

Computing JSJ decompositions of hyperbolic groups

Benjamin Barrett

March 28, 2018

Abstract

We present an algorithm that computes Bowditch's canonical JSJ decomposition of a given one-ended hyperbolic group over its virtually cyclic subgroups. The algorithm works by identifying topological features in the boundary of the group. As a corollary we also show how to compute the JSJ decomposition of such a group over its virtually cyclic subgroups with infinite centre. We also give a new algorithm that determines whether or not a given one-ended hyperbolic group is virtually fuchsian. Our approach uses only the geometry of large balls in the Cayley graph and avoids Makanin's algorithm.

0 Introduction

When studying a group it is natural and often useful to try to cut it into simpler pieces by means of amalgamated free products and HNN extensions over particularly simple subgroups. Sometimes this can be done in a canonical way analogous to the characteristic submanifold decomposition of Jaco, Shalen and Johannson [24, 25], in which the family of embedded tori along which the 3-manifold is cut is unique up to isotopy. Such JSJ decompositions were introduced to group theory by Sela [32] to answer questions about rigidity and the isomorphism problem for torsion-free hyperbolic groups. In [8] Bowditch developed a related type of decomposition for hyperbolic groups possibly with torsion. This decomposition is built from the structure of local cut points in the boundary of the group and is therefore an automorphism invariant of the group; this property of the Bowditch JSJ was used in Levitt's work [27] on outer automorphism groups of one-ended hyperbolic groups. For more general constructions of JSJ decompositions of groups see [30, 16, 18, 23].

The above results describe and prove the existence of various types of JSJ decompositions but do not give an algorithm to construct them. Gerasimov [20] proved that there exists an algorithm that determines whether or not the Gromov boundary of a given hyperbolic group is connected. This algorithm is unpublished; see also [14]. The connectedness of the boundary is determined by the so-called double-dagger condition of Bestvina and Mess [2]; it is this condition that Gerasimov showed to be computable. 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 the computability of Bowditch's JSJ decomposition over virtually cyclic subgroups.

In this paper we present an algorithm that computes Bowditch's decomposition. 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 [15] 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 [29] 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 [28].

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 [12]. Their result is based on Touikan's algorithm [33], 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.

The existence of a splitting of a one-ended hyperbolic group over a virtually cyclic subgroup is reflected in the existence of certain topological features in its Gromov boundary by results of Bowditch [8, 4, 3]; in this paper we show that these topological features can be detected algorithmically.

Theorem 0.1. *There is an algorithm that takes as input a presentation for a one-ended hyperbolic group and returns the graph of groups associated to the three following JSJ decompositions:*

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

The main content of Theorem 0.1 is the computability of the \mathcal{VC} -JSJ; it is shown in [15] 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 [15].

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 Γ be a rigid hyperbolic group (such as the fundamental group of a closed hyperbolic 3-manifold) and let g be an element of Γ that is not a proper power. Let $k > 1$ and consider the group $\Gamma' = \Gamma *_{g=t^k} \langle t \rangle$ obtained by adjoining a k th root of g to Γ . In this case the \mathcal{Z}_{\max} decomposition computed by Dahmani and Guirardel is trivial while the JSJ decomposition over virtually cyclic subgroups with infinite centre is not.

Central to the algorithm of Theorem 0.1 is an algorithm that determines whether or not a given hyperbolic group with a (possibly empty) finite collection of virtually cyclic subgroups admits a proper splitting as an amalgamated product or HNN extension over a virtually cyclic subgroup, relative to that collection of subgroups. Recall that a *cut pair* in a connected topological space S is a pair of points p and q such that $S - \{p, q\}$ is disconnected. It is shown in [8] 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 [3]. 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 show to be computable.

To detect the presence of a cut pair in the Bowditch boundary of a hyperbolic group relative to a given family of subgroups we first show that the connectivity of the complement of a pair of points in the boundary is equivalent to the connectivity of a thickened cylinder around a geodesic connecting that pair of points in the cusped space defined in [22]. Then, supposing that there is a cut pair in the boundary, we use a pumping lemma argument to 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 [10] to control cut pairs in the decomposition space of a line pattern in a free group. This is sufficient to detect a cut pair: the existence of such a periodic geodesic can be detected in finite time by searching a large finite ball in the Cayley graph.

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. Therefore we prove the following theorem, which is interesting in its own right, and does not seem to appear in the literature. Recall that the Convergence Group Theorem of Tukia, Casson, Jungreis and Gabai [34, 11, 19] implies that the Gromov boundary of a hyperbolic group is homeomorphic to a circle if and only if the group surjects with finite kernel onto the fundamental group of a compact hyperbolic orbifold.

Theorem 0.2. *There is an algorithm that takes as input a hyperbolic group Γ and a (possibly empty) collection of virtually cyclic subgroups \mathcal{H} and returns an answer to the question “is $\partial(\Gamma, \mathcal{H})$ homeomorphic to S^1 ?”*

The algorithm of Theorem 0.2 is similar to the algorithm that detects the presence of a cut pair in the Bowditch boundary: it follows from a result in point-set topology that the boundary is homeomorphic to a circle if and only if every pair of points in the boundary is a cut pair. We show that if there is a non-cut pair in the Bowditch boundary then there is a non-cut pair connected in the cusped space by a local geodesic with bounded period.

In section 1 we first review the definition of the cusped space and the Bowditch boundary. We then recall some important properties: the computability of the hyperbolicity constant of the cusped space, the existence of a visual metric on the Bowditch boundary, and most importantly the so-called double-dagger condition, which will be vital in linking the connectivity of the boundary to that of subsets of the cusped space. We then define a thickened cylinder around a geodesic in the cusped space and show that its connectivity determines the connectivity of the complement of the limit set of that geodesic in the boundary.

Section 2 contains the main technical results of the paper we describe the algorithms that determine whether or not the Bowditch boundary of a hyperbolic group contains the three topological features of interest to us: cut points, cut pairs and non-cut pairs.

In section 3 we deal with the special case of a group with circular Bowditch boundary. 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 precisely conjugacy class representatives of the fundamental groups of the boundary components of that orbifold. In [23] a complete list of such orbifolds that do not admit such a splitting is described; we use this to complete this special case.

In section 4 we first record the general definition of a JSJ decomposition and a description of Bowditch’s canonical JSJ decomposition over virtually cyclic subgroups. We then recall the theorem of [8] that links the topology of the Gromov boundary of a hyperbolic group to the existence of a proper splitting of that group and extend it to a relative version. We then show how to use the algorithms described so far to compute a maximal splitting of a one-ended hyperbolic group over virtually cyclic subgroups.

In section 5 we complete the proof of Theorem 0.1 by describing the processes that 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.

It seems plausible that the techniques of this paper 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 [14] 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 [12], so we do not introduce additional technical complexity by trying to give a new proof of this result here.

Acknowledgements

I thank my supervisor Henry Wilton for suggesting this problem and for his guidance towards its solution. I am also grateful to the anonymous reviewer for pointing out several problems with the clarity and accuracy of the earlier draft.

1 The cusped space and cut pairs

1.1 The cusped space and the boundary

In this section we recall some technical results about the geometry of relatively hyperbolic groups. Fix a group Γ with a finite generating set S and a finite collection \mathcal{H} of subgroups of Γ such that $H \cap S$ is a generating set for H for each H in \mathcal{H} . For now we may allow the groups in \mathcal{H} to be arbitrary, although in our application they will be virtually cyclic.

In [22] the cusped space X associated to the triple (Γ, S, \mathcal{H}) is defined; we recall the definition here. See [22] for more details and properties of the cusped space. We use only the 1-skeleton of the cusped space so we omit the 2-cells from the definition.

Definition 1.1. For C a 1-complex define the combinatorial horoball \widehat{C} based on C to be the 1-complex with vertex set $C^{(0)} \times \{0 \cup \mathbb{N}\}$ and three types of edges:

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. For $k \geq 1$ a horizontal edge from (v, k) to (w, k) whenever $0 < d(v, w) \leq 2^k$.
3. A vertical edge from (v, k) to $(v, k+1)$ for each $v \in C^{(0)}$ and each $k \in \mathbb{Z}_{\geq 0}$.

We define a height function $h: \widehat{C} \rightarrow \mathbb{R}_{\geq 0}$ sending a vertex (v, k) to k and interpolating linearly along edges.

Definition 1.2. If Γ , S and \mathcal{H} are as above then the cusped space X associated to this triple is defined as follows. For H in \mathcal{H} let T_H be a left transversal of H in Γ . For H in \mathcal{H} and t in T_H let $C_{H,t}$ be the full subgraph of $\text{Cay}(\Gamma, S)$ containing tH ; note that this is isomorphic to $\text{Cay}(H, H \cap S)$. Then let X be the union

$$\text{Cay}(\Gamma, S) \cup \bigcup_{H \in \mathcal{H}, t \in T_H} \widehat{C_{H,t}}$$

where we identify $C_{H,t} \subset \text{Cay}(\Gamma, S)$ with the part of $\widehat{C_{H,t}}$ with height 0. Then the height function defined on each horoball extends to a height function $h: X \rightarrow \mathbb{R}_{\geq 0}$.

For $k \geq 0$ the k -thick part X_k of X is defined to be $h^{-1}[0, k]$. We say that a path γ in X is vertical if $h \circ \gamma$ is strictly monotonic and horizontal if $h \circ \gamma$ is constant.

Definition 1.3. A (quasi-)isometric embedding α from a subinterval of \mathbb{R} to X is called a (quasi-)geodesic. If that interval is closed and bounded then α is a segment. If the interval is $[0, \infty)$ then α is a ray. If the interval is all of \mathbb{R} then α is bi-infinite.

X is called a geodesic space if for any x and y in X there exists a geodesic segment $\alpha: [a, b]$ with $\alpha(a) = x$ and $\alpha(b) = y$.

Definition 1.4. A geodesic metric space X is δ -hyperbolic if every edge of any geodesic triangle in X is contained in a δ -neighbourhood of the union of the other two edges.

Definition 1.5. Γ is hyperbolic relative to \mathcal{H} if the cusped space associated to (Γ, S, \mathcal{H}) is hyperbolic for some (equivalently for any) generating set S in which $S \cap H$ is a generating set for H for each $H \in \mathcal{H}$.

This definition of relative hyperbolicity is shown to be equivalent to several other standard definitions in [22].

We recall a lemma from [7] that characterises hyperbolicity of the pair (Γ, \mathcal{H}) when Γ is itself a hyperbolic group. Recall that a collection \mathcal{H} of subgroups of Γ is *almost malnormal* if, for H_1 and H_2 in \mathcal{H} and g in Γ , $H_1 \cap gH_2g^{-1}$ is finite unless $H_1 = H_2$ and $g \in H_1$.

Lemma 1.6. [7, Theorem 7.11] Let Γ be a non-elementary hyperbolic group and let \mathcal{H} be a finite set of subgroups of Γ . Then Γ is hyperbolic relative to \mathcal{H} if and only if \mathcal{H} is almost malnormal in Γ and each element of \mathcal{H} is quasi-convex in Γ .

In the cases of interest to us all groups in \mathcal{H} are virtually cyclic. 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 Γ . Putting this together with Lemma 1.6 we obtain the following:

Lemma 1.7. If Γ is hyperbolic and each group in \mathcal{H} is maximal virtually cyclic then Γ is hyperbolic relative to \mathcal{H} .

Fix X as in Definition 1.2 and assume that it is δ -hyperbolic. We shall require the following two lemmas describing the geometry of X . The first is [14, Lemma 2.11] and the second is the Morse lemma for a hyperbolic metric space.

Lemma 1.8. *Let $C = 3\delta$. For any v in X_0 and x in X there is a geodesic ray starting at v that passes within a distance C of x .*

Lemma 1.9. *There exists a (computable) function $D(\lambda, \epsilon, \delta)$ such that whenever γ is a geodesic and γ' is a (λ, ϵ) -quasi-geodesic with the same end points in $X \cup \partial X$ the Hausdorff distance from γ to γ' is at most D .*

Fixing δ so that X is δ -hyperbolic we fix the constant $C = 3\delta$ and function $D = D(\lambda, \epsilon, \delta)$ as in Lemmas 1.8 and 1.9 for the remainder of this and the following section.

We shall need an algorithm to compute the constant δ with respect to which the cusped space is δ -hyperbolic. This is dealt with by the following results.

Proposition 1.10. *[13, Prop. 2.3] There is an algorithm that takes as input a presentation for a group Γ , the generators in Γ for a finite set of subgroups of Γ with respect to which Γ relatively hyperbolic, and a solution to the word problem in Γ and returns the constant of a linear relative isoperimetric inequality satisfied by the given presentation of Γ .*

In the case of interest here, Γ is hyperbolic. Hyperbolic groups have uniformly solvable word problem, so the requirement that the algorithm be given a solution to the word problem in Γ 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. The definition of this space is [22, Definition 2.47].

In [22] the cusped 2-complex is defined; this is a simply connected 2-complex, the 1-skeleton of which is the cusped 1-complex defined in Definition 1.2. 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 Γ .

Theorem 1.11. *[22, Theorem 3.24] Suppose that the coned-off Cayley complex of Γ with respect to S and \mathcal{H} satisfies a linear combinatorial isoperimetric inequality with constant K . Then the cusped two-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 Γ and a set of generators for each group H in \mathcal{H} is therefore completed by the following proposition:

Proposition 1.12. *[22, Prop. 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.*

1.2 The Bowditch boundary

Definition 1.13. *Let X be a hyperbolic geodesic metric space. The Gromov boundary ∂X of X is defined to be the quotient of the set of geodesic rays starting at some chosen base point in which two rays are identified if they are a finite Hausdorff distance apart. It is endowed the quotient of the compact-open topology.*

This definition is quasi-isometry invariant and independent of the chosen base point up to homeomorphism.

Definition 1.14. *Let Γ be a group that is hyperbolic relative to a collection \mathcal{H} of subgroups. The Bowditch boundary $\partial(\Gamma, \mathcal{H})$ is defined to be the Gromov boundary of the cusped space associated to Γ and \mathcal{H} with some generating set.*

This definition of the Bowditch boundary is shown to be equivalent to other definitions in [7]. If each group in \mathcal{H} is maximal virtually cyclic then we have the following description of the Bowditch boundary from [17]:

Lemma 1.15. *If Γ is hyperbolic and each group in \mathcal{H} is maximal virtually cyclic then $\partial(\Gamma, \mathcal{H})$ is the quotient of $\partial(\Gamma, \emptyset)$ by the equivalence relation in which $x \sim y$ if and only if either $x = y$ or $\{x, y\} = g \cdot \Lambda H$ for some g in Γ and H in \mathcal{H} . The topology on $\partial(\Gamma, \mathcal{H})$ is the quotient of the topology on $\partial(\Gamma, \emptyset)$.*

The boundary of a hyperbolic metric space has a natural quasi-conformal class of metrics. Recall the definition of the Gromov product of a pair of points p and q in X with respect to a base point v in X :

$$(p \cdot q)_v = \frac{1}{2}(d(v, p) + d(v, q) - d(p, q)).$$

This definition is extended to allow p and q to be in $X \cup \partial X$ by

$$(p \cdot q)_v = \sup \liminf_{i, j \rightarrow \infty} (p_i \cdot q_j)_v.$$

Here the supremum is taken over pairs of sequences (p_i) and (q_j) in X converging to p and q respectively.

By [9, III.H.3.21] ∂X admits a visual metric at any base point v ; that is, a metric d_v on ∂X satisfying

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

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

1.3 The double-dagger condition

If the boundary of X does not contain a cut point then its local connectivity is controlled by the so-called double-dagger condition. This was defined first in the absolute case in [2] and later in the relative case in [14]. We now record the definition of the condition and some important properties. Fix a base point v in X_0 .

Let $M = 6(C + 45\delta) + 2\delta + 3$. For $\epsilon \geq 0$ we say that a pair of points x and y in X satisfy \star_ϵ if $|d(v, x) - d(v, y)| \leq \epsilon$ and $d(x, y) \leq M$.

For $n \geq 0$ we say that a pair of points x and y satisfying \star_ϵ satisfy $\ddagger(\epsilon, n)(x, y)$ if there exists a path of length at most n from x to y that avoids the ball of radius $m - C - 45\delta + 3\epsilon$ centred at v , where $m = \min\{d(v, x), d(v, y)\}$. After this section we will not need to allow ϵ to be non-zero; we will say that X satisfies \ddagger_n if $\ddagger(0, n)(x, y)$ holds for all x and y in X satisfying \star_0 . However, the full definition is required in the proof of Proposition 1.16.

We require the following proposition, which is essentially Proposition 5.1 of [14].

Proposition 1.16. *Suppose that $\partial(\Gamma, \mathcal{H})$ is connected and without a cut point. Then there exists n such that X satisfies \ddagger_n . Furthermore, there is an algorithm that computes such an n if $\partial(\Gamma, \mathcal{H})$ is connected and without a cut point and that does not terminate if it is not connected.*

For clarity we recall the proof of this proposition from [14] here.

Proof. We begin by recalling some constants from [14]. In this paper these constants will only appear in this proof. Let $k = 2M$, let $K = 3(2^{2M+3}) + M + 3$ and for any n let $R(n) = 4(n + M) + 3k + 50\delta + 3$. For each $n \geq K$ in turn, check whether $\ddagger(10\delta, n)(x, y)$ holds for all pairs of vertices (x, y) in $\text{Ball}_{R(n)}(v) \cap X_k$ satisfying $\star_{10\delta}$.

If $\partial(\Gamma, \mathcal{H})$ is connected and does not contain a cut point then such an n exists by [14, Lemma 4.2]. It is commented that $\ddagger(10\delta, n)(x, y)$ holds for all pairs of vertices (x, y) in $\text{Ball}_{R(n)}(v)$ satisfying $\star_{10\delta}$ with $x \notin X_k$ in the first paragraph of the proof of [14, Lemma 2.16]. Then [14, Corollary 4.6] says that X satisfies \ddagger_n ; note that this corollary applies because the conclusion of [14, Lemma 2.16] holds in the case of virtually cyclic peripheral subgroups.

If $\partial(\Gamma, \mathcal{H})$ is disconnected then there is no n such that the condition \ddagger_n holds by [14, Lemma 4.1]. \square

1.4 Cut points and pairs

We now investigate cut points and pairs in ∂X under the assumption that X satisfies \ddagger_n ; we assume that this condition holds for the remainder of this section. We relate the existence of such a point or pair to the connectedness of thickened cylinders around quasi-geodesics in X .

Let γ be a bi-infinite (λ, ϵ) -quasi-geodesic in X or a (λ, ϵ) -quasi-geodesic ray passing through X_0 in X . Fix a point v in $\gamma \cap X_0$. Since γ is a quasi-geodesic, its limit set, which we shall denote $\Lambda\gamma$, is either a pair of points $\gamma(\pm\infty)$ or a single point $\gamma(\infty)$ depending on whether γ is bi-infinite or a ray.

Recall the definitions $C = 3\delta$ and $D = D(\lambda, \epsilon, \delta)$ as in Lemma 1.9.

We define a subset of X , the connectivity of which will be seen to reflect the connectivity of $\partial X - \Lambda\gamma$. 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 $A_{r,R,K}(\gamma)$ be the union of those connected components of $N_{r,R}(\gamma)$ that meet $C_K(\gamma)$. Note that $A_{r,R,K}(\gamma)$ contains $N_{K,R}(\gamma)$.

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

Lemma 1.17. *Let $r > D$. Then $\bigcup_U \mathcal{S}U = \partial X - \Lambda\gamma$, where the union is taken over the set of connected components of $A_{r,\infty,K}(\gamma)$. Furthermore, $\mathcal{S}U \cap \mathcal{S}V = \emptyset$ for distinct components U and V .*

Proof. The statement that shadows are disjoint is clear from the definition. The assumption that $r > D$ ensures that when U is a subset of $A_{r,\infty,K}(\gamma)$, $\mathcal{S}U$ does not contain a point in $\Lambda\gamma$ by Lemma 1.9.

If p is any point in $\partial X - \Lambda\gamma$ then any geodesic ray α from v to p diverges arbitrarily far from γ , so $d(\alpha(t), \gamma) \geq r + D$ for t at least some number t_0 . Let

U be the component of $A_{r,\infty,K}(\gamma)$ that contains $\alpha(t)$ for $t \geq t_0$. If α' is another geodesic ray from v to p then for any t there exists t' such that $d(\alpha'(t), \alpha(t')) \leq D$ and then $|t - t'|$ is guaranteed to be at most D . If $t \geq t_0 + D$ then $t' \geq t_0$, so $\alpha'(t)$ is in the same component of $A_{r,\infty,K}(\gamma)$ as $\alpha(t')$, which is in U . Therefore $\alpha'(t) \in U$ for $t \geq t_0 + D$. It follows that $p \in \mathcal{S}U$. \square

Lemma 1.18. *Let U be a component of $A_{r,\infty,K}(\gamma)$. Then $\mathcal{S}U$ is non-empty as long as $K \geq r + D + \delta + C$ and $r > D$.*

Proof. Let $x \in C_K(\gamma) \cap U$. Then by Lemma 1.8 there exists a geodesic ray α from v and $t \geq 0$ such that $d(\alpha(t), x) \leq C$. Any geodesic segment in X from x to $\alpha(t)$ is contained in U , so $\alpha(t) \in U$. Also $d(\alpha(t), \gamma) \geq r + D + \delta$; it follows from Lemma 1.9, hyperbolicity of X and the assumption that $r > D$ that $d(\alpha(t'), \gamma) \geq r$ for $t' \geq t$, and therefore that $\alpha(t') \in U$ for $t' \geq t$. As in the proof of Lemma 1.17 it follows from this and the fact that $r > D$ that $\alpha(\infty) \in \mathcal{S}U$. \square

Lemma 1.19. *If U is a component of $A_{r,\infty,K}(\gamma)$, $\mathcal{S}U$ is closed and open in $\partial X - \Lambda\gamma$ as long as $r > D$.*

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

Then there exists t_0 such that for $t \geq t_0$, $d(\alpha(t), \gamma) \geq r + 2D + 7\delta + 1$. We claim that $V_{t_0}(\alpha) \subset \mathcal{S}U$ for t_0 as defined in the previous paragraph. To see this, let $q \in V_{t_0}(\alpha)$ and let β be a geodesic ray from v to q , so $d(\beta(t_0), \alpha(t_0)) < 2\delta + 1$. Let β' be another geodesic ray from v with $\beta'(\infty) = q$, so $d(\beta'(t_0), \beta(t_0)) < 4\delta$. Suppose that there exists $t \geq t_0$ such that $\beta'(t) \notin U$. Then there exists $t' \geq t_0$ such that $d(\beta'(t'), \gamma) \leq r$; without loss of generality assume that $d(\beta'(t'), \gamma|_{[0,\infty)}) \leq r$. Let γ' be a geodesic ray from v to $\gamma(\infty)$, so the Hausdorff between $\gamma|_{[0,\infty)}$ and γ' is at most D . Then $d(\beta'(t'), \gamma') \leq r + D$ and $d(\beta'(t_0), \gamma') \leq r + D + \delta$. Putting these inequalities together, $d(\alpha(t_0), \gamma) \leq r + 2D + 7\delta + 1$, which is a contradiction. Hence $\beta(t) \in U$ for $t \geq t_0$, and so $q \in \mathcal{S}U$.

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

The proof of the following lemma is based on the proof of [2, Proposition 3.2].

Lemma 1.20. *If U is a connected component of $A_{r,\infty,K}(\gamma)$ then $\mathcal{S}U$ is contained in one connected component of $\partial X - \Lambda\gamma$ as long as r satisfies the following inequality.*

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

Proof. Let p and q be points in $\mathcal{S}U$ and let α_1 and α_2 be geodesic rays from v 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 v 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 [2, Prop. 3.2], we show that $r_i(\infty)$ and $r_{i+1}(\infty)$ can be connected in $\partial X - \Lambda\gamma$ for each i ; for notational convenience we prove it for $i = 0$.

Using the condition \ddagger_n , 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:

$$d(\beta_{j/n^k}(m_i + k), \beta_{j+1/n^k}(m_i + k)) \leq M \text{ for each } j \text{ with } 0 \leq j \leq n^k - 1.$$

Note that the first step of the induction holds since M is at least $2C + 1$. The triangle inequality gives the following lower bound on the Gromov product of these points.

$$\begin{aligned} (\beta_{j/n^k}(\infty) \cdot \beta_{j+1/n^k}(\infty))_v &\geq \liminf_{n_1, n_2} (\beta_{j/n^k}(n_1) \cdot \beta_{j+1/n^k}(n_2))_v \\ &\geq (\beta_{j/n^k}(m_0 + k) \cdot \beta_{j+1/n^k}(m_0 + k))_v \\ &= m_0 + k - M/2 \end{aligned}$$

Let d_v be a visual metric on ∂X with base point v , visual parameter a and multiplicative constants k_1 and k_2 . We obtain:

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

Inductively applying the triangle inequality we arrive at the following inequality

$$d_v(\beta_0(\infty), \beta_t(\infty)) \leq \frac{k_2(n-1)a^{-m_0 + M/2}}{1 - a^{-1}} \text{ for each } n\text{-adic rational } t \in [0, 1].$$

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 $t \rightarrow \beta_t(\infty)$ defined on the n -adic rationals, which is established by equation (*). 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 v to $\gamma(\infty)$, so the Hausdorff distance between γ and γ' is at most D . By [9, III.H.3.17],

$$(\beta_0(\infty) \cdot \gamma'(\infty))_v \leq \liminf_{n_1, n_2} (\beta_0(n_1) \cdot \gamma'(n_2))_v + 2\delta.$$

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 , and 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.

$$\begin{aligned}
(\beta_0(n_1) \cdot \gamma'(n_2))_v - (\beta_0(m_0) \cdot \gamma'(m_0))_v &= d(\beta_0(n_1), \beta_0(m_0)) - d(\beta_0(n_1), p) \\
&\quad + d(\beta_0(m_0), \gamma'(m_0)) - d(p, q) \\
&\quad + d(\gamma'(m_0), \gamma'(n_2)) - d(q, \gamma'(n_2)) \\
&\leq \delta + 2\delta + \delta = 4\delta
\end{aligned}$$

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

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

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

Lemma 1.21. *The inclusion map $A_{r,R,K}(\gamma) \hookrightarrow A_{r,\infty,K}(\gamma)$ induces a bijection between the sets of connected components of those subspaces of X as long as $R \geq 4\delta + D + \max\{r + 4\delta + 1, K\}$.*

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 $A_{r,\infty,K}(\gamma)$. We show that the shortest path from x to y in $N_{r,\infty}(\gamma)$ stays within a distance R of γ . 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 $t \in [t_0, t_1]$.

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 [9, III.H.1.13] it is contained in a 2δ -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 $d(\phi(s), \gamma) \leq 4\delta + \max\{r + 8\delta + 2, K\} + D$, which is a contradiction. \square

From the results of this section we conclude the following:

Proposition 1.22. *The map that sends a component U of $A_{r,R,K}(\gamma)$ to the shadow (with respect to some base point $v \in \gamma$) of the component of $A_{r,\infty,K}(\gamma)$ containing U is a well defined bijection between the set of connected components of $A_{r,R,K}(\gamma)$ 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 1.17, 1.18, 1.19, 1.20 and 1.21. \square*

Remark 1.23. *The conditions on r , R and K depend only on δ , n , λ and ϵ and suitable values can be computed from these data.*

We end this section with the following lemma, which shows that connectedness of $A_{r,R,K}(\gamma)$ can be detected locally.

Lemma 1.24. *Suppose that γ is a bi-infinite (λ, ϵ) -quasi-geodesic and that neither point in $\Lambda\gamma$ is a cut point. Let r and R be chosen to satisfy lemmas 1.20 and 1.21. Let T be at least $\log_a(2k_1/k_2) + 3D + 2\delta + K$. Then every component of $A_{r,R,K}(\gamma)$ meets $C_K(\gamma) \cap \text{Ball}_T(\gamma(t))$ for any t such that $\gamma(t)$ is in X_0 .*

Proof. Fix the base point $v = \gamma(t)$. Let U be a component of $A_{r,R,K}(\gamma)$ and let U' be the component of $A_{r,\infty,K}(\gamma)$ containing U . Then it is sufficient to show that U' meets $C_K(\gamma) \cap \text{Ball}_T(v)$ since $U' \cap C_K(\gamma) = U \cap C_K(\gamma)$ by Lemma 1.21.

Suppose that U' does not meet $C_K(\gamma) \cap \text{Ball}_T(v)$. Let p be a point in SU and let α be a geodesic ray from v 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 $d(\gamma'(0), v) \leq D$. Then $d(\alpha(s), \gamma') \leq K + D$; let $d(\alpha(s), \gamma'(s')) \leq K + D$. This implies that $d(\gamma'(s'), v) \geq T - D - K$. We therefore have

$$(\alpha(s) \cdot \gamma'(s'))_v \geq T - D - K.$$

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

$$\begin{aligned} (p \cdot \gamma'(\infty))_v &\geq \liminf_{m,n \rightarrow \infty} (\alpha(m) \cdot \gamma'(n))_v \\ &\geq (\alpha(s) \cdot \gamma'(s'))_v - D \\ &\geq T - 2D - K. \end{aligned}$$

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 SU' 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 $(\gamma(t+s) \cdot \gamma(t-s))_v \leq D$. Then $(\gamma(\infty) \cdot \gamma(-\infty))_v \leq 2\delta + D$, 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 Lemma 1.20 SU' is connected, so is contained in one of these two balls, say in the ball around $\gamma(\infty)$. But then SU 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

2 Detecting cut points and pairs

We now use the results of the previous section to prove some computability results concerning topological features of the Bowditch boundary of a hyperbolic group under the assumption that the cusped space satisfies a double dagger condition. These are the main technical results of this paper. The idea is to identify the topological feature of the boundary with a combinatorial feature of the cusped space of bounded size, so that the existence of that feature can be determined by looking at only a finite part of the the cusped space.

Let Γ be a group hyperbolic relative to a finite set \mathcal{H} of maximal virtually cyclic subgroups. Let S be a generating set for Γ such that $S \cap H$ generates H for each H in \mathcal{H} . Let X be the cusped space associated to (Γ, \mathcal{H}, S) .

2.1 Geodesics in horoballs

First we must understand the connectivity of neighbourhoods of geodesics in the thin part of X . We assume that the peripheral subgroups are virtually cyclic, so the geometry of the cusps of X is relatively simple. For notational convenience we initially restrict to the case in which X consists of a single cusp. Then the vertex set of X can be identified with $H \times \mathbb{Z}_{\geq 0}$ and for any k there is an inclusion of the Cayley graph $\text{Cay}(H, S) \hookrightarrow h^{-1}(k)$ mapping a vertex $h \in \text{Cay}(H, S)$ to (h, k) ; for $k = 0$ this inclusion is an isomorphism of graphs.

Let d_H be the word metric in $\text{Cay}(H, S)$ and let $\text{Cay}(H, S)$ be δ_H -hyperbolic with respect to this word metric. Let α be a bi-infinite geodesic in $\text{Cay}(H, S)$ with respect to d_H ; then any point in $\text{Cay}(H, S)$ 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 $\gamma(0) = (\alpha(0), 0)$. Then for any $k \in \mathbb{Z}_{\geq 0}$, the vertex set of $h^{-1}(k) \cap N_{r,R}(\gamma)$ is

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

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

Assume now that $k \geq \log_2(2\delta_H + 1)$. 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

$$\begin{aligned} \alpha|_{[2^{k+r-1}, 2^{k+R-1}] \times \{k\}} &\subset Y_k^+, \\ \alpha|_{[-2^{k+R-1}, -2^{k+r-1}] \times \{k\}} &\subset Y_k^-, \end{aligned}$$

and each of these components meets $C_K(\gamma)$. Therefore each of the sets Y_k^\pm is a subset of a component of $A_{r,R,K}(\gamma)$. Any vertex in the complement of these two components of Y_k is contained in

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

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, vertices that are in Y_{k+1} but not in Y_{k+1}^\pm are adjacent to vertices in Y_k .

Thus, if $k \geq \log_2(2\delta_H + 1)$ then $A_{r,R,K}(\gamma) \cap h^{-1}[k, \infty)$ 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

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

To make precise the consequences of this description of $A_{r,R,K}(\gamma)$, we make the following definition, now allowing X to consist of more than a single cusp. Let $\gamma: [a, b] \rightarrow X$ be a geodesic segment such that $h(\gamma(a)) = h(\gamma(b)) = k > R$ such that $h \circ \gamma$ is decreasing at a and increasing at b . Let $\hat{\gamma}: (-\infty, \infty) \rightarrow X$ be the path obtained by concatenating γ with two vertical geodesic rays. Note that this is a k -local-geodesic. Let $A'_{r,R,K}(\gamma)$ be

$$A_{r,R,K}(\hat{\gamma}) - h^{-1}[k, \infty) \cap N_R(\hat{\gamma}(-\infty, a] \cup \hat{\gamma}[b, \infty)).$$

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 2.1. *Let γ and $\hat{\gamma}$ be as in the previous paragraph. Let $\delta_{\mathcal{H}}$ be such that each $H \in \mathcal{H}$ is $\delta_{\mathcal{H}}$ -hyperbolic with respect to the generating set $H \cap S$. Suppose that $h(\gamma(a)) = h(\gamma(b)) \geq \min\{R, \log_2(2\delta_{\mathcal{H}} + 1)\}$. Then the inclusion $A'_{r,R,K}(\gamma) \hookrightarrow A_{r,R,K}(\hat{\gamma})$ induces a bijection between the sets of connected components of those spaces. \square*

2.2 Cut points

We now show that there is an algorithm that determines whether or not ∂X contains a cut point under the assumption that X satisfies a double dagger condition.

Proposition 2.2. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ with generating set S , a list of subsets of S generating a collection \mathcal{H} of maximal virtually cyclic subgroups of Γ and integers δ and n such that the cusped space is δ -hyperbolic and satisfies \ddagger_n and returns the answer to the question “does $\partial(\Gamma, \mathcal{H})$ contain a cut point?”*

Proof. It is shown in [4, Theorem 0.2] that any cut point in $\partial(\Gamma, \mathcal{H})$ must be the limit point of gHg^{-1} for some $H \in \mathcal{H}$ and $g \in \Gamma$. Therefore it is sufficient to check whether or not ΛH is a cut point for each $H \in \mathcal{H}$. Choose r, R and K to simultaneously satisfy the conditions of lemmas 1.17, 1.18, 1.19, 1.20 and 1.21 with $\lambda = 1$ and $\epsilon = 0$. Let $\delta_{\mathcal{H}}$ be large enough that each H in \mathcal{H} is $\delta_{\mathcal{H}}$ -hyperbolic with respect to the generating set $S \cap H$ and let $k \geq \log_2(2\delta_{\mathcal{H}} + 1)$.

Then for each vertical geodesic ray γ starting at the identity element $1 \in X_0$ check whether or not $A_{r,R,K}(\gamma) \cap h^{-1}([0, k])$ is connected; as in Lemma 2.1, it is connected if and only if $A_{r,R,K}(\gamma)$ is connected. Also, $A_{r,R,K}(\gamma)$ is disconnected if and only if $\Lambda\gamma = \Lambda H$ is a cut point by lemmas 1.17, 1.18, 1.19, 1.20 and 1.21, where H is the element of \mathcal{H} such that the combinatorial horoball based on $H \subset \text{Cay}(X, S)$ contains γ . This check can be completed in finite time. \square

2.3 Cut pairs

We now assume that X is δ -hyperbolic and satisfies \ddagger_n , that ∂X contains no cut point and that the Cayley graph of H with respect to its generating set $S \cap H$ is $\delta_{\mathcal{H}}$ -hyperbolic for each H in \mathcal{H} . Then all results of sections 1.4 and 2.1 can be applied.

First we show that the existence of a cut pair in ∂X is equivalent to the existence of a feature in X of known bounded size. Then by searching for such a feature one can determine whether or not ∂X contains a cut pair. 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 quasi-geodesic with bounded period that also joins the points of a (possibly different) cut pair.

Before stating the proposition we define some constants. Take λ and ϵ so that any $(8\delta + 1)$ -local-geodesic is a (λ, ϵ) -quasi-geodesic; for example let $\lambda = (12\delta + 1)/(5\delta + 1)$ and let $\epsilon = 2\delta$. Fix r, R and K to simultaneously satisfy the

conditions of propositions 1.17, 1.18, 1.19, 1.20 and 1.21 and fix T to satisfy the conditions of Lemma 1.24 with this choice of λ and ϵ . Also let

$$\begin{aligned} k &= \max\{8\delta + 1, \log_2(2\delta\mathcal{H} + 1), T + R\}, \\ \rho &= (2R + \epsilon)\lambda^2 + \epsilon + R \text{ and} \\ \eta &= \max\{(8\delta + 1)/2, \lambda(T + K) + \lambda\epsilon, \lambda(R + r) + \lambda\epsilon, \lambda(R + \rho) + \lambda\epsilon\}. \end{aligned}$$

Let B 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

$$\begin{aligned} N_{\min} &= \max\{8\delta + 1, \lambda(2R + 1) + \lambda\epsilon + 1\} \text{ and} \\ N_{\max} &= N_{\min} (k + R + 1) B^{2\eta} 2^V + 1. \end{aligned}$$

Proposition 2.3. *∂X contains a cut pair if and only if X contains one of the following two features:*

1. *A short period geodesic at shallow depth in X : a geodesic segment $\gamma: [a - \eta, b + \eta] \rightarrow X$ contained in X_{k+R} such that*

- (a) $N_{\min} \leq b - a \leq N_{\max}$,
- (b) $h(\gamma(a)) = h(\gamma(b))$, so there exists $g \in \Gamma$ such that $\gamma(b) = g \cdot \gamma(a)$,
- (c) $\gamma|_{[b-\eta, b+\eta]} = g \cdot \gamma|_{[a-\eta, a+\eta]}$,
- (d) *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 $C_K(\gamma) \cap \text{Ball}_T(\gamma(c))$ for some $c \in [a, b]$, and*
- (e) *the partition on the vertices of $N_{r,R}(\gamma) \cap \text{Ball}_\rho(\gamma(b))$ 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 $N_{r,R}(\gamma) \cap \text{Ball}_\rho(\gamma(a))$ obtained by the restriction of \mathcal{P} ; note that $N_{r,R}(\gamma) \cap \text{Ball}_\rho(\gamma(b))$ is equal to $g \cdot N_{r,R}(\gamma) \cap \text{Ball}_\rho(\gamma(a))$ by condition 1c.*

2. *A short horseshoe-shaped geodesic: a geodesic segment $\gamma: [a, b] \rightarrow X$ such that*

- (a) $b - a \leq N_{\max} - 2R + 2\eta$
- (b) $h(\gamma(a)) = h(\gamma(b)) \geq k$,
- (c) $h \circ \gamma$ is decreasing at a and increasing at b ,
- (d) $A'_{r,R,K}(\gamma)$ is disconnected.

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

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

for $m \in \mathbb{Z}$ and $t \in [0, b - a]$. γ' is an $(8\delta + 1)$ -local-geodesic by condition 1c since $\eta \geq (8\delta + 1)/2$. It is therefore a (λ, ϵ) -quasi-geodesic by [9]. We now aim to show that $A_{r,R,K}(\gamma')$ 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 $\eta \geq \lambda(R+r) + \lambda\epsilon$, $N_r(\gamma') \cap N_R(\gamma|_{[a,b]})$ is a subset of $N_r(\gamma)$, so is equal to $N_r(\gamma) \cap N_R(\gamma|_{[a,b]})$. Therefore $N_{r,R}(\gamma')$ decomposes as a union

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

Since $b-a > \lambda(2R+1) + \epsilon\lambda$, $g^m \cdot N_R(\gamma|_{[a,b]})$ 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 $g^m \cdot \text{Ball}_\rho(\gamma(b))$ since $\rho \geq (2R+\epsilon)\lambda^2 + \epsilon + R$.

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 1e. Its restriction to $A_{r,R,K}(\gamma')$ is non-trivial: $C_K(\gamma')$ contains $C_K(\gamma) \cap \text{Ball}_T(\gamma(c))$ since $\eta \geq \lambda(T+K) + \lambda\epsilon$ and this set meets both sets in \mathcal{P}' by condition 1d; therefore $A_{r,R,K}(\gamma')$ is disconnected and therefore $\Lambda\gamma'$ is a cut pair by the results of section 1.4.

Now suppose that the second type of feature exists in X . Let $\hat{\gamma}$ be the $(8\delta+1)$ -local-geodesic obtained by concatenating γ with vertical geodesic rays. Then $A_{r,R,K}(\hat{\gamma})$ is disconnected by Lemma 2.1 and $\Lambda\hat{\gamma}$ is a cut pair by the results of section 1.4.

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 $N_{max} + 2\eta$, say $[a-R, b+R]$ with $h(\gamma'(a-R)) = h(\gamma'(b+R)) = k+R$, so $h(\gamma'(a)) = h(\gamma'(b)) = k$. Let $\gamma = \gamma'|_{[a,b]}$. Let $c \in [a, b]$ such that $h(\gamma'(c)) = 0$. $A_{r,R,K}(\gamma')$ is disconnected and each component meets $C_K(\gamma') \cap \text{Ball}_T(\gamma'(c))$ by Lemma 1.24. Since $k \geq T$ this is a subset of $A'_{r,R,K}(\gamma)$, and $A'_{r,R,K}(\gamma)$ is a subset of $A_{r,R,K}(\gamma')$, so $A'_{r,R,K}(\gamma)$ 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 $g_i \in \Gamma$, 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 1.4 since $\Lambda\gamma'$ is a cut pair. Since $\eta \geq \lambda(R+\rho) + \lambda\epsilon$, $N_{r,R}(\gamma') \cap \text{Ball}_\rho(\gamma(a_0))$ is equal to $g_i^{-1}N_{r,R}(\gamma') \cap \text{Ball}_\rho(\gamma(a_i))$ for all i . This set contains at most V vertices, so there exist $0 \leq i < j \leq 2^V$ such that

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

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, 1c and 1e are satisfied by definition of the a_i . Let $c \in [a, b]$ such that $\gamma'(c) \in X_0$. Then $C_K(\gamma) \cap \text{Ball}_T(\gamma(c))$ is equal to $C_K(\gamma') \cap \text{Ball}_T(\gamma'(c))$ since $\eta \geq \lambda(T+K) + \lambda\epsilon$ and Lemma 1.24 guarantees that both sets in \mathcal{P}' meet this set, so condition 1d is satisfied, too. \square

The existence of a geodesic segment with the properties described in the statement of Proposition 2.3 can be checked by looking at just a finite ball

in X . Such a ball can be computed from a solution to the word problem in Γ , which exists since Γ is hyperbolic. Therefore we immediately obtain the following corollary:

Corollary 2.4. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ with generating set S , a list of subsets of S generating a collection \mathcal{H} of maximal virtually cyclic subgroups of Γ and integers δ , $\delta_{\mathcal{H}}$ and n such that the cusped space associated to (Γ, \mathcal{H}, S) is δ -hyperbolic and satisfies \ddagger_n and such that the Cayley graph of each element of \mathcal{H} with respect to its given generating set is $\delta_{\mathcal{H}}$ -hyperbolic and returns the answer to the question “does $\partial(\Gamma, \mathcal{H})$ contain a cut pair?” \square*

2.4 Non-cut pairs

By a similar argument, we now show that there is an algorithm that determines whether or not ∂X contains a non-cut pair. For the detection of non-cut pairs we will need the following lemma.

Lemma 2.5. *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. ∂X is locally connected by the main theorem of [3]. ∂X is assumed not to contain a cut point, so the results of sections 2 and 3 of [8] can be applied. We recall some definitions from that paper. 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 $x \cong y$ if $\text{val}(x) = \text{val}(y)$ and $\partial X - \{x, y\}$ has exactly $\text{val}(x)$ components.

Recall [8, Lemma 3.8]: if $x \cong y$ and $x \cong z$ then $y \cong z$. 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 [8, Lem. 3.2] σ is a cyclically separating set, and so is the closure of σ by [8, Lem. 2.2], which implies that $\{x_{-\infty}, x_{\infty}\}$ is a cut pair as required. \square

Let $\lambda, \epsilon, r, R, K, T, k, \rho, \eta, B$ and V be as defined in section 2.3. Let N_1, N_2 and N_3 be given by

$$\begin{aligned} N_1 &= 2(V-1)((k+R+1)B^{2\eta}V^{V+1} + 2\eta) + 2\eta + 2((k+R+1)B^{2\eta} + 1), \\ N_2 &= (k+R+1)B^{2\eta} + 1, \\ N_3 &= 2(k+R+1)B^{2\eta}V^{V+1} + 4\eta. \end{aligned}$$

Proposition 2.6. *∂X contains a non-cut pair if and only if X contains one of the following two features:*

1. *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 $a_3 = b_2$ such that*
 - (a) $1 \leq b_i - a_i \leq N_1$ for $i = 1, 3$,
 - (b) $1 \leq b_2 - a_2 \leq N_2$,
 - (c) $\gamma_i|_{[b_i - \eta, b_i + \eta]} = \gamma_{i+1}|_{[a_i - \eta, a_i + \eta]}$ for $i = 1, 2$,

- (d) $\gamma_i|_{[b_i-\eta, b_i+\eta]} = g_i \cdot \gamma_i|_{[a_i-\eta, b_i+\eta]}$ for $i = 1, 3$, and
- (e) $h(\gamma_i(a_i)) = h(\gamma_i(b_i))$ for $i = 1, 3$, so there exist $g_i \in \Gamma$ such that $\gamma_i(b_i) = g_i\gamma_i(a_i)$,
- (f) all vertices of $C_K(\gamma_2) \cap \text{Ball}_T(\gamma_2(c))$ 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 geodesic segment $\gamma: [a, b] \rightarrow X$ with image in X_k such that

- (a) $b - a \leq N_3$,
- (b) $h(\gamma(a)) = h(\gamma(b)) = k$ with γ descending vertically at a and ascending vertically at b , and
- (c) $A'_{r,R,K}(\gamma)$ 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:

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

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 conditions 1c and 1d since $\eta \geq 8\delta + 1$ and is therefore a (λ, ϵ) -quasi-geodesic.

$C_K(\gamma') \cap \text{Ball}_T(\gamma'(c))$ is equal to $C_K(\gamma_2) \cap \text{Ball}_T(\gamma_2(c))$ since $\gamma'|_{[a_2-\eta, b_2+\eta]} = \gamma_2$ and $\eta \geq \lambda(T + K) + \lambda\epsilon$. Furthermore, $\eta \geq \lambda(R + r) + \lambda\epsilon$, 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 $C_K(\gamma') \cap \text{Ball}_T(\gamma'(c))$ lies in a single component of $A_{r,R,K}(\gamma')$ by condition 1f, so $A_{r,R,K}(\gamma')$ is connected by Lemma 1.24, so $\Lambda\gamma'$ is a non-cut pair by the results of section 1.4.

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 $k \geq 8\delta + 1$. $A_{r,R,K}(\hat{\gamma})$ is connected by Lemma 2.1, so $\Lambda\hat{\gamma}$ is a non-cut pair by the results of section 1.4.

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 $\gamma'(0) \in X_0$. $C_K(\gamma') \cap \text{Ball}_T(\gamma'(0))$ contains at most V vertices and lies in a single component of $N_{r,R}(\gamma')$. Define n_V^\pm to be $\pm\lambda(K + T + \epsilon)$ 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 $C_K(\gamma') \cap \text{Ball}_T(\gamma'(0))$ meets at most l components of $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_l^-, n_l^+]})$ and $[n_{l+1}^-, n_{l+1}^+] \subset [n_l^-, n_l^+]$. Note that the condition that $|n_l^\pm| \geq \lambda(K + T) + \lambda\epsilon$ ensures that $C_K(\gamma') \cap \text{Ball}_T(\gamma'(0))$ is a subset of $N_{r,R}(\gamma') \cap N_R(\gamma'|_{[n_l^-, n_l^+]})$.

Suppose that $n_{l-1}^+ - n_l^+ > (k + R + 1)B^{2\eta}V^{V+1} + 2\eta$. Let \mathcal{Q} be the partition of $C_K(\gamma') \cap \text{Ball}_T(\gamma'(0))$ 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:

$$\begin{aligned} Y_t &= N_{r,R}(\gamma') \cap \text{Ball}_\rho(\gamma'(t)) \\ Z_t &= N_{r,R}(\gamma') \cap \left(\text{Ball}_\rho(\gamma'(t)) \cup N_R(\gamma'|_{[n_{l-1}^-, t]}) \right) \end{aligned}$$

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 $C_K(\gamma') \cap \text{Ball}_T(\gamma'(0))$ in Z_t . As in the proof of Proposition 2.3 there exist $s_1 < s_2$ in $[n_l^+ + \eta, n_{l-1}^+ - \eta]$ such that such that

1. $h(\gamma'(s_1)) = h(\gamma'(s_2))$, so $\gamma'(s_2) = g\gamma'(s_1)$ for some g in Γ ,
2. $\gamma'|_{[s_2-\eta, s_2+\eta]} = g \cdot \gamma'|_{[s_1-\eta, s_1+\eta]}$, which implies that $Y_{s_2} = g \cdot Y_{s_1}$, and
3. $\mathcal{P}_{s_2} = g \cdot \mathcal{P}_{s_1}$.

Then replace γ' by another path defined by

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

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 $\eta \geq \lambda(R + r + \epsilon)$,

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

and it follows that $C_K(\gamma'') \cap \text{Ball}_T(\gamma''(0))$ meets at most $l - 1$ components of $N_{r,R}(\gamma'') \cap N_R(\gamma''|_{[n_{l-1}^-, n_{l-1}^+ - (s_2 - s_1)]})$. 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 $(k + R + 1)B^{2\eta}V^{V+1} + 2\eta$ for all l , and therefore that

$$|n_1^\pm| \leq \lambda(K + T + \epsilon) + (V - 1)((k + R + 1)B^{2\eta}V^{V+1} + 2\eta).$$

There exist $a_1 \leq b_1 \leq n_1^- - \eta$ with $n_1^- - b_1$ and $b_1 - a_1$ both at most $(k + R + 1)B^{2\eta} + 1$ such that

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

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, 1b, 1e, 1c and 1d clearly hold by construction and condition 1f holds because the condition that $\eta \geq \lambda(R + r) + \lambda\epsilon$ 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 2.1. 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 2.5 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 $b_i - a_i \leq 4\eta + 2(k + R + 1)B^{2\eta}V^{V+1}$; this yields a feature of the second type described in the proposition. \square

From this we deduce the following corollary:

Corollary 2.7. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ with generating set S , a list of subsets of S generating a collection \mathcal{H} of maximal virtually cyclic subgroups of Γ and integers δ , $\delta_{\mathcal{H}}$ and n such that the cusped space associated to (Γ, \mathcal{H}, S) is δ -hyperbolic and satisfies \ddagger_n and such that the Cayley graph of each element of \mathcal{H} with respect to its given generating set is $\delta_{\mathcal{H}}$ -hyperbolic and returns the answer to the question “does $\partial(\Gamma, \mathcal{H})$ contain a non-cut pair?” \square*

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.

3.1 Circular boundary

Let Γ be a hyperbolic group with a finite collection \mathcal{H} of subgroups such that Γ is hyperbolic relative to \mathcal{H} and $\partial(\Gamma, \mathcal{H})$ is homeomorphic to S^1 . Then Γ acts as a discrete convergence group on $\partial(\Gamma, \mathcal{H})$ by [5], so by the Convergence Group Theorem of Tukia, Casson, Jungreis and Gabai [34, 11, 19] there is a properly discontinuous action of Γ by isometries on \mathbb{H}^2 and a Γ -equivariant homeomorphism $\partial(\Gamma, \mathcal{H}) \rightarrow \partial\mathbb{H}^2$.

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

$$1 \longrightarrow K \longrightarrow \Gamma \longrightarrow \Gamma' \longrightarrow 1$$

and the quotient Γ' acts faithfully on \mathbb{H}^2 . Let \mathcal{H}' be the set of images of elements of \mathcal{H} in Γ' .

Lemma 3.1. *With Γ and Γ' as above, Γ splits non-trivially over a virtually cyclic subgroup relative to \mathcal{H} if and only if Γ' admits such a splitting relative to \mathcal{H}' .*

Proof. One direction is clear: if Γ' acts minimally on a non-trivial tree T with virtually cyclic edge stabilisers then the quotient map $\Gamma \rightarrow \Gamma'$ induces a minimal Γ -action on T with edge stabilisers finite extensions of the edge stabilisers of the Γ' -action.

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

It follows that the Γ -action descends to a Γ' action, and the edge stabilisers are quotients of the original edge stabilisers by finite subgroups. Furthermore, elements of \mathcal{H}' are elliptic if elements of \mathcal{H} are. \square

Any finite normal subgroup of Γ is contained in the ball of radius $4\delta + 2$ centred at the identity by [9], so by checking all finite subsets of this ball using a solution to the word problem, the set of finite normal subgroups of Γ can be computed. K is the unique maximal finite normal subgroup of Γ and is therefore computable.

As described above, Γ' 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 Γ' . The action of Γ' on $\partial(\Gamma', \mathcal{H}')$ is minimal, so the limit set of the action is S^1 . It follows that Γ' is a Fuchsian group of the first kind: \mathbb{H}^2/Γ' 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 Γ' as the fundamental group of a compact hyperbolic orbifold such that \mathcal{H}' is a choice of conjugacy class representatives for the boundary subgroups.

Definition 3.2. *A group is bounded Fuchsian if it acts properly discontinuously and convex cocompactly (i.e. cocompactly on a convex subset) on \mathbb{H}^2 . Then a group is bounded Fuchsian if and only if it surjects with finite kernel onto the fundamental group of a compact two-dimensional hyperbolic orbifold. The peripheral subgroups of a bounded Fuchsian group are the fundamental groups of the boundary components of the associated orbifold.*

Therefore we have shown that Γ' is a bounded Fuchsian group and \mathcal{H}' is a collection of representatives of its conjugacy classes of peripheral subgroups. This proves the following lemma:

Lemma 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 Γ' with a set \mathcal{H}' of virtually cyclic subgroups such that $\partial(\Gamma', \mathcal{H}')$ is homeomorphic to $\partial(\Gamma, \mathcal{H})$ and Γ splits over a virtually cyclic subgroup relative to \mathcal{H} if and only if Γ' splits over a virtually cyclic subgroup relative to \mathcal{H}' . Furthermore, if $\partial(\Gamma', \mathcal{H}')$ is homeomorphic to a circle then Γ' 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.*

3.2 Orbifolds

Given the results of the previous section, the problem of determining whether or not a given group Γ 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 Γ is a bounded Fuchsian group and \mathcal{H} is a set of conjugacy class representatives of its peripheral subgroups. Let $\Gamma = \pi_1 Q$ where Q is a compact hyperbolic orbifold, so \mathcal{H} is a set of conjugacy class representatives of the boundary subgroups of Q .

We recall some terminology to describe orbifolds. The universal cover of a hyperbolic orbifold (possibly with boundary) is a convex subspace \tilde{Q} of \mathbb{H}^2 . The orbifold fundamental group $\pi_1 Q$ acts properly discontinuously and by isometries on \tilde{Q} . The *underlying topological surface* Q_{top} of Q is defined to be the topological quotient of \tilde{Q} by this action. The *topological boundary* $\partial_{\text{top}} Q$ of Q is the boundary of Q_{top} , while the *orbifold boundary* ∂Q of Q is the image of the boundary of \tilde{Q} in Q_{top} under the quotient projection. A *mirror* in Q is the image in Q_{top} of the set of fixed points of an order 2 orientation reversing isometry of \tilde{Q} in $\pi_1 Q$. All mirrors in Q are contained in $\partial_{\text{top}} Q$ and any mirror is homeomorphic either to a circle or to an interval. In the latter case each end point of the interval is contained in either ∂Q or in another mirror. The intersection of two mirrors is called a *corner reflector*; the stabiliser in $\pi_1 Q$ of the preimage in \tilde{Q} of a corner reflector is a finite dihedral group. If x is a point in \tilde{Q} whose stabiliser in $\pi_1 Q$ is non-trivial, cyclic and orientation preserving then the image of x in Q_{top} is called a *cone point*. The *singular locus* of Q is the union of all mirrors, corner reflectors and cone points. If y is not contained in the singular locus of Q then any preimage of y in \tilde{Q} has trivial stabiliser in $\pi_1 Q$.

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 $\pi_1 Q$. 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 $\partial_{\text{top}} Q$. For more information about orbifolds see [31].

The theory of splittings of fundamental groups of orbifolds relative to their boundary subgroups is developed in [23]; we recall the following results:

Lemma 3.4. [23, Corollary 5.6] Γ splits non-trivially relative to \mathcal{H} if and only if Q contains an essential simple closed geodesic.

Definition 3.5. We call a compact orbifold without an essential simple closed geodesic *small*.

Proposition 3.6. [23, Proposition 5.12] A hyperbolic 2-orbifold Q is small if and only if it is one of the following.

1. A sphere with three cone points, so $\pi_1 Q \cong \langle a, b | a^p, b^q, ab^r \rangle$ where $p^{-1} + q^{-1} + r^{-1} < 1$ and the peripheral structure is empty.
2. A triangle, all three edges of which are mirrors, so $\pi_1 Q$ has presentation $\langle a, b, c | a^2, b^2, c^2, (ab)^p, (bc)^q, (ca)^r \rangle$ where $p^{-1} + q^{-1} + r^{-1} < 1$ and the peripheral structure is empty.

3. A disc with two cone points, so $\pi_1 Q \cong \langle a, b | a^p, b^q \rangle$ where $p, q > 1$ and the peripheral structure is $\{\langle ab \rangle\}$.
4. A cylinder with one cone point, so $\pi_1 Q \cong \langle a, b | (ab)^p \rangle$ where $p > 1$ and the peripheral structure is $\{\langle a \rangle, \langle b \rangle\}$.
5. A pair of pants, so $\pi_1 Q$ is the free group $\langle a, b \rangle$ and the peripheral structure is $\{\langle a \rangle, \langle b \rangle, \langle c \rangle\}$.
6. A disc with one cone point with edge consisting of an interval boundary component and a mirror, so $\pi_1 Q \cong \langle a, t | a^2, t^p \rangle$ where $p > 1$ and the peripheral structure is $\{\langle a, tat^{-1} \rangle\}$.
7. A square with an interval boundary component and three mirrors edges, so $\pi_1 Q \cong \langle a, b, c | a^2, b^2, c^2, (ab)^p, (bc)^q \rangle$ where $p + q \geq 1$ and the 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 $\pi_1 Q \cong \langle a, t | a^2 \rangle$ with peripheral structure $\{\langle a, tat^{-1} \rangle\}$.
9. A pentagon with two non-adjacent interval boundary components and three mirrors as edges, so $\pi_1 Q \cong \langle a, b, c | a^2, b^2, c^2, (ab)^p \rangle$ with 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 $\pi_1 Q \cong \langle a, b, c | a^2, b^2, c^2 \rangle$ and the peripheral structure is $\{\langle a, b \rangle, \langle b, c \rangle, \langle c, a \rangle\}$.

Lemma 3.7. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ and a collection \mathcal{H} of maximal virtually cyclic peripheral subgroups and terminates if and only if $\partial(\Gamma, \mathcal{H})$ is homeomorphic to a circle and Γ does not split relative to \mathcal{H} over a virtually cyclic subgroup.*

Proof. First use the algorithm of Lemma 3.3 and let the output of that algorithm be (Γ', \mathcal{H}') ; then $\partial(\Gamma, \mathcal{H})$ is homeomorphic to a circle and Γ does not split over a virtually cyclic subgroup relative to \mathcal{H} if and only if there is an isomorphism from Γ' to one of the groups with peripheral structure listed in proposition 3.6 that maps elements of \mathcal{H}' to conjugates of elements of that peripheral structure.

Enumerate the groups and peripheral structures described in Proposition 3.6 and, in parallel, enumerate all homomorphisms from these groups to Γ' and homomorphisms from Γ' 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 Γ' and all groups listed in Proposition 3.6 are hyperbolic. The algorithm then terminates when an inverse pair of such maps that preserve the peripheral structures (up to conjugacy) is found. \square

4 Maximal splittings

We now apply the technical results of sections 2 and 3 to the problem of finding a maximal splitting of a one-ended hyperbolic group.

4.1 JSJ decompositions

We begin by recalling the definition a JSJ decomposition. For a more detailed account see [23]. Let Γ be a group and let \mathcal{A} be a collection of subgroups of Γ that is closed under conjugation and taking subgroups. Let \mathcal{H} be another collection of subgroups of Γ . An $(\mathcal{A}, \mathcal{H})$ -tree is a tree with a Γ -action with no edge inversions in which the stabiliser of each edge is in \mathcal{A} and each element of \mathcal{H} is *elliptic*, that is, it fixes a point in the tree. An $(\mathcal{A}, \mathcal{H})$ -tree T *dominates* another such tree T' if there is a Γ -equivariant morphism from T to T' . T is *elliptic with respect to T'* if the stabiliser of each edge in T is elliptic in T' . For v a vertex of T we denote by Γ_v the stabiliser of v , and similarly if e is an edge of T we let Γ_e denote the stabiliser of e . We will sometimes refer to a $(\mathcal{A}, \mathcal{H})$ -tree as a *splitting of Γ over \mathcal{A} relative to \mathcal{H}* . A splitting is *trivial* if Γ is elliptic. We will frequently implicitly move between the languages of group actions on trees and graphs of groups.

Fixing \mathcal{A} and \mathcal{H} , a subgroup of Γ is *universally elliptic* if it is elliptic with respect to any $(\mathcal{A}, \mathcal{H})$ -tree. An $(\mathcal{A}, \mathcal{H})$ -tree is *universally elliptic* if the stabiliser of each of its edges is universally elliptic. An $(\mathcal{A}, \mathcal{H})$ tree is a *JSJ tree* if it is universally elliptic and is maximal for domination among universally elliptic $(\mathcal{A}, \mathcal{H})$ trees.

4.2 The Bowditch JSJ

In [8] Bowditch defines a canonical JSJ tree Σ for a one-ended hyperbolic group Γ where \mathcal{A} is the set of virtually cyclic subgroups of Γ and \mathcal{H} is empty. (To ensure that \mathcal{A} is closed under taking subgroups we should strictly speaking also include the finite subgroups of Γ in \mathcal{A} . However, since Γ is assumed to be one-ended it does not split over any finite subgroup, so this change is not important.) We recall a description of the tree Σ here.

Σ is a tree with a three-colouring: its vertex set $V(\Sigma)$ admits a partition $V_1(\Sigma) \sqcup V_2(\Sigma) \sqcup V_3(\Sigma)$ preserved by the action of Γ such that no two vertices in either $V_1(\Sigma)$ or $V_2(\Sigma) \cup V_3(\Sigma)$ are adjacent.

The stabiliser of a vertex in $V_1(\Sigma)$ is a maximal virtually cyclic subgroup of Γ and therefore contains the stabiliser of each incident edge at finite index.

To describe the vertices of the second type we require the following definition. Recall definition 3.2 of a bounded Fuchsian group.

Definition 4.1. *A hanging Fuchsian subgroup Q of Γ is a subgroup of Γ that is the stabiliser of a vertex in some finite splitting of Γ over virtually cyclic subgroups such that Q admits an isomorphism with a bounded Fuchsian group that maps the stabilisers of incident edges precisely to the peripheral subgroups.*

Stabilisers of vertices in $V_2(\Sigma)$ are precisely the maximal hanging Fuchsian subgroups of Γ and the stabilisers of incident edges at such a vertex in the JSJ decomposition are precisely the peripheral subgroups referred to in the definition.

The stabiliser of a vertex in $V_3(\Sigma)$ is not virtually cyclic and is not a hanging Fuchsian subgroup.

4.3 Splittings and the topology of the boundary

In this section we recall and extend some results linking the existence of a non-trivial splitting of a hyperbolic group relative to a collection of virtually cyclic subgroups to the topology of a particular Bowditch boundary.

For the remainder of Section 4.3, we fix a hyperbolic group Γ and a collection \mathcal{H} of virtually cyclic subgroups. If H is any virtually cyclic subgroup of Γ , let \widehat{H} be the unique maximal virtually cyclic subgroup of Γ containing H . Then let $\widehat{\mathcal{H}}$ be $\{\widehat{H} \mid H \in \mathcal{H}\}$. Recall Lemma 1.7: since each group in $\widehat{\mathcal{H}}$ is maximal virtually cyclic, Γ is hyperbolic relative to $\widehat{\mathcal{H}}$. We may therefore study the boundary $\partial(\Gamma, \widehat{\mathcal{H}})$.

4.3.1 Boundaries with cut points

We first recall results that deal with the case in which $\partial(\Gamma, \widehat{\mathcal{H}})$ contains a cut point. First recall Bowditch's definition [6] of a *peripheral splitting*.

Definition 4.2. *If Γ is a group and \mathcal{H} is a finite collection of subgroups of Γ , a peripheral splitting of Γ with respect to \mathcal{H} is a representation of Γ as a finite bipartite graph of groups such that each vertex group of one colour is conjugate to an element of \mathcal{H} , and each element of \mathcal{H} is conjugate to a vertex group of that colour.*

As for any splitting, we say that a peripheral splitting is trivial if some vertex group is equal to Γ .

Note that, in our setting, a peripheral splitting of Γ with respect to $\widehat{\mathcal{H}}$ is a splitting of Γ over virtually cyclic subgroups relative to \mathcal{H} .

Then the following proposition, which we obtain by putting together two theorems of Bowditch, completes our treatment of the case in which the boundary contains a cut point.

Proposition 4.3. *Suppose that $\partial(\Gamma, \widehat{\mathcal{H}})$ is connected and contains a cut point. Then the pair Γ admits a non-trivial splitting over virtually cyclic subgroups relative to \mathcal{H} .*

Proof. By [4, Theorem 0.2] the global cut point of $\partial(\Gamma, \widehat{\mathcal{H}})$ is a parabolic fixed point. By [3, Theorem 1.2] Γ admits a non-trivial peripheral splitting with respect to $\widehat{\mathcal{H}}$. The edge groups in this splitting are automatically virtually cyclic, since each edge meets a vertex with vertex group conjugate into $\widehat{\mathcal{H}}$, and the splitting is automatically relative to \mathcal{H} : in fact it is relative to $\widehat{\mathcal{H}}$. \square

4.3.2 Boundaries with cut pairs

In the absence of cut points, the existence of a relative splitting is reflected in the existence of cut pairs in the boundary. In the absolute case, recall the following theorem of [8].

Theorem 4.4. [8, Theorem 6.2] *Let Γ be a one-ended hyperbolic group such that $\partial(\Gamma, \emptyset)$ contains a cut pair and is not homeomorphic to S^1 . Then Γ admits a non-trivial splitting over a virtually cyclic subgroup.*

We require a relative version of this theorem. We only require such a result in the case when Γ arises as a vertex group in a splitting of a larger group

over virtually cyclic subgroups, and \mathcal{H} is the collection of edge groups incident at that vertex. In this section we show that in this simple case the relative result follows from Theorem 4.4, and one can avoid \mathbb{R} -trees machinery. For a discussion in greater generality, see [21].

Proposition 4.5. *Let v be a vertex in a minimal Γ -tree T with virtually cyclic edge groups where Γ is hyperbolic and one-ended. Let $\text{Inc } v$ be a set of representatives of Γ_v -conjugacy classes of stabilisers of edges in T incident at v . Suppose that $\partial(\Gamma_v, \text{Inc } v)$ 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 Γ_v admits a non-trivial splitting over virtually cyclic subgroups of Γ_v relative to $\text{Inc } v$.*

First we need the following lemma. Recall that for a vertex v in a Γ -tree, we defined $\text{Inc } v$ to be a set of conjugacy class representatives of the stabilisers of the edges of the Γ -tree incident at v , and let $\widehat{\text{Inc } v}$ be the set of maximal virtually cyclic subgroups of Γ_v that contain the elements of $\text{Inc } v$.

Lemma 4.6. *Let $f: T_1 \rightarrow T_2$ be an equivariant map of Γ -trees with virtually infinite cyclic edge stabilisers such that the action of Γ on T_1 is cocompact and the action of Γ on T_2 is minimal. (This means that there is no proper Γ -invariant subtree of T_2 .) Let v be a vertex of T_2 such that $\partial(\Gamma_v, \widehat{\text{Inc } v})$ is not a single point, is connected and does not contain a cut point. Then the action of Γ_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 $\Gamma_w \cap \Gamma_v$ is non-elementary. If this is not the case then the stabiliser of each edge of T_1 with respect to the action of Γ_v on T_1 is either finite or commensurable with the stabilisers of its end points. Therefore the action induces a splitting of Γ_v with finite edge groups and virtually cyclic vertex groups. By minimality of the action of Γ on T_2 each edge e incident at v is $f(e')$ for some edge e' of T_1 , and then $\Gamma_{e'}$ is a finite index subgroup of Γ_e , since each is virtually infinite cyclic. In particular, $\widehat{\Gamma_e}$ is elliptic, and the splitting is relative to $\widehat{\text{Inc } v}$. But $\partial(\Gamma_v, \widehat{\text{Inc } v})$ 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 $\Gamma_{w'} \cap \Gamma_v$ 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 [8] there exists a partition of $\partial(\Gamma, \emptyset) - \Lambda\Gamma_e$ as $U_1 \sqcup U_2$. The intersection of Γ_e with Γ_v is either finite or commensurable with a conjugate of an element of $\widehat{\text{Inc } v}$, so the images of $U_1 \cap \Lambda\Gamma_v$ and $U_2 \cap \Lambda\Gamma_v$ under the projection map $\Lambda\Gamma_v \rightarrow \partial(\Gamma, \widehat{\text{Inc } v})$ cover all but at most a point of $\partial(\Gamma, \widehat{\text{Inc } v})$. These sets are disjoint, so one must be empty, say U_2 . But $\Lambda\Gamma_w$ and $\Lambda\Gamma_{w'}$ each contain more than two points, so must both be contained in $U_1 \cup \Lambda\Gamma_e$. This implies that w and w' are on the same side of e .

Therefore the action of Γ_v on T_1 fixes S , for any element of Γ_v must send w to a vertex of S . \square

Proof of Proposition 4.5. Let Σ be Bowditch's JSJ tree for Γ . Σ is then elliptic with respect to T , so by [23, Proposition 2.2] there exists a Γ -tree $\widehat{\Sigma}$ and maps $p: \widehat{\Sigma} \rightarrow \Sigma$ and $f: \widehat{\Sigma} \rightarrow T$ such that:

1. p is a collapse map. (i.e. 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.

Let $S \subset \widehat{\Sigma}$ be the component of $f^{-1}(v)$ fixed by the action of Γ_v constructed in Lemma 4.6. Suppose that a vertex w of S is fixed by the Γ_v -action.

If e is any edge of $\widehat{\Sigma}$ that is adjacent to but not contained in S then $\Gamma_e \leq \Gamma_{f(e)}$ 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 $\partial(\Gamma_v, \widehat{\text{Inc } w})$ is homeomorphic to $\partial(\Gamma_v, \widehat{\text{Inc } v})$. We assumed that $\partial(\Gamma_v, \widehat{\text{Inc } v})$ was neither a point nor homeomorphic to a circle, so pw is not in $V_1(\Sigma)$ or $V_2(\Sigma)$, and is therefore in $V_3(\Sigma)$.

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

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

Then by Theorem 4.4 there is a type 1 or type 2 vertex u of Σ such that $\Lambda\Gamma_u$ contains $\{\tilde{x}, \tilde{y}\}$. Then $\{\tilde{x}, \tilde{y}\} = \Lambda\Gamma_u \cap \Lambda\Gamma_v$, so $\{\tilde{x}, \tilde{y}\}$ is the limit set of an edge incident at v . Therefore $\{x, y\} \subset \partial(\Gamma_v, \widehat{\text{Inc } v})$ is a single point, which is a contradiction because $\partial(\Gamma_v, \widehat{\text{Inc } v})$ was assumed not to contain a cut point. Hence the action of Γ_v on S does not fix any vertex and therefore gives rise to a non-trivial splitting of Γ_v relative to $\text{Inc } v$. □

4.3.3 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 Propositions 4.3 and 4.5.

Proposition 4.7. *Let Γ be a hyperbolic group and let \mathcal{H} be a finite set of virtually cyclic subgroups of Γ such that $\partial(\Gamma, \widehat{\mathcal{H}})$ is connected. Suppose that Γ admits a non-trivial splitting over a virtually cyclic subgroup relative to \mathcal{H} . Then $\partial(\Gamma, \widehat{\mathcal{H}})$ contains a cut point or pair.*

Proof. Let T be the Γ -tree associated to such a non-trivial splitting. Without loss of generality assume that the action of Γ 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 4.6 we obtain a partition of $\partial(\Gamma, \emptyset) - \Lambda\Gamma_e$ as $U_1 \sqcup U_2$ where U_i are open sets given by

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

If a subgroup H of Γ is in \mathcal{H} , $H \leq \Gamma_w$ for some vertex $w \in \Sigma$ and either $\Lambda H = \Lambda\Gamma_e$ or $\Lambda H \cap \Lambda\Gamma_e = \emptyset$. In the latter case either $\Lambda H \subset U_1$ or $\Lambda H \subset U_2$.

Let $\pi: \partial(\Gamma, \emptyset) \rightarrow \partial(\Gamma, \widehat{\mathcal{H}})$ be the quotient projection of Lemma 1.15. It follows that the images of U_1 and U_2 under π in the complement of $\pi(\Lambda\Gamma_e)$ in $\partial(\Gamma, \widehat{\mathcal{H}})$ are disjoint open sets; they are non-empty by the minimality of the action of Γ on T . The image of $\Lambda\Gamma_e$ is either one or two points, and therefore $\partial(\Gamma, \widehat{\mathcal{H}})$ contains a cut point or pair. \square

4.4 Virtually cyclic subgroups and finding splittings

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

Lemma 4.8. *[15, Lemma 2.8] There is an algorithm that, when given a presentation for a hyperbolic group Γ and a finite subset $S \subset \Gamma$, returns an answer to the question “is $\langle S \rangle \leq \Gamma$ virtually cyclic?” If the answer is “yes” then the algorithm also determines*

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_∞)).*
4. *a generating set for the maximal virtually cyclic subgroup of Γ containing $\langle S \rangle$.*

The proof of this lemma in [15] 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 paper.

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 [15]. First compute δ with respect to which Γ is δ -hyperbolic.

Use the algorithm of [26, Proposition 4] to search for a constant K with respect to which $\langle S \rangle$ is K -quasi-convex in Γ . 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. If the algorithm terminates use K and δ to compute δ' such that $\langle S \rangle$ is δ' -hyperbolic. Then all finite subgroups 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 S can be computed using a solution to the word and conjugacy problems in Γ . Once this is computed the algorithm of [15] can be used to determine whether or not $\langle S \rangle$ is virtually cyclic.

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

If $\langle S \rangle$ is not quasi-convex then it contains a free group on two generators, so a pair (g, h) as in the previous paragraph 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. \square

Lemma 4.9. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ and a collection \mathcal{H} of peripheral subgroups and either returns the graph of groups associated to a non-trivial splitting of Γ relative to \mathcal{H} over virtually cyclic subgroups or does not terminate if no such splitting exists.*

Proof. Suppose that Γ admits a proper splitting relative to \mathcal{H} as an amalgamated product over a virtually cyclic subgroup; the case of an HNN extension is similar. Then Γ admits a presentation that makes this splitting explicit in the following sense. There are finite disjoint symmetric sets of symbols S_1, S_2 and S_3 , finite subsets R_1, R_2 and R_3 of the free monoids S_1^*, S_2^* and S_3^* respectively, and maps ι_1 and ι_2 from S_3^* to S_1^* and S_2^* respectively, such that Γ admits an isomorphism to the group $|\mathcal{P}|$ with presentation \mathcal{P} of the form

$$\langle S_1 \cup S_2 \cup S_3 | R_1 \cup R_2 \cup R_3 \cup \{s^{-1}\iota_1(s), s^{-1}\iota_2(s) : s \in S_3\} \rangle$$

where $\{\iota_1(r) | r \in R_3\} \subset R_1$ and $\{\iota_2(r) | r \in R_3\} \subset R_2$. This ensures that ι_1 and ι_2 induce group homomorphisms from $\langle S_3 | R_3 \rangle$ to $\langle S_1 | R_1 \rangle$ and $\langle S_2 | R_2 \rangle$, which we shall denote $\widehat{\iota}_1$ and $\widehat{\iota}_2$. Let $\widehat{\iota}$ be the induced map from $\langle S_3 | R_3 \rangle$ to Γ . Then the assumptions on the nature of the splitting give the following conditions:

1. that $\langle S_3 | R_3 \rangle$ be virtually cyclic,
2. that $\widehat{\iota}_1$ and $\widehat{\iota}_2$ be injective,
3. that $\widehat{\iota}_1$ and $\widehat{\iota}_2$ not be surjective and
4. that for each $H \in \mathcal{H}$ with generating set S_H there exists g_H in Γ such that either S_1 or S_2 contains the image of $g_H S_H g_H^{-1}$ under the isomorphism from Γ to $|\mathcal{P}|$.

We show that there is an algorithm that finds such a presentation for Γ if it exists. Using Tietze transformations, there is an algorithm that enumerates all presentations of Γ and for each presentation gives an explicit isomorphism from Γ to the realisation of that presentation; therefore it is sufficient to show that each of the four conditions above can be checked algorithmically.

The first condition can be checked using the first part of the algorithm of Lemma 4.8.

$\widehat{\iota}_1$ and $\widehat{\iota}_2$ are both injective if and only if $\widehat{\iota}$ is injective: one direction is trivial, the other is a consequence of the normal form theorem for the amalgamated product. To check this condition, compute the maximal finite normal subgroup of $\langle S_3 | R_3 \rangle$ and check the triviality of the image under $\widehat{\iota}$ of each of these elements. Then find an element of $\langle S_3 | R_3 \rangle$ of infinite order and check whether or not the image under $\widehat{\iota}$ of this element has infinite order.

To check the surjectivity of $\widehat{\iota}_1$ first check whether or not $\langle S_1 \rangle$ is virtually cyclic. If it is, check whether or not the maximal finite normal subgroup of $\text{Im } \widehat{\iota}_1$ is equal to the maximal finite normal subgroup of $\langle S_1 \rangle$. If it is, next check whether or not $\langle S_1 \rangle$ and $\langle S_3 \rangle$ are either both of \mathcal{Z} -type or both of D_∞ type. If they are both of the same type, pass to an index 2 subgroup of each if necessary to ensure that they are both of \mathcal{Z} type, then check whether or not the composition of $\widehat{\iota}_1$ with the natural surjection to \mathbb{Z} is surjective. If it is then $\widehat{\iota}_1$ is surjective; if any of these tests produced the opposite answer then $\widehat{\iota}_1$ is not surjective. Repeat this process for $\widehat{\iota}_2$.




Figure 1: The decision process in the algorithm of Proposition 4.11.

The final condition can be checked using a solution to the simultaneous conjugacy problem in Γ .

The existence of such a presentation for Γ guarantees that Γ splits non-trivially as an internal amalgamated product

$$\Gamma \cong \langle S_1 | R_1 \rangle *_{\langle S_3 | R_3 \rangle} \langle S_2 | R_2 \rangle.$$

This splitting is over a virtually cyclic subgroup of Γ and is relative to \mathcal{H} . \square

4.5 Computing a maximal splitting

We will need to be able to determine algorithmically whether or not the boundary of the given hyperbolic group is homeomorphic to a circle. This is achieved using the algorithm of Corollary 2.7 and a theorem from point-set topology, which we note here.

Theorem 4.10. [35, II.2.13] *Any separable, connected, locally connected space containing more than one point that is without a cut point and in which every pair is a cut pair is homeomorphic to S^1 .*

Proposition 4.11. *There is an algorithm that takes as input a presentation for a hyperbolic group Γ with a collection \mathcal{H} of virtually cyclic subgroups such that $\partial(\Gamma, \widehat{\mathcal{H}})$ is connected and Γ appears as the stabiliser of some vertex in the action of a hyperbolic group on a tree and \mathcal{H} is a set of conjugacy class representatives of incident edge groups and returns the answer to the question “does Γ split non-trivially over a virtually cyclic subgroup relative to \mathcal{H} ?”*

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

Next compute δ such that the cusped space X associated to the pair $(\Gamma, \widehat{\mathcal{H}})$ is δ -hyperbolic. Search for a non-trivial splitting of Γ relative to \mathcal{H} using Lemma 4.9 and, in parallel, search for n such that \ddagger_n holds in X ; one of these processes must terminate by Proposition 1.16 and Proposition 4.3.

If a splitting is found then Γ 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 2.4 to check whether or not $\partial(\Gamma, \widehat{\mathcal{H}})$ contains a cut point. If it does then Γ does split properly over a virtually cyclic subgroup by Proposition 4.3

If there is no cut point, use the algorithm of Corollary 2.4 on X to determine whether or $\partial(\Gamma, \widehat{\mathcal{H}})$ contains a cut pair; if it does not then Γ does not split relative to \mathcal{H} by Proposition 4.7.

If there is a cut pair then simultaneously run the algorithms of Lemma 4.9 and Lemma 3.7. If the former terminates then a splitting has been found; if the latter does then no splitting exists. \square

Note that a subprocess of this algorithm, together with the algorithm of [20] that determines whether or not a hyperbolic group is one-ended, provides the algorithm promised in Theorem 0.2.

Proposition 4.12. *There is an algorithm that, when given a presentation for a hyperbolic group, computes the graph of groups associated to a splitting of that group that is maximal for domination.*

Proof. The algorithm iteratively constructs a sequence of Γ -marked graphs of groups G_i . Let G_1 consist of a single vertex with vertex group Γ . Then to obtain G_{i+1} from G_i , use Proposition 4.11 to check whether each vertex group splits non-trivially relative to its incident edge groups. If no vertex group does split then halt the algorithm here. If the group at a vertex v does split then find the non-trivial graph of groups G' with fundamental group Γ_v using Lemma 4.9. Then define G_{i+1} by replacing the vertex v of G_i by the graph G' and connecting edges corresponding to the edges of G_i incident at v to G' in the obvious way.

This process must eventually stabilise: this follows from an accessibility theorem [1] since the associated group actions on trees constructed are minimal and reduced by construction. Let the corresponding Γ -trees stabilise at a Γ -tree T .

Suppose that another Γ -tree T' with virtually cyclic edge stabilisers dominates T . Let v be a vertex of T . If e is an incident edge then Γ_e is elliptic with respect to the action on T' , since it contains the stabiliser of an edge of T' as a subgroup of finite index. Therefore each incident edge subgroup of Γ_v is elliptic with respect to T' . The vertex stabiliser Γ_v does not split over a virtually cyclic subgroup relative to its incident edge groups, so Γ_v is elliptic with respect to T' . The vertex v was arbitrary, so T dominates T' , and so T is maximal for domination. \square

5 JSJ Decompositions

In this section we show that three closely related types of JSJ splittings are computable for hyperbolic groups. Fix a one-ended hyperbolic group Γ and recall that \mathcal{VC} is defined to be the set of all virtually cyclic subgroups of Γ , \mathcal{Z} is defined to be the set of all virtually cyclic subgroups of Γ with infinite centre and \mathcal{Z}_{\max} is defined to be the set of subgroups of Γ in \mathcal{Z} that are maximal for inclusion. We consider the JSJ splittings over groups in these three sets.

5.1 Two-ended edge groups

We now prove the first part of Theorem 0.1.

Theorem 5.1. *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. This decomposition can be taken to be Bowditch's canonical decomposition.*

We first prove the following lemma.

Lemma 5.2. *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 Γ -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 $\Gamma_e \leq \Gamma_{e'}$, which is universally elliptic. But then the image of S in T' contains more than a single vertex. Therefore Γ_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 5.3. *An extended Möbius strip group is a virtually cyclic group of \mathcal{Z} type with peripheral structure consisting of a single index 2 subgroup.*

Definition 5.4. *We say that an edge e connecting vertices v_1 and v_2 of a Γ -tree is a internal surface edge if, for each i , either Γ_{v_i} is a hanging Fuchsian group and Γ_e is maximal among virtually cyclic subgroups of Γ_{v_i} , or Γ_{v_i} is an extended Möbius strip group and $\Gamma_e \leq \Gamma_{v_i}$ is the peripheral subgroup of Γ_{v_i} .*

Lemma 5.5. *If T is reduced (that is, no proper collapse of T dominates T) 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 5.2, 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 [23, 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 . The tree T was assumed to be reduced and v is a hanging Fuchsian vertex so 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 Γ_e is not elliptic with respect to T'' . \square

We now have sufficient tools to prove the computability of a \mathcal{VC} -JSJ for a given hyperbolic group G . In [8], Bowditch defines a canonical JSJ in the class of all \mathcal{VC} -JSJs of a given hyperbolic group. In the language of [23] this is the decomposition corresponding to the *tree of cylinders* of any other \mathcal{VC} -JSJ.

Definition 5.6. *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 Γ . Also denote by \sim the equivalence relation on the set of edges of T defined by letting $e \sim e'$ if and only if $\Gamma_e \sim \Gamma_{e'}$. A cylinder is a subset $Y \subset T$ that is the union of all edges in a \sim -equivalence class.*

Definition 5.7. *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 $v \in Y$.*

Proof of Theorem 5.1. First compute a maximal splitting of the group over virtually cyclic subgroups by Theorem 4.12. 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 4.8. For each edge e connecting vertices v_1 and v_2 of the graph of groups T/Γ determine whether Γ_e is maximal in Γ using the algorithm of Lemma 4.8 and whether the two vertex groups Γ_{v_1} and Γ_{v_2} have circular boundary relative to their incident edge groups by Theorem 0.2. Check also whether each of Γ_{v_1} and Γ_{v_2} is virtually cyclic of \mathcal{Z} -type, and, if it is, whether or not Γ_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; collapse all edges where this is the case.

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 4.8. This is result the content of [15, 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. \square

5.2 \mathcal{Z} edge groups

We now prove the second part of Theorem 0.1.

Theorem 5.8. *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 [15] 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 4.8 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 [15].

Definition 5.9. *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 Γ is any hyperbolic group with a collection \mathcal{H} of virtually cyclic subgroups such that $\partial(\Gamma, \widehat{\mathcal{H}})$ is homeomorphic to a circle then the mirrors splitting of Γ relative to \mathcal{H} is defined to be the splitting induced by the mirrors splitting of the quotient Γ by a maximal finite normal subgroup of Γ as in Proposition 3.3.

Lemma 5.10. *The mirrors splitting of a hyperbolic group Γ with a set \mathcal{H} of virtually cyclic subgroups such that $\partial(\Gamma, \widehat{\mathcal{H}})$ is homeomorphic to a circle is computable.*

Proof. Using Proposition 3.3 it is enough to show that the mirrors splitting is computable in the case where Γ is bounded Fuchsian and \mathcal{H} is the a collection of representatives of peripheral subgroups of Γ . 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 Γ and all homomorphisms from Γ 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

5.3 \mathcal{Z}_{\max} edge groups

In [15] it is shown that the \mathcal{Z}_{\max} -JSJ decomposition can be obtained from a \mathcal{Z} -JSJ decomposition by performing the so-called \mathcal{Z}_{\max} -fold. We note that this can be done algorithmically, which completes the proof of the final part of Theorem 0.1, which we restate here as Theorem 5.11.

Theorem 5.11. *There is an algorithm that takes as input a presentation for a one-ended hyperbolic group and outputs the graph of groups associated to a \mathcal{Z}_{\max} -JSJ decomposition for that group.*

Lemma 5.12. *There is an algorithm that takes as input a graph of groups decomposition of a hyperbolic group Γ over virtually cyclic subgroups and returns the graph of groups associated to the \mathcal{Z}_{\max} -fold of the associated Γ -tree.*

Proof. This can be done by iterating some simple folds described in [15]. We describe this fold at the level of the tree T . Take some edge e in T such that the maximal subgroup $\hat{\Gamma}_e$ of Γ containing Γ_e is a subgroup of $\Gamma_{o(e)}$ where $o(e)$ is the origin of e . Then take the quotient of T by the Γ -equivariant equivalence relation generated by $e \sim \hat{\Gamma}_e \cdot e$. Note that this does not change the underlying graph of the associated graph of groups.

At the level of the graph of groups this fold is achieved by taking an edge e in the graph such that $\hat{\Gamma}_e$ is a subgroup of $\Gamma_{o(e)}$, replacing the group at the edge e by $\hat{\Gamma}_e$ and replacing the group at the terminal vertex $t(e)$ of e by $\langle \hat{\Gamma}_e, \Gamma_{t(e)} \rangle$. Repeat this process until there is no edge e such that $\Gamma_e \neq \hat{\Gamma}_e$. \square

References

- [1] Mladen Bestvina and Mark Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.
- [2] Mladen Bestvina and Geoffrey Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.
- [3] B. H. Bowditch. Boundaries of geometrically finite groups. *Math. Z.*, 230(3):509–527, 1999.
- [4] B. H. Bowditch. Connectedness properties of limit sets. *Trans. Amer. Math. Soc.*, 351(9):3673–3686, 1999.

- [5] B. H. Bowditch. Convergence groups and configuration spaces. In *Geometric group theory down under (Canberra, 1996)*, pages 23–54. de Gruyter, Berlin, 1999.
- [6] B. H. Bowditch. Peripheral splittings of groups. *Trans. Amer. Math. Soc.*, 353(10):4057–4082, 2001.
- [7] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [8] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [9] 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.
- [10] Christopher H. Cashen and Nataša Macura. Line patterns in free groups. *Geom. Topol.*, 15(3):1419–1475, 2011.
- [11] Andrew Casson and Douglas Jungreis. Convergence groups and Seifert fibered 3-manifolds. *Invent. Math.*, 118(3):441–456, 1994.
- [12] F. Dahmani and N. Touikan. Isomorphisms using Dehn fillings: the splitting case. *ArXiv e-prints*, November 2013, 1311.3937.
- [13] François Dahmani. Finding relative hyperbolic structures. *Bull. Lond. Math. Soc.*, 40(3):395–404, 2008.
- [14] François Dahmani and Daniel Groves. Detecting free splittings in relatively hyperbolic groups. *Trans. Amer. Math. Soc.*, 360(12):6303–6318, 2008.
- [15] François Dahmani and Vincent Guirardel. The isomorphism problem for all hyperbolic groups. *Geom. Funct. Anal.*, 21(2):223–300, 2011.
- [16] M. J. Dunwoody and M. E. Sageev. JSJ-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.
- [17] J. Fox Manning. The Bowditch boundary of (G, \mathcal{H}) when G is hyperbolic. *ArXiv e-prints*, April 2015, 1504.03630.
- [18] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.
- [19] David Gabai. Convergence groups are Fuchsian groups. *Ann. of Math. (2)*, 136(3):447–510, 1992.
- [20] V. Gerasimov. Detecting connectedness of the boundary of a hyperbolic group. Unpublished.
- [21] Bradley W. Groff. Quasi-isometries, boundaries and JSJ-decompositions of relatively hyperbolic groups. *J. Topol. Anal.*, 5(4):451–475, 2013.
- [22] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.

- [23] V. Guirardel and G. Levitt. JSJ decompositions of groups. *ArXiv e-prints*, February 2016, 1602.05139.
- [24] 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.
- [25] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.
- [26] 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.
- [27] Gilbert Levitt. Automorphisms of hyperbolic groups and graphs of groups. *Geom. Dedicata*, 114:49–70, 2005.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] Peter Scott. The geometries of 3-manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983.
- [32] 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.
- [33] N. W. M. Touikan. Detecting geometric splittings in finitely presented groups. *ArXiv e-prints*, June 2009, 0906.3902.
- [34] Pekka Tukia. Homeomorphic conjugates of Fuchsian groups. *J. Reine Angew. Math.*, 391:1–54, 1988.
- [35] Raymond Louis Wilder. *Topology of Manifolds*. American Mathematical Society Colloquium Publications, vol. 32. American Mathematical Society, New York, N. Y., 1949.