense of Definition 2.1.13, although the constant δ might be different.*

Note that Definition 2.1.13 is applicable even when X is not assumed to be a geodesic space, although we will not make use of this greater generality in this thesis.

The Morse property

Geodesics in hyperbolic metric spaces satisfy a useful rigidity property.

Definition 2.1.15. Let X be a metric space and let $I \subset \mathbb{R}$ be an interval. Let $\gamma: I \rightarrow X$ be a (possibly discontinuous) map. Then we call γ a λ -*quasi-geodesic* for $\lambda \geq 1$ if it is a λ -quasi-isometric embedding. We call γ a k -*local-geodesic* for $k > 0$ if the restriction to γ to any subinterval of I of length at most k is a geodesic.

In a hyperbolic space both of these weakenings of the definition of a geodesic do not substantially enlarge the class of paths under consideration.

Theorem 2.1.16 (The Morse property for geodesics). *[BH99, Theorem III.H.1.7] Let X be a δ -hyperbolic metric space and let $\lambda \geq 1$. Then there exists a constant $R(\delta, \lambda)$ depending only on δ and λ such that for any λ -quasi-geodesic γ there is a geodesic γ' such that the Hausdorff distance $d_{\text{Haus}}(\gamma, \gamma')$ between γ and γ' is at most R . Furthermore, γ can be taken to have the same end points as γ' .*

Theorem 2.1.17. *[BH99, Theorem III.H.1.13] Let X be a δ -hyperbolic metric space and let $k > 8\delta$. Suppose that $\gamma: [0, l] \rightarrow X$ is a k -local-geodesic. Then γ is a $\max\{2\delta, (k + 4\delta)/(k - 4\delta)\}$ -quasi-geodesic and the Hausdorff distance $d_{\text{Haus}}(\gamma, [\gamma(0), \gamma(l)])$ is at most 3δ .*

2.1.3 Hyperbolic groups

As a consequence of the quasi-isometry invariance of hyperbolicity (Lemma 2.1.9), together with the Švarc-Milnor lemma (Theorem 2.1.6), hyperbolicity is a group-theoretic property.

Definition 2.1.18. [Gro87] A group G is (*Gromov*) *hyperbolic* if some (equivalently any) proper geodesic metric space on which G acts geometrically is hyperbolic.

Example 2.1.19. 1. Any finitely generated free group is hyperbolic: such a group acts geometrically on a tree, which is a 0-hyperbolic space.

2. If M is a closed hyperbolic n -manifold then $\pi_1 M$ acts geometrically on \mathbb{H}^n , so M is hyperbolic.

3. More generally, the fundamental group of a compact Riemannian manifold with strictly negative sectional curvature is hyperbolic.

4. Let $\langle S \mid R \rangle$ be a finite presentation such that for any $r \in R$ all cyclic conjugates of r and r^{-1} are contained in R . If there exist $r_1 \neq r_2$ in R such that the words r_1 and r_2 share a common prefix a then we call a a *piece*. If $\lambda > 0$ is such that for any $r \in R$ with prefix a such that a is a piece we have the inequality $|a| < \lambda|r|$ then we say that the presentation satisfies the *small cancellation condition* $C'(\lambda)$. Then any group admitting a $C'(1/6)$ presentation is hyperbolic. For more details see [LS77].

5. As well as the above groups of classical interest, there is a sense in which most finitely presented groups are hyperbolic. This idea is formalised by the notion of a *random group*; see [Gro03]. In particular, a random group in either the few relator model or the density model at density $< 1/2$ is infinite hyperbolic [Gro87, Gro92].

Subgroups of hyperbolic groups

Subgroups of hyperbolic groups can in general be complicated. However, those subgroups whose geometry is in some way consistent with the geometry of the larger group are far better behaved. Here we introduce Gromov's notion of *quasi-convexity* [Gro87]. For a more detailed introduction to the subgroup structure of hyperbolic groups see [BH99].

Definition 2.1.20. Let X be a geodesic metric space. A subset $A \subset X$ is C -*quasi-convex* (or just *quasi-convex*) if any geodesic in X joining two points in A is contained in a C -neighbourhood of A .

Lemma 2.1.21. [BH99, III.H.3.5–III.H.3.6] *Let G be a hyperbolic group and let $H \leq G$. If H is a quasi-convex subset of $\Gamma(G, S)$ for some (equivalently any) finite generating set for G then H is finitely generated and $H \hookrightarrow G$ is a quasi-isometric embedding with respect to the word metrics.*

Conversely, if H is finitely generated and $H \hookrightarrow G$ is a quasi-isometric embedding with respect to the word metrics then H is a quasi-convex subset of the Cayley graph of G .

From Lemma 2.1.21 and Lemma 2.1.9 we obtain the following corollary.

Corollary 2.1.22. *A quasi-convex subgroup of a hyperbolic group is hyperbolic.*

Note the following proposition.

Proposition 2.1.23. [Gro87] *Any virtually cyclic subgroup of a hyperbolic group is quasi-convex.*

Virtually cyclic subgroups play an additional important role in the subgroup structure of hyperbolic groups, as the following theorem shows.

Theorem 2.1.24. [Gro87] *Let G be a hyperbolic group. Then any subgroup of G is either virtually cyclic or contains a free subgroup of rank 2.*

We call virtually cyclic subgroups of a hyperbolic group *elementary*.

2.2 The Gromov boundary

The space of ends of a graph measures its “connected components at infinity.”

Definition 2.2.1. Let Γ be a locally finite graph. We define the *space of ends* of Γ to be the following inverse limit of finite discrete topological spaces.

$$\text{Ends}(\Gamma) = \varprojlim_K \pi_0(\Gamma - K) \quad (2.4)$$

Here the limit is taken over finite subgraphs of Γ partially ordered by inclusion.

By applying this construction to the Cayley graph of a finitely generated group with respect to a finite generating set we obtain the space of ends of the group: this space is independent of the choice of finite generating set, so we let $\text{Ends}(G) = \text{Ends}(\Gamma(G, S))$ for some finite generating set S .

In the case of hyperbolic groups we can give the ends additional structure: we can make each end into a topological space in a natural way. By viewing \mathbb{H}^n in the Poincaré

ball model, we immediately see that it admits a compactification by a “sphere at infinity”. Hyperbolic spaces admit an analogous compactification by a boundary. This boundary is a large scale property of the space. It is quasi-isometry invariant, and can therefore be thought of as a group invariant. See [KB02] for a more complete overview of the subject.

2.2.1 Definition and topology

Definition 2.2.2. [Gro87] Let X be a hyperbolic metric space. Then the *Gromov boundary* of X is defined to be the following quotient.

$$\partial X = \{\text{geodesic rays}\} / \sim \quad (2.5)$$

where $\gamma_1 \sim \gamma_2$ if the Hausdorff distance from γ_1 to γ_2 is finite. For a geodesic ray γ we denote by $\gamma(\infty)$ the equivalence class of γ in $\partial_\nu X$. The boundary ∂X is endowed with the quotient of the topology of uniform convergence.

It is clear that any group of isometries of a hyperbolic metric space acts naturally on the boundary of that space.

Remark 2.2.3. In view of Theorem 2.1.16, in Definition 2.2.2 we could equivalently have defined ∂X to be a set of equivalence classes of quasi-geodesics.

Proposition 2.2.4. *Let X be a hyperbolic metric space and let $\nu \in X$. Then the natural inclusion*

$$\{\text{geodesic rays } \gamma \mid \gamma(0) = \nu\} / \sim \rightarrow \partial X \quad (2.6)$$

is a homeomorphism. Here $\gamma_1 \sim \gamma_2$ if the Hausdorff distance from γ_1 to γ_2 is finite and the topology is the quotient of the topology of uniform convergence, as in Definition 2.2.2.

Proposition 2.2.4 gives an alternative definition for the boundary of a hyperbolic space, which will usually be more convenient for our purposes.

Remark 2.2.5. The equivalence relation on rays in Proposition 2.2.4 may be replaced by uniformly bounded distance. Let X be δ -hyperbolic, then for geodesic rays γ_1 and γ_2 with $\gamma_1(\infty) = \gamma_2(\infty)$, the geodesic triangle $\Delta(1, \gamma_1(t), \gamma_2(t))$ is δ -slim for all t , from which it follows that $d(\gamma_1(t), \gamma_2(t)) < 2\delta$ for all t .

We claim that the boundary ∂X compactifies a hyperbolic metric space X . We now give two descriptions the topology on the union $X \cup \partial X$, first in terms of convergence of sequences and then in terms of fundamental neighbourhoods.

Definition 2.2.6. Let X be a hyperbolic metric space and let $x \in X \cup \partial X$. Let (x_i) be a sequence of points in $X \cup \partial X$; we define a topology on $X \cup \partial X$ by giving a condition to determine whether or not x_i converges to x .

1. If $x \in X$ then $\lim x_i = x$ if and only if $x_i \in X$ for i sufficiently large and the sequence then converges to x with respect to the metric on X .
2. If $x \in \partial X$ then we say that $\lim x_i = x$ if and only if there exists for each i either a geodesic segment γ_i with endpoints v and x_i (in the case when $x_i \in X$) or a geodesic ray γ_i with $\gamma_i(0) = v$ and $\gamma_i(\infty) = x_i$ (in the case when $x_i \in \partial X$) such that every subsequence of γ_i admits a subsequence that converges uniformly on compact subsets to a geodesic ray γ with $\gamma(\infty) = x$.

Alternatively, we can explicitly describe a fundamental system of neighbourhoods for points in $X \cup \partial X$.

Definition 2.2.7. Fix $k > 2\delta$. Then consider the following set of subsets of $X \cup \partial X$.

1. For each point $x \in X$ and each $r > 0$, the open ball in X with centre x and radius r .
2. For each point $\gamma(\infty) \in \partial X$ and each $n \geq 0$, the set $V_n(\gamma)$ of equivalence classes of geodesic rays and end points of geodesic segments γ' of length at least n such that $\gamma'(0) = v$ and $d(\gamma'(n), \gamma(n)) < k$.

This set of subsets of $X \cup \partial X$ is a fundamental system of (not necessarily open) subsets of $X \cup \partial X$, and therefore defines a topology.

Proposition 2.2.8. [BH99, III.H.3.6] *The topologies on $X \cup \partial X$ given in Definitions 2.2.6 and 2.2.7 are equal. This topology is independent of base point. The inclusion $X \hookrightarrow X \cup \partial X$ is a homeomorphism onto its image and $X \subset X \cup \partial X$ is open. Finally, $X \cup \partial X$ is compact.*

Proposition 2.2.9. *Let $f: X \rightarrow Y$ be a quasi-isometric embedding of hyperbolic metric spaces. Then the map $f_\partial: \partial X \rightarrow \partial Y$ defined by $f_\partial(\gamma(\infty)) = (f \circ \gamma)(\infty)$ is well defined and continuous. (Note that $f \circ \gamma$ is a quasi-geodesic. See Remark 2.2.3.) If f is a quasi-isometry then f_∂ is a homeomorphism.*

In view of Proposition 2.2.9 and Theorem 2.1.6 we may make the following definition: the boundary of a space on which a hyperbolic group acts is an intrinsic property of the group.

Definition 2.2.10. Let G be a hyperbolic group and let X be proper geodesic metric space on which G acts geometrically. Then we define $\partial G = \partial X$.

It is easy to prove the following proposition relating the boundary of a hyperbolic group to its ends.

Proposition 2.2.11. *Let S be a finite generating set for a hyperbolic group G . Then the map that sends a the equivalence class of a geodesic ray to the end of $\Gamma(G, S)$ containing that ray gives a continuous surjection $\partial G \rightarrow \text{Ends}(G)$ and the preimage of any point is a connected component of ∂G . In particular, the number of ends of G is equal to the number of connected components of ∂G .*

Example 2.2.12. 1. Let F be a free group of rank n . Then F acts geometrically on a $2n$ -regular tree, so ∂F is a Cantor set.

2. If M is a closed hyperbolic n -manifold then $\pi_1 M$ acts geometrically on \mathbb{H}^n , so $\partial \pi_1 M = \partial \mathbb{H}^n \cong S^{n-1}$.
3. More generally, if M is a compact Riemannian manifold with strictly negative sectional curvature then $\partial \pi_1 M \cong S^{n-1}$.
4. In contrast to these groups with spherical boundary, most boundaries of hyperbolic groups are quite complicated topological spaces: Champetier showed [Cha95] that most small cancellation groups are hyperbolic and have boundary homeomorphic to a Menger curve.

For an arbitrary group acting on a hyperbolic space we define the following set.

Definition 2.2.13. Let G be a group acting by isometries on a hyperbolic metric space X . Let $x \in X$. Then the *limit set* ΛG of G is defined as follows.

$$\Lambda G = \{p \in \partial X \mid p = \lim_{n \rightarrow \infty} g_n \cdot x \text{ for some sequence } (g_n) \text{ in } G\} \quad (2.7)$$

It is easy to see that this set is independent of the choice of x .

The limit set is particularly well behaved when G is a quasi-convex subgroup of a group acting geometrically on X .

Theorem 2.2.14. [GdlH90] *Let G be a hyperbolic group with generating set S and let H be a quasi-convex subgroup of G . Then the inclusion $H \hookrightarrow \Gamma(G, S)$ extends to a topological embedding $\partial H \hookrightarrow \partial G$ and the image of this embedding is ΛH .*

2.2.2 The visual metric

The boundary of a hyperbolic space has much additional structure beyond its topology. In particular, it admits a natural family of metrics called visual metrics. These measure something like the angle between geodesics from a fixed base point.

To begin, recall the equivalent formulation of hyperbolicity through Gromov products 2.1.13: fixing a base point $\nu \in X$, X is hyperbolic if and only if there exists δ' such that $(x \cdot y)_\nu \geq \min\{(x \cdot z)_\nu, (z \cdot y)_\nu\} - \delta'$ for any points x, y and z in X . It follows that for $a > 1$, the quantity $a^{-(x \cdot y)_\nu}$ is bounded above by a quantity comparable to the maximum of $a^{-(x \cdot z)_\nu}$ and $a^{-(z \cdot y)_\nu}$. Therefore one might hope to use $a^{-(x \cdot y)_\nu}$ as a measure of the angle between x and y with respect to the base point ν , at least when x and y are both far from ν . Some modification is required, but such a metric can be built.

First we must extend the definition of the Gromov product to points on ∂X .

Definition 2.2.15. Let x and y be points in ∂X and let ν be a base point in X . Then we define the Gromov product of x and y with respect to base point ν as follows.

$$(x \cdot y)_\nu = \sup \liminf_{i, j \rightarrow \infty} (x_i \cdot y_j)_\nu. \quad (2.8)$$

Here the supremum is taken over sequences (x_i) and (y_i) in X converging to x and y respectively.

In calculations the following inequality is often useful.

Lemma 2.2.16. [GdlH90, Remarque 8.9] If X is δ -hyperbolic and (x_i) and (y_i) are sequences in X converging to x and y respectively then

$$\liminf_{i, j \rightarrow \infty} (x_i \cdot y_j)_\nu \leq (x \cdot y)_\nu \leq \liminf_{i, j \rightarrow \infty} (x_i \cdot y_j)_\nu + 2\delta. \quad (2.9)$$

Proposition 2.2.17. [GdlH90, Section 7.3] Let X be a hyperbolic metric space. Then there exists parameter $a > 1$ and $0 < k_1 \leq k_2$ and a metric d_ν on ∂X inducing the topology of Definition 2.2.2 and satisfying the following inequality.

$$k_1 a^{-(x \cdot y)_\nu} \leq d_\nu(x, y) \leq k_2 a^{-(x \cdot y)_\nu}. \quad (2.10)$$

Such a metric is called a visual metric on ∂X and a is the visual parameter of the metric.

Remark 2.2.18. A visual metric on the boundary of a hyperbolic space is not unique, and there is usually no canonical choice. Nor is there a canonical choice of visual parameter. There is, however, a canonical class of metrics, called a *quasi-conformal structure* on ∂X [Pan89].

2.2.3 Local connectedness

Boundaries of hyperbolic groups are compact metric spaces, but as we saw in Example 2.2.12, they can be topologically complicated: for example, they are not generally homotopy equivalent to CW complexes. However, they do possess useful connectivity properties, which we will use later. Here we are particularly interested in the case in which ∂G is connected.

Definition 2.2.19. A *continuum* is a compact connected metric space. A *Peano continuum* is a continuum that is locally connected at each point. That is, it admits a topological basis consisting only of connected sets.

Theorem 2.2.20. *Let G be a hyperbolic group. If ∂G is connected then ∂G is a Peano continuum.*

The substance of Theorem 2.2.20 is that the boundary is locally connected. This deep result was first proved by Bestvina and Mess under the assumption that the boundary does not contain a cut point. (Recall that a cut point in a connected topological space is a point whose complement is disconnected.)

Proposition 2.2.21. *[BM91] Let G be a hyperbolic group. If ∂G is connected and does not contain a cut point then it is locally connected.*

We shall return to this result in Section 3.6.1

The problem of whether or not the boundary of a hyperbolic group can contain a cut point remained open for several years, before being finally settled in full generality by Swarup.

Theorem 2.2.22. *[Swa96] Let G be a hyperbolic group such that ∂G is connected. Then ∂G does not contain a cut point.*

This theorem was originally proved by Bowditch [Bow98a] under the additional assumption that G is strongly accessible. Swarup demonstrated in [Swa96] how to extend these arguments to the general case.

2.3 Relative hyperbolicity

Many of the appealing properties of closed hyperbolic manifolds 3-manifolds are shared by complete hyperbolic manifolds with cusps. Motivated by this, relatively hyperbolic groups are defined to be groups that admit actions on hyperbolic spaces that are not

necessarily cocompact, but satisfy a weaker “geometric finiteness” condition. The theory of relatively hyperbolic groups was introduced by Gromov [Gro87]. Gromov’s concept was later expanded on by Bowditch [Bow12] and Farb [Far98]. Farb’s definition is weaker than Bowditch’s, and today Bowditch’s more restrictive definition is more common. Groups that are relatively hyperbolic in the weaker sense of Farb may be termed *weakly relatively hyperbolic*. In this thesis we adopt the stricter definition.

2.3.1 A dynamical characterisation

There are now many equivalent formulations of relative hyperbolicity. Bowditch [Bow98a] gives a dynamical characterisation, which we adopt as a definition. This characterisation makes the connection with geometrically finite Kleinian groups particularly clear and has its origin in work of Beardon and Maskit [BM74] on limit sets of Kleinian groups.

Given a Kleinian group $G \leq \text{PSL}(2, \mathbb{C})$, define the *limit set* $\Lambda G \subset S^2$ of G to be the set of accumulation points in $S^2 = \partial\mathbb{H}^3$ of an orbit in \mathbb{H}^3 of G . Unless G is finite or virtually cyclic, this is the unique non-empty minimal G -invariant closed subset of S^2 . Beardon and Maskit give a characterisation of geometric finiteness in terms of the action of G on ΛG : G is geometrically finite if every point in ΛG is either a *conical limit point* (roughly, a point approximated sufficiently well by translates of a point in \mathbb{H}^3 , such as a fixed point of a loxodromic element of G) or a *bounded parabolic point* (a point x stabilised by a purely parabolic subgroup of G such that the quotient of $\Lambda G - \{x\}$ by that subgroup is compact.) This result gives a plausible generalisation of the notion of geometric finiteness to group actions on other hyperbolic spaces, at least as long as we can generalise the definitions of conical limit points and bounded parabolic points.

Definition 2.3.1. Let G be a group acting by homeomorphisms on a compact metric space M . Then a point $x \in M$ is:

1. a *conical limit point* if there exist p and q in M and a sequence (g_i) in G such that $g_i \cdot x \rightarrow p$ and $g_i \cdot y \rightarrow q$ for every $y \in M - \{x\}$.
2. a *bounded parabolic point* if the action of $\text{Stab}_G(x)$ on $M - \{x\}$ is cocompact.

Definition 2.3.2. Let G be a group and let \mathcal{H} be a finite collection of finitely generated subgroups of G . Suppose that G acts properly discontinuously and by isometries on a proper hyperbolic metric space X such that ∂X does not contain a closed non-empty G -invariant proper subset and every point in ∂X is either a conical limit point or a bounded parabolic point. Assume further that the stabiliser of each bounded parabolic point is

conjugate to an element of \mathcal{H} and that every element of \mathcal{H} arises in this way. Then G is *hyperbolic relative to \mathcal{H}* .

We call elements of \mathcal{H} *peripheral subgroups*. A subgroup that stabilises a parabolic fixed point is called a *parabolic subgroup*, so the maximal parabolic subgroups are precisely the conjugates of the peripheral subgroups.

From the discussion at the beginning of this section we see that this definition generalises that of a geometrically finite Kleinian group. It also generalises the definition of a hyperbolic group: a group is hyperbolic if and only if it is hyperbolic relative to \emptyset .

Remark 2.3.3. Without the assumption of cocompactness the Švarc-Milnor lemma (Theorem 2.1.6) does not apply, and the quasi-isometry type of the space X in Definition 2.3.2 is not determined by the pair (G, \mathcal{H}) .

Remark 2.3.4. It is possible to phrase this definition purely in terms of the action of G on a compact metric space without the assumption that this action arises as the boundary of a properly discontinuous action on a hyperbolic metric space. If one assumes that the action satisfies the technical condition of being a convergence group action, then the equivalence of these formulations of relative hyperbolicity is a theorem of Bowditch [Bow98c] in the case when $\mathcal{H} = \emptyset$ and of Yaman [Yam04] in the general case.

As mentioned at the beginning of this section, Definition 2.3.2 is stronger than Farb's definition. It is equivalent to Farb's condition "relatively hyperbolic with bounded coset penetration"; for details see [Bow98a].

When G is hyperbolic and \mathcal{H} is a finite collection of finitely generated subgroups of G , Bowditch [Bow12] gives a simple condition to decide whether or not G is hyperbolic relative to \mathcal{H} . Recall that a collection \mathcal{H} of subgroups of G is *almost malnormal* if, for H_1 and H_2 in \mathcal{H} and g in G , $H_1 \cap gH_2g^{-1}$ is finite unless $H_1 = H_2$ and $g \in H_1$.

Theorem 2.3.5. [Bow12, Theorem 7.11] *Let G be a non-elementary hyperbolic group and let \mathcal{H} be a finite set of subgroups of G . Then G is hyperbolic relative to \mathcal{H} if and only if \mathcal{H} is almost malnormal in G and each element of \mathcal{H} is quasi-convex in G .*

2.3.2 The cusped space

A hyperbolic group comes with an easily described geometric model: its Cayley graph with respect to any finite generating set. Definition 2.3.2 does not immediately give a corresponding model for a relatively hyperbolic group. We now introduce the Groves-Manning cusped space [GM08] as a geometric model for a relatively hyperbolic group. This space is a graph that will take the role of the Cayley graph of a hyperbolic group. Like

the Cayley graph of a hyperbolic group, this is a hyperbolic metric space. It is obtained by gluing to the Cayley graph a hyperbolic space along each coset of a peripheral subgroup. These hyperbolic spaces are called combinatorial horoballs and they are designed to look like horoballs in \mathbb{H}^n . This construction is analogous to the thick-thin decomposition of a hyperbolic manifold: the Cayley graph is the thick part of the cusped space, and the group acts cocompactly on this set, while the cusps are the thin part.

To begin, we define the combinatorial horoball; for more details of this construction see [GM08].

Definition 2.3.6. Let C be a graph. Then the *combinatorial horoball based on C* is the graph $\text{Hor}(C)$ with the following description.

The vertex set of $\text{Hor}(C)$ is $C^{(0)} \times \mathbb{Z}_{\geq 0}$.

Each edge of $\text{Hor}(C)$ takes one of the following three forms.

1. One *horizontal edge* from $(v, 0)$ to $(w, 0)$ for each edge from v to w in C . (Note that we allow vertices to be connected by multiple edges.)
2. One *horizontal edge* from (v, k) to (w, k) whenever v and w are connected in C by an edge path of length at most 2^k .
3. One *vertical edge* from (v, k) to $(v, k + 1)$ for each $v \in C^{(0)}$ and $k \in \mathbb{Z}_{\geq 0}$.

We define a *height function* $h: \text{Hor}(C) \rightarrow \mathbb{R}_{\geq 0}$ sending a vertex (v, k) to k and extending to an affine function on each edge of $\text{Hor}(C)$. Denote by $\text{Hor}(C)_k$ the set $h^{-1}[k, \infty)$, which we shall call the *k -thin-part* of the horoball based on C .

Definition 2.3.7. Let G be a group, let \mathcal{H} be a finite set of finitely generated subgroups of G and let S be a finite generating set for G such that $S \cap H$ is a generating set for H for each $H \in \mathcal{H}$. For each $H \in \mathcal{H}$ let T_H be a left transversal of H in G . (Recall that a left transversal of a subgroup H of G is a choice of one element of each left coset of H in G .) For each $H \in \mathcal{H}$ and each $t \in T_H$ let $\Gamma_{H,t}$ be the full subgraph of $\Gamma(G, S)$ with vertex set tH ; note that this is isomorphic to $\Gamma(H, S \cap H)$. Then we define the *cusped space* of the triple (G, \mathcal{H}, S) .

$$X(G, \mathcal{H}, S) = \Gamma(G, S) \cup \bigcup_{H \in \mathcal{H}, t \in T_H} \text{Hor}(\Gamma_{H,t}). \quad (2.11)$$

Here the part of $\text{Hor}(\Gamma_{H,t})$ at height zero is identified with $\Gamma_{H,t} \subset \Gamma(G, S)$. The *height function* h defined on each horoball extends to X : it takes the value 0 on $\Gamma(G, S)$.

Denote by X_k the set $h^{-1}[0, k]$, which we shall call the *k -thick part*.

Note that G acts on the cusped space X properly discontinuously and by isometries.

Theorem 2.3.8. [GM08, Theorem 3.25] *Let G be a group, let \mathcal{H} be a finite set of finitely generated subgroups of G and let S be a finite generating set for G such that $S \cap H$ is a generating set for H for each $H \in \mathcal{H}$. Then G is hyperbolic relative to \mathcal{H} if and only if the cusped space of the triple (G, \mathcal{H}, S) is a hyperbolic metric space.*

2.3.3 The Bowditch boundary

As noted above, a pair (G, \mathcal{H}) where G is hyperbolic relative to \mathcal{H} does not determine a canonical quasi-isometry class of spaces on which G admits a geometrically finite action. However, the cusped space is unique up to quasi-isometry: if S_1 and S_2 are finite generating sets for G such that $S_i \cap H$ generates H for each $H \in \mathcal{H}$ then $X(G, \mathcal{H}, S_1)$ is quasi-isometric to $X(G, \mathcal{H}, S_2)$. Therefore, following Bowditch [Bow98a], we make the following definition. (Note that the cusped space satisfies the conditions on spaces studied in [Bow98a, Section 6].)

Definition 2.3.9. Let G be hyperbolic relative to \mathcal{H} . Then we define the (*Bowditch boundary*) of (G, \mathcal{H}) to be the boundary of the hyperbolic metric space $X(G, \mathcal{H}, S)$ for some finite generating set S for G such that $S \cap H$ generates H for each H in \mathcal{H} .

In fact, the boundary of a space on which a relatively hyperbolic group acts as in Definition 2.3.2 is unique up to homeomorphism even though the space itself is not unique up to quasi-isometry. Note the following theorem of Bowditch.

Theorem 2.3.10. [Bow98a] *Suppose that G is hyperbolic relative to a collection \mathcal{H} of subgroups of G . Let X be a proper hyperbolic metric space on which G acts properly discontinuously and by isometries such that ∂X does not contain a closed non-empty G -invariant proper subset and every point in ∂X is either a conical limit point or a bounded parabolic point. Suppose further that ∂X contains no proper non-empty closed G -invariant subset. Then $\partial(G, \mathcal{H}) \cong \partial X$.*

Unlike in the case of hyperbolic groups, the boundary of a relatively hyperbolic group may contain a cut point. The implication of cut points in the boundary will be discussed in Section 3.5.3.

In the situation discussed at the end of Section 2.3.1 in which G is hyperbolic, there is a simple description of the boundary of $\partial(G, \mathcal{H})$ in terms of ∂G due to Tran [Tra13].

Theorem 2.3.11. *Let G be a hyperbolic group and let \mathcal{H} be a finite collection of finitely generated subgroups of G so that G is hyperbolic relative to \mathcal{H} . (Equivalently, each element of \mathcal{H} is quasi-convex and \mathcal{H} is almost malnormal.) Then $\partial(G, \mathcal{H}) \cong \partial G / \sim$, where $x \sim y$ if and only if $x = y$ or $\{x, y\} \subset g\Lambda H$ for some $g \in G$ and $H \in \mathcal{H}$.*

Chapter 3

Splittings and JSJ decompositions

In this chapter we review some of the theory of splittings of groups. We begin with an overview of Bass-Serre theory, before moving onto the theory of JSJ decompositions.

3.1 Amalgams and HNNs

We begin by recalling the two basic operations from which our decompositions of groups will be built. These operations are motivated geometrically.

First, let X be a CW complex and suppose that $X = U_1 \cup U_2$, where U_1 and U_2 are connected subcomplexes with $U_1 \cap U_2$ connected. Then the Seifert-van Kampen Theorem tells us that this decomposition of X into its subcomplexes induces a decomposition of $\pi_1 X$ as an amalgamated product:

$$\pi_1 X \cong \pi_1 U_1 *_{\pi_1(U_1 \cap U_2)} \pi_1 U_2$$

In this case, $U_1 \cap U_2$ is *separating*: cutting X along $U_1 \cap U_2$ separates it into disjoint pieces. There is a second possibility that results in a different decomposition of the fundamental group. First recall the following definition.

Definition 3.1.1. Let H be a group with $A, B \leq H$. Suppose that there exists an isomorphism $\alpha : A \rightarrow B$. Then the *HNN extension of H relative to α* is defined to be the following group:

$$H *_{\alpha} = \langle H, t \mid tat^{-1} = \alpha(a) \text{ for } a \in A \rangle.$$

Now suppose that X is a CW complex and U is a connected subcomplex that is non-separating. In other words, there exists a CW complex Y with disjoint subcomplexes U_1

and U_2 such that U_1 and U_2 are isomorphic, and the quotient of Y obtained by identifying U_1 with U_2 is isomorphic to X , where this isomorphism identifies each of U_1 and U_2 with U . Then another application of the Seifert-van Kampen Theorem gives the following decomposition of $\pi_1 X$:

$$\pi_1 X \cong \pi_1 Y *_{\alpha}$$

where α is the isomorphism $\pi_1 U_1 \rightarrow \pi_1 U_2$ determined by the identification of U_1 with U_2 .

We see that the two ways in which a CW complex can be cut (along separating and non-separating subcomplexes) correspond at the level of fundamental groups to amalgamated products and HNN extensions.

3.2 More general splittings

In this section we develop the general theory of group decompositions that can be built from amalgamated products and HNN extensions. For a full account see [Ser77].

3.2.1 Graphs of groups

The decompositions of groups that can be obtained by repeatedly decomposing as amalgamated products and HNN extensions are described through the formalism of *graphs of groups*. First, let us recall Serre's definition of an oriented graph.

Definition 3.2.1. A *graph* Y is a pair of sets $V = V(Y)$ and $E = E(Y)$ equipped with a map $\iota: E \rightarrow V$ and a fixed-point-free involution $E \rightarrow E$, which we denote by $e \mapsto \bar{e}$. We call elements of V *vertices* and elements of E (*oriented*) *edges*. We think of the map ι as sending an oriented edge to its initial vertex and the map $e \mapsto \bar{e}$ as sending an oriented edge to the same edge with the opposite orientation. We then define a map $\tau: E \rightarrow V$ by $\tau(e) = \iota(\bar{e})$, sending an edge to its terminal vertex.

Y is connected if for any vertices $v_1 \neq v_2$ in $V(Y)$, there exist edges e_1, \dots, e_n such that $\iota(e_1) = v_1$, $\tau(e_n) = v_2$ and $\iota(e_i) = \tau(e_{i-1})$ for $i = 2, \dots, n$.

We now recall the definition of a graph of groups, which is intended to record how a collection of groups should be glued together to obtain a larger group.

Definition 3.2.2. A *graph of groups* \mathcal{Y} is a connected graph Y together with the following data:

1. a group \mathcal{Y}_v for each vertex $v \in V(Y)$;

2. a group \mathcal{Y}_e for each edge $e \in E(Y)$ such that $\mathcal{Y}_e = \mathcal{Y}_{\bar{e}}$ for each edge e ;
3. an injective homomorphism $\partial_-^e : \mathcal{Y}_e \rightarrow \mathcal{Y}_{t_e}$ for each edge e .

We then define injective homomorphisms $\partial_+^e : \mathcal{Y}_e \rightarrow \mathcal{Y}_{\tau_e}$ by $\partial_+^e = \partial_-^{\bar{e}}$.

We still need to describe how to use the data recorded in a graph of groups to obtain a group by a process of attaching the vertex groups together. This can be done combinatorially (this is the approach taken by Serre [Ser77]), but we instead follow the topological approach of Scott and Wall [SW79].

3.2.2 Graphs of spaces

We follow Scott and Wall [SW79] in defining graphs of spaces analogous to graphs of groups.

Definition 3.2.3. A *graph of spaces* \mathbb{Y} is a connected graph Y together with the following data:

1. a CW complex \mathbb{Y}_v for each vertex $v \in V(Y)$;
2. a CW complex \mathbb{Y}_e for each edge $e \in E(Y)$ such that $\mathbb{Y}_e = \mathbb{Y}_{\bar{e}}$ for each edge e ;
3. a π_1 -injective continuous map $\partial_-^e : \mathbb{Y}_e \rightarrow \mathbb{Y}_{t_e}$ for each edge e .

We then define π_1 -injective continuous maps $\partial_+^e : \mathbb{Y}_e \rightarrow \mathbb{Y}_{\tau_e}$ by $\partial_+^e = \partial_-^{\bar{e}}$.

It is now easy to describe a gluing process through which one can define a CW complex from a graph of spaces.

Definition 3.2.4. Given a graph of spaces \mathbb{Y} , define *the space obtained by gluing* $X_{\mathbb{Y}}$ to be the following CW complexes:

$$X_{\mathbb{Y}} = \left(\bigsqcup_{v \in V(Y)} \mathbb{Y}_v \sqcup \bigsqcup_{e \in E(Y)} (\mathbb{Y}_e \times [0, 1]) \right) / \sim, \quad (3.1)$$

where for each $e \in E(Y)$ and each $y \in \mathbb{Y}_e$, $(y, 0) \sim \partial_-^e(y)$.

Using this construction we can now describe an analogous process for gluing together graphs of groups. Note that one can obtain a graph of groups from a graph of spaces \mathbb{Y} by defining $\mathcal{Y}_v = \pi_1 \mathbb{Y}_v$ and $\mathcal{Y}_e = \pi_1 \mathbb{Y}_e$ and defining edge homomorphisms using functoriality of π_1 . On the other hand, one can go the other way by replacing each group in a graph of groups by a topological space with that fundamental group; for uniqueness and to ensure that the edge maps may be defined we do this by applying the functor $K(-, 1)$.

Definition 3.2.5. Given a graph of groups \mathcal{Y} , let \mathbb{Y} be the associated graph of spaces. We then define the *fundamental group* of \mathcal{Y} by $\pi_1\mathcal{Y} = \pi_1X_{\mathbb{Y}}$.

If $G \cong \pi_1\mathcal{Y}$ then we call \mathcal{Y} , together with the isomorphism $\pi_1\mathcal{Y} \rightarrow G$, a *splitting of G* . Let \mathcal{A} be a set of subgroups of G closed under conjugation. If the isomorphism $\pi_1\mathcal{Y} \rightarrow G$ sends each edge group to an element of \mathcal{A} then we say that \mathcal{Y} is a *splitting of G over \mathcal{A}* .

Remark 3.2.6. Note that the image of a vertex group \mathcal{Y}_v in the fundamental group $\pi_1X_{\mathbb{Y}}$ is only determined up to conjugacy: it is dependent on the choice of a path in $X_{\mathbb{Y}}$ from the base point to the base point of \mathbb{Y}_v . Therefore we will only refer to this image when either a preferred conjugacy class representative has been chosen, or we are discussing a property of subgroups of $\pi_1X_{\mathbb{Y}}$ that is conjugacy invariant.

Remark 3.2.7. The requirement in Definition 3.2.2 that the edge homomorphisms in a graph of groups be injective (and, analogously in Definition 3.2.3, that the edge maps in a graph of spaces be π_1 -injective) seems artificial. Certainly, it is not required in the constructions above. We impose this restriction for two reasons: firstly so that we may obtain Bass and Serre’s relationship between graphs of groups and actions on trees in Section 3.3, and secondly so that Lemma 3.2.8 below is applicable.

In any case, this requirement is not particularly restrictive: if \mathbb{Y} is a “graph of spaces” in which the edge maps are not assumed to be π_1 -injective, a graph of spaces with π_1 -injective edge maps can be obtained by the following process. For each edge e and each loop γ in \mathbb{Y}_e such that either $\partial_-^e \circ \gamma$ or $\partial_+^e \circ \gamma$ bounds a disk, attach a disk to \mathbb{Y}_e along γ and attach disks to $\mathbb{Y}_{\iota(e)}$ and $\mathbb{Y}_{\tau(e)}$ along $\partial_-^e \circ \gamma$ and $\partial_+^e \circ \gamma$. Then extend ∂_-^e and ∂_+^e in the obvious way. Note that this process does not change the fundamental group of the space obtained by gluing, $X_{\mathbb{Y}}$. This process might have created new loops in edge spaces that do not bound disks in the edge spaces, but do in the adjacent vertex spaces. Repeat this process and take a direct limit to ensure that the edge maps are π_1 -injective.

A similar process is possible for graphs of groups: this is most easily described by applying the process described for graphs of spaces to the graph of spaces associated to the given graph of groups.

With the assumption that the edge maps are injective, the following lemma holds; this is a generalisation of Britton’s lemma for HNN extensions.

Lemma 3.2.8. [Ser77, Corollary 1] *Let \mathcal{Y} be a graph of groups. Then for each vertex v the map $\mathcal{Y}_v \rightarrow \pi_1\mathcal{Y}$ is injective. The same holds for the edge groups.*

Rather than trying to study all splittings of a given group G , we frequently take some particularly nice class \mathcal{A} of subgroups G (for example, finite subgroups, or virtually cyclic

subgroups) and restrict to studying splittings of G over \mathcal{A} . We will also sometimes be interested in studying a group *relative to* a preferred subgroup or collection of subgroups. In this case we will be interested in splittings with the following property.

Definition 3.2.9. Let \mathcal{H} be a collection of subgroups of G . Then a splitting \mathcal{Y} of G is *relative to* \mathcal{H} if every element of \mathcal{H} is conjugate in G to a subgroup of the image in G of a vertex group \mathcal{Y}_v .

3.2.3 Domination and refinement

A given group may admit many splittings, some of which will encode more information about the structure of the group than others. The following definitions formalise these relationships between splittings. Some of these definitions are more clearly stated in the language of group actions on trees, which will be discussed in the next section. For more details see [GL17]; we adopt their terminology here.

Throughout this thesis, all maps between graphs will send vertices to vertices and edges to edge paths.

Definition 3.2.10. Let \mathcal{Y}_1 and \mathcal{Y}_2 be splittings of a group G with underlying graphs Y_1 and Y_2 . We say that \mathcal{Y}_1 *dominates* \mathcal{Y}_2 if for each $v \in V(Y_1)$ there exists $w \in V(Y_2)$ such that $(\mathcal{Y}_1)_v$ is conjugate in G to a subgroup of $(\mathcal{Y}_2)_w$.

Sometimes we will be interested in a particularly nice type of domination, called refinement.

Definition 3.2.11. Let \mathcal{Y}_1 and \mathcal{Y}_2 be splittings of a group G with underlying graphs Y_1 and Y_2 . We say that \mathcal{Y}_1 *refines* \mathcal{Y}_2 if there is a *collapse map* $p: Y_1 \rightarrow Y_2$ (i.e. a map obtained by collapsing certain edges of Y_1 to points) such that the following conditions are satisfied.

1. For each $v \in V(Y_2)$, $(\mathcal{Y}_2)_v$ is the fundamental group of the graph of groups induced on $p^{-1}(v)$.
2. For each edge $e \in E(Y_2)$ with preimage $e' \in E(Y_1)$, $(\mathcal{Y}_2)_e$ may be identified with $(\mathcal{Y}_1)_{e'}$ such that the map ∂_-^e is conjugate in $(\mathcal{Y}_2)_{\iota(e)}$ to $\partial_-^{e'}$, where we identify $(\mathcal{Y}_1)_{\iota(e')}$ with a subgroup of $(\mathcal{Y}_2)_{\iota(e)}$.

Remark 3.2.12. A refinement is obtained by blowing up certain vertex groups as graphs of groups: a refinement of a graph of groups \mathcal{Y} is equivalent to a choice of splitting of \mathcal{Y}_v relative to $\{\partial_-^e(\mathcal{Y}_e) \mid \iota(e) = v\}$ for each vertex v : see [GL17, Lemma 4.12]. We quote Guirardel and Levitt's precise statement of this result later in the setting of group actions on trees.

Definition 3.2.13. A splitting \mathcal{Y} of a group G is *trivial* if there exists a vertex v in the underlying graph of \mathcal{Y} such that $\mathcal{Y}_v = G$.

Note that any group admits a trivial splitting with one vertex, which is refined by any other splitting.

Also note that any splitting is refined in a trivial manner by subdividing the edges in the splitting, or by adding redundant edges and vertices to the graph. We will sometimes wish to rule out such refinements, so we make the following definition.

Definition 3.2.14. Let \mathcal{Y} be a graph of groups with underlying graph Y . Then \mathcal{Y} is *minimal* if there is no proper subgraph $Y' \subset Y$ with the following properties, where \mathcal{Y}' is the graph of groups with underlying graph Y' and the same vertex and edge groups as \mathcal{Y} .

1. The inclusion $Y' \hookrightarrow Y$ is a homotopy equivalence.
2. For each edge $e \in Y - Y'$ oriented such that $\iota(e)$ is closer to Y' than $\tau(e)$, ∂_+^e is surjective.

A vertex $v \in V(\mathcal{Y})$ is a *redundant vertex* if it is of degree 2 and for each edge e of Y with $o(e) = v$, ∂_-^e is an isomorphism.

3.3 Group actions on trees

In [Ser77], Serre demonstrated a relationship between the splittings of a group and the actions that group admits on trees. This gives an alternative viewpoint on splittings, which we will find useful later on.

Definition 3.3.1. A *tree* is a graph with no cycles. A group G is said to act on a tree T *without edge inversions* if, for any $g \in G$ and any edge $e \in E(T)$, $g \cdot e \neq \bar{e}$. For a vertex $v \in V(T)$ we denote by G_v the stabiliser of v , and for an edge $e \in E(T)$ we denote by G_e the stabiliser of e .

Theorem 3.3.2. [Ser77] *Suppose that a group G acts on a tree T without edge inversions. Then G is isomorphic to the fundamental group of a graph of groups \mathcal{Y} with underlying graph $G \backslash T$. Let v be a vertex of T and let $[v]$ be its image in $G \backslash T$. Then under the identification of $\pi_1 \mathcal{Y}$ with G the group $\mathcal{Y}_{[v]}$ is conjugate to G_v . Similarly, when e is an edge of T we have that $\mathcal{Y}_{[e]}$ is conjugate to G_e .*

Conversely, if $G = \pi_1 \mathcal{Y}$ where \mathcal{Y} is a graph of groups then G admits an action on a tree T without edge inversions such that the stabiliser of each vertex of T is conjugate to a vertex group in \mathcal{Y} and the stabiliser of each edge of T is conjugate to an edge group in \mathcal{Y} .

These operations give a one-to-one correspondence between the actions of a group on trees and the graph of groups decompositions of that group.

As a consequence of this theorem, we could have defined a splitting of a group G over a class \mathcal{A} of subgroups of G to be an action of G on a tree T without edge inversions such that for each edge $e \in E(T)$, $G_e \in \mathcal{A}$. Note that the splitting is relative to a family \mathcal{H} of subgroups of G if and only if each element of \mathcal{H} fixes a point in the action of G on T ; in this case the elements of \mathcal{H} are said to be *elliptic*. Throughout this thesis we shall pass freely between these two notions of splittings.

The following lemma reframes the concepts of Section 3.2.3 in the language of actions on trees.

Lemma 3.3.3. *Let G act on trees T_1 and T_2 without edge inversions. Then:*

1. *the splitting T_1 dominates T_2 if and only if there is a G -equivariant map $T_1 \rightarrow T_2$;*
2. *the splitting T_1 refines T_2 if and only if there is a G -equivariant collapse map $T_1 \rightarrow T_2$.*

Let G act on a tree T without edge inversions. Then:

1. *the splitting is trivial if and only if G is elliptic;*
2. *the splitting is minimal if and only if T contains no proper G -invariant subtree;*
3. *the splitting is without redundant vertices if, for any vertex v in T of degree two, there is an element of G that swaps the two edges with initial vertex v .*

Proof. The conditions for domination and refinement are noted by Guirardel and Levitt in [GL17, Section 1.4.1–1.4.2].

The condition for triviality is clear.

For minimality, note that if T contains a proper G -invariant subtree T' then $G \setminus T' \subset G \setminus T$ satisfies the conditions of Definition 3.2.14, since T admits a G -equivariant deformation retract onto T' , and for any edge $e \in T - T'$, G_e fixes the (unique) vertex of T' closest to e . Therefore the splitting is not minimal.

Conversely, suppose that the splitting is not minimal and let \mathcal{Y} be the graph of groups associated to T , so there is a subgraph $Y' \subset Y$ satisfying the conditions of Definition 3.2.14. Then $\pi_1 \mathcal{Y}'$ is naturally isomorphic to G and \mathcal{Y}' dominates \mathcal{Y} . Let T' be the G -tree associated to \mathcal{Y}' . Then there is a G -equivariant map $T' \rightarrow T$. This map is not surjective, so its image is a proper G -invariant subtree of T .

The condition for redundant vertices is clear. □

The following lemma tells us how to build refinements of splittings by splitting the vertex groups relative to their incident edge groups.

Lemma 3.3.4. [GL17, Lemma 4.12] *Let T be the tree associated to a splitting of G . Let v be a vertex of G . Suppose that we have a splitting of G_v relative to $\{G_e \mid \iota(e) = v\}$; let S be the associated tree. Then T admits a refinement $p : T' \rightarrow T$ such that $p^{-1}(v)$ is G_v -equivariantly isomorphic to S .*

3.4 Splittings over finite subgroups

In this section we recall some results about splittings of finitely presented groups over finite subgroups. In this setting each finitely presented group G admits a maximal splitting that encodes all splittings of G over finite subgroups.

Definition 3.4.1. A group G is *accessible* if for any sequence $(\mathcal{Y}_i)_{i=1}^{\infty}$ of finite splittings of G over its finite subgroups such that each \mathcal{Y}_i is minimal and without redundant vertices, and such that \mathcal{Y}_{i+1} refines \mathcal{Y}_i for each i , there exists $I \in \mathbb{N}$ such that \mathcal{Y}_i is equal to \mathcal{Y}_{i+1} for all $i \geq I$.

The first result on accessibility follows from the Grushko theorem [Gru40]: this shows that all finitely presented torsion-free groups are accessible. Dunwoody showed that the assumption on torsion is unnecessary:

Theorem 3.4.2. [Dun85] *All finitely presented groups are accessible.*

Remark 3.4.3. Note that not all groups are accessible. Dunwoody constructed an example in [Dun93] of a finitely generated (but infinitely presented) inaccessible group.

It follows that for any finitely presented group G there exists a splitting over finite subgroups that is maximal for refinement; we shall call such a splitting a *Dunwoody decomposition of G* .

Remark 3.4.4. Note that this splitting dominates all finite splittings of G : let T be the corresponding tree and let S be the tree corresponding to some other splitting of G over its finite subgroups. Then for any edge e of T , G_e is elliptic with respect to the action on S , so for any vertex v of T , G_v must be elliptic with respect to the action on S , since otherwise the splitting over finite subgroups corresponding to T can be non-trivially refined at $[v]$ using Lemma 3.3.4. So there exists a G -equivariant map from T to S .

Remark 3.4.5. This maximal splitting is not unique: for example, if G is a finitely generated free group then the set of its maximal splittings over the trivial subgroup is closely related




Figure 3.1: Cutting this surface along the red curves yields a maximal splitting of the fundamental group over its cyclic subgroups. However, this splitting does not dominate all others: in particular, it does not dominate the splitting obtained by cutting the surface along the blue curve.

to the Culler-Vogtmann Outer Space [CV86]. However, the vertex groups in a maximal splitting over finite subgroups are canonically determined: they are precisely the maximal zero- or one-ended subgroups of G . Furthermore, if T_1 and T_2 are trees corresponding to two maximal splittings over finite subgroups then there exist G -equivariant maps $T_1 \rightarrow T_2$ and $T_2 \rightarrow T_1$.

The first example of the splittings of a group being reflected in the large scale geometry of that group is then given by the following theorem of Stallings.

Theorem 3.4.6. [Sta68, Sta71] *Let G be a finitely generated group. Then G has more than one end if and only if G admits a non-trivial splitting over its finite subgroups.*

If G is accessible then a maximal decomposition has the property that each vertex group is zero- or one-ended. Recall that if G is hyperbolic then the ends of G are in bijection with the connected components of ∂G , so in the setting of hyperbolic groups each vertex group in the Dunwoody decomposition has connected boundary.

3.5 JSJ decompositions

We now leave behind the problem of classifying splittings of a group over its finite subgroups: we suppose that the given group is one-ended and consider splittings over other families of subgroups. In this case, Dunwoody's accessibility theorem is replaced by a theorem of Bestvina and Feighn. Recall that a finitely generated group H is *small* if for any minimal non-trivial action of H on a tree, H maps either a ray or a bi-infinite geodesic into itself. In particular, any virtually cyclic group is small.

Theorem 3.5.1. [BF91] *Given a finitely presented group G , there exists a bound on the number of vertices in a minimal graph of groups decomposition of G without redundant vertices with small edge groups.*

We immediately encounter a difficulty that was not present in the case of splittings over finite groups: there is not necessarily a splitting with the property of Remark 3.4.4. This non-existence of a splitting dominating all others is essentially due to the presence of surface pieces in the group: see Figure 3.1. JSJ decompositions try to work around this difficulty, drawing inspiration from the following theorem.

Theorem 3.5.2. [JS79, Joh79] *Let M be an irreducible orientable closed 3-manifold. Then there exists a collection Σ of disjointly embedded incompressible tori in M such that each component of $M - \Sigma$ is either atoroidal or Seifert fibred. Furthermore, if Σ_1 and Σ_2 are minimal collections of tori with this property then Σ_1 is isotopic in M to Σ_2 .*

Essentially, the theorem guarantees the existence and uniqueness of a collection Σ of disjointly tori in M so that cutting M along Σ leaves only pieces that cannot be cut further along tori (the atoroidal case) and pieces that can contain very many different intersecting tori (the Seifert fibred case.) Although it is then possible to decompose M further by cutting the Seifert fibred pieces along embedded tori, we choose not to, since doing so would require the choice of a collection of *disjoint* tori on which to cut, violating uniqueness of the decomposition.

A JSJ decomposition of a group attempts to capture this uniqueness property. JSJ decompositions were introduced to group theory by Sela [Sel97] to answer questions about rigidity and the isomorphism problem for torsion-free hyperbolic groups. In [Bow98b] Bowditch developed a related type of decomposition for hyperbolic groups possibly with torsion; we return to study this decomposition in Section 3.5.2. Bowditch's decomposition is built from the structure of local cut points in the boundary of the group and is therefore invariant under automorphisms of the group. For other constructions of JSJ decompositions of groups see [RS97, DS99, FP06].

3.5.1 Guirardel and Levitt

In [GL17] Guirardel and Levitt provide a general definition of the JSJ decomposition of a group. This definition unifies and simplifies many of the preexisting ideas on this subject and will be very useful later on, so we recall their definition here. We work in the language of group actions on trees. As before, let G be a group and let \mathcal{A} be a collection of subgroups of G .

Definition 3.5.3. [GL17] An \mathcal{A} -tree is a tree equipped with a G -action without edge inversions in which all edge stabilisers are in \mathcal{A} , so \mathcal{A} -trees correspond to splittings of G over \mathcal{A} .

A subgroup $H \leq G$ is *universally elliptic* if it is elliptic with respect to any \mathcal{A} -tree. We denote by \mathcal{A}_{ell} the set of all universally elliptic elements of \mathcal{A} .

A *JSJ-tree of G over \mathcal{A}* is a \mathcal{A}_{ell} -tree that dominates any other \mathcal{A}_{ell} -tree. We call the splitting associated to such a tree a *JSJ decomposition*.

Guirardel and Levitt similarly define a relative JSJ decomposition. Here let \mathcal{H} be a finite collection of finitely generated subgroups of G .

Definition 3.5.4. [GL17] An $(\mathcal{A}, \mathcal{H})$ -tree is an \mathcal{A} -tree in which every element of \mathcal{H} is elliptic.

A subgroup $H \leq G$ is *universally elliptic relative to \mathcal{H}* (or just universally elliptic if \mathcal{H} is understood) if it is elliptic with respect to any $(\mathcal{A}, \mathcal{H})$ -tree. We denote by \mathcal{A}_{ell} the set of all universally elliptic elements of \mathcal{A} relative to \mathcal{H} .

A *JSJ-tree of G over \mathcal{A} relative to \mathcal{H}* is a $(\mathcal{A}_{\text{ell}}, \mathcal{H})$ -tree that dominates any other $(\mathcal{A}_{\text{ell}}, \mathcal{H})$ -tree.

Theorem 3.5.5. [GL17, Theorem 2.16 or Theorem 2.20] *If G is finitely presented then a JSJ tree over \mathcal{A} (relative to \mathcal{H} if applicable) exists.*

If \mathcal{Y} is a JSJ decomposition of G then for each vertex v in the underlying graph, there are two possibilities. In the first case, which is comparable to the atoroidal case of the classical JSJ decomposition of a 3-manifold, the decomposition \mathcal{Y} cannot be refined at v . Then \mathcal{Y}_v does not admit a non-trivial decomposition relative to the incident edge groups by Lemma 3.3.4. In this case we say that v is *rigid*. Alternatively, and analogously to the Seifert fibred case, there may be very many different ways to refine \mathcal{Y} at v ; in this case we say that v is *flexible*.

In many cases, Guirardel and Levitt are able to describe the flexible vertices. Essentially, these all come from surfaces pieces in G , so Figure 3.1 is descriptive of flexible vertices in considerable generality. We do not need the full strength of this description in a general setting, although in Section 3.5.2 we will observe that this description of flexible vertex groups holds for Bowditch's decomposition.

3.5.2 Bowditch's canonical decomposition

Before Guirardel and Levitt's general definition of JSJ decompositions, Bowditch defined a canonical decomposition of a one-ended hyperbolic group that coincides with a JSJ

decomposition over virtually cyclic subgroups in the sense of Guirardel and Levitt. The canonicity of this decomposition follows from the way it is constructed: the decomposition is built from the structure of local cut points in the Gromov boundary of the group, and is therefore preserved by automorphisms of the group. The definition of this decomposition motivates our interest in detecting local cut points in the boundary of a hyperbolic group, which we discuss in Chapter 4. We use these results to prove the computability of Bowditch's JSJ decomposition in Chapter 5. This computability theorem is the main result of this thesis.

In this section we describe Bowditch's JSJ decomposition of a one-ended hyperbolic group. We begin by recalling Bowditch's terminology to describe the structure of local cut points.

Let ∂G be the Gromov boundary of a one-ended hyperbolic group G ; note then that ∂G is compact, connected, locally connected and has no cut point (i.e. no $x \in \partial G$ such that $\partial G - \{x\}$ is disconnected). Then, for x in ∂G , denote by $\text{val}(x)$ the number of ends of $\partial G - \{x\}$, the *valency* of x . If a point x in ∂G has valency at least 2 then we call x a *local cut point* of ∂G . Let $M(n)$ be $\{x \in \partial G \mid \text{val}(x) = n\}$ and $M(n+)$ be $\{x \in \partial G \mid \text{val}(x) \geq n\}$. For x and y in ∂G , let $N(x, y)$ be the number of components of $\partial G - \{x, y\}$; note that $N(x, y) \leq \max\{\text{val}(x), \text{val}(y)\}$. Define a relation \sim on $M(2)$ by letting $x \sim y$ if and only if $x = y$ or $N(x, y) = 2$. Define a second relation \approx on $M(3+)$ by letting $x \approx y$ if and only if $N(x, y) = \text{val}(x) = \text{val}(y)$.

Lemma 3.5.6. [Bow98b, Lemma 3.1, Lemma 3.8, Proposition 5.13]. *The relation \sim is an equivalence relation. The relation \approx partitions $M(3+)$ into pairs.*

Let T_1 be the set of \approx -pairs and let T_2 be the set of \sim -classes. Bowditch shows that $T_1 \sqcup T_2$ can be embedded as a subset of the vertex set of a tree T by adding a third class of vertices T_3 . Under this embedding the "betweenness" of elements of $T_1 \sqcup T_2$ is preserved. Elements of T_3 are the subsets F of $T_1 \sqcup T_2$ satisfying the following two conditions. Firstly, for any x, z in F and y in $T_1 \sqcup T_2$, no pair of points in y separates a point in x from a point in z . Secondly, for any $x \notin F$, there exist y, z in F such that some pair of points in y does separate a point in x from a point in z . For more details of the construction of the tree T from $T_1 \sqcup T_2$, see [Bow98b, Section 2].

As in Guirardel and Levitt's JSJ decompositions discussed in Section 3.5.1, some vertices in Bowditch's decomposition correspond to surface pieces in G . To describe these vertices, Bowditch makes the following definition.

Definition 3.5.7. Recall that a *bounded Fuchsian group* is a non-elementary group Q that admits a properly discontinuous and convex cocompact action on the hyperbolic

plane \mathbb{H}^2 . The *peripheral subgroups* of a bounded Fuchsian group are the stabilisers of the components of the boundary of the minimal non-empty closed convex subset of \mathbb{H}^2 on which Q acts.

Definition 3.5.8. A subgroup $H \leq G$ is *hanging Fuchsian* if there exists a splitting of G over virtually cyclic subgroups such that H occurs as a vertex stabiliser, and H admits an isomorphism with a bounded Fuchsian group such that the stabilisers of incident edges are mapped precisely to the peripheral subgroups of the bounded Fuchsian group.

The following theorem summarises the properties of Bowditch's construction.

Theorem 3.5.9. [Bow98b, Theorem 5.28] *Let G be a one-ended hyperbolic group. Then G acts minimally and without edge inversions on a tree T such that all edge stabilisers are virtually cyclic. The vertex set of T is $T_1 \sqcup T_2 \sqcup T_3$, and these vertices have the following properties.*

1. *If $v \in T_1$ then its stabiliser G_v is a maximal virtually cyclic subgroup of G .*
2. *If $v \in T_2$ then its stabiliser G_v is a maximal hanging Fuchsian subgroup of G .*
3. *If $v \in T_3$ then its stabiliser is neither virtually cyclic nor hanging Fuchsian.*

Remark 3.5.10. It is convenient to modify the vertex set T_1 by adding a degree 2 vertex at the midpoint of each edge in T connecting a vertex in T_2 to a vertex in T_3 . Then T becomes a bipartite tree with vertex set partitioned into T_1 and $T_2 \sqcup T_3$. It is this decomposition that we will refer to as *Bowditch's JSJ decomposition*.

Theorem 3.5.9 has the following important corollary.

Corollary 3.5.11. [Bow98b, Theorem 6.2] *Let G be a one-ended hyperbolic group. Suppose that G is not cocompact Fuchsian. Then G admits a non-trivial splitting over its virtually cyclic subgroups if and only if ∂G contains a local cut point.*

Remark 3.5.12. The assumption that G is not cocompact Fuchsian is essential. If G is cocompact Fuchsian then ∂G is homeomorphic to a circle, so certainly contains a local cut point. However, not all cocompact Fuchsian groups admit non-trivial splittings over virtually cyclic subgroups: all torsion free cocompact Fuchsian groups do, but the hyperbolic triangle groups and their index 2 orientation preserving subgroups do not. This exception to Corollary 3.5.11 will cause us difficulties in computing JSJ decompositions, and is discussed further in Section 5.3.




Figure 3.2: The (graph of spaces corresponding to a) JSJ decomposition of a hyperbolic group. Vertices in T_1 can be seen as cylinders, vertices in T_2 as surfaces and vertices in T_3 as squares.

Remark 3.5.13. Bowditch's JSJ decomposition is in fact a JSJ decomposition of G in the sense of Guirardel and Levitt with \mathcal{A} equal to the set of virtually cyclic subgroups of G . This result is stated in [GL17]. For completeness, we give a full proof in Chapter 5: see Proposition 5.5.6. In the language of Guirardel and Levitt, vertices in T_2 will be seen to be flexible unless the quotient of \mathbb{H}^2 by the corresponding Fuchsian group is a particularly simple orbifold, such as a pair of pants, while vertices in $T_1 \cup T_3$ are rigid.

Geometrically, Bowditch's JSJ decomposition allows us to decompose a hyperbolic group into surface pieces and rigid pieces connected by cylinders. (In the presence of torsion the surface pieces are replaced by finite extensions of orbifold pieces and the cylinders are replaced by finite extensions products of closed one dimensional orbifolds with an interval.) Such a decomposition is shown in Figure 3.2.

3.5.3 Peripheral splittings

In [Bow01, Bow99a], Bowditch defines a second decomposition built from the topology of the boundary of a group: the peripheral splitting of a relatively hyperbolic group. The boundary of a relatively hyperbolic group *can* contain a cut point, and it is the structure of cut points in the boundary that determines the decomposition.

Definition 3.5.14. Let G be a group and let \mathcal{H} be a collection of subgroups of G . A *peripheral splitting* of G relative to \mathcal{H} is a representation of G as the fundamental group of a graph of groups with bipartite underlying graph so that, up to conjugacy, the vertex groups of one colour are precisely the conjugates of elements of \mathcal{H} .

A peripheral splitting of G relative to \mathcal{H} is a special kind of splitting of G over subgroups of conjugates of elements of \mathcal{H} relative to \mathcal{H} . It is clear that any peripheral splitting is indeed such a splitting and the following proposition provides a converse to this statement.

Proposition 3.5.15. *[Bow01, Proposition 5.1] The group G admits a non-trivial peripheral splitting relative to \mathcal{H} if and only if G admits a non-trivial splitting over subgroups of elements of \mathcal{H} relative to \mathcal{H} .*

The relationship between peripheral splittings and the boundary of a relatively hyperbolic group is demonstrated by the following two theorems.

Theorem 3.5.16. *[Bow01, Theorem 1.2] Let G be a group and \mathcal{H} a finite set of finitely presented subgroups of G such that G is hyperbolic relative to \mathcal{H} . Suppose that $\partial(G, \mathcal{H})$ is connected. If G admits a non-trivial peripheral splitting relative to \mathcal{H} then $\partial(G, \mathcal{H})$ contains a cut point.*

Putting together two theorems of Bowditch gives the following converse.

Theorem 3.5.17. *[Bow99b, Theorem 0.2][Bow99a, Theorem 1.2] Let G be a group and \mathcal{H} a finite set of finitely presented subgroups of G such that G is hyperbolic relative to \mathcal{H} . Suppose further that every element of \mathcal{H} is one- or two-ended and no element of \mathcal{H} contains an infinite torsion subgroup. If $\partial(G, \mathcal{H})$ is connected and contains a cut point then G admits a non-trivial peripheral splitting relative to \mathcal{H} .*

3.5.4 Cut pairs in the relative setting

To prove the results of this thesis we will only need to understand the implications of cut pairs in the Bowditch boundary of a relatively hyperbolic group in the very special case in which all peripheral subgroups are virtually cyclic. Elementary proofs of these results will be provided in Chapter 5. However, some elements of Bowditch's Theorem 3.5.9 have recently been shown to have analogues in the relative case. For the sake of completeness, we include this result here.

Theorem 3.5.18. *[Hau17] Let G be a group that is hyperbolic relative to a finite set \mathcal{H} of subgroups such that every element of \mathcal{H} is one- or two-ended and no element of \mathcal{H} contains an infinite torsion subgroup. If $\partial(G, \mathcal{H})$ contains a non-parabolic local cut point (that is, a point p that is not fixed by any conjugate of any element of \mathcal{H} such that $\partial(G, \mathcal{H}) - p$ has more than one end) then G admits a splitting over a virtually cyclic subgroup. Conversely, if G admits a non-trivial splitting over a non-parabolic virtually cyclic subgroup then $\partial(G, \mathcal{H})$ contains a non-parabolic local cut point.*

Remark 3.5.19. Note that if $\{p, q\}$ is a cut pair and neither of p and q is a cut point then p and q are local cut points.

3.6 Detecting splittings over finite groups

In this section we give an overview of Dahmani and Groves's result [DG08] that there is an algorithm that computes a maximal splitting of a hyperbolic group over its finite subgroups. This result will use Bestvina and Mess's double dagger condition: this is a geometric property of the Cayley graph of a hyperbolic group that turns out to be equivalent to the connectedness of the boundary of that group. This condition gives a quantitative local connectedness property for the boundary, which will be used in the proof of Proposition 4.1.8, a key step in our proof of Theorem 1.1.1.

In fact, Dahmani and Groves work in the greater generality of a certain class of relatively hyperbolic groups. Although our application of their results is to Theorem 4.3.2, which concerns only hyperbolic groups, our methods require that we allow virtually cyclic peripheral subgroups, since the Bowditch boundary of the group relative to this family determines whether or not the group admits a splitting relative to this family.

Theorem 3.6.1. [DG08] *There is an algorithm that takes as input a presentation for a group G and generators for a collection \mathcal{H} of abelian subgroups of G such that G is hyperbolic relative to \mathcal{H} , and returns a Dunwoody decomposition of G .*

Note that it is not difficult to find splittings of G over finite subgroups algorithmically when they exist: using moves called Tietze transformations one may enumerate all presentations of G . Assuming that G does admit a splitting, there is a presentation that clearly demonstrates that G has the structure of an amalgamated product or HNN extension; it remains to check that the edge group in this splitting is finite. The existence of an algorithm that checks this is a standard result in the theory of hyperbolic groups. A similar approach will be applied to the problem of computing splittings over virtually cyclic subgroups in Chapter 5. Following this argument, Dahmani and Groves obtain the following proposition.

Proposition 3.6.2. [DG08, Proposition 5.2] *There is an algorithm that takes as input a presentation for a group G together with generators for a finite collection \mathcal{H} of abelian subgroups of G such that G is hyperbolic relative to \mathcal{H} and terminates if and only if that group admits a non-trivial splitting over a finite subgroup relative to those abelian subgroups. (In other words, the property of admitting a non-trivial splitting over a finite subgroup*

is recursively enumerable among hyperbolic groups.) In the case when it terminates, the algorithm returns the non-trivial relative splitting.

In light of this result, to prove Theorem 3.6.1 Dahmani and Groves required a way of certifying that the group does not admit a non-trivial splitting over a finite subgroup in cases when it does not. Recall Theorem 3.4.6: this is equivalent to certifying that the boundary of the group is connected.

3.6.1 The double dagger condition

The so-called double dagger condition was introduced by Bestvina and Mess in [BM91] to study the local connectedness of the boundary of a hyperbolic group. The condition was later modified by Dahmani and Groves to guarantee its computability; we adopt their definition of the condition here. Let X be the Cayley graph of a hyperbolic group G with respect to a finite generating set, or, more generally, the cusped space associated to (G, \mathcal{H}, S) where \mathcal{H} is a finite collection of subgroups of G such that G is hyperbolic relative to \mathcal{H} , and S is a finite generating set for G such that $S \cap H$ generates H for each $H \in \mathcal{H}$.

Fix a base point $1 \in X$ corresponding to the identity element of G and fix an integer $\delta \geq 0$ such that X is δ -hyperbolic.

Lemma 3.6.3. [DG08, Lemma 2.11] *For any x in X , there is a geodesic ray starting at 1 that passes within a distance 3δ of x .*

For the remainder of this section let $C = 3\delta$. Also let $M = 6(C + 45\delta) + 2\delta + 3$.

Definition 3.6.4. [DG08] Let $\epsilon \geq 0$. A pair of points x and y in X satisfy \star_ϵ if $d(x, y) \leq M$ and $|d(1, x) - d(1, y)| \leq \epsilon$.

Let $n \in \mathbb{Z}_{\geq 0}$. For a pair (x, y) of points in X satisfying \star_ϵ we say that the condition $\ddagger(\epsilon, n)(x, y)$ holds if there is a path from x to y in X of length at most n that avoids the open ball $B_1(m - C - 45\delta + 3\epsilon)$, where $m = \min\{d(1, x), d(1, y)\}$.

We shall say that X satisfies $\ddagger(n)$ if for any pair of points x and y in X satisfying \star_0 , the condition $\ddagger(0, n)(x, y)$ holds.

Remark 3.6.5. Note that the parameter n in $\ddagger(n)$ differs from the corresponding definition of the double dagger definition of Bestvina and Mess [BM91]; the definition given here is more convenient for our purposes.

The following lemma relates the double dagger condition to the connectedness of ∂X .




Figure 3.3: The points x and y , which we assume to satisfy \star_ϵ , can be joined by a (blue) path of length at most n outside the (red) open ball of radius $m - C - 45\delta + 3\epsilon$ around 1 , so $\ddagger(\epsilon, n)(x, y)$ is satisfied.

Proposition 3.6.6. [BM91, Proposition 3.2] *Suppose that X satisfies $\ddagger(n)$ for some n . Then ∂X is path-connected and locally path-connected.*

This proposition is proved by noting that the double dagger condition allows one to “push” paths in X out onto ∂X . This construction will be used in Chapter 4 and is key to relating the existence of cut pairs in the boundary to the geometry of the cusped space. See the proof of Proposition 4.1.8 for a detailed example of how to use the double dagger condition in this manner.

A converse to Proposition 3.6.6 was proved by Bestvina and Mess [BM91, Proposition 3.3] in the absolute case (that is, $\mathcal{H} = \emptyset$) under the assumption that ∂X does not contain a cut point. Recall from Section 2.2 that this assumption is now known to hold for all hyperbolic groups. More generally, Dahmani and Groves [DG08] show that this converse holds for relatively hyperbolic groups as long as we place some restrictions on the peripheral subgroups, still with the assumption that ∂X does not contain a cut point. Dahmani and Groves provide a complete proof of this converse in the case that \mathcal{H} consists of abelian subgroups of rank at least 2 of G and comment on how to extend their arguments to abelian subgroups of rank 1. Most of their methods apply equally well when the elements of \mathcal{H} are virtually cyclic but not necessarily abelian, and it is this case that will interest us in Section 4, so we now collate the results of [DG08] and fill in the gaps in the virtually cyclic case. We divide into cases according to whether x and y are in the thick or the thin part of X .

The following proposition deals with the thick part of X , and holds without assumption on the nature of the peripheral subgroups.

Proposition 3.6.7. [DG08, Lemma 4.2] *Let $k \in \mathbb{Z}_{\geq 0}$. If ∂X is connected and does not contain a cut point then there exists $n \in \mathbb{Z}_{\geq 0}$ such that $\ddagger(10\delta, n)(x, y)$ holds for all x and y in X_k satisfying $\star_{10\delta}$.*

Note that although this proposition was originally stated for $k = 2M$, the proof given in [DG08] works for any value of k .

The following proposition deals with cusps corresponding to abelian parabolic subgroups of G .

Proposition 3.6.8. [DG08, Lemma 2.16] *Let $K = 3(2^{2M+3}) + M + 3$. Let $H \in \mathcal{H}$ be abelian of rank at least 2 and let $t \in T_H$ be an element of the left transversal of H in G . Then consider the horoball $\text{Hor}(\Gamma_{H,t}) \subset X$. Let $\text{Hor}(\Gamma_{H,t})_k$ be the set of points in the horoball of height at least k . Let x and y be points in $\text{Hor}(\Gamma_{H,t})_k$ satisfying $\star_{20\delta}$. Then x and y are connected by a path in X of length at most K that does not meet the open ball $B_1(\min\{d(1, x), d(1, y)\})$.*

As noted in the proof of [DG08, Lemma 2.16], a similar result holds for rank 1 abelian peripheral subgroups. The authors of that paper indicate that this follows from the fact that the corresponding horoballs are quasi-isometric to horoballs in \mathbb{H}^2 . We now give a precise statement and an alternative proof of this result, generalised to allow all virtually cyclic parabolics.

Proposition 3.6.9. *Let $L = 48 \cdot 2^{2M} + 5M + 20$. Let $H \in \mathcal{H}$ be virtually cyclic and let $t \in T_H$ be an element of the left transversal of H in G . Fix δ_H such that $\Gamma_{H,t}$ is δ_H -hyperbolic with respect to the path metric on $\Gamma_{H,t} \subset X$. Let $k \geq \min\{\log_2(2\delta_H + 1), \delta\}$. Let x_1 and x_2 be points in $\text{Hor}(\Gamma_{H,t})_k$ such that $d(x_1, x_2) \leq M$ and $d(1, x_i) \geq m$ for each i . Then there is a path from x_1 to x_2 in $\text{Hor}(\Gamma_{H,t})_k$ of length at most L that does not meet $B_1(m - 20\delta - 7)$.*

Proof. First note that $\Gamma_{H,t} \times \{0\} \subset \text{Hor}(\Gamma_{H,t})$ contains a bi-infinite geodesic $\alpha \times \{0\}$ with respect to the path metric. Furthermore, $t \mapsto \alpha(2^h t)$ defines the integer points of a geodesic in $\Gamma_{H,t} \times \{h\}$ with respect to the path metric. Also, since $\Gamma_{H,t} \times \{0\}$ is δ_H -hyperbolic, α is $(2\delta + 1)$ -quasi-surjective, and therefore every point in $\Gamma_{H,t} \times \{h\}$ is within a distance 1 of a point on $\alpha \times \{h\}$ as long as $h \geq \log_2(2\delta + 1)$. It follows that, at the cost of changing some constants by 2, we may assume that $t \mapsto \alpha(t) \times \{0\} \subset \Gamma_{H,t} \times \{0\}$ is an isometry from \mathbb{R} .

We therefore assume that this is the case and that x_1 and x_2 are points in $\text{Hor}(\Gamma_{H,t})_k$ such that $d(x_1, x_2) \leq M + 2$ and $d(1, x_i) \geq m - 2$ for each i . We must show that x_1 and x_2 are connected by a path in $\text{Hor}(\Gamma_{H,t})_k$ of length at most $L - 2$ that does not meet the open

ball $B_1(m - 20\delta - 7)$. To construct this path we must first understand the intersection $B_1(m - 20\delta - 7) \cap \text{Hor}(\Gamma_{H,t})_k$. Let $v \in \text{Hor}(\Gamma_{H,t})_k$ be a point closest to 1 in $\text{Hor}(\Gamma_{H,t})_k$ and let y be an arbitrary point in $B_1(m - 20\delta - 7) \cap \text{Hor}(\Gamma_{H,t})_k$. Then we apply the formulation of hyperbolicity given in Definition 2.1.10: the geodesic triangle $\Delta(1, v, y)$ is 4δ -thin by Lemma 2.1.11. Take points i_1, i_v and i_y on edges $[v, y]$, $[1, y]$ and $[1, v]$ respectively such that $\{i_1, i_v, i_y\}$ has diameter at most 4δ . Then we have the following inequality.

$$\begin{aligned} d(1, y) &= d(1, i_v) + d(i_v, y) \geq d(1, i_y) - 4\delta + d(i_1, y) - 4\delta \\ &= d(1, v) + d(v, y) - 8\delta - d(i_y, v) - d(v, i_1). \end{aligned}$$

Furthermore, $\text{Hor}(\Gamma_{H,t})_k$ is convex as long as $k > \delta$ by [GM08, Lemma 3.26], so $[v, y] \subset \text{Hor}(\Gamma_{H,t})_k$ and therefore $d(i_y, v) \leq d(i_z, i_1) \leq 4\delta$. Also, $d(v, i_1) \leq d(v, i_z) + d(i_z, i_1) \leq 8\delta$. It follows from these inequalities that we have the following inclusion.

$$(B_1(m - 20\delta - 7) \cap \text{Hor}(\Gamma_{H,t})_k) \subset (B_v(m - d(1, v) - 7) \cap \text{Hor}(\Gamma_{H,t})_k).$$

Note that since $\text{Hor}(\Gamma_{H,t})_k$ is convex in X the second intersection is equal to the ball in $\text{Hor}(\Gamma_{H,t})_k$ with centre v and radius $m - d(1, v) - 7$ with respect to the path metric.

We now describe this set in terms of coordinates in $\text{Hor}(\Gamma_{H,t})_k$. By reparametrising α we may assume that $v = (\alpha(0), k)$. Consider a geodesic segment from v to a point $(\alpha(t), h) \in \text{Hor}(\Gamma_{H,t})_k$ of length at most $m - d(1, v) - 7$. Immediately we see that $k \leq h \leq k + m - d(1, v) - 7$. We may assume that this geodesic initially ascends vertically, then travels horizontally a distance of at most 3, then descends vertically. The maximum height of any point on this geodesic is at most $(k + h + m - d(1, v) - 7)/2$, so t is bounded above:

$$|t| \leq 3 \cdot 2^{(k+h-7+m-d(1,v))/2}.$$

It follows that $B_1(m - 20\delta - 5) \cap \text{Hor}(\Gamma_{H,t})_k$ is contained in the following set, which we shall call B .

$$B = \left\{ (\alpha(t), h) \mid k \leq h \leq k + m - d(1, v) - 5, |t| \leq 3 \cdot 2^{(k+h-7+m-d(1,v))/2} \right\}$$

We now show that we have not enlarged the ball too much: we show that x_1 and x_2 are not in B . Let $x_i = (\alpha(t_i), h_i)$. It follows from the triangle inequality and the assumption that $d(1, x_i) \geq m - 2$ that $d(v, x_i) \geq m - d(1, v) - 2$. There is a geodesic from v to x_i comprising an ascending vertical segment, followed by a horizontal segment of length at most 3, followed by a descending vertical segment. If the descending segment is trivial then

$h_i > k + m - d(1, v) - 4$, so $x_i \notin B$. Otherwise the horizontal segment has length at least 2, and the height of the apex of this path is at least $(k + h_i - 3 + m - d(1, v) - 2)/2$, so $|t_i| > 2^{(k+h_i-5+m-d(1,v))/2}$. It again follows that $x_i \notin B$.

We shall construct the path from x_1 to x_2 to avoid B . We divide into cases according to the configuration of x_1 and x_2 with respect to B . First deal with the case in which t_1 and t_2 have opposite signs; in this case we show that both x_1 and x_2 must be higher than or almost as high as the highest point in B . In this case, for each i we have $d(x_i, \{\alpha(0)\} \times [k, \infty)) \leq M+3$ and $d(v, x_i) \geq m - d(1, v) - 2$. By considering a geodesic from v to x_i it follows that $h_i \geq k + m - d(1, v) - 2M - 8$. Therefore if either x_1 or x_2 lies underneath B then this can be rectified by moving that point along a path following α of length at most $24 \cdot 2^{2M}$ that avoids B . Still avoiding B , if either point has height less than the maximum height of B then this can be rectified by moving that point upwards by a distance of at most $2M + 5$. The two points can then be connected by a geodesic of length at most $6 + (M + 2)$. The lowest point of this geodesic is higher than the highest point of B , so the geodesic does not intersect with B . Therefore x_1 can be joined to x_2 by a path avoiding B of length at most $48 \cdot 2^{2M} + 5M + 18 \leq L - 2$.

Now assume that t_1 and t_2 have the same sign, say $0 \leq t_1 \leq t_2$. Note that $t_2 - t_1 \leq 2^{M+2+h_1}$ and $|h_2 - h_1| \leq M + 2$. If x_1 and x_2 both have height less than the highest point in B then it follows that x_1 can be connected to x_2 by a path of length at most $2^{M+2} + M + 2 \leq L - 2$ avoiding B ; this path consists of a horizontal segment from x_1 until x_1 lies directly over or directly under x_2 , followed by a vertical segment.

If x_1 and x_2 are both at least as high as the highest point in B then it follows that the geodesic from x_1 to x_2 avoids B ; this has length at most $M + 2 \leq L - 2$.

In the remaining case, say that x_1 has height less than that of the highest point in B and x_2 has height at least that of the highest point in B . Then $h_1 \geq k + m - d(1, v) - 4 - M - 2$, so if x_1 lies underneath B then this can be rectified by moving x_1 along a path following α of length at most $3 \cdot 2^{M+2}$ that avoids B . Append to this path an ascending vertical path ending at height equal to the that of the highest point in B ; this path has length at most $M + 2$. The end point of this path is within a distance $M + 5$ of x_2 , so x_1 and x_2 can be joined by a path of length at most $3 \cdot 2^{M+2} + 2M + 7 \leq L - 2$. \square

Putting together Propositions 3.6.7, 3.6.8 and 3.6.9, we obtain the following.

Proposition 3.6.10. *Suppose that every element of \mathcal{H} is virtually cyclic or abelian. Suppose also that ∂X is connected and does not contain a cut point. Then there exists n such that $\ddagger(10\delta, n)(x, y)$ holds for all x and y in X .*

3.6.2 The computability of \ddagger

In Section 3.6.1 we recalled the definition of the double dagger condition on the cusped space of a relatively hyperbolic group; this is a quantitative local connectedness condition for the boundary of that group. In order to apply this condition in algorithms we require a method to compute the parameters with which the condition holds.

We now recall results of Dahmani and Groves [DG08] that accomplish this goal. To incorporate the possibility of virtually cyclic parabolic subgroups we must modify the constants of these results: Dahmani and Groves treat separately the parts of X that lie in X_{2M} and the parts that lie in $\text{Hor}(\Gamma_{H,t})_M$ for some peripheral subgroup $H \in \mathcal{H}$ and t in the left transversal T_H . In the presence of virtually cyclic subgroups we still treat the thick and thin parts of X separately, but we must allow the cut-off between these two regimes to depend on the distortion of the virtually cyclic parabolic subgroups as in Proposition 3.6.9. In spite of this minor modification the proofs of the following results can be taken from [DG08] almost verbatim.

Proposition 3.6.11. *[DG08, Lemma 4.4] Let $k \in \mathbb{Z}_{\geq 0}$ and let $R(n) = 4(n + M) + 7k + 50\delta + 3$. Suppose that n is given such that $\ddagger(10\delta, n)(x, y)$ holds for all pairs of vertices x and y satisfying $\star_{10\delta}$ in $B_1(R(n))$. Then $\ddagger(0, n)(x, y)$ holds for all pairs of vertices x and y satisfying \star_0 in X_k .*

Theorem 3.6.12. *[DG08, Proposition 5.1] There is an algorithm that takes as input a presentation for a group G and finite generating sets for a finite collection \mathcal{H} of abelian and virtually cyclic subgroups of G such that G is hyperbolic relative to \mathcal{H} and terminates if and only if $\ddagger(n)$ holds for some $n \in \mathbb{Z}_{\geq 0}$. If the algorithm terminates it returns this value of n .*

The following proof is modified from [DG08].

Proof. First compute δ so that the cusped space is δ -hyperbolic and $\delta_{\mathcal{H}}$ so that $\Gamma(H, S \cap H)$ is hyperbolic for each $H \in \mathcal{H}$. Algorithms to accomplish this will be discussed in Section 5.1.1.

Take $k \geq 2M$ and $\geq \log(2\delta_H) + M$ for every virtually cyclic group $H \in \mathcal{H}$. Initially set n to be the maximum of the constants K and L defined in Propositions 3.6.8 and 3.6.9. Then check whether or not $\ddagger(10\delta, n)(x, y)$ holds for all pairs of points x and y satisfying $\star_{10\delta}$ in $B_1(R(n))$. If it does then $\ddagger(n)$ holds by Propositions 3.6.8, 3.6.9 and 3.6.11. The algorithm now terminates and returns n . If the condition does not hold for some pair of points then increment n by 1 and repeat this process. \square

Chapter 4

Cut points and cut pairs

Our goal in this chapter is to relate the existence of cut points and cut pairs in the boundary of a hyperbolic group (or a relatively hyperbolic group with virtually cyclic peripheral subgroups) to the geometry of some large ball in the Cayley graph or cusped space associated to that group. The novel material in this chapter is taken from [Bar18, Section 1–2]. Refer to Chapter 2 for definitions and key properties of the objects used in this chapter.

We begin the chapter by relating the connectedness of the complement of a pair of points (or a single point) in the boundary to the connectedness of a cylindrical region around a geodesic connecting those points (respectively a geodesic ray ending at that point.) Importantly, this result is quantitative: we give explicit bounds on the size of the cylinder. This is the goal of Section 4.1.

Then, in Section 4.2 we study the geometry of this cylindrical region when the geodesic penetrates deep into a horoball in the cusped space. Since we assume that the peripheral subgroups are virtually cyclic, the geometry of the cusps is particularly easy to understand.

Finally, in Section 4.3 we shall apply an argument based on the pumping lemma to show that the existence of a cut point or cut pair in the boundary of a group is determined by the geometry of a finite ball in the cusped space of known size. Supposing that there is a cut pair in the boundary, we show that the geodesic γ connecting the points in a cut pair may be assumed to be periodic and with bounded period: we take a short subsegment $\gamma|_{[a,b]}$ of that geodesic such that both the geodesic and the components of the thickened cylinder are identical in small neighbourhoods of a and b and form a new (local) geodesic that also connects the two points in a (possibly different) cut pair by concatenating infinitely many copies of $\gamma|_{[a,b]}$. A similar method is used in [CM11] to control cut pairs in the decomposition space associated to a line pattern in a free group. By characterising circles as Peano continua in which every pair of points is a cut pair, we do the same for the property of having circular boundary: we show that the boundary of a relatively hyperbolic

group is homeomorphic to a circle if and only if the cusped space contains a periodic local geodesic with bounded period such that the limit set of that geodesic is not a cut pair. This completes the proof of Theorem 1.1.1 and is a key step in the proofs of Theorems 1.2.1 and 1.2.2.

4.1 Cylinders, cut points and cut pairs

4.1.1 Spaces and constants

Let G be a group and let \mathcal{H} be a finite set of virtually cyclic subgroups of G . We work under the assumption that G is hyperbolic relative to \mathcal{H} . A virtually cyclic subgroup of a hyperbolic group is always quasi-convex, and is almost malnormal if and only if it is maximal among virtually cyclic subgroups of G . Then applying Theorem 2.3.5 gives the following lemma.

Lemma 4.1.1. *Suppose that G is hyperbolic and that \mathcal{H} is a finite set of virtually cyclic subgroups of G . Suppose further that each group in \mathcal{H} is maximal virtually cyclic and no two elements of \mathcal{H} are conjugate. Then G is hyperbolic relative to \mathcal{H} .*

Let S be a finite generating set for G such that $S \cap H$ is a generating set for H for each $H \in \mathcal{H}$. Then let $X = X(G, \mathcal{H}, S)$ be the cusped space associated to the triple (G, \mathcal{H}, S) as in Definition 2.3.7. Note that this is a hyperbolic metric space. Let $1 \in X$ be the vertex corresponding to the identity element of G .

We make the additional assumption that X satisfies a double dagger condition. Note that by Proposition 3.6.10 this is guaranteed if the boundary $\partial(G, \mathcal{H})$ is connected and does not contain a cut point, and by Proposition 3.6.6 this implies that $\partial(G, \mathcal{H})$ is connected. We now fix some constants for the remainder of this chapter.

1. Fix δ such that X is δ -hyperbolic.
2. Let $C = 3\delta$ as in Lemma 3.6.3.
3. Let D be chosen so that for any $(8\delta + 1)$ -local-geodesic γ and any geodesic γ' with the same endpoints in $X \cup \partial X$ as γ , the Hausdorff distance between γ and γ' is at most D . (See Theorem 2.1.17.)
4. Let $M = 6(C + 45\delta) + 2\delta + 3$ and let n be such that X satisfies $\ddagger(n)$.

5. Let d_v be a visual metric on ∂G and fix a , k_1 and k_2 such that it satisfies the following inequality. (See Proposition 2.2.17.)

$$k_1 a^{-(p \cdot q)_1} \leq d_v(p, q) \leq k_2 a^{-(p \cdot q)_1}. \quad (4.1)$$

In fact, a , k_1 and k_2 can be taken to be $2^{1/4\delta}$, $3 - 2\sqrt{2}$ and 1 respectively.

4.1.2 Cylinders in the cusped space

We now relate the existence of cut points and cut pairs in the boundary $\partial(G, \mathcal{H})$ to the connectedness of thickened cylinders around local geodesics in X . Let γ be an $(8\delta + 1)$ -local-geodesic in X with either one or both end points in ∂X . (In other words, γ is either a local geodesic ray or a bi-infinite local geodesic.) Note that the end points of γ at infinity are well defined by Remark 2.2.3. We shall denote by $\Lambda\gamma$ the limit set of γ in ∂X , which is either one or two points. By translating by an element of G we may assume that the point $1 \in X$ is $\gamma(0)$.

We define a subset of X , the connectivity of which will be seen to reflect the connectivity of $\partial X - \Lambda\gamma$.

Definition 4.1.2. For $R \geq 0$ let $N_R(\gamma)$ be the closed R -neighbourhood of γ and for $0 \leq r \leq R \leq \infty$ let $N_{r,R}(\gamma)$ be $\{x \in X : r \leq d(x, \gamma) \leq R\}$. For $K \geq 0$ let $C_K(\gamma)$ be $\{x \in X : d(x, \gamma) = K\}$. Finally, for $0 \leq r \leq K \leq R \leq \infty$ let $\text{Cyl}_{r,R,K}(\gamma)$ be the union of those connected components of $N_{r,R}(\gamma)$ that meet $C_K(\gamma)$. Note that $\text{Cyl}_{r,R,K}(\gamma)$ contains $N_{K,R}(\gamma)$. See Figure 4.1 for an illustration of such a subset of X .

Definition 4.1.3. For a component U of $\text{Cyl}_{r,\infty,K}(\gamma)$ define its *shadow* $\mathcal{S}U$ to be the set of points p in ∂X such that for any geodesic ray α from 1 to p , $\alpha(t)$ is in U for t sufficiently large.

The condition of Definition 4.1.3 that all geodesic rays with a particular end point are eventually contained in a component U of $\text{Cyl}_{r,\infty,K}$ is equivalent to a seemingly weaker condition.

Lemma 4.1.4. *Let α be a geodesic ray with $\alpha(0) = 1$. Suppose that $t \geq 0$ and $d(\alpha(t), \gamma) \geq 3\delta + \max\{r, D\}$. Let U be the component of $N_{r,\infty}(\gamma)$ containing $\alpha(t)$. Then U meets $C_K(\gamma)$ for any $K \geq r$ and $\alpha(\infty)$ is contained in $\mathcal{S}U$.*

Proof. Let α' be another geodesic ray with $\alpha'(0) = 1$ and $\alpha'(\infty) = \alpha(\infty)$. Let s be greater than $t + d_{\text{Haus}}(\alpha, \alpha')$. Then the triangle $\Delta(1, \alpha(s), \alpha'(s))$, two edges of which are subsegments of α and α' , is δ -slim. It follows that $d(\alpha'(t), \alpha) \leq \delta$, and then $d(\alpha'(t), \alpha(t)) \leq 2\delta$, so $d(\alpha'(t), \gamma) \geq \delta + \max\{r, D\}$ and $\alpha'(t) \in U$.

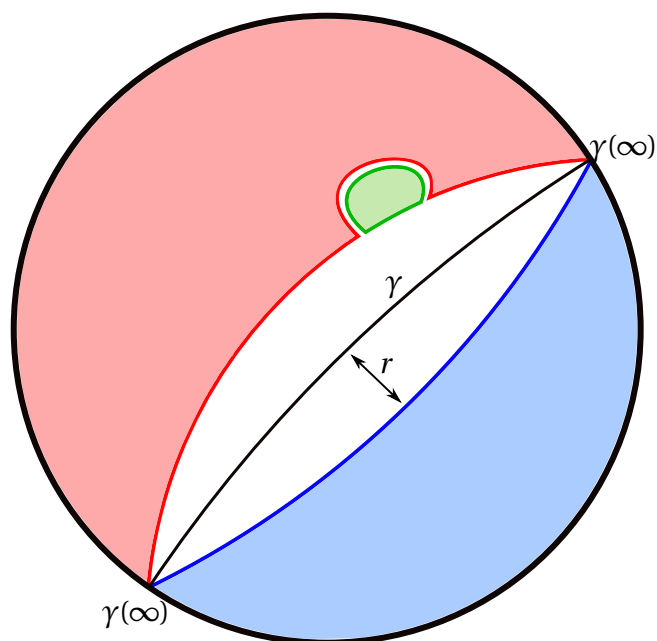


Figure 4.1: The three connected components of $N_{r,\infty}(\gamma)$ are shaded red, blue and green. The small green component does not correspond to a component of $\partial X - \Lambda\gamma$ and is discarded when $N_{r,\infty}(\gamma)$ is replaced by $\text{Cyl}_{r,\infty,K}(\gamma)$. The union of the red and blue components is $\text{Cyl}_{r,\infty,K}(\gamma)$.

We now show that $\alpha'(s) \in U$ for $s \geq t$. If not then there exists $s \geq t$ such that $d(\alpha(s), \gamma) < r$; let $d(\alpha(s), \gamma(s')) < r$. Then the triangle $\Delta(1, \alpha'(s), \gamma(s'))$ is δ -slim, so $\alpha'(t)$ is within a distance δ of $[1, \gamma(s')] \cup [\gamma(s'), \alpha(s)]$. But $[1, \gamma(s')]$ is contained in a D -neighbourhood of γ , while $[\gamma(s'), \alpha(s)]$ is contained in an r -neighbourhood of γ , so this is impossible.

Therefore $\alpha'(s) \in U$ for $s \geq t$ for any geodesic ray α' with $\alpha'(0) = 1$ and $\alpha'(\infty) = \alpha(\infty)$. In particular, this is true for $\alpha' = \alpha$ and α diverges arbitrarily far from γ , so U meet $C_K(\gamma)$. Furthermore, $\alpha(\infty) \in \mathcal{S}(U)$. \square

Lemma 4.1.5. *Let $r > D$ and let $K \geq r$. Then ∂X is covered by shadows:*

$$\bigcup_U \mathcal{S}U = \partial X - \Lambda\gamma, \quad (4.2)$$

where the union is taken over the set of connected components of $\text{Cyl}_{r, \infty, K}(\gamma)$. Furthermore, $\mathcal{S}U \cap \mathcal{S}V = \emptyset$ for distinct components U and V of $\text{Cyl}_{r, \infty, K}(\gamma)$.

Proof. The statement that shadows are disjoint is clear from the definition of the shadow. The assumption that $r > D$ ensures that when U is a component of $\text{Cyl}_{r, \infty, K}(\gamma)$, $\mathcal{S}U$ does not contain a point in $\Lambda\gamma$, since any geodesic ray from 1 to a point in $\Lambda\gamma$ is contained in a D -neighbourhood of γ .

If p is any point in $\partial X - \Lambda\gamma$ then any geodesic ray α from 1 to p diverges arbitrarily far from γ , so $d(\alpha(t), \gamma) \geq 3\delta + \max\{r, D\}$ for some $t \geq 0$. Let U be the component of $N_{r, \infty}(\gamma)$ that contains $\alpha(t)$. Then U is a component of $\text{Cyl}_{r, \infty, K}(\gamma)$ and $p \in \mathcal{S}U$ by Lemma 4.1.4. \square

Lemma 4.1.6. *Let U be a component of $\text{Cyl}_{r, \infty, K}(\gamma)$. Then $\mathcal{S}U$ is non-empty as long as $K \geq C + 3\delta + \max\{r, D\}$.*

Proof. Let $x \in C_K(\gamma) \cap U$. Then by Lemma 3.6.3 there exists a geodesic ray α with $\alpha(0) = 1$ and $t \geq 0$ such that $d(\alpha(t), x) \leq C$. Then $\alpha(t) \in U$ and $d(\alpha(t), \gamma) \geq 3\delta + \max\{r, D\}$, so $\alpha(\infty) \in \mathcal{S}U$ by Lemma 4.1.4. \square

Lemma 4.1.7. *Let $r > D$, let $K \geq r$ and let U be a component of $\text{Cyl}_{r, \infty, K}(\gamma)$. Then $\mathcal{S}U$ is closed and open in $\partial X - \Lambda\gamma$.*

Proof. Let p be a point in $\mathcal{S}U$ and let $p = \alpha(\infty)$ where α is a geodesic ray with $\alpha(0) = 1$. For $t \geq 0$ let $V_t(\alpha)$ be the set of end points of geodesic rays β from 1 such that $d(\beta(t), \alpha(t)) < 2\delta + 1$. From Proposition 2.2.8, the collection $\{V_t(\alpha) \mid t \in \mathbb{N}\}$ of such sets forms a fundamental system of neighbourhoods of $p \in \partial X$.

Then there exists t_0 such that for $t \geq t_0$, $d(\alpha(t), \gamma) \geq 5\delta + 1 + \max\{r, D\}$, and therefore $\alpha(t_0) \in U$ by Lemma 4.1.4. We claim that $V_{t_0}(\alpha) \subset \mathcal{S}U$. Let β be a geodesic ray with

$\beta(0) = 1$ and $d(\beta(t_0), \alpha(t_0)) \leq 2\delta + 1$. Then $\beta(t_0) \in U$ and $d(\beta(t_0), \gamma) \geq 3\delta + \max\{r, D\}$, so $\beta(\infty) \in \mathcal{S}U$ by Lemma 4.1.4.

Since p was arbitrary in $\mathcal{S}U$, it follows that $\mathcal{S}U$ is open. As U ranges over the connected components of $\text{Cyl}_{r,\infty,K}(\gamma)$, $\mathcal{S}U$ ranges over a cover of $\partial X - \Lambda\gamma$ by disjoint open subsets by Lemma 4.1.5 since $r > D$, so each is also closed. \square

To complete our proof of the fact that connectedness of $\partial X - \Lambda\gamma$ is equivalent to connectedness of $\text{Cyl}_{r,\infty,K}(\gamma)$ for appropriate values of r and K we prove the following proposition, in which we show that the double dagger condition can be used to push paths in a component U of $\text{Cyl}_{r,\infty,K}(\gamma)$ onto $\mathcal{S}U$. The proof of this Proposition is similar to on the proof of [BM91, Proposition 3.2].

Proposition 4.1.8. *Assume that r satisfies the following inequality.*

$$r > 2 \log_a \left(\frac{k_2}{k_1} \frac{n-1}{1-a^{-1}} \right) + M + 12\delta + D. \quad (4.3)$$

If U is a connected component of $\text{Cyl}_{r,\infty,K}(\gamma)$ then $\mathcal{S}U$ is contained in one connected component of $\partial X - \Lambda\gamma$.

Proof. Let p and q be points in $\mathcal{S}U$ and let α_1 and α_2 be geodesic rays from 1 to p and q respectively. Then there exist t_1 and t_2 such that $\alpha_1(t_1)$ and $\alpha_2(t_2)$ are in U . Let $\phi: [0, \ell] \rightarrow X$ be a path in U parametrised by arc length connecting the two points $\alpha_1(t_1)$ and $\alpha_2(t_2)$.

For each integer i in $[0, \ell]$ let z_i be a point within a distance C of $\phi(i)$ so that there is a geodesic ray β_i from 1 with $\beta_i(m_i) = z_i$. We can assume that $\beta_0 = \alpha_1$ and $\beta_\ell = \alpha_2$. Then, following the argument of [BM91, Proposition 3.2], we show that $\beta_i(\infty)$ and $\beta_{i+1}(\infty)$ can be connected in $\partial X - \Lambda\gamma$ for each i . For notational convenience we prove it for $i = 0$.

Note first that $d(\beta_0(m_0), \beta_1(m_1)) \leq 2C + 1$, which implies that $|m_0 - m_1| \leq 2C + 1$, and therefore $d(\beta_0(m_0), \beta_1(m_0)) \leq 4C + 2 \leq M$. Then, using the condition $\ddagger(n)$, choose a sequence of points $w_0 = \beta_0(m_0), w_1, \dots, w_n = \beta_1(m_0)$ such that $d(w_j, w_{j+1}) \leq 1$ and $d(1, w_j) \geq m_0 - C - 45\delta$ for each j . Choose geodesic rays $\beta_{1/n}, \beta_{2/n}, \dots, \beta_{n-1/n}$ such that $d(w_j, \beta_{j/n}) \leq C$.

We claim that $d(\beta_{j/n}(m_0 + 1), \beta_{j+1/n}(m_0 + 1)) \leq M$. To see this, choose s_j such that $d(\beta_{j/n}(s_j), w_j) \leq C$, so $d(\beta_{j/n}(s_j), \beta_{j+1/n}(s_{j+1})) \leq 2C + 1$. If $s_j \geq m_0 + 1$ then, since the triangle $\Delta(1, \beta_{j/n}(s_j), \beta_{j+1/n}(s_{j+1}))$ is δ -slim, $d(\beta_{j/n}(m_0 + 1), \beta_{j+1/n}(m_0 + 1)) \leq \delta + 2C + 1$, so $d(\beta_{j/n}(m_0 + 1), \beta_{j+1/n}(m_0 + 1)) \leq 2(\delta + 2C + 1) \leq M$. Alternatively, if $s_j \leq m_0 + 1$, then $|s_j - (m_0 + 1)| \leq 2C + 45\delta + 1$, so by the triangle inequality $d(\beta_{j/n}(m_0 + 1), \beta_{j+1/n}(m_0 + 1)) \leq 2C + 1 + 2(2C + 45\delta + 1) \leq M$.

Now proceed by induction: using the condition $\ddagger(n)$ as above, define geodesic rays β_t for each n -adic rational t in $[0, 1]$ inductively on the power k of the denominator of t to satisfy the following inequality for each j with $0 \leq j \leq n^k - 1$.

$$d\left(\beta_{j/n^k}(m_i + k), \beta_{(j+1)/n^k}(m_i + k)\right) \leq M. \quad (4.4)$$

The triangle inequality gives the following lower bound on the Gromov product of these points.

$$\left(\beta_{j/n^k}(\infty) \cdot \beta_{(j+1)/n^k}(\infty)\right)_1 \geq \liminf_{n_1, n_2} \left(\beta_{j/n^k}(n_1) \cdot \beta_{(j+1)/n^k}(n_2)\right)_1 \quad (4.5)$$

$$\geq \left(\beta_{j/n^k}(m_0 + k) \cdot \beta_{(j+1)/n^k}(m_0 + k)\right)_1 \quad (4.6)$$

$$= m_0 + k - M/2 \quad (4.7)$$

Recall that we defined a visual metric d_v on ∂X with base point 1, visual parameter a and multiplicative constants k_1 and k_2 . We obtain:

$$d_v\left(\beta_{j/n^k}(\infty), \beta_{(j+1)/n^k}(\infty)\right) \leq k_2 a^{-m_0 - k + M/2}. \quad (4.8)$$

Inductively applying the triangle inequality, we arrive at the following inequality for each n -adic rational $t \in [0, 1]$.

$$d_v(\beta_0(\infty), \beta_t(\infty)) \leq \frac{k_2(n-1)a^{-m_0 + M/2}}{1 - a^{-1}} \quad (4.9)$$

Define a path $\psi: [0, 1] \rightarrow \partial X$ with $\psi(t) = \beta_t(\infty)$ for each n -adic rational t in $[0, 1]$.

This extends continuously to a path from $\beta_0(\infty)$ to $\beta_1(\infty)$ by the uniform continuity of the map $t \rightarrow \beta_t(\gamma)$ defined on the n -adic rationals, which is established by applying the triangle inequality and Equation 4.8 repeatedly. This path is contained in the ball of radius $k_2(n-1)a^{-m_0 + M/2}/(1 - a^{-1})$ around $\beta_0(\infty)$.

We now bound below the distance $d_v(\beta_0(\infty), \Lambda\gamma)$. Let γ' be a geodesic ray from 1 to $\gamma(\infty)$, so the Hausdorff distance between γ and γ' is at most D . By Lemma 2.2.16, we have the following inequality.

$$\left(\beta_0(\infty) \cdot \gamma'(\infty)\right)_1 \leq \liminf_{n_1, n_2} \left(\beta_0(n_1) \cdot \gamma'(n_2)\right)_1 + 2\delta. \quad (4.10)$$

Let n_1 and n_2 each be at least m_0 . Certainly $d(\beta_0(m_0), \gamma') > \delta$ since $r > \delta + D$, so there exists a point p on $[\beta_0(n_1), \gamma'(n_2)]$ within a distance δ of $\beta_0(m_0)$. In fact, $d(\beta_0(m_0), \gamma) > 2\delta + D$,

so $d(\beta_0, \gamma'(m_0)) > 2\delta$. Therefore there exists a point q on $[\beta_0(n_1), \gamma'(n_2)]$ within a distance δ of $\gamma'(m_0)$.

Suppose that q is closer to $\beta_0(n_1)$ than p . Then by considering the geodesic triangle with vertices $\beta_0(n_1)$, $\beta_0(m_0)$ and p we see that q is within distance 2δ of β_0 . Therefore the distance from $\gamma'(m_0)$ to $\beta_0(m_0)$ is at most 6δ . But we assumed that $r > 6\delta + D$, which gives a contradiction. This implies that $d(\beta_0(n_1), \gamma'(n_2))$ is equal to the sum of the distances $d(\beta_0(n_1), p)$, $d(p, q)$ and $d(q, \gamma'(n_2))$.

Then we have the following inequality.

$$(\beta_0(n_1) \cdot \gamma'(n_2))_1 - (\beta_0(m_0) \cdot \gamma'(m_0))_1 = d(\beta_0(n_1), \beta_0(m_0)) - d(\beta_0(n_1), p) \quad (4.11)$$

$$+ d(\beta_0(m_0), \gamma'(m_0)) - d(p, q) \quad (4.12)$$

$$+ d(\gamma'(m_0), \gamma'(n_2)) - d(q, \gamma'(n_2)) \quad (4.13)$$

$$\leq \delta + 2\delta + \delta = 4\delta \quad (4.14)$$

This implies a lower bound on the distance from $\beta_0(\infty)$ to $\gamma(\infty)$ with respect to the visual metric.

$$d_v(\beta_0(\infty), \gamma(\infty)) \geq k_1 a^{-m_0 + (r-D)/2 - 6\delta}. \quad (4.15)$$

In the case that γ is bi-infinite, $d_v(\beta_0(\infty), \gamma(-\infty))$ similarly satisfies the same bound. Therefore, by the assumption on r the path constructed from $\beta_0(\infty)$ to $\beta_1(\infty)$ avoids $\Lambda\gamma$. \square

This completes the proof of the assertion that the connectedness of $\text{Cyl}_{r,\infty,K}(\gamma)$ determines the connectedness of $\partial X - \Lambda\gamma$. We now show that this connectedness can be detected without looking too far from γ : we show that for R at least some predetermined constant, $\text{Cyl}_{r,R,K}(\gamma)$ has the same number of components as $\text{Cyl}_{r,\infty,K}(\gamma)$.

Proposition 4.1.9. *Let $R \geq 2\delta + 2D + \max\{r + 4\delta + 1, K\}$. Then the inclusion map $\text{Cyl}_{r,R,K}(\gamma) \hookrightarrow \text{Cyl}_{r,\infty,K}(\gamma)$ induces a bijection between the sets of connected components of those subspaces of X .*

Proof. Surjectivity is clearly guaranteed by the fact that $R \geq K$. For injectivity, let x and y be points in $C_K(\gamma)$ that lie in the same connected component of $\text{Cyl}_{r,\infty,K}(\gamma)$. We show that the shortest path from x to y in $N_{r,\infty}(\gamma)$ stays within a distance R of γ . This argument is illustrated in Figure 4.2. Let $\phi: [0, \ell] \rightarrow X$ be such a shortest path parametrised by arc length. Suppose that $d(\phi(s), \gamma) > R$. Let $[t_0, t_1]$ be a maximal subinterval of $[0, \ell]$ containing s such that $d(\phi(t), \gamma) \geq r + 4\delta + 1$ for $t \in [t_0, t_1]$.




Figure 4.2: The configuration of points and paths in the proof of Proposition 4.1.9. We show that the shortest path from x to y that avoids $N_r(\gamma)$ is contained in $N_R(\gamma)$, so the cusped space cannot contain a large “hole” obstructing such paths as shown in the diagram.

Then for $t \in [t_0, t_1]$, $\phi|_{[t-4\delta-1, t+4\delta+1] \cap [t_0, t_1]}$ has image in $N_{r+4\delta+1, \infty}(\gamma)$. Therefore any geodesic segment from $\phi(\min\{t-4\delta-1, t_0\})$ to $\phi(\max\{t+4\delta+1, t_1\})$ is contained in $N_{r, \infty}(\gamma)$, so by minimality of the length of ϕ , $\phi|_{[t-4\delta-1, t+4\delta+1] \cap [t_0, t_1]}$ is a geodesic. This means that $\phi|_{[t_0, t_1]}$ is an $(8\delta + 2)$ -local geodesic. Therefore by Theorem 2.1.17 it is contained in a D -neighbourhood of any geodesic from $\phi(t_0)$ to $\phi(t_1)$.

By maximality of $[t_0, t_1]$, either $d(\phi(t_0), \gamma) = r + 4\delta + 1$ or $t_0 = 0$, so certainly $d(\phi(t_0), \gamma) \leq \max\{r + 4\delta + 1, K\}$, and similarly $d(\phi(t_1), \gamma)$ satisfies the same inequality. By δ -hyperbolicity applied to the geodesic quadrilateral with vertices $\phi(t_0)$, $\phi(t_1)$ and the points $\gamma(s_0)$ and $\gamma(s_1)$ on γ minimising the distances to $\phi(t_0)$ and $\phi(t_1)$, any geodesic from $\phi(t_0)$ to $\phi(t_1)$ is contained in a $2\delta + \max\{r + 4\delta + 1, K\}$ neighbourhood of a geodesic from $\gamma(s_0)$ to $\gamma(s_1)$, so is a subset of $N_{2\delta + \max\{r + 4\delta + 1, K\} + D}(\gamma)$. Hence $d(\phi(s), \gamma) \leq 2\delta + \max\{r + 4\delta + 2, K\} + 2D$, which is a contradiction. \square

From the results of this section we draw the following conclusion, which was main goal of this section.

Proposition 4.1.10. *The map that sends a component U of $\text{Cyl}_{r, R, K}(\gamma)$ to the shadow (with respect to the base point $1 \in \gamma$) of the component of $\text{Cyl}_{r, \infty, K}(\gamma)$ containing U is a well defined bijection between the set of connected components of $\text{Cyl}_{r, R, K}(\gamma)$ and the set of*

connected components of $\partial X - \Lambda\gamma$ as long as r , R and K are taken to simultaneously satisfy the conditions of Lemmas 4.1.5, 4.1.6, 4.1.7 and Propositions 4.1.8 and 4.1.9. \square

Remark 4.1.11. The conditions on r , R and K depend only on δ , n and λ . Suitable values for r , R and K can be computed from these data.

We end this section with the following proposition, which shows that it is possible to detect which components of $N_{r,R}(\gamma)$ meet $C_K(\gamma)$ just by looking at a ball of a bounded size in X . Therefore it is possible to construct bounded subsets of $A_{r,R,K}(\gamma)$ algorithmically.

Proposition 4.1.12. *Suppose that γ is a bi-infinite λ -quasi-geodesic and that neither point in $\Lambda\gamma$ is a cut point. Let r and R be chosen to satisfy Propositions 4.1.8 and 4.1.9. Let T be at least $\log_a(2k_1/k_2) + 3D + 2\delta + K$. Then every component of $\text{Cyl}_{r,R,K}(\gamma)$ meets $C_K(\gamma) \cap B_T(\gamma(t))$ for any t such that $\gamma(t)$ is in X_0 .*

Proof. Without loss of generality assume that $\gamma(t) = 1$. Let U be a component of $\text{Cyl}_{r,R,K}(\gamma)$ and let U' be the component of $\text{Cyl}_{r,\infty,K}(\gamma)$ containing U . Then it is sufficient to show that U' meets $C_K(\gamma) \cap B_T(1)$ since $U' \cap C_K(\gamma) = U \cap C_K(\gamma)$ by Proposition 4.1.9.

Suppose that U' does not meet $C_K(\gamma) \cap B_T(1)$. We show then that $\mathcal{S}U$ must then be clustered around the two points in $\Lambda\gamma$. Let p be a point in $\mathcal{S}U$ and let α be a geodesic ray from 1 to p . Then $\alpha(s) \in U \cap C_K(\gamma)$ for some s ; by assumption $d(\alpha(s), v) \geq T$.

Let γ' be a geodesic connecting the points of $\Lambda\gamma$ so that the Hausdorff distance between γ and γ' is at most D . Parametrise γ' so that $d(\gamma'(0), 1) \leq D$. Then $d(\alpha(s), \gamma') \leq K + D$; let $d(\alpha(s), \gamma'(s')) \leq K + D$. This implies that $d(\gamma'(s'), 1) \geq T - D - K$. We therefore have the following inequality.

$$(\alpha(s) \cdot \gamma'(s'))_1 \geq T - D - K. \quad (4.16)$$

Assume that $s' \geq 0$; this implies that

$$(p \cdot \gamma'(\infty))_1 \geq \liminf_{m,n \rightarrow \infty} (\alpha(m) \cdot \gamma'(n))_1 \quad (4.17)$$

$$\geq (\alpha(s) \cdot \gamma'(s'))_1 - D \quad (4.18)$$

$$\geq T - 2D - K. \quad (4.19)$$

So $d_v(p, \gamma'(\infty)) \leq k_2 a^{-T+2D+K}$. Similarly, if $s' \leq 0$, $d_v(p, \gamma'(-\infty)) \leq k_2 a^{-T+2D+K}$. Therefore $\mathcal{S}U'$ is contained in a $k_2 a^{-T+2D+K}$ neighbourhood of $\Lambda\gamma$. Also, for any s , the geodesic from $\gamma(t+s)$ to $\gamma(t-s)$ passes within a distance D of $\gamma(t)$, so $(\gamma(t+s) \cdot \gamma(t-s))_1 \leq D$. Then $(\gamma(\infty) \cdot \gamma(-\infty))_1 \leq 2\delta + D$, and so $d_v(\gamma(\infty), \gamma(-\infty)) \geq k_1 a^{-(2\delta+D)}$. It follows by the inequality satisfied by T that the closed balls of radius $k_2 a^{-T+2D+K}$ around $\gamma(\infty)$ and

$\gamma(-\infty)$ are disjoint. By Proposition 4.1.8 $\mathcal{S}U'$ is connected, so is contained in one of these two balls, say in the ball around $\gamma(\infty)$. But then $\mathcal{S}U$ is a non-empty proper subset of $\partial X - \{\gamma(\infty)\}$ that is closed and open, so $\gamma(\infty)$ is a cut point, which is a contradiction. \square

4.2 Cylinders in the thin part of the cusped space

As in the previous section, let G be a hyperbolic group and let \mathcal{H} be a finite set of pairwise non-conjugate maximal virtually cyclic subgroups of G , so that G is hyperbolic relative to \mathcal{H} . Let S be a generating set for G such that $S \cap H$ generates H for each $H \in \mathcal{H}$. Let X be the cusped space associated to (G, \mathcal{H}, S) . Let γ be a $(8\delta + 1)$ -local-geodesic in X with one or both end points in ∂X .

In this section we study the connectedness of the intersection of $\text{Cyl}_{r,R,K}(\gamma)$ with the thin part of X . Because of our assumption that the peripheral subgroups are virtually cyclic, the geometry of the cusps of X is quite simple. We begin by restricting to the case in which γ is a vertical geodesic ray contained in a horoball corresponding to a subgroup $H \in \mathcal{H}$. Then the vertex set of $\text{Hor}(\Gamma(H, S \cap H))$ can be identified with $H \times \mathbb{Z}_{\geq 0}$ and for any k there is an inclusion of the Cayley graph $\Gamma(H, S \cap H) \hookrightarrow h^{-1}(k) \cap \text{Hor}(\Gamma(H, S \cap H))$ mapping a vertex $g \in \Gamma(H, S)$ to (g, k) . For $k = 0$ this inclusion is an isomorphism of graphs.

Let d_H be the path metric in $\Gamma(H, S \cap H)$ and let $\Gamma(H, S \cap H)$ be δ_H -hyperbolic with respect to this metric. Let α be a bi-infinite geodesic in $\Gamma(H, S)$ with respect to d_H . Then any point in $\Gamma(H, S)$ is within a distance of at most $2\delta_H + 1$ of α .

Let $\gamma: [0, \infty) \rightarrow X$ be a vertical geodesic ray with $\gamma(0) = (\alpha(0), 0)$. The δ -thin part of the horoball is convex by [GM08, Lemma 3.26], so for any point x in $\text{Hor}(\Gamma(H, S \cap H))$ with $h(x) \geq \delta$, the distance from x to γ is minimised by a geodesic comprising a vertical segment followed by a single horizontal edge. Then for any $k \geq \delta$, the vertex set of $h^{-1}(k) \cap N_{r,R}(\gamma)$ is

$$\{g \in H: 2^{k+r-2} < d_H(g, \alpha(0)) \leq 2^{k+R-1}\} \times \{k\}. \quad (4.20)$$

We will denote by Y_k the set $h^{-1}(k) \cap N_{r,R}(\gamma)$.

Assume now that $k \geq \max\{\log_2(2\delta_H + 1), \delta\}$. Then every vertex in $h^{-1}(k)$, and therefore every vertex in Y_k , is adjacent to a vertex in $\alpha \times \{k\}$. Therefore Y_k contains connected

components Y_k^+ and Y_k^- with

$$\alpha|_{[2^{k+r-1}, 2^{k+R-1}] \times \{k\}} \subset Y_k^+, \quad (4.21)$$

$$\alpha|_{[-2^{k+R-1}, -2^{k+r-1}] \times \{k\}} \subset Y_k^-, \quad (4.22)$$

$$(4.23)$$

and each of these components meets $C_K(\gamma)$. Therefore each of the sets Y_k^\pm is a subset of a component of $\text{Cyl}_{r,R,K}(\gamma)$. Any vertex in the complement of these two components of Y_k is adjacent to a point in $\{(\alpha(\pm 2^{k+r-2}), k)\}$, so is contained in the following set.

$$\{g \in H: d_H(g, \alpha(0)) \leq 2^{k+r-1}\} \times \{k\}. \quad (4.24)$$

Therefore only those vertices of Y_k that are in $Y_k^+ \cup Y_k^-$ are adjacent in X to vertices of Y_{k+1} . Furthermore, Y_k^+ is adjacent to Y_{k+1}^+ and not to Y_{k+1}^- and likewise for Y_k^- . Finally, any vertex that is in Y_{k+1} but not in $Y_{k+1}^+ \cup Y_{k+1}^-$ is adjacent to a vertex in $Y_k^+ \cup Y_k^-$.

Therefore, if $k \geq \max\{\log_2(2\delta_H + 1), \delta\}$ then $\text{Cyl}_{r,R,K}(\gamma) \cap h^{-1}[k, \infty) \cap \text{Hor}(\Gamma(H, S \cap H))$ contains two unbounded components $Y_{\geq k}^+$ and $Y_{\geq k}^-$ containing $\cup_{l \geq k} Y_l^+$ and $\cup_{l \geq k} Y_l^-$ respectively and the complement of these two components is contained in

$$\{g \in H: d_H(g, \alpha(0)) \leq 2^{k+r}\} \times \{k\}. \quad (4.25)$$

To make precise the consequences of this description of $\text{Cyl}_{r,R,K}(\gamma)$, we make the following definition, now allowing γ to be any geodesic segment, geodesic ray or bi-infinite geodesic in X .

Definition 4.2.1. Let γ be a geodesic segment, geodesic ray or bi-infinite geodesic in X . We call an end point of γ a *loose end* if it is in X (as opposed to ∂X), and this end point is a local maximum of the function $h \circ \gamma$. (In other words, if the end point is at the start of γ then γ is descending vertically at that point, while if the end point is at the end of γ then γ is ascending vertically at that point.) Then let $\hat{\gamma}$ be the bi-infinite path obtained by concatenating γ with geodesic rays at any loose ends of γ . Note that this can be done in a unique way, since the loose ends of γ have positive height.

Note that if the height of the loose ends of γ is at least $8\delta + 1$, $\hat{\gamma}$ is a $(8\delta + 1)$ -local-geodesic.

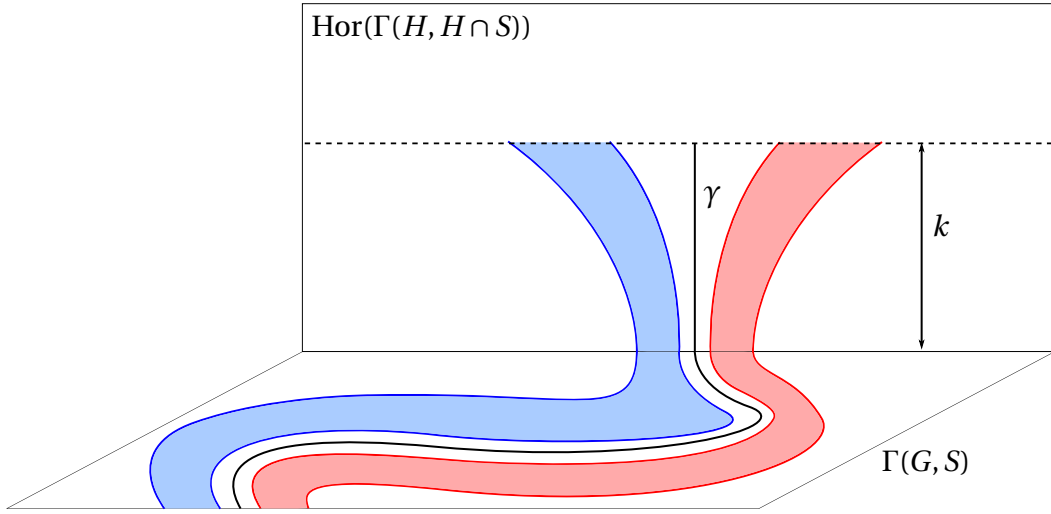


Figure 4.3: The union of the red and blue components of $\text{Cyl}'_{r,R,K}(\gamma)$, where γ is a geodesic with a loose end at height k . Note that extending γ by a geodesic ray (to obtain $\hat{\gamma}$) does not change the number of connected components of the cylinder.

With γ as in Definition 4.2.1, let $\text{Cyl}'_{r,R,K}(\gamma)$ be the space obtained by truncating $\text{Cyl}_{r,R,K}(\hat{\gamma})$ at the loose ends of γ :

$$\text{Cyl}_{r,R,K}(\hat{\gamma}) - (h^{-1}[k, \infty) \cap N_R(\hat{\gamma} - \gamma)). \quad (4.26)$$

This space is shown in Figure 4.3.

The results of this section together give the following lemma, which we shall use to control the depth to which geodesics in X connecting cut pairs in ∂X penetrate into the thin part of X .

Lemma 4.2.2. *Let γ be a geodesic segment, geodesic ray or bi-infinite geodesic in X such that any ends of γ not in ∂X are loose ends of γ and have height at least $\min\{R, \log_2(2\delta_{\mathcal{H}} + 1), 8\delta + 1\}$. Let $\delta_{\mathcal{H}}$ be such that each $H \in \mathcal{H}$ is $\delta_{\mathcal{H}}$ -hyperbolic with respect to the generating set $H \cap S$. Then the inclusion $\text{Cyl}'_{r,R,K}(\gamma) \hookrightarrow \text{Cyl}_{r,R,K}(\hat{\gamma})$ induces a bijection between the sets of connected components of those spaces. \square*

4.3 The geometry of finite balls in the cusped space

In this section we apply the results of Sections 4.1 and 4.2 to prove some of the results referred to in the introduction. We show that certain topological properties of the boundary of the pair (G, \mathcal{H}) are determined by the geometry of a large ball in the cusped space. All

of these results require the assumption that the cusped space satisfies a double dagger condition.

4.3.1 Cut points

The first topological feature we treat is a cut point. As discussed in Sections 2.2.3 and 2.3.3 there is never a cut point in the boundary in the Gromov boundary of a hyperbolic group, but cut points are possible in the Bowditch boundary of relatively hyperbolic groups. Note also that we require a double dagger condition, which is not guaranteed if the boundary does contain a cut point.

Proposition 4.3.1. *Let G be a group and let \mathcal{H} be a finite set of virtually cyclic subgroups of G so that G is hyperbolic relative to \mathcal{H} . Let S be a finite generating set for G such that $S \cap H$ generates H for each $H \in \mathcal{H}$. Suppose that the cusped space X associated to the triple (G, \mathcal{H}, S) is δ -hyperbolic and satisfies $\ddagger(n)$. Fix r, R and K satisfying Proposition 4.1.10. Let $\delta_{\mathcal{H}}$ be large enough that $\Gamma(H, S \cap H)$ is $\delta_{\mathcal{H}}$ -hyperbolic for each $H \in \mathcal{H}$ and let $k \geq \log_2(2\delta_{\mathcal{H}} + 1)$. Then $\partial(G, \mathcal{H}, S)$ contains a cut point if and only if there is a vertical geodesic segment $\gamma: [0, k] \rightarrow X$ with $\gamma(0) = 1$ such that $\text{Cyl}'_{r,R,K}(\gamma)$ is disconnected.*

Proof. It is shown in [Bow99b, Theorem 0.2] that any cut point in $\partial(G, \mathcal{H})$ must be the limit point of gHg^{-1} for some $H \in \mathcal{H}$ and $g \in G$. Therefore $\partial(G, \mathcal{H})$ contains a cut point if and only if ΛH is a cut point for some $H \in \mathcal{H}$. For $H \in \mathcal{H}$ let γ_H be the vertical geodesic ray contained in $\text{Hor}(\Gamma(H, S \cap H))$ with $\gamma_H(0) = 1$. Note that $\gamma_H(\infty) = \Lambda H$.

By Proposition 4.1.10, $\gamma_H(\infty)$ is a cut point if and only if $\text{Cyl}_{r,R,K}(\gamma_H)$ is disconnected. By Lemma 4.2.2 $\text{Cyl}_{r,R,K}(\gamma_H)$ is disconnected if and only if $\text{Cyl}'_{r,R,K}(\gamma_H|_{[0,k]})$ is disconnected. \square

4.3.2 Cut pairs

We now assume that $\partial(G, \mathcal{H})$ does not contain a cut point and give criteria for the existence of a cut pair in $\partial(G, \mathcal{H})$. In this situation all results of Section 4.1.2 are applicable. We use a pumping lemma argument: we aim to replace an arbitrary geodesic joining the two points in a cut pair in ∂X with a periodic local geodesic with bounded period that also joins the points of a (possibly different) cut pair.

Before stating the theorem we define some constants. Let the cusped space X be δ -hyperbolic and satisfy the double dagger condition $\ddagger(n)$. Fix $\delta_{\mathcal{H}}$ so that $\Gamma(H, S \cap H)$ is $\delta_{\mathcal{H}}$ -hyperbolic for each $H \in \mathcal{H}$. Take $\lambda = \max\{(12\delta + 1)/(4\delta + 1), 2\delta\}$ so that any $(8\delta + 1)$ -local-geodesic is a λ -quasi-geodesic as in Theorem 2.1.17 and fix D so that any $(8\delta + 1)$ -local geodesic is contained in a D -neighbourhood of a geodesic with the same end points.

Fix r , R and K to simultaneously satisfy the conditions of the results of Section 4.1.2 as summarised in Proposition 4.1.10. Fix T to satisfy the condition of Proposition 4.1.12. We also make the following definitions.

$$k = \max\{8\delta + 1, \log_2(2\delta_{\mathcal{H}} + 1), T + R\}. \quad (4.27)$$

$$\rho = (2R + \lambda + 1)\lambda^2 + \lambda + R. \quad (4.28)$$

$$\eta = \max\{(8\delta + 1)/2, \lambda(T + K + \lambda), \lambda(R + r + \lambda), \lambda(R + \rho + \lambda), \lambda(2R + \lambda + 1)\}. \quad (4.29)$$

To summarise, we take $r \ll K \ll R \ll \rho \ll \eta$ and $k \gg R$.

Let Val be the maximum valence of any vertex in X_{k+R} and let V be the maximum number of vertices in any ball of radius ρ around any vertex in X_{k+R} . Then define N_{\min} and $N_{\max} \gg \eta$ as follows.

$$N_{\min} = \max\{8\delta + 1, \lambda(2R + \lambda + 1) + 1\} \quad (4.30)$$

$$N_{\max} = N_{\min} (k + R + 1) \text{Val}^{2\eta} 2^V + 1 \quad (4.31)$$

Theorem 4.3.2. *There is a cut pair in $\partial(G, \mathcal{H})$ if and only if and only if X contains one of the following two features:*

1. *A short coarsely separating geodesic segment at shallow depth in X : a geodesic segment $\gamma: [a - \eta, b + \eta] \rightarrow X$ contained in X_{k+R} satisfying the following conditions. Such a feature is shown in Figure 4.4.*
 - (a) *The segment is short: $N_{\min} \leq b - a \leq N_{\max}$.*
 - (b) *The ends of the segment match: $h(\gamma(a)) = h(\gamma(b))$, so there exists $g \in G$ such that $\gamma(b) = g \cdot \gamma(a)$, and $\gamma|_{[b-\eta, b+\eta]} = g \cdot \gamma|_{[a-\eta, a+\eta]}$.*
 - (c) *The segment coarsely separates X : there is a partition \mathcal{P} of the vertices of $N_{r,R}(\gamma) \cap N_R(\gamma|_{[a,b]})$ into two subsets such that adjacent vertices lie in the same subset and each of the sets meets $C_K(\gamma) \cap B_T(\gamma(c))$ for some $c \in [a, b]$ such that $\gamma(c) \in X_0$.*
 - (d) *The ends of the partition match: the partition on the vertices of $N_{r,R}(\gamma) \cap B_\rho(\gamma(b))$ induced by the restriction of \mathcal{P} to that subset is the same as the translate by g of the partition on the vertices of $N_{r,R}(\gamma) \cap B_\rho(\gamma(a))$ obtained by restricting \mathcal{P} . Note that $N_{r,R}(\gamma) \cap B_\rho(\gamma(b))$ is equal to $g \cdot N_{r,R}(\gamma) \cap B_\rho(\gamma(a))$ by Condition 1b.*
2. *A short coarsely separating horseshoe-shaped geodesic: a geodesic segment $\gamma: [a, b] \rightarrow X$ satisfying the following conditions.*

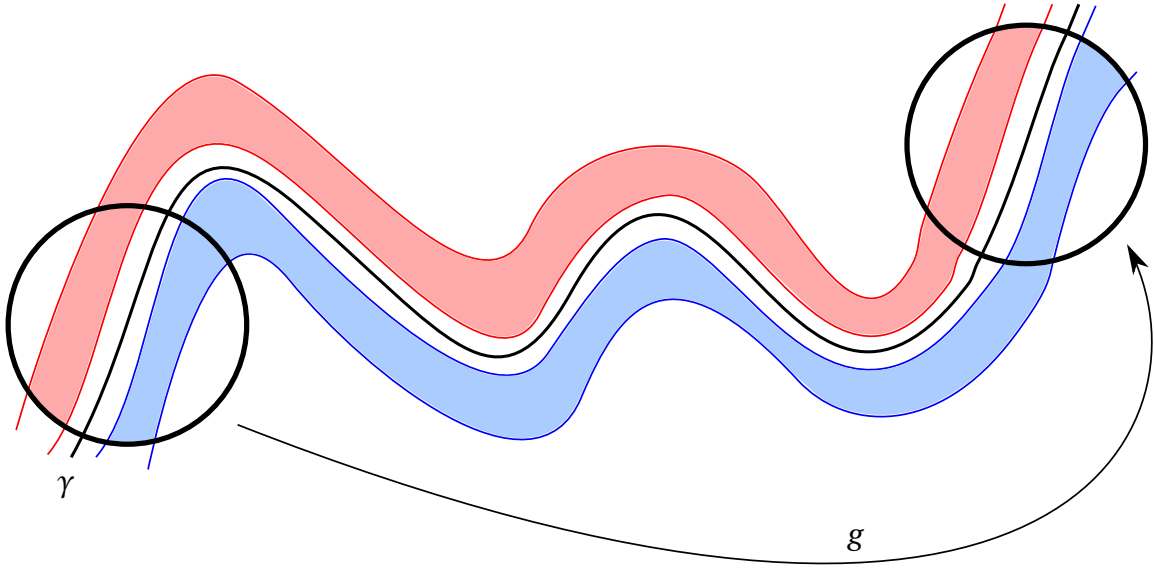


Figure 4.4: A short coarsely separating geodesic segment in X . The partition \mathcal{P} consists of the red set and the blue set. Note that the two ends of the segment, and the partitions around the two ends, match under translation by g .

- (a) *The segment is short: $b - a \leq N_{max} - 2R + 2\eta$.*
- (b) *Both ends are loose ends: $h(\gamma(a)) = h(\gamma(b)) \geq k$ and $h \circ \gamma$ is decreasing at a and increasing at b .*
- (c) *The segment coarsely separates X : $\text{Cyl}'_{r,R,K}(\gamma)$ is disconnected.*

Proof. First suppose that the first type of feature exists in X . Define a path γ' in X by concatenating translates of γ :

$$\gamma'((b-a)m + t) = g^m \gamma(a + t) \quad (4.32)$$

for $m \in \mathbb{Z}$ and $t \in [0, b-a]$. Then γ' is an $(8\delta + 1)$ -local-geodesic by Condition 1b. (Here we use the condition that $\eta \geq (8\delta + 1)/2$.) It is therefore a λ -quasi-geodesic by Theorem 2.1.17. We now show that $\text{Cyl}_{r,R,K}(\gamma')$ is disconnected.

Note that $N_R(\gamma')$ is a union $\bigcup_{m \in \mathbb{Z}} g^m \cdot N_R(\gamma|_{[a,b]})$ of translates of neighbourhoods of γ . Since $\eta \geq \lambda(R + r + \lambda)$, $N_r(\gamma') \cap N_R(\gamma|_{[a,b]})$ is a subset of $N_r(\gamma)$, so is equal to $N_r(\gamma) \cap N_R(\gamma|_{[a,b]})$. Therefore $N_{r,R}(\gamma')$ is also a union:

$$N_{r,R}(\gamma') = \bigcup_{m \in \mathbb{Z}} g^m \cdot (N_{r,R}(\gamma) \cap N_R(\gamma|_{[a,b]})) \quad (4.33)$$

Since $b - a > \lambda(2R + \lambda + 1)$, $g^m \cdot N_R(\gamma|_{[a,b]})$ and $g^l \cdot N_R(\gamma|_{[a,b]})$ contain no adjacent vertices for $|m - l| \geq 2$. Furthermore, if $l = m + 1$ then any pair of adjacent vertices in these two sets is contained in $g^m \cdot B_\rho(\gamma(b))$ since $\rho \geq (2R + \lambda + 1)\lambda^2 + \lambda + R$.

For each set $U \in \mathcal{P}$ define a set U' of vertices of $N_{r,R}(\gamma')$ by letting $u \in U'$ if $g^m u \in U$ for some $m \in \mathbb{Z}$. This gives a well defined partition \mathcal{P}' of the vertices of $N_{r,R}(\gamma')$ such that adjacent vertices lie in the same set by condition 1d. Its restriction to $\text{Cyl}_{r,R,K}(\gamma')$ is non-trivial: $C_K(\gamma')$ contains $C_K(\gamma) \cap B_T(\gamma(c))$ since $\eta \geq \lambda(T + K + \lambda)$ and this set meets both sets in \mathcal{P}' by condition 1c. Therefore $\text{Cyl}_{r,R,K}(\gamma')$ is disconnected and therefore $\Lambda\gamma'$ is a cut pair by the results of Section 4.1.

Now suppose that the second type of feature exists in X . As in Section 4.2, let $\hat{\gamma}$ be the $(8\delta + 1)$ -local-geodesic obtained by concatenating γ with vertical geodesic rays. Then $\text{Cyl}_{r,R,K}(\hat{\gamma})$ is disconnected by Lemma 4.2.2 and $\Lambda\hat{\gamma}$ is a cut pair by the results of Section 4.1.

Conversely, suppose that ∂X does contain a cut pair. Let γ' be a geodesic in X such that $\Lambda\gamma'$ is a cut pair. Assume first that some connected component of $\gamma'^{-1}h^{-1}[0, k + R]$ is an interval of length less than $N_{max} + 2R$, say $[a - R, b + R]$ with $h(\gamma'(a - R)) = h(\gamma'(b + R)) = k + R$, so $h(\gamma'(a)) = h(\gamma'(b)) = k$. Shortening γ if necessary, we may assume that $h \circ \gamma$ is decreasing at a and increasing at b . Let $\gamma = \gamma'|_{[a,b]}$. Let $c \in [a, b]$ such that $h(\gamma'(c)) = 0$. $\text{Cyl}_{r,R,K}(\gamma')$ is disconnected and each component meets $C_K(\gamma') \cap B_T(\gamma'(c))$ by Proposition 4.1.12. Since $k \geq T$ this is a subset of $\text{Cyl}'_{r,R,K}(\gamma)$, and $\text{Cyl}'_{r,R,K}(\gamma)$ is a subset of $\text{Cyl}_{r,R,K}(\gamma')$, so $\text{Cyl}'_{r,R,K}(\gamma)$ is disconnected. Therefore γ' is a feature of the second kind described in the proposition.

On the other hand, suppose that some interval $[-\eta, N_{max} + \eta]$ is a subset of $\gamma'^{-1}h^{-1}[0, k + R]$. Then there exist $a_0 < a_1 < \dots < a_{2^V}$ in $[0, N_{max}]$ such that $h(a_i) = h(a_j)$ for all i and j , so $a_i = g_i a_0$ for some $g_i \in G$, such that $\gamma'|_{[a_i - \eta, a_i + \eta]} = g_i \cdot \gamma'|_{[a_0 - \eta, a_0 + \eta]}$, and such that $a_i - a_{i-1} \geq N_{min}$. Let \mathcal{P}' be a partition of the vertices of $N_{r,R}(\gamma')$ into two subsets such that adjacent vertices are in the same set and so that both sets meet $C_K(\gamma')$. Such a partition exists by the results of section 4.1.2 since $\Lambda\gamma'$ is a cut pair. Since $\eta \geq \lambda(R + \rho + \lambda)$, $N_{r,R}(\gamma') \cap B_\rho(\gamma(a_0))$ is equal to $g_i^{-1}N_{r,R}(\gamma') \cap B_\rho(\gamma(a_i))$ for all i . This set contains at most V vertices, so there exist $0 \leq i < j \leq 2^V$ such that

$$g_i^{-1}\mathcal{P}'|_{N_{r,R}(\gamma') \cap B_\rho(\gamma(a_i))} = g_j^{-1}\mathcal{P}'|_{N_{r,R}(\gamma') \cap B_\rho(\gamma(a_j))}. \quad (4.34)$$

Let $a = a_i$ and $b = a_j$ and let $\gamma = \gamma'|_{[a-\eta, b+\eta]}$. We claim that γ is then a feature of the first kind described in the proposition. Setting $g = g_j g_i^{-1}$ and $\mathcal{P} = \mathcal{P}'|_{N_{r,R}(\gamma') \cap N_R(\gamma|_{[a,b]})}$, conditions 1a, 1b, and 1d are satisfied by definition of the a_i . Let $c \in [a, b]$ such that $\gamma'(c) \in X_0$. Then $C_K(\gamma) \cap B_T(\gamma(c))$ is equal to $C_K(\gamma') \cap B_T(\gamma'(c))$ since $\eta \geq \lambda(T + K + \lambda)$ and Proposition 4.1.12 guarantees that both sets in \mathcal{P}' meet this set, so Condition 1c is satisfied, too. \square

4.3.3 Circular boundaries

We now show that the boundary is homeomorphic to a circle if and only if a large ball in the cusped space satisfies a geometric condition. To identify circles we require the following theorem from point-set topology.

Theorem 4.3.3. [Wil49, II.2.13] *Let X be a separable, connected, locally connected space containing more than one point that is without a cut point. If every pair of points in X is a cut pair then X is homeomorphic to S^1 .*

Recall a theorem of Bowditch [Bow99a, Theorem 0.1]: if G is hyperbolic relative to \mathcal{H} where every element of \mathcal{H} is finitely presented, one- or two-ended and contains no infinite torsion subgroup then $\partial(G, \mathcal{H})$ is locally connected whenever it is connected. The boundary of a relatively hyperbolic group is a compact metric space, so it is separable. Therefore we see that the boundary of a relatively hyperbolic group is homeomorphic to a circle if and only if it is connected, does not contain a cut point and every pair of points is a cut pair. We call a pair of points that is not a cut pair a *non-cut pair*. This reduces the problem of identifying circular boundaries to the problem of determining whether or not $\partial(G, \mathcal{H})$ contains a non-cut pair.

By a similar argument to that of Section 4.3.2 we now show that the existence of a non-cut pair is equivalent to the presence in a ball of large radius in the cusped space of a short coarsely non-separating geodesic segment. First we prove the following lemma.

Lemma 4.3.4. *Suppose that ∂X does not contain a cut point. Let $(x_n)_{n \in \mathbb{Z}}$ be a sequence of points in ∂X with $x_n \rightarrow x_{\pm\infty}$ as $n \rightarrow \pm\infty$. Suppose that each pair $\{x_n, x_{n+1}\}$ is a cut pair. Then so is $\{x_{-\infty}, x_{\infty}\}$.*

Proof. The boundary ∂X is locally connected by the main theorem of [Bow99a]. By assumption ∂X does not contain a cut point, so the results of [Bow98b, Sections 2 and 3] can be applied. We recall some definitions from that paper; these definitions were previously discussed in Section 3.5.2. For $x \in \partial X$ we define $\text{val}(x) \in \mathbb{N}$ to be the number of ends of $\partial X - \{x\}$. Then we let $M(n) = \{x \in \partial X : \text{val}(x) = n\}$ and $M(n+) = \{x \in \partial X : \text{val}(x) \geq n\}$. For x and y in $M(2)$ we write $x \sim y$ if $x = y$ or $\{x, y\}$ is a cut pair; this defines an equivalence relation. For x and y in $M(3+)$ we write $x \approx y$ if $\text{val}(x) = \text{val}(y)$ and $\partial X - \{x, y\}$ has exactly $\text{val}(x)$ components.

Recall [Bow98b, Lemma 3.8]: if $x \approx y$ and $x \approx z$ then $y \approx z$. Therefore $x_n \in M(2)$ for all n , so $\{x_n\}_{n \in \mathbb{Z}}$ is a subset of a \sim -equivalence class σ . By [Bow98b, Lemma 3.2] σ is a cyclically separating set, and so is the closure of σ by [Bow98b, Lem. 2.2], which implies that $\{x_{-\infty}, x_{\infty}\}$ is a cut pair as required. \square

Let $\lambda, r, R, K, T, k, \rho, \eta, \text{Val}$ and V be constants as defined in Section 4.3.2. Let N_1, N_2 and N_3 be given by the following expressions.

$$N_1 = 2(V-1)((k+R+1)\text{Val}^{2\eta}V^{V+1} + 2\eta) + 2\eta + 2((k+R+1)\text{Val}^{2\eta} + 1), \quad (4.35)$$

$$N_2 = (k+R+1)\text{Val}^{2\eta} + 1, \quad (4.36)$$

$$N_3 = 2(k+R+1)\text{Val}^{2\eta}V^{V+1} + 4\eta. \quad (4.37)$$

Theorem 4.3.5. *The boundary $\partial(G, \mathcal{H})$ is homeomorphic to a circle if and only if X contains one of the following two features.*

1. *A short non-separating geodesic: geodesic segments $\gamma_i: [a_i - \eta, b_i + \eta] \rightarrow X$ with image in X_k for $i = 1, 2, 3$ with $a_2 = b_1$ and $a_3 = b_2$ satisfying the following conditions.*

(a) *The segments are short: $1 \leq b_i - a_i \leq N_1$ for $i = 1$ and for $i = 3$, and $1 \leq b_2 - a_2 \leq N_2$.*

(b) *The ends of the segments match: $\gamma_i|_{[b_i - \eta, b_i + \eta]} = \gamma_{i+1}|_{[a_i - \eta, a_i + \eta]}$ for $i = 1, 2$. Also, $h(\gamma_i(a_i)) = h(\gamma_i(b_i))$ for $i = 1, 3$, so there exist $g_i \in G$ such that $\gamma_i(b_i) = g_i\gamma_i(a_i)$. Finally, $\gamma_i|_{[b_i - \eta, b_i + \eta]} = g_i \cdot \gamma_i|_{[a_i - \eta, b_i + \eta]}$ for $i = 1, 3$.*

(c) *The segment does not coarsely separate: all vertices of $C_K(\gamma_2) \cap B_T(\gamma_2(c))$ lie in the same connected component of $N_{r,R}(\gamma_2) \cap N_R(\gamma_2|_{[a_2, b_2]})$ for some $c \in [a_2, b_2]$ such that $\gamma_2(c) \in X_0$.*

2. *A horseshoe shaped non-separating geodesic: a geodesic segment $\gamma: [a, b] \rightarrow X$ with image in X_k satisfying the following conditions.*

(a) *The segment is short: $b - a \leq N_3$.*

(b) *Both ends are loose ends: $h(\gamma(a)) = h(\gamma(b)) = k$ with γ descending vertically at a and ascending vertically at b .*

(c) *The segment does not coarsely separate: $\text{Cyl}'_{r,R,K}(\gamma)$ is connected.*

Proof. First suppose that X contains a feature of the first kind described in the proposition. Then define a path γ' in X as follows:

$$\gamma'(t) = \begin{cases} g_1^m \cdot \gamma_1(t') & \text{if } t = m(b_1 - a_1) + t' \text{ for } m \in \mathbb{Z}_{\leq 0} \text{ and } t' \in [a_1, b_1] \\ \gamma_2(t) & \text{if } t \in [a_2, b_2] \\ g_3^m \cdot \gamma_3(t') & \text{if } t = m(b_3 - a_3) + t' \text{ for } m \in \mathbb{Z}_{\geq 0} \text{ and } t' \in [a_3, b_3] \end{cases} \quad (4.38)$$

That is, γ' is obtained by concatenating infinitely many translates of γ_1 , then a copy of γ_2 , then infinitely many translates of γ_3 . Note that this is an $(8\delta + 1)$ -local-geodesic by Condition 1b since $\eta \geq 8\delta + 1$ and is therefore a λ -quasi-geodesic.

Since $\gamma'|_{[a_2-\eta, b_2+\eta]} = \gamma_2$ and $\eta \geq \lambda(T + K + \lambda)$, $C_K(\gamma') \cap B_T(\gamma'(c))$ is equal to $C_K(\gamma_2) \cap B_T(\gamma_2(c))$. Furthermore, $\eta \geq \lambda(R + r + \lambda)$, which similarly guarantees that $N_{r,R}(\gamma_2) \cap N_R(\gamma_2|_{[a_2, b_2]})$ is a subset of $N_{r,R}(\gamma')$. Therefore $C_K(\gamma') \cap B_T(\gamma'(c))$ lies in a single component of $\text{Cyl}_{r,R,K}(\gamma')$ by Condition 1c, so $\text{Cyl}_{r,R,K}(\gamma')$ is connected by Proposition 4.1.12, so $\Lambda\gamma'$ is a non-cut pair by the results of Section 4.1.2.

Now suppose that a feature of the second type exists in X . Let $\hat{\gamma}$ be the path obtained by concatenating γ with vertical geodesic rays. Then $\hat{\gamma}$ is an $(8\delta + 1)$ -local-geodesic since $k \geq 8\delta + 1$. Also, $\text{Cyl}_{r,R,K}(\hat{\gamma})$ is connected by Lemma 4.2.2, so $\Lambda\hat{\gamma}$ is a non-cut pair by the results of Section 4.1.2.

Conversely, suppose that ∂X contains a non-cut pair. Let γ' be an $(8\delta + 1)$ -local-geodesic in X such that $\Lambda\gamma'$ is such a pair. Assume first that γ' is contained in X_{k+R} and reparametrise γ' so that $\gamma'(0) \in X_0$. Then $C_K(\gamma') \cap B_T(\gamma'(0))$ contains at most V vertices and lies in a single component of $N_{r,R}(\gamma')$. Define n_V^\pm to be $\pm\lambda(K + T + \lambda)$ and then for l decreasing from $V - 1$ to 1 let n_l^+ and n_l^- be chosen to minimise $n_l^+ - n_l^-$ among pairs such that $C_K(\gamma') \cap B_T(\gamma'(0))$ meets at most l components of

$$N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_l^-, n_l^+]}) \quad (4.39)$$

and $[n_{l+1}^-, n_{l+1}^+] \subset [n_l^-, n_l^+]$. Note that the condition that $|n_l^\pm| \geq \lambda(K + T + \lambda)$ ensures that $C_K(\gamma') \cap B_T(\gamma'(0))$ is a subset of $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_l^-, n_l^+]})$.

Suppose that $n_{l-1}^+ - n_l^+ > (k + R + 1)\text{Val}^{2\eta} V^{V+1} + 2\eta$. Let \mathcal{Q} be the partition of $C_K(\gamma') \cap B_T(\gamma'(0))$ into l non-empty subsets induced by connectivity in $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_{l-1}^-, n_{l-1}^+ - 1]})$. For each $t \in [n_l + \eta, n_{l-1} - \eta]$ define the following sets:

$$Y_t = N_{r,R}(\gamma') \cap B_\rho(\gamma'(t)) \quad (4.40)$$

$$Z_t = N_{r,R}(\gamma') \cap \left(B_\rho(\gamma'(t)) \cup N_R(\gamma'|_{[n_{l-1}^-, t]}) \right) \quad (4.41)$$

and let \mathcal{P}_t be the partition of the vertices of Y_t into $l + 1$ subsets: l corresponding to the l sets in \mathcal{Q} by connectivity in Z_t and one containing the part of Y_t not connected to $C_K(\gamma') \cap B_T(\gamma'(0))$ in Z_t . As in the proof of Theorem 4.3.2 there exist $s_1 < s_2$ in $[n_l^+ + \eta, n_{l-1}^+ - \eta]$ such that the following conditions hold.

1. The geodesic γ' matches at s_1 and s_2 : $h(\gamma'(s_1)) = h(\gamma'(s_2))$, so $\gamma'(s_2) = g\gamma'(s_1)$ for some g in G , and $\gamma'|_{[s_2-\eta, s_2+\eta]} = g \cdot \gamma'|_{[s_1-\eta, s_1+\eta]}$, which implies that $Y_{s_2} = g \cdot Y_{s_1}$.

2. The partitions matches at s_1 and s_2 : the partition \mathcal{P}_{s_2} is equal to the translation of \mathcal{P}_{s_1} by the group element g .

Then replace γ' by another path defined by

$$\gamma''(t) = \begin{cases} \gamma'(t) & \text{if } t \leq s_1 \\ g^{-1} \cdot \gamma'(t + (s_2 - s_1)) & \text{if } t \geq s_2 \end{cases} \quad (4.42)$$

This is an $(8\delta + 1)$ -local-geodesic since $\eta \geq 8\delta + 1$. By definition of ρ , the intersection of Z_{s_2} and $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[s_2, n_{l-1}^+]})$ is contained in Y_{s_2} . Then by definition of n_{l-1}^+ , two distinct sets in \mathcal{P}_{s_2} meet $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[s_2, n_{l-1}^+]})$. Since $\eta \geq \lambda(R + r + \lambda)$,

$$g^{-1}N_{r,R}(\gamma') \cap N_R(\gamma'|_{[s_2, n_{l-1}^+]}) = N_{r,R}(\gamma'') \cap N_R(\gamma''|_{[s_1, n_{l-1}^+ - (s_2 - s_1)]}), \quad (4.43)$$

and it follows that $C_K(\gamma'') \cap B_T(\gamma'(0))$ meets at most $l - 1$ components of

$$N_{r,R}(\gamma'') \cap N_R(\gamma''|_{[n_{l-1}^-, n_{l-1}^+ - (s_2 - s_1)]}). \quad (4.44)$$

This implies that the process of replacing γ' by γ'' leaves unchanged $n_{l'}^+$ for all $l' < l$ and $n_{l'}^-$ for all l' and strictly reduces n_l^+ . Therefore by repeating this process we can assume that γ' was chosen to ensure that $|n_{l-1}^\pm - n_l^\pm|$ is at most $(k + R + 1)\text{Val}^{2\eta}V^{V+1} + 2\eta$ for all l , and therefore that

$$|n_1^\pm| \leq \lambda(K + T + \lambda) + (V - 1)((k + R + 1)\text{Val}^{2\eta}V^{V+1} + 2\eta). \quad (4.45)$$

There exist $a_1 \leq b_1 \leq n_1^- - \eta$ with $n_1^- - b_1$ and $b_1 - a_1$ both at most $(k + R + 1)\text{Val}^{2\eta} + 1$ such that the following matching conditions hold.

1. Firstly, $h(\gamma'(b_1)) = h(\gamma'(a_1))$, so $\gamma'(b_1) = g_1 \cdot \gamma'(a_1)$ for some $g_1 \in G$.
2. Secondly, $\gamma'|_{[b_1 - \eta, b_1 + \eta]} = g_1 \cdot \gamma'|_{[a_1 - \eta, a_1 + \eta]}$.

Then let $\gamma_1 = \gamma'|_{[a_1 - \eta, b_1 + \eta]}$. Similarly define $b_3 \geq a_3 \geq n_1^+$ and let $\gamma_3 = \gamma'|_{[a_3 - \eta, b_3 + \eta]}$. Let $a_2 = b_1$ and $b_2 = a_3$ and let $\gamma_2 = \gamma'|_{[a_2 - \eta, b_2 + \eta]}$; note that $b_2 - a_2 \leq N_1$. Then the triple $(\gamma_1, \gamma_2, \gamma_3)$ is a feature in X of the first kind listed in the proposition: conditions 1a and 1b clearly hold by construction and Condition 1c holds because the condition that $\eta \geq \lambda(R + r + \lambda)$ ensures that $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_1^-, n_1^+]})$ is a subset of $N_{r,R}(\gamma_2) \cap N_R(\gamma_2|_{[a_2, b_2]})$.

If γ' is not contained in X_{k+R} let $\gamma'^{-1}h^{-1}[0, k + R]$ be a (possibly infinite) union of (possibly infinite) intervals $\cup_{i \in I} [a_i - R, b_i + R]$ where we order the intervals so that $a_{i+1} > b_i$ for all i . Then $h(\gamma'(a_i)) = h(\gamma'(b_i)) = k$ for all i and we can apply the results of Section 4.2.

For $i \in I$ let $\gamma'_i = \gamma'|_{[a_i, b_i]}$. Let $\hat{\gamma}'_i$ be the bi-infinite $(8\delta + 1)$ -local-geodesic obtained by concatenating γ'_i with either one or two vertical geodesic rays.

Suppose that $\Lambda\hat{\gamma}'_i$ is a cut pair for all i . Then Lemma 4.3.4 tells us that $\Lambda\gamma'$ is a cut pair, which is a contradiction. Therefore there exists i such that $\Lambda\hat{\gamma}_i$ is a non-cut pair.

Arguing as before, the geodesic $\hat{\gamma}_i$ can be altered to ensure that $b_i - a_i \leq 4\eta + 2(k + R + 1)\text{Val}^{2\eta}V^{V+1}$; this yields a feature of the second type described in the proposition. \square

4.4 Generalisations

It seems possible that these methods could work for a wider class of relatively hyperbolic groups. The assumption that the peripheral subgroups are virtually cyclic is used in two essential steps. Firstly, we need the cusped space to satisfy a double dagger condition. As discussed in Section 3.6.1, the cusped space is known by work of Dahmani and Groves [DG08] to satisfy this condition when the peripheral subgroups are finitely generated abelian, and we showed in Proposition 3.6.9 that this result may be extended to allow virtually cyclic peripheral subgroups. The arguments used make use of fairly explicit descriptions of paths in the Cayley graphs of the peripheral subgroups, so any generalisation to a substantially broader class of peripheral subgroups would require new methods.

Secondly, in Section 4.2 we assumed that the peripheral subgroups are virtually cyclic in our analysis of the geometry of cylinders in the thin part of the cusped space. Lemma 4.2.2 followed from the two-endedness of the peripheral subgroups. It seems likely that a similar result will hold for one-ended peripheral subgroups: the cylinder $\text{Cyl}_{r,R,K}(\gamma)$ should be connected whenever γ strays too deep into a horoball based on a one-ended subgroup. However, without the double dagger condition this fact does not seem to be particularly useful.

Chapter 5

Algorithmic consequences

In this chapter we prove some algorithmic consequences of Theorems 4.3.2 and 4.3.5. Our ultimate aim is to prove Theorem 1.2.2: the computability of some JSJ decompositions of a one-ended hyperbolic group, including Bowditch's JSJ decomposition. The novel results of this chapter are taken from [Bar18, Sections 3–5].

Central to the algorithm of Theorem 1.2.2 is an algorithm that determines whether or not a given hyperbolic group with a (possibly empty) finite collection of virtually cyclic subgroups admits a non-trivial splitting as an amalgamated product or HNN extension over a virtually cyclic subgroup, relative to that collection of subgroups. One of the consequences of [Bow98b] is that in the absolute case (that is, if the collection of subgroups is empty), a one-ended hyperbolic group admits such a splitting if and only if its Gromov boundary contains a cut pair, at least as long as its boundary is not homeomorphic to a circle. In the case of interest here we obtain a relative version of this statement by replacing the Gromov boundary with the Bowditch boundary of the group relative to the given family of subgroups. Unlike the Gromov boundary, the Bowditch boundary might contain a cut point, in which case the group admits a relative splitting. This is the peripheral splitting in the sense of Bowditch [Bow99a]. In the absence of a cut point the existence of a relative splitting is determined by the existence of a cut pair, as in the absolute case. It is the presence of these topological features of the boundary that we use the results of Chapter 4 to show to be computable.

A maximal splitting is obtained from a JSJ decomposition by refining at the flexible vertices. Conversely, to obtain a JSJ decomposition we must decide which edges of the maximal splitting should be collapsed to reassemble the flexible vertices in the JSJ decomposition. In Bowditch's JSJ decomposition the stabilisers of flexible vertex groups are the maximal hanging fuchsian subgroups. These are those subgroups that occur as a vertex

stabiliser in some splitting such that the Bowditch boundary of the subgroup relative the stabilisers of the incident edges is homeomorphic to a circle.

In Section 5.1 we apply the results of Chapter 4: we prove that it is possible to detect algorithmically whether or not the boundary $\partial(G, \mathcal{H})$ of a hyperbolic group relative to a collection of virtually cyclic subgroups contains a cut point or cut pair, and whether or not $\partial(G, \mathcal{H})$ is homeomorphic to a circle.

In Section 5.2 we give a one-sided algorithm to compute non-trivial splittings of a hyperbolic group over its virtually cyclic subgroups. The algorithm eventually outputs every splitting, but we still need a way to tell whether or not a given splitting is maximal, so that we can tell when to stop the process. This algorithm requires tools for dealing with virtually cyclic subgroups of hyperbolic groups.

In Section 5.3 we deal with the special case of determining whether or not a group with circular boundary admits a non-trivial relative splitting over a virtually cyclic subgroup. We first recall some results that reduce the problem of determining whether or not such a group admits a proper splitting relative to its given virtually cyclic subgroups to the case in which the group is the fundamental group of a compact two-dimensional hyperbolic orbifold and the given subgroups are conjugacy class representatives of the fundamental groups of the boundary components of that orbifold. In [GL17] a complete list of such orbifolds that do not admit such a splitting is described; we use this list to complete this special case.

In Section 5.4 we show how the topology of the boundary determines whether or not a group admits a non-trivial relative splitting, and then use the results of Sections 5.1, 5.2 and 5.3 to show how to compute a maximal splitting of a given one-ended hyperbolic group over its virtually cyclic subgroups.

Finally, in Section 5.5 we put together these results to complete the proof of the computability of various JSJ decompositions. We describe a processes to convert a maximal splitting of a one-ended hyperbolic group over virtually cyclic subgroups into its \mathcal{VC} -JSJ, \mathcal{Z} -JSJ and \mathcal{Z}_{\max} -JSJ. The \mathcal{VC} -JSJ can be taken to be Bowditch's canonical decomposition.

5.1 Detecting cut points, cut pairs and circular boundaries

5.1.1 Computing the constants

To begin, we review some results that prove the computability of the various constants on which the criteria of Theorems 4.3.2 and 4.3.5 depend. Most importantly, we need an algorithm that computes a constant δ such that the cusped space is δ -hyperbolic. For

the Cayley graph of a hyperbolic group, the existence of such an algorithm is originally due to Gromov [Gro87]. A more detailed account of an alternative (and more practical) algorithm to accomplish this is due to Papasoglu [Pap95].

For relatively hyperbolic groups, the computability of δ such that the cusped space is δ -hyperbolic is essentially due to Dahmani. Dahmani's result proves that the constant of a linear isoperimetric inequality satisfied by the group is computable. We recall this result, then describe results of Groves and Manning that translate this linear isoperimetric inequality into a quantitative hyperbolicity result for the cusped space.

Proposition 5.1.1. *[Dah08, Prop. 2.3] There is an algorithm that takes as input a presentation for a group G , the generators in G for a finite set of subgroups of G with respect to which G is relatively hyperbolic, and a solution to the word problem in G and returns the constant of a linear relative isoperimetric inequality satisfied by the given presentation of G .*

In the case of interest here, G is hyperbolic. Hyperbolic groups have uniformly solvable word problem, so the requirement that the algorithm be given a solution to the word problem in G is no restriction to its applicability.

The linear relative isoperimetric inequality satisfied by the group is closely related to a linear combinatorial isoperimetric inequality satisfied by the coned-off Cayley complex. Refer to [GM08, Definition 2.47] for a definition of this space.

In [GM08] the cusped 2-complex is defined; this is a simply connected 2-complex, the 1-skeleton of which is the cusped graph defined in Definition 2.3.7. The length of the attaching map of each 2-cell in this complex is bounded above by the maximum of 5 and the length of the longest relator in the given presentation for G .

Theorem 5.1.2. *[GM08, Theorem 3.24] Suppose that the coned-off Cayley complex of G with respect to S and \mathcal{H} satisfies a linear combinatorial isoperimetric inequality with constant K . Then the cusped 2-complex X associated to the triple satisfies a linear combinatorial isoperimetric inequality with constant $3K(2K + 1)$.*

The computation of δ from a presentation for G and a set of generators for each group H in \mathcal{H} is therefore completed by the following proposition:

Proposition 5.1.3. *[GM08, Proposition 2.23] Suppose that a 2-complex X is simply connected, that each attaching map has length at most M and that X satisfies a linear combinatorial isoperimetric inequality. Then the 1-skeleton $X^{(1)}$ of X is δ -hyperbolic for some δ and this δ is computable from M and the constant of the isoperimetric inequality.*

We also need to be able to compute the constant $\delta_{\mathcal{H}}$ defined in Section 4.2. Recall that $\delta_{\mathcal{H}}$ is defined so that for each $H \in \mathcal{H}$, the Cayley graph $\Gamma(H, S \cap H)$ is $\delta_{\mathcal{H}}$ -hyperbolic.

Note that it is not enough to simply apply Papasoglu's algorithm [Pap95] to each $H \in \mathcal{H}$, since Papasoglu invokes an algorithm of Epstein, Paterson, Cannon, Holt, Levy and Thurston [Eps92, Theorem 5.2.4] that computes an automatic structure for a given automatic group. This requires a *presentation* for the group, which we are not given for the peripheral subgroups. Since we can compute large balls in the Cayley graph $\Gamma(H, S \cap H)$ using hyperbolicity of G , this deficiency might be fixable, at least with our assumption that H is a virtually cyclic subgroup of a hyperbolic group. However, rather than delve into the inner workings of this algorithm, we adopt an alternative approach.

Proposition 5.1.4. *There is an algorithm that takes as input a presentation for a hyperbolic group G with finite generating set S and a set of subsets of S generating a set \mathcal{H} of virtually cyclic subgroups of G , and returns an integer $\delta_{\mathcal{H}}$ such that $\Gamma(H, S \cap H)$ is $\delta_{\mathcal{H}}$ -hyperbolic for each $H \in \mathcal{H}$.*

Proof. Using Papasoglu's algorithm [Pap95] compute δ such that $\Gamma(G, S)$ is hyperbolic. Let $H \in \mathcal{H}$. Since H is virtually cyclic, it is quasi-convex in G . The quasi-convexity constant may be computed from the given data using an algorithm of Kapovich [Kap96, Proposition 4]. Equivalently, the inclusion $\Gamma(H, S \cap H) \hookrightarrow \Gamma(G, S)$ is a λ -quasi-isometric embedding by Lemma 2.1.21, where λ can be computed from the given data. If Δ is a geodesic triangle in $\Gamma(H, S \cap H)$ then it is a λ -quasi-geodesic triangle in $\Gamma(G, S)$, so by Theorem 2.1.16 is close to a geodesic triangle in $\Gamma(G, S)$, where the distance depends only on λ . Therefore geodesic triangles in $\Gamma(H, S \cap H)$ are uniformly slim with respect to the metric on $\Gamma(G, S)$, where the thinness constant depends only on δ and λ . But the inclusion $\Gamma(H, S \cap H) \hookrightarrow \Gamma(G, S)$ is a λ -quasi-isometric embedding, so geodesic triangles in $\Gamma(H, S \cap H)$ are uniformly thin with respect to the metric on $\Gamma(H, S \cap H)$, where the thinness depends in a computable way on δ and λ . Therefore there is a computable constant δ_H such that $\Gamma(H, S \cap H)$ is δ_H -hyperbolic.

Repeat this process for each $H \in \mathcal{H}$ and let $\delta_{\mathcal{H}} = \max\{\delta_H \mid H \in \mathcal{H}\}$. □

Finally, note also that there is an algorithm that computes n such that X satisfies $\ddagger(n)$: see Theorem 3.6.12.

5.1.2 Algorithms

Let G be a group and let \mathcal{H} be a finite collection of virtually cyclic subgroups of G such that G is hyperbolic relative to \mathcal{H} . Let S be a finite generating set for G such that $S \cap H$ is a generating set for H for each $H \in \mathcal{H}$. Note that the values of δ such that the associated cusped space X is δ -hyperbolic, n such that X satisfies $\ddagger(n)$, $\delta_{\mathcal{H}}$ such that $\Gamma(H, S \cap H)$ is $\delta_{\mathcal{H}}$

hyperbolic for each $H \in \mathcal{H}$, and $|S|$ together determine the values of all constants defined in Chapter 4, and that these constants can be computed from these data. Note also that the geometric features in X described in Proposition 4.3.1, Theorem 4.3.2 and Theorem 4.3.5 all have bounded size, so their existence or non-existence can be determined by looking only at some finite ball in X , where the radius of that ball depends on the constants defined in Chapter 4. Finally, note that for any given radius, the ball in X of that radius can be computed from a presentation for G with generating set S and for each $H \in \mathcal{H}$ the set $S \cap H$ of generators for H : this follows immediately from the uniform solvability of the word problem in hyperbolic groups.

From this discussion we deduce the following corollaries.

Corollary 5.1.5. *There is an algorithm that takes as input a presentation for a hyperbolic group G with generating set S and a list of subsets of S generating a collection \mathcal{H} of pairwise non-conjugate maximal virtually cyclic subgroups of G such that the cusped space associated to the triple (G, \mathcal{H}, S) satisfies a double dagger condition and returns the answer to the question “does $\partial(G, \mathcal{H})$ contain a cut point?”* \square

Corollary 5.1.6. *There is an algorithm that takes as input a presentation for a hyperbolic group G with generating set S and a list of subsets of S generating a collection \mathcal{H} of pairwise non-conjugate maximal virtually cyclic subgroups of G such that $\partial(G, \mathcal{H})$ is connected and does not contain a cut point and returns the answer to the question “does $\partial(G, \mathcal{H})$ contain a cut pair?”* \square

Corollary 5.1.7. *There is an algorithm that takes as input a presentation for a hyperbolic group G with generating set S and a list of subsets of S generating a collection \mathcal{H} of pairwise non-conjugate maximal virtually cyclic subgroups of G such that $\partial(G, \mathcal{H})$ is connected and does not contain a cut point and returns the answer to the question “is $\partial(G, \mathcal{H})$ homeomorphic to a circle?”* \square

Note that Corollaries 5.1.6 and 5.1.7 together prove Theorem 1.2.1.

5.2 Computing splittings

In this section we show that there is an algorithm that, when given as input a presentation for a hyperbolic group G , returns all splittings of G over its virtually cyclic subgroups. The algorithm does not terminate. First, we clarify exactly what we mean when we say that an algorithm returns a splitting of G . The algorithm should return the following data.

1. A finite connected graph Y in the sense of Serre; see Definition 3.2.1.

2. For each vertex $v \in Y$ a finite alphabet S_v (with $S_{v_1} \cap S_{v_2} = \emptyset$ for distinct vertices v_1 and v_2) and a finite set R_v of words in S_v .
3. For each edge $e \in Y$ a finite alphabet S_e and a finite set R_e of words in S_e such that $S_{\bar{e}} = S_e$ and $R_{\bar{e}} = R_e$, and other alphabets are disjoint.
4. For each edge $e \in Y$ a map $(\partial')^e_- : S_e \rightarrow S_{t(e)}$ such that $(\partial')^e_-$ extends to an injective group homomorphism $\partial^e_- : \langle S_e \mid R_e \rangle \rightarrow \langle S_{t(e)} \mid R_{t(e)} \rangle$.
5. For each vertex $v \in Y$ a map $\theta'_v : S_v \rightarrow G$ such that θ'_v extends to a group homomorphism $\theta_v : \langle S_v \mid R_v \rangle \rightarrow G$ satisfying the following condition. Let \mathcal{Y} be the graph of groups given by the data of items 1–4. Then we require that there exists an isomorphism $\theta : \pi_1 \mathcal{Y} \rightarrow G$ and a choice of inclusion $\langle S_v \mid R_v \rangle \hookrightarrow \pi_1 \mathcal{Y}$ for each vertex v as discussed in Remark 3.2.6 such that the composition of this inclusion with θ is equal to θ_v .

In Section 5.2.1 we provide useful algorithms for solving various problems regarding virtually cyclic subgroups of a hyperbolic group. These algorithms will be used several times throughout the rest of this Chapter.

Then, in Section 5.2.2 we show that there is an algorithm that enumerates all splittings of a hyperbolic group over its virtually cyclic subgroups in the sense described above.

5.2.1 Virtually cyclic subgroups

We will need the following lemma, which allows us to do various computations related to virtually cyclic subgroups of hyperbolic groups.

Lemma 5.2.1. *[DG11, Lemma 2.8] There is an algorithm that, when given a presentation for a hyperbolic group G and a finite subset $S \subset G$, returns an answer to the question “is $\langle S \rangle \leq G$ virtually cyclic?” If the answer is “yes” then the algorithm also determines the following.*

1. *The (unique) maximal finite normal subgroup of $\langle S \rangle$.*
2. *A presentation for $\langle S \rangle$.*
3. *Whether $\langle S \rangle$ is of type \mathcal{Z} or D_∞ . (Recall that we say that a virtually cyclic group of type \mathcal{Z} (respectively D_∞) if it surjects onto \mathbb{Z} (respectively D_∞), and that any virtually cyclic group has exactly one of these types.)*
4. *A generating set for the maximal virtually cyclic subgroup of G containing $\langle S \rangle$.*

The proof of this lemma in [DG11] uses Makanin's algorithm for solving equations in hyperbolic groups. We modify that part of the argument to use only elementary methods in keeping with the themes of this thesis.

Proof. We give an alternative method to determine whether or not $\langle S \rangle$ is virtually cyclic and to produce a maximal finite normal subgroup of $\langle S \rangle$ in the case that it is; the rest of the argument can be copied verbatim from [DG11]. First compute δ with respect to which G is δ -hyperbolic.

Use Kapovich's algorithm [Kap96, Proposition 4] to search for a constant K with respect to which $\langle S \rangle$ is K -quasi-convex in G . This algorithm finds such a constant if it exists and does not terminate if $\langle S \rangle$ is not quasi-convex; note that if $\langle S \rangle$ is virtually cyclic then it is guaranteed to be quasi-convex.

In parallel to Kapovich's algorithm, search for a pair of elements g and h in $\langle S \rangle$ such that the commutator $[g^2, h^2]$ has infinite order. This can be checked since the order of an element of G of finite order is bounded above by the number of elements of G in the ball of radius $4\delta + 2$.

If Kapovich's algorithm terminates, use K and δ to compute δ' such that $\langle S \rangle$ is δ' -hyperbolic, as in the proof of Proposition 5.1.4. Then any finite subgroup of $\langle S \rangle$ can be conjugated into a ball of radius at most $4\delta' + 2$ with respect to the word metric in $\langle S \rangle$, so all finite normal subgroups of $\langle S \rangle$ can be computed using solutions to the word and conjugacy problems in G . Once this is computed the algorithm of [DG11, Lemma 2.8] can be used to determine whether or not $\langle S \rangle$ is virtually cyclic.

If $\langle S \rangle$ is not quasi-convex then it contains a free group on two generators, so a pair (g, h) such that $[g^2, h^2]$ has infinite order certainly exists. Conversely, if such a pair exists then $\langle S \rangle$ cannot be virtually cyclic, since any virtually cyclic group contains a subgroup of index two that surjects onto \mathbb{Z} with finite kernel. But for any g and h in a virtually cyclic group, $[g^2, h^2]$ is contained in this index two subgroup and is mapped to the identity in \mathbb{Z} , so is contained in this finite kernel. \square

We also give an algorithm that checks whether or not given maximal virtually cyclic subgroups of a hyperbolic group are conjugate.

Lemma 5.2.2. *There is an algorithm that, when given a presentation for a hyperbolic group G and finite subsets $S_1, S_2 \subset G$ such that $\langle S_1 \rangle$ and $\langle S_2 \rangle$ are maximal virtually cyclic subgroups of G , returns an answer to the question "is $\langle S_1 \rangle$ conjugate to $\langle S_2 \rangle$ in G ?"*

Proof. Use Lemma 5.2.1 to check that $\langle S_1 \rangle$ and $\langle S_2 \rangle$ are either both of type \mathcal{Z} or both of type D_∞ ; if they are not of the same type then clearly they cannot be conjugate.

Compute presentations of $\langle S_1 \rangle$ and $\langle S_2 \rangle$, together with maximal finite normal subgroups K_i of those groups. Note that $\langle S_i \rangle / K_i$ is isomorphic to either \mathbb{Z} or D_∞ according to the types of the groups.

If both are of type \mathcal{Z} , find the $2|K_i|$ elements of $\langle S_i \rangle$ that map to a generator of \mathbb{Z} and use a solution to the conjugacy problem in G (which exists due to Gromov [Gro87] since G is hyperbolic) to check whether there is an element of $\langle S_1 \rangle$ with this property that is conjugate in G to an element of $\langle S_2 \rangle$ with this property. If there is such a pair then the groups are conjugate since they are maximal. If there is no such pair then they cannot be conjugate.

If both are of type D_∞ find the $|K_i|$ pairs of elements of $\langle S_i \rangle$ that map to some prechosen pair of reflections that generate D_∞ . Use a solution to the simultaneous conjugacy problem [BH05, BH13] to check whether some pair of pairs are conjugate in G up to reordering. If there is then the groups are conjugate, again by maximality. If there is no such pair then they cannot be conjugate since any pair of (unordered) generating sets for D_∞ consisting of two reflections are conjugate in D_∞ . \square

5.2.2 Enumerating the splittings

Lemma 5.2.3. *There is an algorithm that takes as input a presentation for a hyperbolic group G and outputs all minimal splittings of G over virtually cyclic subgroups. (The algorithm does not terminate.)*

Proof. At the beginning of Section 5.2 we described a finite set of data that together determine a splitting of a group G . Given this data, one can use the Seifert-van Kampen theorem find an isomorphism θ from a presentation $\langle S \mid R \rangle$ to G where there is a finite connected graph Y with a maximal subtree Y' such that $\langle S \mid R \rangle$ takes the following form.

1. The generating set S can be partitioned as

$$S = \bigsqcup_{v \in V(Y)} S_v \sqcup \bigsqcup_{e \in E(Y)} S_e \sqcup \{t_e \mid e \in E(Y) - E(Y')\}$$

and there is a bijection from $S_e \rightarrow S_{\bar{e}}$ for each edge e , which we denote $s \mapsto \bar{s}$.

2. For each $e \in E(Y)$ there is a map $(\partial')_-^e S_e \rightarrow S_{t(e)}$. We then denote by $(\partial')_+^e$ the map $(\partial')_-^{\bar{e}}$.

3. The set of relators R can be partitioned as

$$R = \left(\bigcup_{v \in V(Y)} R_v \right) \cup \left(\bigcup_{e \in E(Y)} R_e \right) \cup \left(\bigcup_{e \in E(Y')} \{s^{-1} \bar{s} \mid s \in S_e\} \right) \cup \left(\bigcup_{e \in E(Y)} \{s^{-1} (\partial')_-^e(s) \mid s \in R_e\} \right) \\ \cup \{t_e t_{\bar{e}} \mid e \in E(Y) - E(Y')\} \cup \left(\bigcup_{e \in E(Y) - E(Y')} \{t_e s t_e^{-1} \bar{s}^{-1} \mid s \in R_e\} \right)$$

where relators in R_v only involve generators in S_v for $v \in V(Y)$ and relators in R_e only involve generators in S_e for $e \in E(Y)$, and so that for each relator $r \in R_e$ for each edge $e \in E(Y)$, the word in the free monoid generated $S_{\bar{e}}$ obtained by replacing each letter in s appearing in r by \bar{s} is in $R_{\bar{e}}$.

4. For each $e \in E(Y)$ and each $r \in R_e$, the word in the free monoid generated by $S_{t(e)}$ obtained by applying $(\partial')_+^e$ to each letter of r is in $R_{t(e)}$. It follows that $(\partial')_-^e$ extends to a homomorphism $\partial_-^e : \langle S_e \mid R_e \rangle \rightarrow \langle S_v \mid R_v \rangle$, and similarly for $(\partial')_+^e$.
5. For each $e \in E(Y)$ the group $\langle S_e \mid R_e \rangle$ is virtually cyclic.
6. For each $e \in E(Y)$ the homomorphism $\partial_-^e : \langle S_e \mid R_e \rangle \rightarrow \langle S_v \mid R_v \rangle$ is injective.

It is clear that when given such a presentation one can find a corresponding splitting of G ; we shall say that such a presentation *exhibits a splitting of G over virtually cyclic subgroups*. Note that Condition 6 and Lemma 3.2.8 guarantee that for any vertex $v \in V(Y)$, $\langle \theta(S_v) \rangle$ is isomorphic to $\langle S_v \mid R_v \rangle$, and the same statement holds for edges.

Using Tietze transformations (see [LS77, Section 2.2]) one can enumerate all presentations $\langle S \mid R \rangle$ for G together the isomorphisms $\langle S \mid R \rangle \rightarrow G$. It remains to show that the conditions for a presentation to exhibit a splitting of G over virtually cyclic subgroups can be checked, and that one can check whether or not that splitting is minimal. The processes to check Conditions 1–4 are clear. Condition 5 can be checked using the algorithm of Lemma 5.2.1.

If conditions 1–5 hold, then condition 6 holds if and only if the restriction of θ to $\langle S_e \mid R_e \rangle$ is injective for each edge e . To check this condition for an edge e , compute the maximal finite normal subgroup of $\langle S_e \mid R_e \rangle$ and check the injectivity of the restriction of θ to this subgroup. If it is injective, next find an element of $\langle S_e \mid R_e \rangle$ of infinite order and check whether or not its image under θ has infinite order. If it does then the restriction of θ to $\langle S_e \mid R_e \rangle$ is injective.

To check minimality, it is enough to be able to check if a boundary map is surjective, which we do by applying Lemma 5.2.1. To check whether ∂_-^e is surjective, first check whether or not $\langle S_{t(e)} \mid R_{t(e)} \rangle$ is virtually cyclic. If it is, check whether or not the boundary

map sends the maximal finite normal subgroup of $\langle S_e \mid R_e \rangle$ surjectively to the maximal finite normal subgroup of $\langle S_{i(e)} \mid R_{i(e)} \rangle$. If it does, check whether or not $\langle S_e \mid R_e \rangle$ and $\langle S_{i(e)} \mid R_{i(e)} \rangle$ are both of type \mathcal{Z} or both of type D_∞ . If they are both of the same type, pass to an index two subgroup of each if necessary to ensure that they are both of type \mathcal{Z} . Now check whether or not the composition of ∂_-^e with the natural surjection to \mathbb{Z} is surjective. If it is then ∂_-^e is surjective; if any of these tests produced the opposite answer then ∂_-^e is not surjective. \square

5.3 Splittings of groups with circular boundary

The question of the existence of a relative splitting of a hyperbolic group cannot be answered by consideration of the topology of its boundary alone if the boundary is homeomorphic to a circle: some groups with circular boundary split and some do not. In this section we deal with this special case.

5.3.1 Groups with circular boundary

Relatively hyperbolic groups with circular boundary are closely related to fundamental groups of compact orbifolds of dimension two, with peripheral subgroups corresponding to boundary components. We now make this relationship explicit, so that we can answer the question of when such a group admits a non-trivial splitting over a virtually cyclic subgroup by analysing the structure of the orbifold.

Let G be a hyperbolic group with a finite collection \mathcal{H} of subgroups such that G is hyperbolic relative to \mathcal{H} and $\partial(G, \mathcal{H})$ is homeomorphic to S^1 . Then G acts as a discrete convergence group on $\partial(G, \mathcal{H})$ by [Bow99c], so we can apply the Convergence Group Theorem of Tukia, Casson, Jungreis and Gabai:

Theorem 5.3.1. [Tuk88, CJ94, Gab92] *Let G act as a convergence group on S^1 . Then there is a properly discontinuous action of G by isometries on \mathbb{H}^2 and a G -equivariant homeomorphism $\partial(G, \mathcal{H}) \rightarrow \partial\mathbb{H}^2$.*

Let K be the (finite) kernel of the action of G on $\partial(G, \mathcal{H})$; note that this is the same as the kernel of the extension of the action to \mathbb{H}^2 . Then G is an extension

$$1 \longrightarrow K \longrightarrow G \longrightarrow G' \longrightarrow 1$$

where the quotient G' acts faithfully on \mathbb{H}^2 . Let \mathcal{H}' be the set of images of elements of \mathcal{H} in G' . Note that \mathcal{H}' is a set of conjugacy class representatives of maximal parabolic

subgroups of $G' \leq \text{Isom}\mathbb{H}^2$, so each element of \mathcal{H}' is either infinite cyclic or infinite dihedral depending on whether or not it preserves orientation.

Lemma 5.3.2. *With G and G' as above, G splits non-trivially over a virtually cyclic subgroup relative to \mathcal{H} if and only if G' splits non-trivially over a virtually cyclic subgroup relative to \mathcal{H}' .*

Proof. One direction is clear: if G' acts minimally on a non-trivial tree T with virtually cyclic edge stabilisers then the quotient map $G \rightarrow G'$ induces a minimal action of G on T . If e is an edge of T then G_e is an extension of G'_e by K , which is a finite group, so is itself virtually cyclic. Note also that elements of \mathcal{H} act elliptically on T , since elements of \mathcal{H}' do.

Conversely, suppose that G acts minimally on a non-trivial tree T with virtually cyclic edge stabilisers. The kernel K is finite, so the restriction of the action to K fixes a vertex v . Also, K is normal, so its action fixes the orbit $G \cdot v$ pointwise. Any point in T lies on a geodesic path connecting points in $G \cdot v$, since the union of all such paths is a G -invariant subtree of T and the action was assumed to be minimal. Therefore K acts trivially on T .

It follows that the G -action descends to a G' action. If e is an edge of T then G'_e is the quotient G_e/K , so G'_e is virtually cyclic. Furthermore, elements of \mathcal{H}' are elliptic if elements of \mathcal{H} are. \square

Any finite normal subgroup of G is contained in the ball of radius $4\delta + 2$ centred at the identity by [BH99, Theorem 3.2], so by checking all finite subsets of this ball using a solution to the word problem, the set of finite normal subgroups of G can be computed. Any finite subgroup of G fixes a point in \mathbb{H}^2 , and any finite normal subgroup of \mathbb{H}^2 fixes the G -orbit of that point pointwise and is therefore in the kernel of the action of G on $\partial(G, \mathcal{H})$, so K is the unique maximal finite normal subgroup of G . Therefore the kernel K of the action of G on \mathbb{H}^2 can be computed from a presentation of G . By adjoining all elements of K to the set of relators of G one obtains a presentation for G' .

As described above, G' can be realised as a discrete subgroup of $\text{Isom}\mathbb{H}^2$. The conjugacy classes of elements of \mathcal{H}' are then precisely the conjugacy classes of maximal parabolic subgroups of G' . The action of G' on $\partial(G', \mathcal{H}')$ is minimal, so the limit set of the action is S^1 . It follows that G' is a Fuchsian group of the first kind: \mathbb{H}^2/G' is a finite volume orbifold and \mathcal{H}' is a choice of conjugacy class representatives for the cusp subgroups. Truncating the cusps of the orbifold, we realise G' as the fundamental group of a compact hyperbolic orbifold such that \mathcal{H}' is a choice of conjugacy class representatives for the boundary subgroups. (See the following section for the definition of an orbifold.)

Recall Definition 3.5.7 of a *bounded Fuchsian* group. We shall call a choice of conjugacy class representatives of peripheral subgroups of a bounded Fuchsian group a *boundary peripheral structure* on that group. Therefore G' is a bounded Fuchsian group and \mathcal{H}' is a boundary peripheral structure on G' . The results of Section 5.3.1 are summarised by the following lemma.

Lemma 5.3.3. *There is an algorithm that, when given a presentation for a hyperbolic group with a set \mathcal{H} of virtually cyclic peripheral subgroups, returns a presentation for another hyperbolic group G' with a set \mathcal{H}' of virtually cyclic subgroups such that $\partial(G', \mathcal{H}')$ is homeomorphic to $\partial(G, \mathcal{H})$ and G splits over a virtually cyclic subgroup relative to \mathcal{H} if and only if G' splits over a virtually cyclic subgroup relative to \mathcal{H}' . Furthermore, if $\partial(G', \mathcal{H}')$ is homeomorphic to a circle then G' is the fundamental group of a compact two-dimensional hyperbolic orbifold and \mathcal{H}' is a set of representatives of fundamental groups of components of the boundary of the orbifold.*

5.3.2 Orbifolds

Given the results of the previous section, the problem of determining whether or not a given group G with circular boundary splits non-trivially over a virtually cyclic subgroup relative to a collection \mathcal{H} of peripheral subgroups reduces to the case in which G is a bounded Fuchsian group and \mathcal{H} is a boundary peripheral structure on G . Equivalently, we may assume that $G = \pi_1^{\text{orb}} Q$ where Q is a compact hyperbolic orbifold, so \mathcal{H} is a set of conjugacy class representatives of orbifold fundamental groups of boundary components of Q .

In this section we review some basic terminology for describing compact two dimensional orbifolds. We then recall a result of Guirardel and Levitt [GL17], which tells us precisely when a bounded Fuchsian group admits a non-trivial splitting over a virtually cyclic subgroup relative to its boundary peripheral structure.

In particular, we concentrate on compact hyperbolic orbifolds. For more information about orbifolds see [Sco83]. Hyperbolic orbifolds are particularly easy to deal with, since they are covered by a closed convex subset of \mathbb{H}^2 . (This property is in contrast to a general two dimensional orbifold: such an orbifold need not be covered by any manifold.) This allows us to make the following definition.

Definition 5.3.4. A compact hyperbolic orbifold Q is a compact underlying surface, which we denote Q_{top} , together with a map from a closed convex subset \tilde{Q} of \mathbb{H}^2 to Q that is a quotient by a cocompact properly discontinuous group action by isometries. The *orbifold*

fundamental group $\pi_1^{\text{orb}}Q$ is the group of deck transformations, so $\pi_1^{\text{orb}}Q$ acts cocompactly, properly discontinuously and by isometries of \tilde{Q} with quotient Q .

The orbifold Q can be described without reference to the covering map $\tilde{Q} \rightarrow Q_{\text{top}}$ by endowing the compact surface Q_{top} with additional structure.

Definition 5.3.5. A point in Q is a *singular point* if it is the image of a point in \tilde{Q} with non-trivial stabiliser. Otherwise it is a *regular point*. The *singular locus* of Q is the set of all singular points.

We now classify the singular points of Q . Note that we have the following classification of elliptic subgroups of $\text{Isom } \mathbb{H}^2$: any elliptic subgroup is either cyclic acting by rotations, cyclic of order two generated by a reflection or dihedral generated by reflections through intersecting geodesics in \tilde{Q} .

Definition 5.3.6. The image in Q_{top} of a point in \tilde{Q} with a cyclic stabiliser consisting of rotations is a *cone point*. The image of the set of fixed points in \tilde{Q} of a reflection in $\pi_1^{\text{orb}}Q$ is a *mirror*. The image of a point in \tilde{Q} with dihedral stabiliser is a *corner reflector*.

Now we consider the boundary of Q_{top} . A point in the boundary of Q_{top} is either contained in a mirror or is in the image of the boundary $\partial\tilde{Q}$ of \tilde{Q} . Note however that not every point in the image of $\partial\tilde{Q}$ is in ∂Q_{top} , since two components of $\partial\tilde{Q}$ can be identified in the quotient.

Definition 5.3.7. The *topological boundary* $\partial_{\text{top}}Q$ of Q is the boundary of Q_{top} , while the *orbifold boundary* ∂Q of Q is the intersection of $\partial_{\text{top}}Q$ with image in Q_{top} of the boundary of \tilde{Q} under the quotient projection. Each component of the orbifold boundary is either a component of the topological boundary homeomorphic to S^1 , in which case we call it a *circular boundary component*, or is a proper subset of a component of $\partial_{\text{top}}Q$, in which case it is homeomorphic to an interval and we call it an *interval boundary component*.

Note that the singular locus is the union of all cone points and mirrors in Q_{top} , and the orbifold boundary ∂Q is the closure of the set of regular points in $\partial_{\text{top}}Q$.

This description of the singular locus of an orbifold allows us to draw a picture of any two dimensional hyperbolic orbifold as follows. Start by drawing the underlying surface Q_{top} . Then mark each cone point in Q_{top} and annotate each one with the order of the cyclic stabiliser of one of its preimages. Next, identify each component of the boundary of Q_{top} with a polygon so that each edge is either a mirror or a component of the orbifold boundary ∂Q and mark those edges that are mirrors. Finally, wherever two mirror edges of a boundary polygon meet at a vertex, annotate that vertex with the order of the dihedral




Figure 5.1: Three pictures of orbifolds. Orbifold boundary components are coloured red. The first has underlying surface of genus one with two boundary components, one of which is a mirror and the other is an orbifold boundary component. The second has underlying surface homeomorphic disc. Its topological boundary is identified with a hexagon, two edges of which are orbifold boundary components and the other four are mirrors. Where mirrors meet there are corner reflectors of order depending on the angle at which the mirrors meet. Mirrors meet orbifold boundary components at right angles. The third also has underlying surface homeomorphic to a disc. Here there is a single cone point of order 3. The total angle around the cone point is 120° . Note that this is not actually a hyperbolic orbifold: it is covered by a disc.

stabiliser of one of its preimages. Some pictures of this type—illustrating all possible types of singular points in orbifolds—are displayed in Figure 5.1.

Orbifolds satisfy the Seifert-van Kampen Theorem. See for example [BMP03, Corollary 2.3]. Using this, one can recover the fundamental group $\pi_1^{\text{orb}} Q$ from the picture described in the previous paragraph.

Our description of $\pi_1^{\text{orb}} Q$ from its picture is completed by the observation that one can recover the boundary peripheral structure. The boundary peripheral structure consists of a set of conjugacy class representatives of stabilisers of orbifold boundary components of \tilde{Q} . The boundary peripheral structure contains a subgroup of $\pi_1^{\text{orb}} Q$ isomorphic to \mathbb{Z} for each circular boundary component of Q and a subgroup of $\pi_1^{\text{orb}} Q$ isomorphic to D_∞ for each interval boundary component of Q . Therefore the picture of an orbifold contains sufficient information about that orbifold for our purposes.

Definition 5.3.8. A *geodesic* in Q is a curve γ that is the image of a geodesic $\tilde{\gamma}$ in \tilde{Q} under the quotient projection. A geodesic γ is *closed* if it is compact and *simple* if it is closed and furthermore $g \cdot \tilde{\gamma}$ is either equal to or disjoint from $\tilde{\gamma}$ for each g in $\pi_1^{\text{orb}} Q$. A simple closed geodesic is homeomorphic to either a circle or an interval with each of its end points contained in a mirror in Q . A simple closed geodesic is *essential* if it is not contained in $\partial_{\text{top}} Q$.

5.3.3 Splittings of fundamental groups of orbifolds

The theory of splittings of fundamental groups of orbifolds relative to their boundary peripheral structures is developed in [GL17]; we recall the following results:

Lemma 5.3.9. [GL17, Corollary 5.6] *The fundamental group of a compact orbifold Q splits non-trivially over a virtually cyclic subgroup relative to its boundary peripheral structure if and only if Q contains an essential simple closed geodesic.*

Definition 5.3.10. We call a compact orbifold without an essential simple closed geodesic *small*.

Any sufficiently complicated orbifold contains an essential simple closed geodesic, so there are not very many small orbifolds. Guirardel and Levitt provide a complete list.

Proposition 5.3.11. [GL17, Proposition 5.12] *A hyperbolic orbifold Q of dimension two is small if and only if it is one of the following.*

1. *A sphere with three cone points, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b \mid a^p, b^q, ab^r \rangle$ where $p^{-1} + q^{-1} + r^{-1} < 1$ and the boundary peripheral structure is empty.*
2. *A triangle, all three edges of which are mirrors, so $\pi_1^{\text{orb}}Q$ has presentation*

$$\langle a, b, c \mid a^2, b^2, c^2, (ab)^p, (bc)^q, (ca)^r \rangle$$

where $p^{-1} + q^{-1} + r^{-1} < 1$ and the boundary peripheral structure is empty.

3. *A disc with two cone points, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b \mid a^p, b^q \rangle$ where $p, q > 1$ and the boundary peripheral structure is $\{\langle ab \rangle\}$.*
4. *A cylinder with one cone point, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b \mid (ab)^p \rangle$ where $p > 1$ and the boundary peripheral structure is $\{\langle a \rangle, \langle b \rangle\}$.*
5. *A pair of pants, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b \mid \rangle$ and the boundary peripheral structure is $\{\langle a \rangle, \langle b \rangle, \langle ab \rangle\}$.*
6. *A disc with one cone point with edge consisting of an interval boundary component and a mirror, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, t \mid a^2, t^p \rangle$ where $p > 1$ and the boundary peripheral structure is $\{\langle a, tat^{-1} \rangle\}$.*

7. A square with an interval boundary component and three mirrors edges, so $\pi_1^{\text{orb}}Q$ has presentation

$$\langle a, b, c \mid a^2, b^2, c^2, (ab)^p, (bc)^q \rangle$$

where $p + q \geq 1$ and the boundary peripheral structure is $\{\langle a, c \rangle\}$.

8. An annulus in which one edge comprises an interval boundary component and a mirror and the other is a circular boundary component, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, t \mid a^2 \rangle$ with boundary peripheral structure $\{\langle a, tat^{-1} \rangle\}$.
9. A pentagon with two non-adjacent interval boundary components and three mirrors as edges, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b, c \mid a^2, b^2, c^2, (ab)^p \rangle$ with boundary peripheral structure $\{\langle b, c \rangle, \langle c, a \rangle\}$.
10. A hexagon, the six edges of which are alternately interval boundary components and mirrors, so $\pi_1^{\text{orb}}Q$ has presentation $\langle a, b, c \mid a^2, b^2, c^2 \rangle$ and the boundary peripheral structure is

$$\{\langle a, b \rangle, \langle b, c \rangle, \langle c, a \rangle\}.$$

5.3.4 Detecting whether or not an orbifold is small

In this section we complete the objective of Section 5.3: we show that there is an algorithm that, when given a presentation for the fundamental group of a small orbifold together with the boundary peripheral structure of that orbifold, eventually certifies that the orbifold is small.

Lemma 5.3.12. *There is an algorithm that takes as input a presentation for a hyperbolic group G and a collection \mathcal{H} of maximal virtually cyclic peripheral subgroups and terminates if and only if $\partial(G, \mathcal{H})$ is homeomorphic to a circle and G does not split relative to \mathcal{H} over a virtually cyclic subgroup.*

Proof. First use the algorithm of Lemma 5.3.3 and let the output of that algorithm be (G', \mathcal{H}') . Then $\partial(G, \mathcal{H})$ is homeomorphic to a circle and G does not split over a virtually cyclic subgroup relative to \mathcal{H} if and only if there is an isomorphism from G' to one of the groups with boundary peripheral structure listed in Proposition 5.3.11 that maps elements of \mathcal{H}' to conjugates of elements of the given boundary peripheral structure.

The algorithm runs infinitely many processes in parallel using a standard diagonal argument, so that it first runs the first step of the first process, then the second step

of the first process, then the first step of the second process, and so on. Enumerate the groups and peripheral structures described in Proposition 5.3.11 and, in parallel, enumerate all homomorphisms from these groups to G' and homomorphisms from G' to these groups. Note that this is possible since one can test whether or not a map defined on the generators of a group extends to a homomorphism using a solution to the word problem in the codomain of the map, and G' and all groups listed in Proposition 5.3.11 are hyperbolic. Whenever an inverse pair of homomorphisms is found, start a search for a set of conjugating elements to check that the map $G' \rightarrow \pi_1^{\text{orb}} Q$ maps each element of \mathcal{H}' to a conjugate of an element of the boundary peripheral structure for Q .

If an inverse pair of homomorphisms and a set of conjugating elements is found, the algorithm terminates. \square

5.4 Maximal splittings

We now apply the results of Sections 5.1, 5.2 and 5.3 to the problem of finding a maximal splitting of a one-ended hyperbolic group over its virtually cyclic subgroups.

First we fix some notation.

Definition 5.4.1. Given a virtually cyclic subgroup H of a hyperbolic group G we shall denote by \widehat{H} the unique maximal virtually cyclic subgroup of G containing H . Given a set \mathcal{H} of virtually cyclic subgroups of G let $\widehat{\mathcal{H}}$ be a set of subgroups of G obtained by choosing one representative of each conjugacy class that meets the set $\{\widehat{H} \mid H \in \mathcal{H}\}$. Given a vertex v in the underlying graph of a graph of groups \mathcal{Y} , let $\text{Inc } v$ be the set $\{\partial_-^e(\mathcal{Y}_e) \mid \iota(e) = v\}$ of subgroups of \mathcal{Y}_v .

Remark 5.4.2. Given a hyperbolic group G and a finite set of elements generating a virtually cyclic subgroup $H \leq G$, Lemma 5.2.1 provides an algorithm that computes a generating set for \widehat{H} . Given a finite set of subsets of G generating a set \mathcal{H} of virtually cyclic subgroups of G one can compute finite generating sets for the elements of $\widehat{\mathcal{H}}$ by first applying the algorithm of Lemma 5.2.1 and then throwing away any duplicate subgroups up to conjugacy using the algorithm of Lemma 5.2.2.

Now note the following proposition, which allows us to iteratively apply the results of Section 5.1.

Proposition 5.4.3. *Let \mathcal{Y} be a graph of groups with virtually cyclic edge groups underlying graph Y such that $\pi_1 \mathcal{Y}$ is hyperbolic and $\partial \pi_1 \mathcal{Y}$ is connected. Then for each vertex $v \in Y$, \mathcal{Y}_v is hyperbolic relative to $\widehat{\text{Inc } v}$ and $\partial(\mathcal{Y}_v, \widehat{\text{Inc } v})$ is connected.*

Proof. First note that \mathcal{Y}_e is virtually cyclic so is quasi-convex in $\pi_1\mathcal{Y}$ for each edge e in Y . Therefore \mathcal{Y}_v is quasi-convex by [Bow98b, Proposition 1.2], and therefore \mathcal{Y}_v is hyperbolic. The elements of $\widehat{\text{Inc } v}$ are maximal non-conjugate virtually cyclic subgroups of a hyperbolic group, so they are certainly malnormal. Therefore hyperbolicity of \mathcal{Y}_v relative to $\widehat{\text{Inc } v}$ follows from Theorem 2.3.5.

If $\partial(\mathcal{Y}_v, \widehat{\text{Inc } v})$ is disconnected then \mathcal{Y}_v admits a non-trivial splitting over a finite group relative to $\widehat{\text{Inc } v}$ by [Bow12, Proposition 10.1]. This splitting is then also relative to $\text{Inc } v$. Therefore the graph of groups \mathcal{Y} admits a refinement \mathcal{Y}' including an edge e mapping to v under the collapse map $Y' \rightarrow Y$ such that \mathcal{Y}'_e is finite and neither ∂_+^e nor ∂_-^e is surjective. By collapsing $Y - e$ to either one or two points (depending on whether or not e is separating) one obtains a non-trivial splitting of $\pi_1\mathcal{Y}$ over finite subgroups, contracting the assumption that $\partial\pi_1\mathcal{Y}$ is connected. \square

5.4.1 Boundaries with cut pairs

Recall Corollary 3.5.11 of Bowditch relating splittings of a hyperbolic group over its virtually cyclic subgroups to the existence of cut pairs in its boundary. We require a relative version of this theorem. Such a theorem was recently proved in considerable generality by Haulmark [Hau17], but we only need the special case in which G arises as a vertex group in a splitting of a larger hyperbolic group over virtually cyclic subgroups, and \mathcal{H} is the collection of edge groups incident at that vertex. In this section we give an elementary proof that this simple case follows from Corollary 3.5.11. In particular we can avoid \mathbb{R} -trees machinery. For a discussion in greater generality, see [Gro13].

Proposition 5.4.4. *Let v be a vertex in a minimal G -tree T with virtually cyclic edge groups where G is hyperbolic and one-ended. Let $\text{Inc } v$ be a set of representatives of G_v -conjugacy classes of stabilisers of edges in T incident at v . Suppose that $\partial(G_v, \widehat{\text{Inc } v})$ is not a single point, does not contain a cut point and is not homeomorphic to a circle but does contain a cut pair. Then G_v admits a non-trivial splitting over virtually cyclic subgroups of G_v relative to $\text{Inc } v$.*

First we need the following lemma.

Lemma 5.4.5. *Let $f: T_1 \rightarrow T_2$ be an equivariant map of G -trees with virtually infinite cyclic edge stabilisers such that the action of G on T_1 is cocompact and the action of G on T_2 is minimal. Let v be a vertex of T_2 such that $\partial(G_v, \widehat{\text{Inc } v})$ is not a single point, is connected and does not contain a cut point. Then the action of G_v on T_1 fixes a component of $f^{-1}(v)$.*

Proof. First we show that there is a vertex $w \in T_1$ such that $G_w \cap G_v$ is non-elementary. If this is not the case then the stabiliser of each edge of T_1 with respect to the action of G_v on T_1 is either finite or commensurable with the stabilisers of its end points. Therefore the action induces a splitting of G_v with finite edge groups and virtually cyclic vertex groups. By minimality of the action of G on T_2 each edge e incident at v is $f(e')$ for some edge e' of T_1 , and then $G_{e'}$ is a finite index subgroup of G_e , since each is virtually infinite cyclic. In particular, \widehat{G}_e is elliptic, and the splitting is relative to $\widehat{\text{Inc } v}$. But $\partial(G_v, \widehat{\text{Inc } v})$ was assumed to be connected and not a single point, which is a contradiction.

Then $f(w) = v$, otherwise any edge separating $f(w)$ from v in T_2 has non-elementary stabiliser. Let S be the component of $f^{-1}(v)$ containing w . We now show that any other vertex w' of T_1 such that $G_{w'} \cap G_v$ is non-elementary is also in S . Suppose that e is an edge of T_1 that is not in $f^{-1}(v)$. As in Section 1 of [Bow98b] there exists a partition of $\partial(G, \emptyset) - \Lambda G_e$ as $U_1 \sqcup U_2$. The intersection of G_e with G_v is either finite or commensurable with a conjugate of an element of $\widehat{\text{Inc } v}$, so the images of $U_1 \cap \Lambda G_v$ and $U_2 \cap \Lambda G_v$ under the projection map $\Lambda G_v \rightarrow \partial(G, \widehat{\text{Inc } v})$ cover all but at most a point of $\partial(G, \widehat{\text{Inc } v})$. These sets are disjoint, so one must be empty, say U_2 . But ΛG_w and $\Lambda G_{w'}$ each contain more than two points, so must both be contained in $U_1 \cup \Lambda G_e$. This implies that w and w' are on the same side of e .

Therefore the action of G_v on T_1 fixes S , for any element of G_v must send w to a vertex of S . □

Proof of Proposition 5.4.4. Let Σ be Bowditch's JSJ tree for G . Σ is then elliptic with respect to T , so by [GL17, Proposition 2.2] there exists a G -tree $\widehat{\Sigma}$ and maps $p: \widehat{\Sigma} \rightarrow \Sigma$ and $f: \widehat{\Sigma} \rightarrow T$ such that:

1. p is a collapse map. (That is, a map given by collapsing some edges of $\widehat{\Sigma}$ to points.)
2. For $w \in \Sigma$, the restriction of f to $p^{-1}(w)$ is injective.

By Lemma 5.4.5 there is a component S of $f^{-1}(v)$ that is fixed by the action of G_v . Suppose that a vertex w of S is fixed by the G_v -action.

If e is any edge of $\widehat{\Sigma}$ that is adjacent to but not contained in S then $G_e \leq G_{f(e)}$ and the subgroup is necessarily of finite index. Conversely the stabiliser of any edge incident at v contains the stabiliser of an edge adjacent to S at finite index. The stabiliser of any edge adjacent to S is then commensurable with the stabiliser of an edge incident at w , and vice versa. Therefore $\partial(G_v, \widehat{\text{Inc } w})$ is homeomorphic to $\partial(G_v, \widehat{\text{Inc } v})$. We assumed that $\partial(G_v, \widehat{\text{Inc } v})$ was neither a point nor homeomorphic to a circle, so pw is not in T_1 or T_2 , and is therefore in T_3 , where here we recall the partition $T_1 \sqcup T_2 \sqcup T_3$ of vertices of Σ .

Let x and y be points in $\partial(G_\nu, \widehat{\text{Inc } \nu})$ and choose preimages \tilde{x} and \tilde{y} in $\partial(G_\nu, \emptyset)$, which we identify with $\Lambda G_\nu \subset \partial(G, \emptyset)$. The set of components of $\partial(G_\nu, \widehat{\text{Inc } \nu}) - \{x, y\}$ is in bijection with the set of those components of $\partial(G, \emptyset) - \{\tilde{x}, \tilde{y}\}$ that meet ΛG_ν .

Suppose then that $\{x, y\}$ is a cut pair in $\partial(G_\nu, \widehat{\text{Inc } \nu})$, so at least two components of $\partial(G, \emptyset) - \{\tilde{x}, \tilde{y}\}$ meet ΛG_ν .

Then by Corollary 3.5.11 there is a vertex u of Σ in T_1 or T_2 such that ΛG_u contains $\{\tilde{x}, \tilde{y}\}$. Then $\{\tilde{x}, \tilde{y}\} = \Lambda G_u \cap \Lambda G_\nu$, so $\{\tilde{x}, \tilde{y}\}$ is the limit set of an edge incident at ν . Therefore $\{x, y\} \subset \partial(G_\nu, \widehat{\text{Inc } \nu})$ is a single point, which is a contradiction because $\partial(G_\nu, \widehat{\text{Inc } \nu})$ was assumed not to contain a cut point. Hence the action of G_ν on S does not fix any vertex and therefore gives rise to a non-trivial splitting of G_ν relative to $\text{Inc } \nu$. □

5.4.2 Boundaries without cut points or pairs

Our description of the relationship between the existence of splittings and the topology of the boundary is completed by the following proposition, which serves as a converse to Theorem 3.5.17 and Proposition 5.4.4.

Proposition 5.4.6. *Let G be a hyperbolic group and let \mathcal{H} be a finite set of virtually cyclic subgroups of G such that $\partial(G, \widehat{\mathcal{H}})$ is connected. Suppose that G admits a non-trivial splitting over a virtually cyclic subgroup relative to \mathcal{H} . Then $\partial(G, \widehat{\mathcal{H}})$ contains a cut point or pair.*

Proof. Let T be the G -tree associated to such a non-trivial splitting. Without loss of generality assume that the action of G on T is minimal. Let e be any edge in T . Let T_1 and T_2 be the two components of the complement of the interior of e in T . Then as in the proof of Lemma 5.4.5 we obtain a partition of $\partial(G, \emptyset) - \Lambda G_e$ as $U_1 \sqcup U_2$ where U_i are open sets given by

$$U_i = \partial T_i \cup \bigcup_{w \in T_i} (\Lambda G_w - \Lambda G_e)$$

If a subgroup H of G is in \mathcal{H} , $H \leq G_w$ for some vertex $w \in \Sigma$ and either $\Lambda H = \Lambda G_e$ or $\Lambda H \cap \Lambda G_e = \emptyset$. In the latter case either $\Lambda H \subset U_1$ or $\Lambda H \subset U_2$. Let $q: \partial(G, \emptyset) \rightarrow \partial(G, \widehat{\mathcal{H}})$ be the quotient projection of Theorem 2.3.11. It follows that the images of U_1 and U_2 under q in the complement of $q(\Lambda G_e)$ in $\partial(G, \widehat{\mathcal{H}})$ are disjoint open sets; they are non-empty by the minimality of the action of G on T . The image of ΛG_e is either one or two points, and therefore $\partial(G, \widehat{\mathcal{H}})$ contains a cut point or pair. □

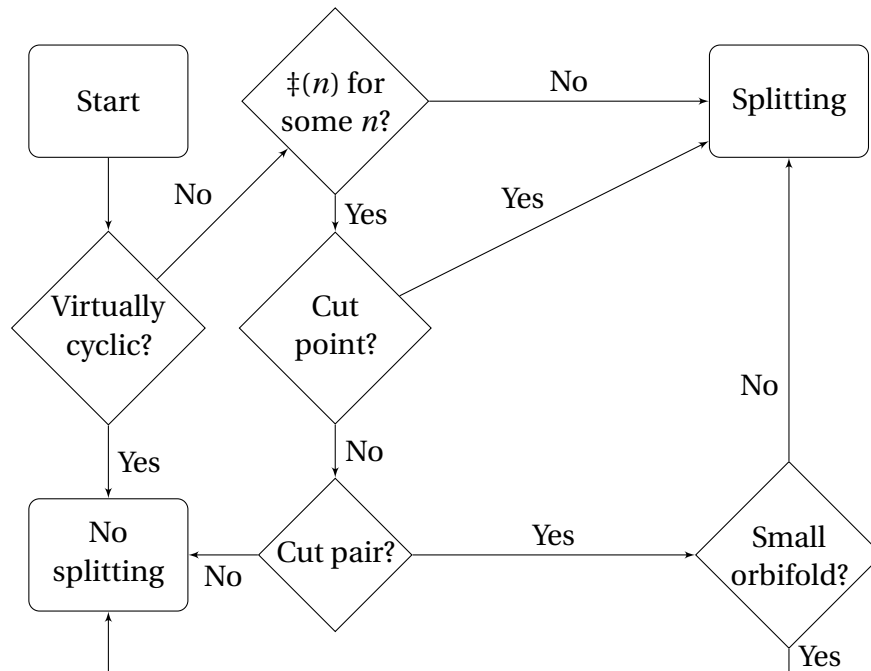


Figure 5.2: The decision process in the algorithm of Proposition 5.4.7.

5.4.3 Computing a maximal splitting

Proposition 5.4.7. *There is an algorithm that takes as input a presentation for a hyperbolic group G with a collection \mathcal{H} of virtually cyclic subgroups of G such that G appears as a vertex group in a minimal splitting of a one-ended hyperbolic group and \mathcal{H} is the set of incident edge groups and returns the answer to the question “does G split non-trivially over a virtually cyclic subgroup relative to \mathcal{H} ?”*

Proof. Let the given group be G and the peripheral structure be \mathcal{H} . First check whether or not G is virtually cyclic. If it is then G does not split properly over a virtually cyclic subgroup. If it is not then compute $\widehat{\mathcal{H}}$; then $\partial(G, \widehat{\mathcal{H}})$ contains more than a single point and the results of this section can be applied.

Next use the results described in Section 5.1.1 to compute δ such that the cusped space X associated to the pair $(G, \widehat{\mathcal{H}})$ is δ -hyperbolic. Search for a non-trivial splitting of G relative to \mathcal{H} using Lemma 5.2.3 and, in parallel, search for n such that $\ddagger(n)$ holds in X ; one of these processes must terminate by Proposition 3.6.10 and Theorem 3.5.17.

If a splitting is found then G does split non-trivially over a virtually cyclic subgroup relative to \mathcal{H} , and the algorithm can return “yes”. If X satisfies $\ddagger(n)$ then use the algorithm of Corollary 5.1.6 to check whether or not $\partial(G, \widehat{\mathcal{H}})$ contains a cut point. If it does then G does split properly over a virtually cyclic subgroup by Theorem 3.5.17

If there is no cut point, use the algorithm of Corollary 5.1.6 on X to determine whether or $\partial(G, \widehat{\mathcal{H}})$ contains a cut pair; if it does not then G does not split relative to \mathcal{H} by Proposition 5.4.6.

If there is a cut pair then simultaneously run the algorithms given in Lemmas 5.2.3 and 5.3.12. If the former terminates then a splitting has been found; if the latter does then no splitting exists. \square

Proposition 5.4.8. *There is an algorithm that, when given a presentation for a one-ended hyperbolic group, computes the graph of groups associated to a splitting of that group that is maximal for domination.*

Proof. Apply the algorithm of Lemma 5.2.3 to enumerate all minimal splittings of G over virtually cyclic subgroups. By Bestvina and Feighn's accessibility theorem (Theorem 3.5.1), this algorithm will eventually output a splitting that is maximal for domination. By Lemma 3.3.4 no vertex group in this maximal splitting admits a non-trivial splitting over its incident edge groups.

Conversely, let T be the tree associated to a minimal splitting of G over its virtually cyclic subgroups such that no vertex stabiliser admits a non-trivial splitting relative to the stabilisers of the incident edges and let T' be a minimal G -tree that dominates T . Let v be a vertex of T and let e be an edge of T with $\iota(e) = v$. Since T is minimal, the map $T' \rightarrow T$ is surjective, so there exists an edge e' of T' such that the image of e' contains e . Then $G_{e'} \leq G_e$. By assumption, T' is minimal and G is one-ended, so $G_{e'}$ is infinite, so it is a finite index subgroup of G_e , so the orbit of e' under the action of G_e is finite. It follows that G_e is elliptic with respect to the action on T' . Therefore G_v is elliptic with respect to the action on T' , since otherwise this action gives a non-trivial splitting of G_v relative to $\{G_e \mid \iota(e) = v\}$. This argument applies to every vertex of T , so T dominates T' by Lemma 3.3.3 and it follows that T is maximal for domination.

For each splitting returned by the algorithm of Lemma 5.2.3 use the algorithm of Proposition 5.4.7 to test whether each vertex group splits relative to its incident edge groups. When a splitting is found such that no vertex group splits relative to its incident edge groups, return that splitting. \square

5.5 JSJ decompositions

In this section we complete the proof of Theorem 1.2.2: we show that three closely related types of JSJ splittings are computable for hyperbolic groups. Fix a one-ended hyperbolic group G . We define three classes of virtually cyclic subgroups of G .

1. Let \mathcal{VC} be the set of virtually cyclic subgroups of G .
2. Let \mathcal{Z} be the set of virtually cyclic subgroups of G with infinite centre. Equivalently, \mathcal{Z} is the set of virtually cyclic subgroups of G that surject onto \mathbb{Z} .
3. Let \mathcal{Z}_{\max} be the set of subgroups of G in \mathcal{Z} that are maximal among subgroups of \mathcal{Z} .

We consider JSJ decompositions over groups in these three classes of subgroups. Note that here we follow [DG11] in defining the \mathcal{Z}_{\max} -JSJ: since a proper subgroup of a \mathcal{Z}_{\max} group is not \mathcal{Z}_{\max} , this does not fit into Guirardel and Levitt's framework for JSJ theory. We call a splitting over \mathcal{Z}_{\max} a \mathcal{Z}_{\max} -JSJ if it is \mathcal{Z} -universally elliptic and is maximal for domination among \mathcal{Z} -universally elliptic splittings over \mathcal{Z}_{\max} .

5.5.1 Virtually cyclic edge groups and Bowditch's decomposition

In this section we show that a JSJ decomposition over \mathcal{VC} is computable. We then show that Bowditch's canonical decomposition is a \mathcal{VC} -JSJ and show how to obtain this special JSJ algorithmically. This completes the proof of the first part of Theorem 1.2.2.

We first prove the following lemma.

Lemma 5.5.1. *The tree obtained from the tree associated to a maximal splitting by collapsing each edge whose stabiliser is not universally elliptic is a \mathcal{VC} -JSJ tree.*

Proof. Let T be the tree associated to a maximal splitting and let T' be the tree obtained by collapsing each edge of T that is not universally elliptic. Then certainly T' is universally elliptic, so it is sufficient to show that if Σ is another universally elliptic G -tree then T' dominates Σ .

The tree Σ can be refined to dominate T , so there exists a map $f: T \rightarrow \Sigma$. Let v be a vertex of T' and let S be a component of its preimage in T . Then $f|_S$ is constant: if an edge e in S is mapped into an edge e' in Σ then $G_e \leq G_{e'}$, which is universally elliptic. But then the image of S in T' contains more than a single vertex. Therefore G_v fixes the vertex $f(S)$ in Σ , so is elliptic with respect to Σ . This shows that T' dominates Σ . \square

We must now identify the edges in the tree associated to the maximal splitting that are not maximally elliptic. We make the following definitions.

Definition 5.5.2. An *extended Möbius strip group* is a virtually cyclic group H of type \mathcal{Z} with peripheral structure consisting of the preimage of $2\mathbb{Z}$ under the canonical surjection $H \rightarrow \mathbb{Z}$. Equivalently, it is an extension of the pair $(\mathbb{Z}, \{2\mathbb{Z}\})$ by a finite group.

Recall Definition 3.5.8 of a hanging Fuchsian group.

Definition 5.5.3. We say that an edge e connecting vertices v_1 and v_2 of a G -tree is a *internal surface edge* if, for each i , either G_{v_i} is a hanging Fuchsian group and G_e is maximal among virtually cyclic subgroups of G_{v_i} , or G_{v_i} is an extended Möbius strip group and $G_e \leq G_{v_i}$ is the peripheral subgroup of G_{v_i} .

Lemma 5.5.4. *Suppose that T is a maximal \mathcal{VC} -tree. Then the edges of T that are not universally elliptic are precisely the internal surface edges.*

Proof. Let T' be the tree obtained by collapsing each edge of T that is not universally elliptic as in Lemma 5.5.1, so T' is a JSJ tree and there is a collapse map from T to T' . The edges of T that are not universally elliptic are precisely those edges that are mapped to flexible vertices of T' under the collapse map; by [GL17, Theorem 6.2] all flexible vertices of T' are hanging Fuchsian vertices.

Any splitting of a hanging Fuchsian group is dual to a family of curves on the associated orbifold, so any edge in such a splitting is an internal surface edge, so all edges that are not universally elliptic are internal surface edges.

Conversely, let e be an internal surface edge. Let T' be the tree obtained by collapsing each edge in the orbit of e . Let v be the vertex of T' in the image of e . Then T is obtained from T' by refining at v . This refinement is dual to an essential simple closed curve ℓ on the associated orbifold Q . Then Q contains another essential simple closed curve ℓ' that is not homotopic to a curve disjoint from ℓ . Refine T' at v dual to ℓ' to obtain a tree T'' ; then G_e is not elliptic with respect to T'' . \square

Theorem 5.5.5. *There is an algorithm that takes as input a presentation for a one-ended hyperbolic group and returns as output the graph of groups associated to a \mathcal{VC} -JSJ decomposition for that group.*

Proof. First compute a maximal splitting of the group over virtually cyclic subgroups by Theorem 5.4.8. Let T be the associated tree. By construction T is reduced; in any case, T can easily be made reduced using the processes of Lemma 5.2.1. For each edge e connecting vertices v_1 and v_2 of the graph of groups T/G determine whether G_e is maximal in G using the algorithm of Lemma 5.2.1 and whether the two vertex groups G_{v_1} and G_{v_2} have circular boundary relative to their incident edge groups by Corollary 5.1.7. Check also whether each of G_{v_1} and G_{v_2} is virtually cyclic of \mathcal{Z} -type, and, if it is, whether or not G_e has index 2 in that group. One of these possibilities is the case if and only if e is not universally elliptic by Lemma 5.5.4; collapse all edges where this is the case. \square

Recall the description of Bowditch's canonical decomposition from Section 3.5.2. We now show that this decomposition of a one-ended hyperbolic group is computable. First we show that it is a \mathcal{VC} -JSJ decomposition.

Proposition 5.5.6. *Bowditch's canonical decomposition of a one-ended hyperbolic group is a \mathcal{VC} -JSJ decomposition.*

Proof. Let T be the tree associated to Bowditch's canonical decomposition of a one-ended hyperbolic group G with vertex set partitioned as $T_1 \sqcup T_2 \sqcup T_3$. First note that stabilisers of edges in T are virtually cyclic.

By [GL17, Corollary 9.20] G admits a JSJ decomposition over its virtually cyclic subgroups in which all flexible vertices are hanging Fuchsian. Therefore the stabiliser of any flexible vertex in the JSJ tree stabilises a point in T_2 by [Bow98b, Proposition 6.1]. Automatically the stabiliser of any rigid vertex in the JSJ tree is elliptic in T , so the JSJ decomposition dominates T and T is universally elliptic.

Therefore every edge in T is elliptic in the JSJ tree. The stabiliser of a vertex in T_1 is virtually cyclic and it contains stabilisers of its incident edge groups at finite index, so it follows that the stabiliser of a vertex in T_1 is elliptic in the JSJ tree. As in the proof of Lemma 5.5.4 no stabiliser of a vertex in T_2 admits a non-trivial universally elliptic splitting over a virtually cyclic subgroup relative to its incident edge groups, so the stabiliser of any vertex in T_2 is elliptic in the JSJ tree. The stabiliser of a vertex in T_3 does not split non-trivially over a virtually cyclic subgroup relative to its incident edge groups (see the remark after [Bow98b, Proposition 5.31]), so the stabiliser of any vertex in T_3 is elliptic in the JSJ tree.

It follows that T dominates the JSJ decomposition, so T is a JSJ tree. \square

In [DG11] Dahmani and Guirardel define an operation on \mathcal{VC} -trees. Their construction is more general, but here we quote the special case that is relevant to Bowditch's decomposition.

Definition 5.5.7. Let T be a \mathcal{VC} -tree. Define the *commensurability* equivalence relation \sim on \mathcal{VC} by letting $A \sim B$ if and only if A and B lie in the same maximal virtually cyclic subgroup of G . Also denote by \sim the equivalence relation on the set of edges of T defined by letting $e \sim e'$ if and only if $G_e \sim G_{e'}$. A *cylinder* is a subset of T that is the union of all edges in a \sim -equivalence class.

A cylinder in a subtree of T : if $e \sim e'$ and e'' is an edge separating e' from e then $G_{e''} \leq G_e \cap G_{e'}$, so $G_{e''}$ is commensurable with G_e and $G_{e'}$.

Definition 5.5.8. Let T be a \mathcal{VC} -tree. The corresponding *tree of cylinders* T_c is a bipartite tree with vertex set $V_1 \sqcup V_2$, where V_1 is the set of vertices of T that lie in at least two cylinders and V_2 is the set of cylinders in T . A vertex $v \in V_1$ is connected by an edge to $Y \in V_2$ if and only if $v \in Y$.

Remark 5.5.9. Guirardel and Levitt then define a *collapsed tree of cylinders* obtained by collapsing all edges of the tree of cylinders that are not in the class of allowed edge groups. This is unnecessary in this special case: let e be an edge contained in some cylinder in T and let g be in the stabiliser of that cylinder. Then $gG_e g^{-1} = G_{g \cdot e} \sim G_e$, so $\overline{gG_e g^{-1}} = \widehat{G_e}$ and therefore g is in the normaliser of $\widehat{G_e}$, which is $\widehat{G_e}$ itself. So stabilisers of cylinders are virtually cyclic, so all edge stabilisers in the tree of cylinders are virtually cyclic.

Lemma 5.5.10. [GL17, Lemma 7.3] *The tree T_c is equal to its own tree of cylinders $(T_c)_c$.*

Lemma 5.5.11. *The G -tree corresponding to Bowditch's canonical decomposition is the tree of cylinders of any \mathcal{VC} -JSJ tree.*

Proof. Using Proposition 5.5.6 and Lemma 5.5.10 it is enough to show that Bowditch's tree is equal to its own tree of cylinders. Recall that Bowditch's tree T is bipartite with vertex set partitioned as $T_1 \sqcup (T_2 \sqcup T_3)$ and the stabiliser of any vertex in T_1 is maximal virtually cyclic. Therefore cylinders in Bowditch's tree are precisely the stars of vertices in T_1 , from which the result follows immediately. \square

Bowditch's canonical decomposition is the graph of cylinders of the decomposition obtained in this way. The operation of replacing a decomposition with the decomposition associated to its tree of cylinders can be done algorithmically using Lemma 5.2.1. This result is the content of [DG11, Lemma 2.34]; note that while the result is stated for a \mathcal{Z} -tree, replacing this with a \mathcal{VC} -tree makes no difference to the proof. Therefore the results of this section give the following corollary.

Corollary 5.5.12. *Bowditch's canonical decomposition of a one-ended hyperbolic groups is computable.*

Remark 5.5.13. Using Lemma 5.2.1 and Corollary 5.1.7 one can algorithmically recover the partition $T_1 \sqcup T_2 \sqcup T_3$ of the vertex set of Bowditch's canonical decomposition.

5.5.2 \mathcal{Z} edge groups

We now prove the second part of Theorem 1.2.2.

Theorem 5.5.14. *There is an algorithm that takes as input a presentation for a one-ended hyperbolic group and returns the graph of groups associated to a \mathcal{Z} -JSJ decomposition for that group.*

In [DG11] it is shown that the \mathcal{Z} -JSJ decomposition is closely related to the \mathcal{VC} -JSJ: a \mathcal{Z} -JSJ tree can be obtained from a \mathcal{VC} -JSJ tree T by first refining T by applying the so-called mirrors splitting to each hanging Fuchsian vertex group and then collapsing each edge with stabiliser of dihedral type. The second of these processes can be done algorithmically using the part of the algorithm of Lemma 5.2.1 that determines whether or not a given virtually cyclic group is of dihedral type. Therefore we must now show that the mirrors splitting is computable.

Recall the definition of the mirrors splitting of the fundamental group of a compact 2-dimensional orbifold Q from [DG11].

Definition 5.5.15. Let N be a regular neighbourhood of the union of the mirrors and D_∞ -boundary components of Q that does not contain any cone point of Q . If $Q - N$ is an annulus or a disc with at most one cone point then the *mirrors splitting* of $\pi_1 Q$ is defined to be trivial; otherwise it is the splitting obtained by cutting Q along each component of ∂N .

If the mirrors splitting is non-trivial then the graph of groups associated to the splitting is a star; the group at the central vertex is the fundamental group of an orbifold with no mirrors and the group at each leaf is the fundamental group of an orbifold with no cone points and underlying surface an annulus, one of whose topological boundary components is a circular orbifold boundary component and the other a union of interval boundary components and at least one mirror. If G is any hyperbolic group with a collection \mathcal{H} of virtually cyclic subgroups such that $\partial(G, \widehat{\mathcal{H}})$ is homeomorphic to a circle then the mirrors splitting of G relative to \mathcal{H} is defined to be the splitting induced by the mirrors splitting of the quotient G by a maximal finite normal subgroup of G as in Proposition 5.3.3.

Lemma 5.5.16. *The mirrors splitting of a hyperbolic group G with a set \mathcal{H} of virtually cyclic subgroups such that $\partial(G, \widehat{\mathcal{H}})$ is homeomorphic to a circle is computable.*

Proof. Using Proposition 5.3.3 it is enough to show that the mirrors splitting is computable in the case where G is bounded Fuchsian and \mathcal{H} is a collection of representatives of peripheral subgroups of G . To do this we enumerate all mirrors splittings: for each non-negative integer k enumerate all fundamental groups of compact orbifolds without mirrors and with at least k boundary components and all k -tuples of fundamental groups of orbifolds homeomorphic to an annulus with no cone points and such that one topological

boundary component of the orbifold is a circular orbifold boundary component. In each case form the graph of groups in which the underlying graph is a k -pointed star, the group at the central vertex is the fundamental group of the orbifold without mirrors, the group at each leaf is the fundamental group of an orbifold homeomorphic to an annulus and the group at each edge is infinite cyclic and is identified with the fundamental group of a circular orbifold boundary component of each of the orbifolds associated to the end points of that edge. Compute the fundamental group of each such graph of groups and record also the peripheral structure consisting of conjugacy class representatives of the fundamental groups of components of the orbifold boundary of the orbifold.

Also enumerate all groups with trivial mirrors splitting, i.e. fundamental groups of orbifolds homeomorphic as topological spaces to a disc with at most one cone point, or homeomorphic to an annulus with no cone points and such that one topological boundary component is a circular orbifold boundary component.

In parallel enumerate all homomorphisms from the fundamental groups of these graphs of groups to G and all homomorphisms from G to the fundamental groups of these graphs of groups. Some such pair of homomorphisms is an inverse pair that preserves the peripheral structure up to conjugacy. On finding this pair the algorithm returns the associated mirrors splitting. \square

5.5.3 \mathcal{I}_{\max} edge groups

The \mathcal{I}_{\max} -JSJ decomposition can be obtained from a \mathcal{I} -JSJ decomposition by performing the so-called \mathcal{I}_{\max} -fold: see [DG11, Proposition 4.11]. Let T be the G -tree corresponding to a \mathcal{I} -JSJ. Dahmani and Guirardel show that the tree corresponding to the \mathcal{I}_{\max} -fold is obtained by iterating the following process. Suppose that there exists an edge e in T such that $G_e \neq \widehat{G}_e$. Then there exists such an edge e such that $\widehat{G}_e \leq G_{l(e)}$. Then quotient T by the G -equivariant equivalence relation generated by $e \sim g \cdot e$ for each $g \in \widehat{G}_e$. Repeat this process until $G_e = \widehat{G}_e$ for each edge e . Note that this process takes at most as many steps as there are edges in $G \setminus T$, which is finite as long as T is minimal.

We now note that this process can be performed algorithmically, which completes the proof of the final part of Theorem 1.2.2.

Lemma 5.5.17. *There is an algorithm that takes as input a graph of groups decomposition of a hyperbolic group G over virtually cyclic subgroups and returns the graph of groups associated to the \mathcal{I}_{\max} -fold of the associated G -tree.*

Proof. Let \mathcal{Y} be the graph of groups with underlying graph Y associated to a \mathcal{I} -JSJ. We describe the folding process described above at the level of the graph of groups \mathcal{Y} .

Unfortunately this description is a little more complicated than the description at the level of the tree. Note that the folds do not change the quotient graph $G \backslash T$, since $e \sim e'$ only when e and e' lie in the same G -orbit.

Take an edge e of Y with the following two properties.

1. Let $v = \iota(e)$. Then $\partial_-^e(G_e)$ is not a maximal virtually cyclic subgroup of \mathcal{Y}_v .
2. For each edge e' with $\iota(e') = v$ with $\partial_-^{e'}(\mathcal{Y}_{e'})$ contained in the maximal virtually cyclic subgroup of \mathcal{Y}_v containing $\partial_-^e(\mathcal{Y}_e)$, $\partial_+^{e'}(\mathcal{Y}_{e'})$ is a maximal virtually cyclic subgroup of $\mathcal{Y}_{\tau(e)}$.

Note that if such an edge exists then it can be found using the algorithm of Lemma 5.2.1.

Then change \mathcal{Y} by making the following two replacements.

1. Replace $\mathcal{Y}_{\tau(v)}$ by the fundamental group of the graph of groups \mathcal{Y}' with a single edge e , $\mathcal{Y}'_e = \mathcal{Y}_e$, $\mathcal{Y}'_{\tau(e)} = \mathcal{Y}_{\tau(e)}$ and $\mathcal{Y}'_{\iota(e)}$ equal to the maximal virtually cyclic subgroup of $\mathcal{Y}_{\iota(v)}$ containing $\partial_-^e(\mathcal{Y}_e)$. Note that this fundamental group is either an amalgamated product or an HNN extension, so this can easily be done algorithmically.
2. Replace \mathcal{Y} by the maximal virtually cyclic subgroup of $\mathcal{Y}_{\iota(v)}$ containing $\partial_-^e(\mathcal{Y}_e)$.

Repeat this process until there is no edge e satisfying the given condition. □

This completes the proof of Theorem 1.2.2.

5.6 Generalisations

In light of Haulmark's theorem [Hau17], it seems possible that the techniques used in this chapter might be applicable to determining whether or not groups in a larger class of relatively hyperbolic groups admit non-trivial splittings over virtually cyclic subgroups. As discussed in Section 4.4, the main impediment to our ability detect cut pairs in the boundary in more general situations is the double dagger condition: except when the peripheral subgroups are either abelian or virtually cyclic, we do not know that the cusped space satisfies a double dagger condition, even when the boundary is connected and without a cut point. A resolution to this problem would answer [DT13, Question 1.8] and would therefore extend the class of relatively hyperbolic groups for which the isomorphism problem is known to be solvable: in light of Dahmani and Touikan's work, the problem of detecting rigidity is the only barrier to the solution for the isomorphism problem for relatively hyperbolic groups with virtually nilpotent peripheral subgroups. The class of

relatively hyperbolic groups with virtually nilpotent peripheral subgroups is particularly natural, since the Margulis lemma tells us that groups that act geometrically finitely on complete simply connected Riemannian manifolds with pinched negative sectional curvature fall into this class.

Chapter 6

The cohomology of decomposition spaces

Epstein asked whether there is an algorithm that, when given a presentation for a hyperbolic group G , computes the Čech cohomology $\check{H}^k(\partial G)$ as a G -module. This problem appears as Question 1.18 in Bestvina's problem list [Bes04]. Throughout this chapter all Čech cohomology groups will have coefficients in \mathbb{Z} , although most of the results given apply to other rings. Our purpose in this chapter is to answer a relative version of this question in the case of Otal's decomposition space, a special case of the Bowditch boundary of a relatively hyperbolic group. We then apply this result to answer Epstein's question for a restricted class of hyperbolic groups: those arising as fundamental groups of graphs of free groups with cyclic edge groups.

Fix a free group F of rank n with free generating set S . Let T be the corresponding Cayley graph containing a vertex 1 corresponding to the identity in F . Fix a finite set $\{w_i\}$ of cyclically reduced words in F . Then, following Otal [Ota92], we make the following definitions.

Definition 6.0.1. The *line pattern* \mathcal{L} associated to the set $\{w_i\}$ is defined to be the set of lines spanned by the points $\{gw_i^k\}_k$ for $g \in F$. Let \mathcal{D} be the associated *decomposition space*: this is the quotient of ∂F by the equivalence relation that identifies the two end points of each line in \mathcal{L} . Let $q: \partial F \rightarrow \mathcal{D}$ be the quotient projection.

If we assume that each of the elements of $\{w_i\}$ is not a proper power and that for distinct i and j the element w_i is not conjugate in F to w_j^\pm then \mathcal{D} is the Bowditch boundary of group F relative to $\langle\langle w_i \rangle\rangle$ by a result of Tran [Tra13]: see Theorem 2.3.11.

We present an algorithm that computes the Čech cohomology of \mathcal{D} as an F -module. The F -module structure is induced by the action of F on ∂F , which descends to an action

of F on \mathcal{D} . The cohomology functor is contravariant, so we use the action of g^{-1} on \mathcal{D} to define the action of g on $\check{H}^k(\mathcal{D})$. Given $g \in F$ let ϕ_g be the corresponding homeomorphism of \mathcal{D} . Then we define the action of g on $\check{H}^k(\mathcal{D})$ by setting $g \cdot x = (\phi_{g^{-1}})^* x$.

The Čech cohomology of \mathcal{D} is defined to be the direct limit over open covers \mathcal{U} of \mathcal{D} that provide successively better combinatorial approximations to \mathcal{D} of the singular cohomology of the nerve of \mathcal{U} . In Section 6.1 we shall see how to associate an open cover of \mathcal{D} to a finite subtree of T . The combinatorics of such a cover will then be recorded in the Whitehead graph associated to that subtree. Therefore Whitehead graphs provide approximations to \mathcal{D} , with larger Whitehead graphs providing better approximations.

These open covers have no triple intersections, so we immediately see that the Čech cohomology is concentrated in dimensions zero and one. This could also be seen as a consequence of the Urysohn–Menger addition theorem [Eng78, 3.1.17]: if X is a compact metric space and A and B are subsets of X with $X = A \cup B$ then $\dim X \leq \dim A + \dim B + 1$ where \dim is the Lebesgue covering dimension. The decomposition space \mathcal{D} can be written as a union of the set of points in ∂F that are not end points of lines in \mathcal{L} with the set of images in \mathcal{D} of end points of lines in \mathcal{L} . The first of these sets is a subset of a Cantor set, so has dimension zero, while the second is countable, so also has dimension zero, thus proving that \mathcal{D} has dimension at most one.

In Sections 6.2 and 6.3 we show how to compute the Čech cohomology in dimensions zero and one respectively. In dimension zero this could be accomplished by computing a maximal free splitting of the group relative to the set $\{\langle w_i \rangle\}$ of subgroups of F and then relating the decomposition space to the tree associated to the splitting; in Section 6.4 we apply an argument along these lines. In Section 6.2 we take an alternative approach: our methods rely on showing that some (large) finite subtree of T contains sufficient information to compute the Čech cohomology. This approach is based on the proof of [CM11, Lemma 4.12].

In Section 6.4 we apply these results to prove the computability of the Čech cohomology of the boundary of a hyperbolic fundamental group of a graph of free groups with virtually cyclic edge groups. This result answers Epstein’s question for this restricted class of hyperbolic groups.

6.1 Whitehead graphs and open covers

Recall the following definition from [CM11].

Definition 6.1.1. Let X be a subtree of T . Then we define the *Whitehead graph* $\text{Wh}(X)$ of \mathcal{L} at X to be the graph with a vertex corresponding to each vertex of T adjacent to (but

not contained in) X and an edge connecting a pair of vertices in $\text{Wh}(X)$ for each line in \mathcal{L} connecting the corresponding vertices in T . The edges of $\text{Wh}(X)$ are labelled by the corresponding lines.

For more information about Whitehead graphs and their applications, see [Man10].

Definition 6.1.2. For $v \in T$ let $\mathcal{S}_1(v) \subset \partial T$ be the *shadow* of v from 1 as defined in [CM11]: the set of boundary points $p \in \partial T$ such that the geodesic $[1, p]$ contains v . These sets are open and closed and the collection of such sets as v varies in T is a basis for the topology on ∂T .

Lemma 6.1.3. *Let X be a finite subtree of T containing 1. Then there is a covering of \mathcal{D} by a collection of open sets U_i in bijection with the vertices a_i of $\text{Wh}(X)$ satisfying the following conditions.*

- *The projection of $\mathcal{S}_1(a_i)$ to \mathcal{D} is a subset of U_i .*
- *The pairwise intersections of sets in the cover can be read from the Whitehead diagram: $U_i \cap U_j \neq \emptyset$ if and only if there is an edge connecting a_i and a_j in $\text{Wh}(X)$.*
- *There are no triple intersections.*

Proof. We aim to construct open sets V_i covering ∂T satisfying the following conditions.

- The shadow $\mathcal{S}_1(a_i)$ is contained in V_i .
- Intersections are determined by the Whitehead graph: $V_i \cap V_j \neq \emptyset$ if and only if there is an edge connecting a_i and a_j in $\text{Wh}(X)$.
- There are no triple intersections.
- For each line l in the line pattern \mathcal{L} , each set V_i contains either both of the end points of l or neither.

Then the projection of these sets in \mathcal{D} satisfies the requirements of the lemma.

We build these inductively. For the first step, take $V_i^0 = \mathcal{S}_1(a_i)$. Then, for each i , there are finitely many lines in the line pattern passing through a_i . Add to V_i^0 an open neighbourhood of the end point not in V_i^0 of each such line to obtain V_i^1 . This can be done in a way that ensures that V_i^1 is the union of V_i^0 and finitely many other shadows, that the open sets added are all disjoint, and that no line in the line pattern has an end in two different added sets. This is possible since if a subset of ∂T is a union of finitely

many shadows then only finitely many lines in the line pattern have exactly one end in that subset.

After each V_i^1 is defined, continue inductively, ensuring that each V_i^k is the union of finitely many shadows, so that if a line in the line pattern has one end in V_i^k then its other end is in V_i^{k+1} . We can do this without introducing any new intersections, so all intersections correspond to lines from $\mathcal{S}_1(a_i)$ to $\mathcal{S}_1(a_j)$ for some i and j , so all intersections correspond to edges in the Whitehead graph and there are no triple intersections.

Then let $U_i = q(\cup_k V_i^k)$; these sets cover \mathcal{D} and have the required properties. \square

For X a finite subtree of T , we shall denote by \mathcal{U}_X a cover of \mathcal{D} associated to X as constructed in Lemma 6.1.3. This cover is not unique, but this non-uniqueness will not cause us a problem, since any two such covers are very similar. In particular, they admit a common refinement, and refinement between different open covers associated to X as in Lemma 6.1.3 induces a natural isomorphism between the singular cochain complexes of the nerves of those covers. If $X \subset X'$ then $\mathcal{U}_{X'}$ can be chosen to be a refinement of \mathcal{U}_X and the refinement map $\mathcal{U}_{X'} \rightarrow \mathcal{U}_X$ is canonical.

It will be convenient to define an open cover associated to the empty subtree of X : this is the trivial covering $\mathcal{U}_\emptyset = \{\mathcal{D}\}$.

Lemma 6.1.4. *Let \mathcal{W} be a finite open cover of \mathcal{D} . Then some refinement of \mathcal{W} is of the form given in Lemma 6.1.3.*

Proof. Let \mathcal{V} be the pullback of \mathcal{W} to ∂T . Consider the set

$$C = \{a \in T \mid \mathcal{S}_1(a) \subset V \text{ for some } V \in \mathcal{V}\}. \quad (6.1)$$

The collection $\{\mathcal{S}_1(x) \mid x \in T\}$ is a basis for the topology on ∂T so sets of the form $\mathcal{S}_1(a)$ for $a \in C$ cover each $V \in \mathcal{V}$. Hence such sets cover ∂T .

The boundary ∂T is compact, so there is a finite set of points a_1, \dots, a_n such that $\{\mathcal{S}_1(a_i)\}$ covers ∂T and each $\mathcal{S}_1(a_i)$ is contained in some $V_{\sigma(i)} \in \mathcal{V}$. Let H be the convex hull of $\{a_i\} \cup \{1\}$. Call vertices of H adjacent to vertices in $T - H$ boundary vertices. If we take $\{a_i\}$ to be minimal with its covering property then the set of boundary points of H is precisely $\{a_i\}$. Let X be the subtree of H obtained by pruning off its boundary vertices.

Let $\mathcal{U}_X = \{U_i\}$ be the finite cover of \mathcal{D} corresponding to X as in Lemma 6.1.3. Define a new set \mathcal{U}' of open subsets of \mathcal{D} by

$$\mathcal{U}' = \{U_i \cap q(V_{\sigma(i)})\}. \quad (6.2)$$

The set \mathcal{U}' is a cover of \mathcal{D} since it covers $q(\mathcal{S}_1(a_i))$ for each i . It is certainly a refinement of \mathcal{W} and it is easy to check that it corresponds to X in the sense of the statement of Lemma 6.1.3. \square

The results of this section together imply the following corollary:

Corollary 6.1.5. *The finite Whitehead diagrams determine the cohomology of \mathcal{D} :*

$$\check{H}^n(\mathcal{D}) = \varinjlim_X \check{H}^n(\mathcal{U}_X) \quad (6.3)$$

where the direct limit is taken over finite subtrees X of T , partially ordered by inclusion. \square

6.2 Computing $\check{H}^0(\mathcal{D})$

For each element $[\sigma] \in \check{H}^0(\mathcal{D})$ there exists a subtree $X \subset T$ such that $[\sigma]$ is represented by some $\sigma \in \check{H}^0(\mathcal{U}_X) = \ker(d^0: \check{C}^0(\mathcal{U}_X) \rightarrow \check{C}^1(\mathcal{U}_X))$. Such an element is an assignment of an integer to each connected component of $\text{Wh}(X)$. In this situation we shall say that $[\sigma]$ is *supported on X* and we shall refer to the minimal such subtree as the *support of $[\sigma]$* . A unique minimal such subtree exists: it is clear that if $[\sigma]$ is supported on X_1 and on X_2 then it is supported on $X_1 \cap X_2$.

As discussed at the beginning of the chapter, F acts on \mathcal{D} by homeomorphisms, giving the Čech cohomology the structure of an F -module with g acting as $\phi_{g^{-1}}^*$ on $\check{H}^0(\mathcal{D})$. In terms of Whitehead graphs, any $g \in F$ induces a map

$$g: \check{H}^0(\mathcal{U}_X) \rightarrow \check{H}^0(g \cdot \mathcal{U}_X) = \check{H}^0(\mathcal{U}_{g(X)}). \quad (6.4)$$

This map takes an element represented by a \mathbb{Z} -labelling of the connected components of $\text{Wh}(X)$ to the element represented by the translate by g of this labelled diagram.

We now aim to find an algorithm that computes a presentation for this F -module. First we describe an algorithm that computes a generating set. The argument is loosely based on the proof of [CM11, Lemma 4.12].

Proposition 6.2.1. *There exists a finite number N , computable from n and \mathcal{L} , such that $\check{H}^0(\mathcal{D})$ is generated as an abelian group by 0-cycles supported on subtrees of T with at most N vertices.*

Proof. Let $[\sigma]$ be a 0-cycle supported on a subtree X of T with more than N vertices, where N is to be chosen later. Then by induction it is sufficient to show that $[\sigma]$ can be written as




Figure 6.1: A partition of a disconnected Whitehead graph into two connected components. An assignment of the integer 1 to the red part and 0 to the blue part represents an element of $\check{H}^0(\mathcal{D})$. In this example F is free on two generators and the line pattern is generated by the word a^2bab .

the sum of 0-cycles supported on strictly smaller subtrees. The idea is that if N is large enough then there will be two vertices in X at which $[\sigma]$ looks similar and cutting out everything between these vertices allows us to split $[\sigma]$ into strictly smaller summands.

$\check{H}^0(\mathcal{Z}_X)$ is generated by 0-cycles represented by Whitehead diagrams at X with one connected component labelled 1 and the others labelled 0. We can assume without loss of generality that $[\sigma]$ is such a 0-cycle. Then $[\sigma]$ can be thought of as a partition of $\text{Wh}(X)$ into a connected component and its complement. An example of such a partition is shown pictorially in Figure 6.1.

Suppose that N is large enough that any subtree of T with more than N vertices is guaranteed to contain an embedded arc of length at least $M + 2$, where M is a computable function of n and \mathcal{L} to be chosen later. Then let v_1, \dots, v_M be the interior vertices of such an embedded arc in X . Traversing this arc in the direction from v_1 to v_M , record for each vertex v_i an ordered pair (s_i, t_i) of elements of S^\pm , where s_i labels the incoming edge at v_i of the embedded arc and t_i labels the outgoing edge.

Suppose that M is large enough that at least K of these pairs are equal. Here K is a computable function of \mathcal{L} to be chosen later. Then let v_{i_1}, \dots, v_{i_K} be vertices with equal associated edge pairs. The edges of each $\text{Wh}(v_{i_j})$ extend to edges in $\text{Wh}(X)$, hence the partition of $\text{Wh}(X)$ associated to $[\sigma]$ gives a partition of the edges of $\text{Wh}(v_{i_j})$ into a subset and its complement.

Treating the v_{i_j} as elements of F , the translation of $\text{Wh}(v_{i_j})$ by $v_{i_j}^{-1}$ gives a partition on the edges of $\text{Wh}(1)$. There is a finite number of such partitions; let K be greater than that number. Then we obtain v and $w = g(v)$ in $\{v_{i_1}, \dots, v_{i_K}\}$ such that the translates of the associated partitions agree.

Now we define two disjoint subsets of X . Let A be the vertices $u \neq v$ of X such that the geodesic in T from w to u passes through v , and let B be the same with the rôles of v




Figure 6.2: Vertices v and w are chosen so that σ induces the same partition on $\text{Wh}(v)$ as on $\text{Wh}(w)$. Then τ is an element of $\check{H}^0(\mathcal{U}_{A \cup g^{-1}B})$ chosen to induce the same partition on $\text{Wh}(A')$ as σ does.

and w reversed. Without loss of generality, A contains at least as many vertices as B does. Then let $A' = A \cup \{v\}$. See Figure 6.2.

We now cancel off the part of $[\sigma]$ supported on A . Let $Y = A' \cup g^{-1}B$. As described above, σ induces a partition on the edges of $\text{Wh}(u)$ for each vertex u in A' and B , and hence by translation on $\text{Wh}(u)$ for each vertex u in Y . The partitions at $v \in A'$ and the vertex of $g^{-1}B$ adjacent to v are consistent with respect to the attaching map between the two Whitehead graphs, so we obtain a partition of the graph $\text{Wh}(Y)$ into its connected components. Hence we can define $\tau \in \check{H}^0(\mathcal{U}_Y)$ to be a cycle represented by assigning the integer one to one component of $\text{Wh}(Y)$ and zero to the others in a way that agrees at the vertices of A' with the labelling of components represented by σ .

Then $[\tau]$ is supported on Y , which contains at most $|A'| + |B| < |X|$ vertices, and $[\sigma] - [\tau]$ is supported on $(X - A') \cup g^{-1}B$, which has fewer vertices than X since A' has more vertices than B . \square

Therefore $\check{H}^0(\mathcal{D})$ admits a finite generating set as an F -module and this generating set can be computed from the line pattern: this is equivalent to finding generators for the kernel of the restriction of d to an abelian group of finite rank, which may be done using Smith Normal Form. Let $[\sigma_1], \dots, [\sigma_k]$ be a generating set for $\check{H}^0(\mathcal{D})$. This is equivalent to a surjection $p: (\mathbb{Z}F)^k \rightarrow \check{H}^0(\mathcal{D})$ of F -modules. Let e_i be the i^{th} basis vector in the free module, and let it be mapped to $[\sigma_i]$ under p . To complete the computation of a presentation for $\check{H}^0(\mathcal{D})$ we need an algorithm that computes a generating set for the kernel of p .

For each $[\sigma_i]$ let X_i be the support of $[\sigma_i]$. A general element $x \in (\mathbb{Z}F)^k$ is of the form

$$x = \sum_j \left(\sum_i n_{ij} g_{ij} \right) e_j, \text{ where } n_{ij} \in \mathbb{Z}, g_{ij} \in F. \quad (6.5)$$

Define the support of x to be the convex hull

$$\text{hull}\left(\bigcup_{i,j} g_{ij}(X_j)\right). \quad (6.6)$$

Note that the support of px is contained in the support of x .

We can now state and prove a theorem that shows that the kernel of p is generated by elements of bounded size, in the same way that Proposition 6.2.1 shows that $\check{H}^0(\mathcal{D})$ is generated by elements of bounded size.

Proposition 6.2.2. *ker p is generated as an abelian group by elements whose supports have at most N vertices, where N is a computable function of \mathcal{L} and n .*

Proof. Our approach here is similar to that in the proof of Proposition 6.2.1: we show that—for sufficiently large (computable) N —an element of $\ker p$ supported on a set with more than N vertices can be written as the sum of two elements of $\ker p$ supported on strictly smaller sets.

Let D be the maximum of the diameters of the X_i . Choose L such that any subtree of T of diameter at most $2D$ contains at most L vertices.

From the proof of Proposition 6.2.1 it is clear that in picking a preimage x under p of an element $[\sigma] \in \check{H}^0(\mathcal{D})$ it might well be necessary for the support of x to be strictly larger than the support of $[\sigma]$. We will need to be able to bound the size of the support of x for $[\sigma]$ supported on a set with diameter at most $2D$. We deal with this first.

With some care, the proof of Proposition 6.2.1 gives an explicit bound. At each step, the cochain is split into two pieces, each supported on a set with strictly fewer vertices. Hence, since $[\sigma]$ is supported on a set with at most L vertices, it can certainly be written as a \mathbb{Z} -linear combination of at most 2^L elements of our generating set. So if each generator has at most M vertices, any $[\sigma]$ supported on a set with diameter at most $2D$ has a preimage supported on a set with at most $2^L M$ vertices. By construction, this set can be taken to be connected.

Let N be large enough that any subtree X of T with at least N vertices contains a vertex v such that the connected components of $X - v$ can be partitioned into two sets with unions A and B , so that each of A and B contains at least $2^L M + L$ vertices. For example, this holds if X is guaranteed to contain an embedded arc of length at least $2(2^L M + L) + 1$. Then suppose that some element $x \in \ker p$ is supported on a subtree $X \subset T$ with at least N vertices. Let v be a vertex as in the definition of N . Then we aim to divide x as the sum of two smaller relators by cutting at v .

The element x is of the form of Equation 6.5 and is such that $g_{ij}X_j \subset X$ for each i and j . Let A and B be the two subsets of $X - v$ as described above, and let C be the ball in X of radius D centred at v . Let $y \in \mathbb{Z}F^k$ be the sum of those summands of x in Equation 6.5 whose supports are contained in A . Then the support of y is a subset of A and the support of $x - y$ is a subset of $B \cup C$.

Roughly, y and $x - y$ will be the two desired smaller relators whose sum is x . However $py \neq 0$, so we shall need to add a small correction term.

Since $py = -p(x - y)$, py is supported on $A \cap (B \cup C) = A \cap C$. This is a subtree of a tree of diameter $2D$, so by assumption py has a preimage w under p that is supported on a set with at most $2^L M$ vertices. Then $p(y - w) = 0$ and $x = (y - w) + (x - y + w)$ so it remains to show that $y - w$ and $x - y + w$ have strictly smaller supports than x . But the support of x has $|A| + |B| + 1$ vertices, while $y - w$ and $x - y + w$ are supported on sets with at most $|A| + 2^L M$ and $|B| + |C| + 2^L M$ vertices respectively. The subtrees $|A|$ and $|B|$ have at least $2^L M + |C|$ vertices, so this completes the proof. \square

Proposition 6.2.3. *There is an algorithm that computes a finite presentation for $\check{H}^0(\mathcal{D})$.*

Proof. Note that F acts on $\mathbb{Z}F^k$ by translation in the sense that if the support of $x \in \mathbb{Z}F^k$ is X then the support of $g \cdot x$ is $g(X)$. Hence if N is as in the statement of Proposition 6.2.2 then that proposition shows that $\ker p$ is generated as an F -module by those of its elements that are supported on a ball of radius N ball centred at 1.

In other words, $\ker p$ is generated by its intersection with the set of those \mathbb{Z} -linear combinations of translates of the $\{e_i\}$ by F whose supports are contained in this ball of radius N . To find all such linear combinations is simply to solve a finite dimensional \mathbb{Z} -linear equation, which can be done algorithmically, for example using Smith Normal Form. \square

6.3 Computing $\check{H}^1(\mathcal{D})$

The first cohomology $\check{H}^1(\mathcal{U}_X)$ is the quotient of $\check{C}^1(\mathcal{U}_X)$ by $d\check{C}^0(\mathcal{U}_X)$, since $\check{C}^2(\mathcal{U}_X)$ is trivial. Since taking direct limits of families of \mathbb{Z} -modules is an exact functor, $\check{H}^1(\mathcal{D})$ is also a quotient: the sequence

$$0 \longrightarrow d\varinjlim_X \check{C}^0(\mathcal{U}_X, \mathbb{Z}) \longrightarrow \varinjlim_X \check{C}^1(\mathcal{U}_X, \mathbb{Z}) \longrightarrow \check{H}^1(\mathcal{D}) \longrightarrow 0$$

is exact. Here we use the fact that for $X \subset X'$ there is a *canonical* refinement map $\mathcal{U}_{X'} \rightarrow \mathcal{U}_X$. As in the previous section, each of these abelian groups can be endowed with the

structure of an F -module so that the homomorphisms in the short exact sequence are homomorphisms of F -modules.

Now finding a presentation for $\check{H}^1(\mathcal{D})$ is equivalent to finding a presentation for $\varinjlim_X \check{C}^1(\mathcal{U}_X)$ and a generating set for $d\varinjlim_X \check{C}^0(\mathcal{U}_X)$. We present an algorithm that does the former in Proposition 6.3.1 and an algorithm that does the latter in Lemma 6.3.3.

As in the previous section, cochains have a convenient representation in terms of the Whitehead graph. A 1-cochain (with respect to an open cover \mathcal{U}) is a map that associates an integer to each pair $U_1, U_2 \in \mathcal{U}$ with $U_1 \cap U_2 \neq \emptyset$. Equivalently, if \mathcal{U}_X is the open cover associated to a Whitehead graph $\text{Wh}(X)$, this is the assignment of an integer to each edge in the Whitehead graph, with the restriction that if two edges connect the same pair of vertices then they are assigned the same integer. Refinement to the open cover associated to a larger Whitehead graph preserves the labelling of the old edges, and assigns the integer zero to each new edge.

Proposition 6.3.1. *The direct limit $\varinjlim_X \check{C}^1(\mathcal{U}_X)$ is isomorphic as an F -module to $\mathbb{Z}\mathcal{L}$ with F -action induced by the action of F on the set \mathcal{L} of lines.*

Proof. We define a map $\theta : \mathbb{Z}\mathcal{L} \rightarrow \varinjlim_X \check{C}^1(\mathcal{U}_X)$. Given a line l , choose a subtree $X \subset T$ large enough that no line in \mathcal{L} except l connects the same two components of $T - X$ that l does. Then let $\theta(l)$ be the element of $\check{C}^1(\mathcal{U}_X)$ represented by labelling the line l in $\text{Wh}(X)$ with the integer one and all other lines with the integer zero. Note that the image of l under this map is independent of the chosen subtree of T satisfying the condition, so we obtain a map $\mathbb{Z}\mathcal{L} \rightarrow \check{C}^1(\mathcal{U}_X)$.

From the description of all elements of $\varinjlim_X \check{C}^1(\mathcal{U}_X)$ in terms of Whitehead graphs we see that this map is surjective. From the description of the refinement maps in terms of Whitehead graphs we see that the map is injective, and it is clearly F -equivariant. \square

Remark 6.3.2. We can give an explicit presentation for the F -module $\mathbb{Z}\mathcal{L}$. Let \mathcal{L} be generated by the words $\{w_i\}_{i=1}^n$ and let \widehat{w}_i be the unique element of F such that w_i is a power of \widehat{w}_i and \widehat{w}_i is not a proper power. Then $\mathbb{Z}\mathcal{L} \cong \bigoplus_i \mathbb{Z}(F/\langle \widehat{w}_i \rangle)$.

Proposition 6.3.1 and Remark 6.3.2 describe the F -module $\varinjlim_X \check{C}^1(\mathcal{U}_X)$ of one dimensional cocycles. The cohomology $\check{H}^1(\mathcal{D})$ is the quotient of this module by the image under the boundary map of $\varinjlim_X \check{C}^0(\mathcal{U}_X)$, so it remains to show that this image has a computable generating set.

Lemma 6.3.3. *The F -module $d\varinjlim_X \check{C}^0(\mathcal{U}_X)$ is generated by $d\check{C}^0(\mathcal{U}_{\{1\}})$. This abelian group is generated by the images under d of the cochains given by assigning the value one to one of the $2|S|$ sets in $\mathcal{U}_{\{1\}}$ and zero to the others.*

Proof. The cochain module $\check{C}^0(\mathcal{U}_X)$ is generated as an abelian group by those elements that are contained in $\check{C}^0(\mathcal{U}_{\{1\}})$ for some vertex v of T , and is therefore generated as an F -module by those of its elements that are contained in $\check{C}^0(\mathcal{U}_{\{1\}})$. It follows that $d\varinjlim_X \check{C}^0(\mathcal{U}_X)$ is generated as an F -module by $d\check{C}^0(\mathcal{U}_{\{1\}})$. The claim concerning the generators of $d\check{C}^0(\mathcal{U}_{\{1\}})$ is clear. \square

Putting the results of this section along with Proposition 6.2.3 we obtain the following Theorem.

Theorem 6.3.4. *There is an algorithm that, when given a free group F and a finite collection of elements of that group, computes presentations for the F -modules $\check{H}^k(\mathcal{D})$, where \mathcal{D} is the associated decomposition space.*

6.4 The cohomology of hyperbolic graphs of free groups

In this section we apply Theorem 6.3.4 to prove Theorem 1.3.1: we prove the computability of the Čech cohomology of Gromov boundaries of those hyperbolic groups that are fundamental groups of graphs of free groups with cyclic edge groups. By Bestvina and Feighn's combination theorem [BF92], these are precisely those fundamental groups of finite graphs of free groups with cyclic edge groups that do not contain a Baumslag–Solitar subgroup. By a theorem of Bestvina and Mess [BM91] this proves the computability of the group cohomology of such a group with coefficients in the group ring. We do this geometrically: we build the boundary of such a group out of the decomposition spaces of the free vertex groups with line patterns generated by the incident edge groups.

In dimension zero an algebraic approach describes the cohomology in greater generality. In dimensions at least two one can easily prove topologically that the cohomology must vanish. The most complicated part of the proof of Theorem 1.3.1 is the computability of the cohomology in dimension one. It has very recently been shown by Manning and Wang [MW18] that the cohomology of the Bowditch boundary of a relatively hyperbolic group is equal to the group cohomology of the group relative to its peripheral subgroups. In light of this, one could alternatively try to relate the cohomology of a hyperbolic fundamental group of a graph of free groups to the cohomology of the decomposition spaces of its vertex groups algebraically using a group cohomological Mayer-Vietoris argument.

6.4.1 Boundaries of fundamental groups of graphs of groups

Let G be a hyperbolic group and let T be a minimal G -tree such that all edge stabilisers are quasi-convex. Then by [Bow98b, Proposition 1.2], for each vertex v of T , G_v is quasi-

convex in G and is therefore finitely generated and hyperbolic. The inclusion $G_\nu \hookrightarrow G$ extends to a map $\partial G_\nu \rightarrow \partial G$ with image ΛG_ν and this map is a homeomorphism onto its image. Then by [Bow98b, Proposition 1.3] ∂G has the following description.

$$\partial G = \bigcup_{\nu \in V(T)} \Lambda G_\nu \sqcup \partial T \quad (6.7)$$

Furthermore, whenever an edge e separates vertices ν_1 and ν_2 in T , $\Lambda G_{\nu_1} \cap \Lambda G_{\nu_2} \subset \Lambda G_e$ with equality when $\nu_1 = \iota(e)$ and $\nu_2 = \tau(e)$. We now aim to give a topological description of ∂G in terms of the subspaces $\{\Lambda G_\nu \mid \nu \in V(T)\}$.

Let X be a finite subtree of T . Now define the subset ∂_X of ∂G as follows.

$$\partial_X = \bigcup_{\nu \in V(X)} \Lambda G_\nu / \sim. \quad (6.8)$$

Here \sim is the equivalence relation generated by letting $x \sim y$ whenever $\{x, y\} \subset \Lambda G_e$ for some edge $e \in E(T)$ with $\iota(e) \in X$ and $\tau(e) \notin X$.

We specialise to two classes of edge stabilisers.

Infinite cyclic edge stabilisers

First consider the case in which all edge stabilisers are infinite cyclic and G is one-ended. Then for any edge $e \in T$, ΛG_e is a cut pair. Therefore there is a natural map $\partial G \rightarrow \partial_X$ obtained by, for each edge e of T with $\iota(e) \in X$ and $\tau(e) \notin X$, collapsing the union of ΛG_e with a component of $\partial G - \Lambda G_e$ to a point: for each such edge e , let T_e be the component of $T - \text{int } e$ containing $\tau(e)$; then we collapse to a point following set.

$$\bigcup_{\nu \in V(T_e)} \Lambda G_\nu \sqcup \partial T_e \quad (6.9)$$

Call this map q_X .

Note that when $X_1 \subset X_2$ are finite subtrees of T , q_{X_1} factors through q_{X_2} : denote by p_{X_1, X_2} the map $\partial_{X_2} \rightarrow \partial_{X_1}$ obtained by collapsing $q_{X_2}(\Lambda G_\nu)$ to a point for each vertex $\nu \in V(X_2) - V(X_1)$. Then $q_{X_1} = p_{X_1, X_2} \circ q_{X_2}$.

From this discussion we see that ∂G admits a natural map to the inverse limit $\varprojlim_X \partial_X$ of the system $(\{\partial_X\}_X, \{p_{X_1, X_2}\})$ partially ordered by inclusion of subtrees. From the decomposition of ∂G described by Equation 6.7 we see that this map is a bijection. Each space is compact Hausdorff, so the map is a homeomorphism.

Trivial edge stabilisers

Secondly, we perform a similar construction in the case in which all edge stabilisers are trivial. Here we must assume that G_v is infinite for each vertex v . In this case, for any finite subtree X of T , ∂_X is simply a union

$$\bigcup_{v \in V(X)} \Lambda G_v. \quad (6.10)$$

This subset of ∂G has the topology of the disjoint union of the sets $\{\Lambda G_v \mid v \in V(X)\}$.

We now define retracts $q_X : X \rightarrow \partial_X$. For each vertex $v \in T$ choose a point $z_v \in \Lambda G_v$. Now for each edge $e \in E(X)$ such that $\iota(e) \in X$ and $\tau(e) \notin X$ let T_e be the component of $T - \text{int } e$ containing $\tau(e)$ and map the following subset of ∂_X to the point $z_{\iota(e)} \in \partial_X$.

$$\bigcup_{v \in V(T_e)} \Lambda G_v \sqcup \partial T_e \quad (6.11)$$

Then, as in the case of infinite cyclic edge stabilisers, we realise ∂G as an inverse limit $\varprojlim_X \partial_X$ of the system $(\{\partial_X\}, \{p_{X_1, X_2}\})$.

In either case, using the fact that Čech cohomology is a continuous contravariant functor we obtain a description of the Čech cohomology group $\check{H}^k(\partial G)$.

Proposition 6.4.1. *If G is a hyperbolic group and T is a minimal G -tree such that either all edge stabilisers are trivial and all vertex stabilisers are infinite, or all edge stabilisers are infinite cyclic, then we have the following isomorphism.*

$$\check{H}^k(\partial G) \cong \varprojlim_X \check{H}^k(\partial_X) \quad (6.12)$$

6.4.2 Computing the cohomology

We now use Proposition 6.4.1 to describe the cohomology of ∂G in terms of the vertex groups.

Groups with more than one end

First, we apply Proposition 6.4.1 in the case when the edge stabilisers are trivial and the vertex stabilisers are infinite. Assume also that G is torsion free. In this case we have observed that $\partial_X \subset \partial G$ is a disjoint union $\bigsqcup_{v \in V(X)} \Lambda G_v$. We split into cases: first we treat the cohomology in dimension zero, then in dimensions at least one. Let T be a Dunwoody decomposition for G . Since G is torsion free this coincides with the Grushko decomposition.

In dimension zero, $\check{H}^0(\partial_X) \cong \bigoplus_{v \in V(X)} \check{H}^0(\Lambda G_v)$, but the inclusion $\check{H}^0(\Lambda G_v) \rightarrow \check{H}^0(\partial_X)$ does not coincide with the map $p_{\{v\}, X}^*$, so we cannot conclude that $\check{H}^0(\partial G)$ is a direct sum indexed by $V(T)$. Instead we describe the cohomology in terms of the edges of T . Let $E^+(T)$ be a G -invariant orientation on T : that is, a choice of one element from each pair $\{e, \bar{e}\}$ such that if $e \in E^+(T)$ then $g \cdot e \in E^+(T)$ for all g in G . We define a map from $\mathbb{Z}E^+(T)$ to $\varinjlim_X \check{H}^k(\partial_X)$. Choose a base point $v \in T$ and let $y = \sum_i n_i e_i \in \mathbb{Z}E^+(T)$. Then for any subtree X of T containing e_i for each i , define the image of y in $\check{H}^0(\partial_X)$ to be given by the following assignment of integers to connected components of ∂_X . For $w \in X$, assign to ΛG_w the sum of the coefficients in y of the edges traversed in the path from v to w in T , with edges traversed in the opposite direction counting negatively. The composition of this map with the map q_X^* is independent of the choice of X , so gives a homomorphism of abelian groups $\mathbb{Z}E^+(T) \rightarrow \check{H}^0(\partial G)$.

The map constructed in the previous paragraph depends on choice of base point v , and therefore cannot provide an isomorphism of G -modules. However, the induced map $\mathbb{Z}E^+(T) \rightarrow \check{H}_1^0(\partial G)$ to the *reduced* cohomology is an isomorphism of abelian groups and does not depend on the base point. Injectivity of this map follows from the fact that for any edge e of T , G_e is an infinite index subgroup of $G_{i(e)}$. It is clear that this map is G -equivariant with respect to the G -action on $\mathbb{Z}E^+(T)$. All edge stabilisers are trivial, so $\mathbb{Z}E^+(T) \cong (\mathbb{Z}G)^{|E(G \setminus T)|}$. Finally, note that the map $\check{H}^0(\partial G) \rightarrow \check{H}_1^0(\partial G)$ splits, so $\check{H}^0(\partial G) \cong \check{H}_1^0(\partial G) \oplus \mathbb{Z}$, where G acts trivially on \mathbb{Z} .

This gives the following proposition.

Proposition 6.4.2. *Let G be a one-ended hyperbolic group. Let T be a tree on which G acts minimally such that all edge stabilisers are trivial and all vertex stabilisers are infinite. Then $\check{H}^0(\partial G) \cong (\mathbb{Z}G)^{|E(G \setminus T)|} \oplus \mathbb{Z}$.*

Remark 6.4.3. Algebraic analogues of this statement are true more generally. See, for example, [Geo08, Section 13.5].

In dimension $k \geq 1$ the situation is simpler: if e is an edge of X and A and B are the two components of $X - \text{int } e$ then $\check{H}^k(\partial_X) \cong \check{H}^k(\partial_A) \oplus \check{H}^k(\partial_B)$ where the inclusions $\check{H}^k(\partial_A) \rightarrow \check{H}^k(\partial_X)$ and $\check{H}^k(\partial_B) \rightarrow \check{H}^k(\partial_X)$ coincide with the maps $p_{A, X}^*$ and $p_{B, X}^*$. It follows that $\check{H}^k(\partial G) \cong \bigoplus_{v \in V(T)} \check{H}^k(\Lambda G_v)$ as abelian groups for $k \geq 1$.

The G -action is clear: if $g \in G_v$ then the action of g fixes $\check{H}^k(\Lambda G_v) \subset \check{H}^k(\partial G)$ and the action on this subspace coincides with the action induced by the action of g^{-1} on ΛG_v . If $g \notin G_v$ then the action of g^{-1} on ∂G maps $g(\Lambda G_v) = \Lambda G_{g^{-1}v}$ homeomorphically onto ΛG_v , so the action of g on $\check{H}^k(\partial G)$ maps $\check{H}^k(\Lambda G_v)$ isomorphically onto $\check{H}^k(\Lambda G_{g^{-1}v})$.

Therefore $\check{H}^k(\partial G)$ is isomorphic as a G -module for $k \geq 1$ to the following direct sum.

$$\bigoplus_{[v] \in V(G \setminus X)} \left(\mathbb{Z}G \otimes_{\mathbb{Z}G_v} \check{H}^k(\partial G_v) \right) \quad (6.13)$$

This argument reduces the problem to the case in which G is one-ended: let \mathcal{Y} be a Dunwoody decomposition of G , which may be computed using [DG08], and suppose that every vertex group in \mathcal{Y} is infinite. Then the cohomology of G in dimension zero is given by Proposition 6.4.2, while the cohomology in higher dimensions is given in terms of the one-ended vertex groups by Equation 6.13.

Groups with one end

Now suppose that G is one-ended, that stabilisers of vertices in T are free and that the stabilisers of edges in T are cyclic. We now describe $\check{H}^k(\partial_X)$ in terms of the limit sets of stabilisers of vertices in X . Let e be an edge of X and let A and B be the two components of $X - \text{int } e$. Let $\widehat{\partial}_A$ be $q_X(\bigcup_{v \in V(A)} \Lambda G_v) \subset \partial_X$, so that ∂_A is obtained from $\widehat{\partial}_A$ by identifying the two points in ΛG_e . Similarly define $\widehat{\partial}_B$, so $\partial_X = \widehat{\partial}_A \cup \widehat{\partial}_B$.

Note that the Mayer-Vietoris exact sequence may be applied to the union $\partial_X = \widehat{\partial}_A \cup \widehat{\partial}_B$. This is non-trivial: $\widehat{\partial}_A$ and $\widehat{\partial}_B$ are not open in ∂_X and $\widehat{\partial}_A \cap \widehat{\partial}_B$ is not a strong neighbourhood deformation retract in either $\widehat{\partial}_A$ or $\widehat{\partial}_B$. The exactness of the Mayer-Vietoris sequence for this union is a special feature of Čech cohomology: Čech cohomology satisfies a very strong excision property with the consequence that any compact triad is a proper triad with respect to Čech cohomology, and therefore there is a Mayer-Vietoris exact sequence for any compact triad (Y, Y_1, Y_2) with $Y = Y_1 \cup Y_2$. For more details see [ES52, Chapter 10].

For dimension $k > 1$ the Mayer-Vietoris theorem tells us that we have the following isomorphism.

$$\check{H}^k(\partial_X) \cong \check{H}^k(\widehat{\partial}_A, \Lambda G_e) \oplus \check{H}^k(\widehat{\partial}_B, \Lambda G_e) \cong \check{H}^k(\partial_A) \oplus \check{H}^k(\partial_B). \quad (6.14)$$

Applying this repeatedly we reduce to a direct sum of groups $\check{H}^k(\partial_{\{v\}})$, where v is a vertex of X . But $\partial_{\{v\}}$ is simply a decomposition space, so the cohomology in dimensions $k > 1$ vanishes. Therefore $\check{H}^k(\partial_X) = 0$ for all finite subgraphs X , so $\check{H}^k(\partial G) = 0$.

In dimension one, the map $\partial_X \rightarrow \partial_A$ maps $\widehat{\partial}_B$ to a point, so its restriction to $\widehat{\partial}_B$ is trivial in first cohomology. Therefore the following diagram commutes.

$$\begin{array}{ccc}
\check{H}^1(\partial_A) & \longrightarrow & \check{H}^1(\widehat{\partial}_A) \\
p_{A,X}^* \downarrow & & \downarrow \\
\check{H}^1(\partial_X) & \longrightarrow & \check{H}^1(\widehat{\partial}_A) \oplus \check{H}^1(\widehat{\partial}_B)
\end{array}$$

Here the map $\check{H}^1(\widehat{\partial}_A) \rightarrow \check{H}^1(\widehat{\partial}_A) \oplus \check{H}^1(\widehat{\partial}_B)$ is the obvious inclusion map.

Note that $\check{H}^1(\partial_A) \cong \check{H}^1(\widehat{\partial}_A, \Lambda G_e)$, so the Mayer-Vietoris exact sequence applied to $\partial_X = \widehat{\partial}_A \cup \widehat{\partial}_B$ and the long exact sequences associated to the pairs $(\widehat{\partial}_A, \Lambda G_e)$ and $(\widehat{\partial}_B, \Lambda G_e)$ give the following commutative diagram, in which rows are exact sequences.

$$\begin{array}{ccccccc}
\check{H}_r^0(\widehat{\partial}_A) & & & \check{H}^1(\partial_A) & \longrightarrow & \check{H}^1(\widehat{\partial}_A) & \\
\downarrow & \searrow & & \downarrow p_{A,X}^* & & \downarrow & \\
\check{H}_r^0(\widehat{\partial}_A) \oplus \check{H}_r^0(\widehat{\partial}_B) & \longrightarrow & \check{H}_r^0(\Lambda G_e) & \longrightarrow & \check{H}^1(\partial_X) & \longrightarrow & \check{H}^1(\widehat{\partial}_A) \oplus \check{H}^1(\widehat{\partial}_B) \longrightarrow 0 \\
\uparrow & \nearrow & & \uparrow p_{B,X}^* & & \uparrow & \\
\check{H}_r^0(\widehat{\partial}_B) & & & \check{H}^1(\partial_B) & \longrightarrow & \check{H}^1(\widehat{\partial}_B) &
\end{array}$$

Let M be the pushout of the maps $\check{H}_r^0(\Lambda G_e) \rightarrow \check{H}^1(\partial_A)$ and $\check{H}_r^0(\Lambda G_e) \rightarrow \check{H}^1(\partial_B)$. The kernel of the map $\check{H}_r^0(\Lambda G_e) \rightarrow M$ is the image of the map $\check{H}_r^0(\widehat{\partial}_A) \oplus \check{H}_r^0(\widehat{\partial}_B) \rightarrow \check{H}_r^0(\Lambda G_e)$. We therefore see that M fits into the following commutative diagram, in which rows are exact sequences.

$$\begin{array}{ccccccc}
& & & M & & & \\
& & \nearrow & \downarrow & \searrow & & \\
\check{H}_r^0(\widehat{\partial}_A) \oplus \check{H}_r^0(\widehat{\partial}_B) & \longrightarrow & \check{H}_r^0(\Lambda G_e) & & \check{H}^1(\widehat{\partial}_A) \oplus \check{H}^1(\widehat{\partial}_B) & \longrightarrow & 0 \\
& & \searrow & \check{H}^1(\partial_X) & \nearrow & &
\end{array}$$

By the five lemma it follows that $\check{H}^1(\partial_X) \cong M$ and the isomorphism from M to $\check{H}^1(\partial_X)$ is induced by the maps $p_{A,X}^*$ and $p_{B,X}^*$.

Since the limit set of an edge stabiliser consists of precisely two points, $\check{H}_r^0(\Lambda G_e) \cong \mathbb{Z}$ for any edge $e \in E(T)$. For each edge $e \in E(T)$ choose a generator for $\check{H}_r^0(\Lambda G_e)$ equivariantly with respect to the G -action on T : if x_e is the image in ∂G of the chosen generator of $\check{H}_r^0(\Lambda G_e)$ then we require that $g \cdot x_e = x_{ge}$ for any $g \in G$. (Here we use the fact that G is torsion free, so G_e acts trivially on ΛG_e . If infinite dihedral edge stabilisers were allowed then we would have to be more careful.)

Therefore we see that, as an abelian group, $\check{H}^1(\partial G)$ is the quotient of $\bigoplus_{v \in V(T)} \check{H}^1(\partial_{\{v\}})$ obtained by identifying the image of the chosen generator of $\check{H}_\Gamma^0(\Lambda G_e)$ in $\check{H}^1(\partial_{\{l(e)\}})$ with its image in $\check{H}^1(\partial_{\{\tau(e)\}})$ for each edge $e \in E(T)$. Denote by $x_{l(e)}$ the image of the generator in $\check{H}^1(\partial_{\{l(e)\}})$ and by $x_{\tau(e)}$ the image of the generator in $\check{H}^1(\partial_{\{\tau(e)\}})$

Then we have the following isomorphism of abelian groups.

$$\check{H}^1(\partial G) \cong \bigoplus_{v \in V(T)} \check{H}^1(\partial_{\{v\}}) / \langle x_{l(e)} - x_{\tau(e)} \mid e \in E(T) \rangle \quad (6.15)$$

We now describe the G -module structure of $\check{H}^1(\partial G)$ in terms of the graph of groups \mathcal{Y} associated to T . Let Y be the underlying graph of \mathcal{Y} . Vertex group in \mathcal{Y} are only identified with stabilisers of vertices in T up to conjugacy in G . To remedy this, let S be a maximal subtree of Y and choose a lift \tilde{S} of S in T . For each vertex and edge of S choose the group labelling that vertex or edge in \mathcal{Y} to be the stabiliser of the lift of that vertex or edge to \tilde{S} . Denote this lift of a vertex v by \tilde{v} and the lift of an edge e by \tilde{e} . For each edge e in $E(Y) - E(S)$ choose a lift \tilde{e} of e to T such that $l(\tilde{e}) \in \tilde{S}$ and label e with the stabiliser \tilde{e} . Also choose an element $g_e \in G$ such that $\tau(g_e(\tilde{e})) \in \tilde{S}$. Note then that $x_{\tau(\tilde{e})} = g_e^{-1} \cdot x_{\tau(g_e(\tilde{e}))}$.

We obtain the following isomorphism of G -modules.

$$\check{H}^1(\partial G) \cong \left(\bigoplus_{v \in V(Y)} \mathbb{Z}G \otimes_{\mathbb{Z}G_{\tilde{v}}} \check{H}^1(\partial_{\{\tilde{v}\}}) \right) / R \quad (6.16)$$

where R is the G -submodule generated by the following finite set.

$$\{1 \otimes (x_{l(\tilde{e})} - x_{\tau(\tilde{e})}) \mid e \in E(S)\} \cup \{1 \otimes x_{l(\tilde{e})} - g_e^{-1} \otimes x_{\tau(g_e(\tilde{e}))} \mid e \in E(Y) - E(S)\} \quad (6.17)$$

In the case of interest to us, in which all vertex stabilisers in T are free, $\partial_{\{v\}}$ is precisely the decomposition space associated to the pair $(G_v, \text{Inc } v)$, and therefore $\check{H}^1(\partial_{\{v\}})$ is computable by Theorem 6.3.4. Furthermore, $x_{l(e)}$ is simply the generator corresponding to the line spanned by $G_e \leq G_{l(e)}$, which can readily be identified in $\check{H}^1(\partial_{\{v\}})$.

Finally, note that given a presentation for a one-ended hyperbolic group that is the fundamental group of a graph of free groups with cyclic edge groups, a decomposition of G with free vertex groups and cyclic edge groups can be computed from that presentation. First compute Bowditch's canonical decomposition for the group using Theorem 1.2.2. If this decomposition is trivial consisting of a single vertex in T_2 then the boundary of the group is a circle and the Čech cohomology of the boundary is \mathbb{Z} in dimensions zero and one, with trivial G -action. Otherwise, the stabiliser of every vertex in T_2 is free.

Bowditch's canonical decomposition can be refined at vertices in T_2 to dominate the decomposition with free vertex groups, so every vertex in $T_1 \cup T_3$ is elliptic with respect to this decomposition and therefore every vertex group is free. It follows that every vertex in Bowditch's canonical decomposition has free stabiliser.

Putting the results of this section together, we obtain the following theorem.

Theorem 6.4.4. *There is an algorithm that takes as input a presentation for a hyperbolic fundamental group G of a graph of groups with free vertex groups and cyclic edge groups and returns presentations for the Čech cohomology G -modules of ∂G .*

Corollary 6.4.5. *The group cohomology of a hyperbolic fundamental group of a graph of groups with free vertex groups and cyclic edge groups with coefficients in the group ring $\mathbb{Z}G$ can be computed from a presentation of that group.*

Bibliography

- [Bar17] Benjamin Barrett. Computing the cech cohomology of decomposition spaces. arXiv:1712.00890, December 2017.
- [Bar18] Benjamin Barrett. Computing JSJ decompositions of hyperbolic groups. *Journal of Topology*, 11(2):527–558, 2018.
- [Bes04] M. Bestvina. Questions in geometric group theory, July 2004. Unpublished.
- [BF91] Mladen Bestvina and Mark Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.
- [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.
- [BH05] Martin R. Bridson and James Howie. Conjugacy of finite subsets in hyperbolic groups. *Internat. J. Algebra Comput.*, 15(4):725–756, 2005.
- [BH13] D. J. Buckley and Derek F. Holt. The conjugacy problem in hyperbolic groups for finite lists of group elements. *Internat. J. Algebra Comput.*, 23(5):1127–1150, 2013.
- [BM74] Alan F. Beardon and Bernard Maskit. Limit points of Kleinian groups and finite sided fundamental polyhedra. *Acta Math.*, 132:1–12, 1974.
- [BM91] Mladen Bestvina and Geoffrey Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.

- [BMP03] Michel Boileau, Sylvain Maillot, and Joan Porti. *Three-dimensional orbifolds and their geometric structures*, volume 15 of *Panoramas et Synthèses [Panoramas and Syntheses]*. Société Mathématique de France, Paris, 2003.
- [Bow98a] B. H. Bowditch. Boundaries of strongly accessible hyperbolic groups. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 51–97. Geom. Topol. Publ., Coventry, 1998.
- [Bow98b] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [Bow98c] Brian H. Bowditch. A topological characterisation of hyperbolic groups. *J. Amer. Math. Soc.*, 11(3):643–667, 1998.
- [Bow99a] B. H. Bowditch. Boundaries of geometrically finite groups. *Math. Z.*, 230(3):509–527, 1999.
- [Bow99b] B. H. Bowditch. Connectedness properties of limit sets. *Trans. Amer. Math. Soc.*, 351(9):3673–3686, 1999.
- [Bow99c] B. H. Bowditch. Convergence groups and configuration spaces. In *Geometric group theory down under (Canberra, 1996)*, pages 23–54. de Gruyter, Berlin, 1999.
- [Bow01] B. H. Bowditch. Peripheral splittings of groups. *Trans. Amer. Math. Soc.*, 353(10):4057–4082, 2001.
- [Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [CDP90] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes*, volume 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990. Les groupes hyperboliques de Gromov. [Gromov hyperbolic groups], With an English summary.
- [Cha95] Christophe Champetier. Propriétés statistiques des groupes de présentation finie. *Adv. Math.*, 116(2):197–262, 1995.
- [CJ94] Andrew Casson and Douglas Jungreis. Convergence groups and Seifert fibered 3-manifolds. *Invent. Math.*, 118(3):441–456, 1994.

- [CM11] Christopher H. Cashen and Nataša Macura. Line patterns in free groups. *Geom. Topol.*, 15(3):1419–1475, 2011.
- [CV86] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.
- [Dah08] François Dahmani. Finding relative hyperbolic structures. *Bull. Lond. Math. Soc.*, 40(3):395–404, 2008.
- [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.
- [DS99] M. J. Dunwoody and M. E. Sageev. JSJ-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.
- [DT13] F. Dahmani and N. Touikan. Isomorphisms using Dehn fillings: the splitting case. *ArXiv e-prints*, November 2013.
- [Dun85] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.
- [Dun93] Martin J. Dunwoody. An inaccessible group. In *Geometric group theory, Vol. 1 (Sussex, 1991)*, volume 181 of *London Math. Soc. Lecture Note Ser.*, pages 75–78. Cambridge Univ. Press, Cambridge, 1993.
- [Eng78] Ryszard Engelking. *Dimension theory*. North-Holland Publishing Co., Amsterdam-Oxford-New York; PWN—Polish Scientific Publishers, Warsaw, 1978. Translated from the Polish and revised by the author, North-Holland Mathematical Library, 19.
- [Eps92] D. B. A. Epstein. *Word Processing in Groups*. A K Peters, November 1992.
- [ES52] Samuel Eilenberg and Norman Steenrod. *Foundations of algebraic topology*. Princeton University Press, Princeton, New Jersey, 1952.
- [Far98] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.
- [FP06] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.

- [Gab92] David Gabai. Convergence groups are Fuchsian groups. *Ann. of Math. (2)*, 136(3):447–510, 1992.
- [GdlH90] Étienne Ghys and Pierre de la Harpe. Espaces métriques hyperboliques. In *Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)*, volume 83 of *Progr. Math.*, pages 27–45. Birkhäuser Boston, Boston, MA, 1990.
- [Geo08] Ross Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*. Springer, New York, 2008.
- [Ger] V. Gerasimov. Detecting connectedness of the boundary of a hyperbolic group. Unpublished.
- [GL17] Vincent Guirardel and Gilbert Levitt. JSJ decompositions of groups. *Astérisque*, (395):vii+165, 2017.
- [GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [Gro92] M. Gromov. *Asymptotic Invariants of Infinite Groups*. Institut des hautes études scientifiques, 1992.
- [Gro03] M. Gromov. Random walk in random groups. *Geometric & Functional Analysis GAFA*, 13(1):73–146, Feb 2003.
- [Gro13] Bradley W. Groff. Quasi-isometries, boundaries and JSJ-decompositions of relatively hyperbolic groups. *J. Topol. Anal.*, 5(4):451–475, 2013.
- [Gru40] I. A. Grushko. On the bases of a free product of groups. *Mat. Sb.*, 8:169–182, 1940.
- [Hau17] Matthew Haulmark. Local cut points and splittings of relatively hyperbolic groups. arXiv:1708.02855, August 2017.
- [Joh79] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.
- [JS79] William Jaco and Peter B. Shalen. Seifert fibered spaces in 3-manifolds. In *Geometric topology (Proc. Georgia Topology Conf., Athens, Ga., 1977)*, pages 91–99. Academic Press, New York-London, 1979.

- [Kap96] Ilya Kapovich. Detecting quasiconvexity: algorithmic aspects. In *Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994)*, volume 25 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 91–99. Amer. Math. Soc., Providence, RI, 1996.
- [KB02] Ilya Kapovich and Nadia Benakli. Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 39–93. Amer. Math. Soc., Providence, RI, 2002.
- [Lev05] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geom. Dedicata*, 114:49–70, 2005.
- [LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin-New York, 1977. *Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89*.
- [Mak82] G. S. Makanin. Equations in a free group. *Izv. Akad. Nauk SSSR Ser. Mat.*, 46(6):1199–1273, 1344, 1982.
- [Man10] Jason Fox Manning. Virtually geometric words and Whitehead’s algorithm. *Math. Res. Lett.*, 17(5):917–925, 2010.
- [MNS99] C. F. Miller, III, Walter D. Neumann, and G. A. Swarup. Some examples of hyperbolic groups. In *Geometric group theory down under (Canberra, 1996)*, pages 195–202. de Gruyter, Berlin, 1999.
- [MW18] Jason Fox Manning and Oliver Wang. Cohomology and the bowditch boundary. arXiv:1806.07074, June 2018.
- [Ota92] Jean-Pierre Otal. Certaines relations d’équivalence sur l’ensemble des bouts d’un groupe libre. *J. London Math. Soc. (2)*, 46(1):123–139, 1992.
- [Pan89] Pierre Pansu. Dimension conforme et sphère à l’infini des variétés à courbure négative. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 14(2):177–212, 1989.
- [Pap95] P. Papasoglu. Strongly geodesically automatic groups are hyperbolic. *Invent. Math.*, 121(2):323–334, 1995.
- [Pau91] Frédéric Paulin. Outer automorphisms of hyperbolic groups and small actions on \mathbf{R} -trees. In *Arboreal group theory (Berkeley, CA, 1988)*, volume 19 of *Math. Sci. Res. Inst. Publ.*, pages 331–343. Springer, New York, 1991.

- [RS95] E. Rips and Z. Sela. Canonical representatives and equations in hyperbolic groups. *Invent. Math.*, 120(3):489–512, 1995.
- [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.
- [Sco83] Peter Scott. The geometries of 3-manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983.
- [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.
- [Ser77] Jean-Pierre Serre. *Arbres, amalgames, SL_2* . Société Mathématique de France, Paris, 1977. Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46.
- [Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.
- [Sta71] John Stallings. *Group theory and three-dimensional manifolds*. Yale University Press, New Haven, Conn.-London, 1971. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969, Yale Mathematical Monographs, 4.
- [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.
- [Swa96] G. A. Swarup. On the cut point conjecture. *Electron. Res. Announc. Amer. Math. Soc.*, 2(2):98–100 (electronic), 1996.
- [Tou18] Nicholas W. M. Touikan. Detecting geometric splittings in finitely presented groups. *Trans. Amer. Math. Soc.*, 370(8):5635–5704, 2018.
- [Tra13] Hung Cong Tran. Relations between various boundaries of relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 23(7):1551–1572, 2013.
- [Tuk88] Pekka Tukia. Homeomorphic conjugates of Fuchsian groups. *J. Reine Angew. Math.*, 391:1–54, 1988.

-
- [Wil49] Raymond Louis Wilder. *Topology of Manifolds*. American Mathematical Society Colloquium Publications, vol. 32. American Mathematical Society, New York, N. Y., 1949.
- [Yam04] Asli Yaman. A topological characterisation of relatively hyperbolic groups. *J. Reine Angew. Math.*, 566:41–89, 2004.