

Deciding if a hyperbolic group splits over a given quasiconvex subgroup

Joseph Paul MacManus

Abstract. We present an algorithm which decides whether a given quasiconvex residually finite subgroup H of a hyperbolic group G is associated with a splitting. The methods developed also provide algorithms for computing the number of filtered ends $\tilde{e}(G, H)$ of H in G under certain hypotheses and give a new straightforward algorithm for computing the number of ends $e(G, H)$ of the Schreier graph of H . Our techniques extend those of Barrett via the use of labelled digraphs, the languages of which encode information on the connectivity of $\partial G - \Lambda H$.

Introduction

The study of decision problems within group theory is almost as old as the definition of an abstract group itself, dating back to Dehn’s classical word, conjugacy, and isomorphism problems. The classical theorems of Novikov–Boone [6, 29] and Adian–Rabin [1, 31] state that these problems – among many, many others – turn out to be undecidable in the class of finitely presented groups. Thus, if we wish to search for effective solutions to group theoretical problems, one must restrict their scope to some “nice” subclass of groups. One such subclass is the class of hyperbolic groups. Introduced by Gromov in his landmark essay [17], these are groups whose Cayley graphs possess geometric properties reminiscent of negative curvature. Indeed, within this class many problems become decidable. For example, the class of hyperbolic groups has uniformly solvable word and conjugacy problems [9, Ch. III.H], and more recently it was shown that one can distinguish isomorphism classes of hyperbolic groups [12, 37].

Another algorithmic problem which has received attention in recent years is that of splitting detection. In the language of Bass–Serre theory [38], a group G *splits* over a subgroup H if G admits a minimal simplicial action on a tree T without inversions, and H stabilises an edge in this action. This problem could be traced back to the algorithm of Jaco–Oertel [20] which decides if a given closed irreducible 3-manifold M is Haken, or equivalently if $\pi_1(M)$ splits over an infinite surface group. Splittings over finite subgroups are called *finite splittings*, and a celebrated theorem of Stallings [39, 40] states that a finitely generated group admits a finite splitting if and only if it has more than one

geometric end (cf. [36]). This provides a powerful link between the coarse geometry of a group and its splitting properties. Returning to the realm of hyperbolic groups, we have the following unpublished result due to Gerasimov¹.

Theorem (Gerasimov). *There is an algorithm which, upon input of a presentation of a hyperbolic group G , will compute the number of ends of G .*

In particular, one can effectively detect finite splittings of hyperbolic groups. This result was later generalised by Diao–Feighn to finite graphs of finitely generated free groups [13], by Dahmani–Groves [11] to relatively hyperbolic groups, and by Touikan [42] to finitely presented groups with a solvable word problem and no 2-torsion. Also worthy of mention is the algorithm by Jaco–Letscher–Rubinstein for computing the prime decomposition of a closed orientable 3-manifold [19], as well as the classical description of the Grushko decomposition of a one-relator groups [26, Prop. II.5.13].

Finite splittings aside, the next logical step is to detect splittings over two-ended (i.e., virtually cyclic) subgroups. This was achieved for (relatively) hyperbolic groups independently by Barrett [3] and Touikan [42] using quite distinct approaches. Note that Touikan’s algorithm here only applies in the torsion-free case.

Theorem (Barrett, Touikan). *There is an algorithm which, upon input of a presentation of a hyperbolic group G , will decide if G splits over a two-ended subgroup.*

The algorithm of Barrett, which is of particular interest to us, makes use of Bowditch’s deep theorem on two-ended splittings of hyperbolic groups [7]. This theorem states that a one-ended hyperbolic group G which is not virtually Fuchsian will admit such a splitting if and only if ∂G contains a cut pair. In fact, Barrett applies this result to effectively construct Bowditch’s canonical JSJ (Jaco–Shalen–Johansson) decomposition of a hyperbolic group.

In this paper, we will aim to expand on the techniques of [3] and apply them to larger splittings. We will restrict our attention to *quasiconvex* subgroups, that is, those subgroups whose inclusion maps are quasi-isometric embeddings, since distorted subgroups exhibit global geometry which is harder to understand on a local scale. If a group G splits over a subgroup commensurable with H , we say H is *associated with a splitting*. Finding sufficient conditions for a subgroup to be associated with a splitting is a problem which has received a great amount of interest (e.g., [27, 28, 33–35]). Applying the results of [27] to the setting of quasiconvex subgroups of hyperbolic groups, we are able to prove the following decidability result.

Theorem (cf. Corollary 4.8). *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasiconvex, residually finite subgroup H . This*

¹V. V. Gerasimov, Detecting connectedness of the boundary of a hyperbolic group. Unpublished, 1999.

algorithm will then decide if H is associated with a splitting, and will output such a splitting if one exists.

It is possible to somewhat weaken the residual finiteness assumption placed on H in the theorem above and give a more general (but more involved) result. We will postpone this more technical statement until Section 4.3 (see Theorems 4.6 and 4.7).

In light of Stallings' theorem, it is a natural generalisation to define the number of ends of a pair of groups (G, H) where $H \leq G$. This definition was first introduced by Houghton [18] and later explored in more depth in the context of discrete groups by Scott [32]. The number of ends of the pair (G, H) , denoted $e(G, H)$, can be identified with the number of geometric ends of the quotient of the Cayley graph of G by the left action of H . This quotient graph is sometimes called the *coset graph* or *Schreier graph* of H . It is not hard to show that if G splits over H then $e(G, H) \geq 2$, but the converse does not hold. Our methods give a new proof of the following theorem, originally due to Vonseel [43].

Theorem (Vonseel, cf. Theorem 3.11). *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasiconvex subgroup H , will output $e(G, H)$.*

There is a competing notion of “ends” of a pair of groups which goes by several names in the literature. This idea was considered independently by Bowditch [8], Kropholler–Roller [25], and Geoghegan [15], who refer to this invariant as *coends*, *relative ends*, and *filtered ends*, respectively. See [35, Ch. 2] for a discussion on the equivalence of these three definitions. In this paper, we will adopt the terminology and notation of Geoghegan and denote the number of filtered ends of the pair (G, H) by $\tilde{e}(G, H)$. This value appears to be more resilient to calculation without extra hypotheses, but nonetheless we have some partial results. Recall that the *generalised word problem* for a finitely generated group H is the problem of, given words $w_0, \dots, w_n$ in the generators of H , deciding whether $w_0 \in \langle w_1, \dots, w_n \rangle_H$. We then have the following statement.

Theorem (cf. Theorems 3.12 and 3.13). *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasiconvex subgroup H . This algorithm will terminate if and only if $\tilde{e}(G, H)$ is finite, and if it terminates, it will output the value of $\tilde{e}(G, H)$.*

Furthermore, if one is also given a solution to the generalised word problem for H , then there is an algorithm which decides whether $\tilde{e}(G, H) \geq N$ for any given $N \geq 0$.

We remark in Section 2.2 that $\tilde{e}(G, H)$ can be identified with the number of components of $\partial G - \Lambda H$. Thus, the above algorithm allows us to decide if $\partial G - \Lambda H$ is disconnected. It is also worth noting that if we know *a priori* that $\tilde{e}(G, H)$ is finite, for example, if H is two-ended, then $\tilde{e}(G, H)$ is fully computable. We will also see that the value of $\tilde{e}(G, H)$ is computable if H is free. It does not seem possible to decide in general if $\tilde{e}(G, H) = \infty$ using our machinery for an arbitrary quasiconvex subgroup, without assuming further hypotheses. We discuss this limitation in Section 3.4.

1. Preliminaries

In this section, we recall the basic notions and tools we require. We begin with a look at almost invariant sets and (filtered) ends of pairs, before briefly turning towards hyperbolic groups and their quasiconvex splittings. Throughout this paper we will assume a working knowledge of Bass–Serre theory, a good reference for which is [38].

1.1. Almost invariant subsets and (filtered) ends

We will need the idea of an almost invariant subset. A very good introduction to the upcoming definitions can be found in [33], which features many helpful examples. The reader should note, however, that this paper contains an error, a correction of which can be found in [35].

First, some notation. In what follows, G will be a finitely generated group and $H \leq G$ a finitely generated subgroup. If Z is a set upon which H acts on the left, then denote by $H \backslash Z$ the quotient of Z by this action.

Definition 1.1. Let U and V be two sets. Denote by $U \Delta V$ the *symmetric difference* of U and V , defined as

$$U \Delta V := (U - V) \cup (V - U).$$

We say that two sets U, V are *almost equal* if $U \Delta V$ is finite.

Definition 1.2. Let G acts on the right on a set Z . We say that $U \subset Z$ is *almost invariant* if for all $g \in G$, Ug is almost equal to U .

Definition 1.3. We say a subset $U \subset G$ is *H-finite* or *small*, if U projects to a finite subset of $H \backslash G$. If U is not *H-finite*, then we say U is *H-infinite*, or *large*.

To ease notation, given a subset $X \subset G$, we will write $X^* := G - X$.

Definition 1.4. We say that a subset $X \subset G$ is *H-almost invariant* if it is invariant under the left action of H , and $H \backslash X$ is almost invariant under the right action of G on $H \backslash G$. We say that X is *non-trivial* if both X and X^* are *H-infinite*.

Let X and Y be two non-trivial *H-almost invariant* subsets of G . We say that X and Y are *equivalent*, if $X \Delta Y$ is *H-finite*.

Definition 1.5. Let X be an *H-almost invariant* subset. Given $g \in G$, we say that gX *crosses* X if all of

$$gX \cap X, \quad gX \cap X^*, \quad gX^* \cap X, \quad gX^* \cap X^*$$

are large. If there exists $g \in G$ such that gX crosses X , then we say that X *crosses itself*. If X does not cross itself, we say it is *almost nested*. If one of the above intersections is empty, we say X is *nested*.

It is easy to see that if X and Y are equivalent H -almost invariant sets, then X crosses itself if and only if Y crosses itself.

Example 1.6. Suppose a group G splits as an amalgam or HNN extension over a subgroup H . Then one can construct a non-trivial nested H -almost invariant subset $X \subset G$ as follows. Let T be the Bass–Serre tree of this splitting. We now construct a G -equivariant map $\phi : G \rightarrow VT$, where VT denotes the set of vertices of T . Given $1 \in G$, set $\phi(1)$ arbitrarily to some $w \in VT$. For each $g \in G$, set $\phi(g) = g\phi(1)$. Since G acts upon itself freely and transitively, ϕ is well defined for all $g \in G$.

Given ϕ as above, let $e \in ET$ be the edge stabilised by H with endpoints u, v . Deleting the interior of e separates T into two components, T_u and T_v containing u and v , respectively. Set $X = \phi^{-1}(VT_u)$, then it is a simple exercise to check that X is a non-trivial nested H -almost invariant subset of G .

We can now state the following key theorem due to Scott–Swarup [34], which is in some sense a converse to Example 1.6. Recall that a subgroup H of G is said to be *associated with a splitting* if G splits over a subgroup commensurable with H .

Theorem 1.7 ([34, Thm. 2.8]). *Let G be a finitely generated group, H a finitely generated subgroup, and X an H -almost invariant subset of G . Suppose that X is almost nested, then H is associated with a splitting.*

There is a generalisation of the above, which will be important to us. Firstly, we must further loosen our requirements for nesting. Denote by $\text{Comm}_G(H)$ the commensurator of H in G . That is,

$$\text{Comm}_G(H) = \{g \in G : |H : H \cap H^g| < \infty, |H^g : H \cap H^g| < \infty\}.$$

Then, we have the following definition.

Definition 1.8. Let $X \subset G$ be H -almost invariant. We say that X is *semi-nested* if $\{g \in G : gX \text{ crosses } X\}$ is contained in $\text{Comm}_G(H)$.

Informally, we relax our definition to allow crossings of X by gX on the condition that gH is “very close” to H . We then have the following useful result, due to Niblo–Sageev–Scott–Swarup [27], which says that this relaxation still produces splittings.

Theorem 1.9 ([27, Thm. 4.2]). *Let G be a finitely generated group and H a finitely generated subgroup. Suppose that there exists a non-trivial H -almost invariant subset $X \subset G$ which is semi-nested. Then G splits over a subgroup commensurable with H .*

This idea of “crossings” of almost invariant sets is much more rich than what is presented here, and pertains to the idea of “compatible” splittings. The interested reader should consult [34], which features many helpful examples, as a starting point.

There is a natural way to “count” these H -almost invariant subsets, which provides a useful integer invariant of the subgroup. Let $\mathcal{P}(H \setminus G)$ denote the power set of $H \setminus G$. Let $\mathcal{F}(H \setminus G)$ denote the set of finite subsets. Under the operation of symmetric difference Δ , $\mathcal{P}(H \setminus G)$ can be seen as a $\mathbb{Z}_2$ -vector space, and $\mathcal{F}(H \setminus G)$ a subspace. The quotient space $\mathcal{E}(H \setminus G) := \mathcal{P}(H \setminus G) / \mathcal{F}(H \setminus G)$ can be identified naturally with the set of H -almost invariant sets of G , modulo equivalence.

Definition 1.10. Let G be a group and $H \leq G$. We define the *number of ends of the pair* (G, H) as the rank of $\mathcal{E}(H \setminus G)$ as a $\mathbb{Z}_2$ -vector space. Denote by $e(G) = e(G, \{1\})$ and say that $e(G)$ is the *number of ends of G* .

There is also the following characterisation of ends of pairs, which will be helpful later. This result motivates the earlier description that $e(G, H)$ “counts” H -almost invariant subsets.

Proposition 1.11 ([32, Lem. 1.6]). *Let G be a group, H a subgroup, and $n \geq 0$. Then $e(G, H) \geq n$ if and only if there exists a collection of n pairwise disjoint H -almost invariant subsets of G .*

There is a more geometric intuition for the above, which can be seen in the coset graph.

Proposition 1.12 ([32, Lem. 1.1]). *Let Γ be a Cayley graph of G , and $H \subset G$. Then $e(G, H)$ is equal to the number of ends of the coset graph $H \setminus \Gamma$.*

From the above it is then clear that this definition generalises the standard geometric notion of ends. This also provides a helpful way of seeing that the number of geometric ends of a finitely generated group is indeed independent of the choice of generating set.

There is another competing, but equally interesting notion of ends of a pair of groups, namely the idea of *filtered ends*. This was considered independently by Geoghegan [15, Ch. 14], Kropholler–Roller [25], and Bowditch [8]. We now summarise the definition as it appears in [15]. We first need the following technical preliminaries relating to filtrations.

Let Y be a connected, locally finite cell complex. A *filtration* $\mathcal{K} = \{K_i\}$ of Y is an ascending sequence of subcomplexes $K_1 \subset K_2 \subset \cdots \subset Y$ such that $\bigcup_i K_i = Y$. We say that this filtration is *finite* if each K_i is finite. We call the pair $(Y, \mathcal{K})$ a *filtered complex*, and a map $f : (Y, \mathcal{K}) \rightarrow (X, \mathcal{L})$ between filtered complexes is called a *filtered map* if the following four conditions hold:

- (1) $\forall i, \exists j$ such that $f(K_j) \subset L_i$,
- (2) $\forall i, \exists j$ such that $f(K_i) \subset L_j$,
- (3) $\forall i, \exists j$ such that $f(Y - K_j) \subset X - L_i$,
- (4) $\forall i, \exists j$ such that $f(Y - K_i) \subset X - L_j$.

We say that a homotopy H_t between two filtered maps is a *filtered homotopy* if H_t is a filtered map for each t . Fix a basepoint $b \in Y$, then a *filtered ray based at b* is a map

$\gamma : [0, \infty) \rightarrow Y$ with $\gamma(0) = b$, which is filtered with respect to the filtration $\{[0, i] : i \in \mathbb{N}\}$ of $[0, \infty)$. We say that two filtered rays based at b are *equivalent* if there is a filtered homotopy between them fixing b . A *filtered end* of $(Y, \mathcal{K})$ is an equivalence class of filtered rays based at b . It is easy to see that the choice of b does not affect the number of filtered ends.

Let G be a finitely generated group and H a finitely generated subgroup. Let X_G be a Cayley complex for G , and let $p : X_G \rightarrow H \backslash X_G$ denote quotient map. Choose a finite filtration $\mathcal{K} = \{K_i\}$ of $H \backslash X_G$, and lift this to a filtration $\mathcal{L} = \{L_i\}$ of X_G , where $L_i = p^{-1}(K_i)$.

Definition 1.13. The *number of filtered ends of the pair (G, H)* , denoted by $\tilde{e}(G, H)$, is defined as the number of filtered ends of the filtered complex $(X_G, \mathcal{L})$.

Though it seems as though this definition depends on a choice of the filtration $\mathcal{K}$, it is a helpful fact that it does not – see [15, Ch. 14] for details. While this definition is a little technical, we will see later on that, in the case of quasiconvex subgroups of hyperbolic groups, these filtered ends can be clarified by looking at the Gromov boundary. We conclude this section with two results relating ends of pairs and filtered ends of pairs.

Proposition 1.14 ([15, Prop. 14.5.3]). *Let G be a group and H be a finitely generated subgroup. Then $e(G, H) \leq \tilde{e}(G, H)$.*

Equality is certainly possible, but not true in general. A simple counterexample is presented in [35, p. 32], which involves a one-sided essential simple closed curve on a non-orientable surface. We finally conclude with the following, which demonstrates the common thread between the two competing definitions.

Proposition 1.15 ([35, Lem. 2.40]). *Let G be a group and H be a finitely generated subgroup of infinite index. Then $\tilde{e}(G, H) > 1$ if and only if there is some subgroup $K \leq H$ such that $e(G, K) > 1$.*

1.2. Hyperbolic groups

Let $\delta \geq 0$, and recall that a geodesic metric space X is δ -hyperbolic if every geodesic triangle in X is δ -slim. That is, if every side of a geodesic triangle is contained in the δ -neighbourhood of the other two sides. We say that a finitely generated group G is *hyperbolic* if some (equivalently, any) Cayley graph of G is δ -hyperbolic for some $\delta \geq 0$.

For the remainder of this section, let Γ be a δ -hyperbolic Cayley graph of the finitely generated group G . We recall some basic facts about hyperbolic groups which will be important later. The first is often known as the *visibility* property.

Lemma 1.16. *For every $g, g' \in \Gamma$, there is a geodesic ray γ based at g' which passes within $C = 3\delta$ of g .*

Viewing geodesic rays as “lines of sight” from the basepoint, we can imagine this theorem as saying that every point in Γ is (nearly) visible from every other point.

The second result we need is the following, commonly known as the *Morse Lemma*. This fact will play a fundamental role in much of the machinery of this paper. Informally, it says that quasi-geodesics must stay uniformly close to geodesics with the same endpoints.

Lemma 1.17. *Let $\lambda \geq 1$, $\varepsilon \geq 0$. Then there is some computable $D = D(\delta, \lambda, \varepsilon) \geq 0$ such that (λ, ε) -quasi-geodesics in Γ are contained within the D -neighbourhood of any geodesic with the same endpoints.*

We now briefly introduce the Gromov boundary of a hyperbolic group. Consider the set $\mathcal{R} = \mathcal{R}(\Gamma, 1)$ of geodesic rays in Γ based at 1. Throughout this paper, we may abuse notation and identify a ray or path with its image in Γ . We say that two such rays $\gamma, \gamma' \in \mathcal{R}$ are *equivalent* and write $\gamma \sim \gamma'$, if their Hausdorff distance as subsets of Γ is finite. Let $\partial\Gamma = \mathcal{R} / \sim$, and call this set the *Gromov boundary* of Γ . Given $\gamma \in \mathcal{R}$, denote by $\gamma(\infty)$ its equivalence class in $\partial\Gamma$. We topologise $\partial\Gamma$ as follows. Given $x, y, z \in \Gamma$, let

$$(x \cdot y)_z = \frac{1}{2}(d(x, z) + d(y, z) - d(x, z))$$

denote the *Gromov product*. For $p \in \partial\Gamma$, $r \geq 0$, let

$$V(p, r) = \{q \in \partial\Gamma : \exists \gamma \in p, \gamma' \in q \text{ such that } \liminf_{t \rightarrow \infty} (\gamma(t) \cdot \gamma'(t))_1 \geq r\}.$$

Given $p \in \partial\Gamma$, we set $\{V(p, r) : r \geq 0\}$ as a basis of neighbourhoods about p . This topologises $\partial\Gamma$. Informally, two rays are “close” in $\partial\Gamma$ if they fellow-travel for a long time. It can be seen that, up to homeomorphism, $\partial\Gamma$ does not depend on a choice of basepoint, or even on the choice of Cayley graph Γ . Thus, we may abuse notation and write ∂G .

One can also define the *sequential boundary* of Γ . Fix a basepoint $z \in \Gamma$. We say that a sequence $(p_n) \in \Gamma^{\mathbb{N}}$ *tends to infinity* if $\liminf_{n, m \rightarrow \infty} (p_n \cdot p_m)_z = \infty$. We say two such sequences $(p_n), (q_m)$ are *equivalent* if $\liminf_{n, m \rightarrow \infty} (p_n \cdot q_m)_z = \infty$. One can then define $\partial\Gamma$ as the set of equivalence classes of sequences which tend to infinity, and topologise it in a similar way. Given a sequence (p_n) which tends to infinity, denote by $\lim_{n \rightarrow \infty} p_n$ its equivalence class in $\partial\Gamma$. This definition is equivalent to the previous [22, Prop. 2.14]. Other equivalent variations of this construction are possible and can be found in [22].

We can extend the definition of the Gromov product to $\partial\Gamma$, by defining

$$(p \cdot q)_z = \sup \left\{ \liminf_{n, m \rightarrow \infty} (p_n \cdot q_m)_z : (p_n) \in p, (q_m) \in q \right\},$$

for $p, q \in \partial\Gamma$, $z \in \Gamma$. Using this extension there is a natural way to put a metric on ∂G , though the metric is not canonical, as it depends on some chosen parameters. In particular, a *visual metric* on $\partial\Gamma$ with parameter a and multiplicative constants k_1, k_2 is a metric ρ satisfying

$$k_1 a^{-(p \cdot q)_1} \leq \rho(p, q) \leq k_2 a^{-(p \cdot q)_1}$$

for every $p, q \in \partial\Gamma$. See [9, p. 434] for how to construct such a metric given suitable parameters.

Note that the number of connected components of the boundary can be identified with the number of ends $e(G)$ of G . In particular, the boundary $\partial\Gamma$ is connected if and only if G is one-ended. Thus, by Stallings' theorem for ends of groups (see, e.g., [36]), we have that ∂G is connected if and only if G admits no finite splitting.

We conclude this section with an important and deep result due to Bestvina–Mess [5], Bowditch [7], and Swarup [41]. This result essentially quantifies the connectedness of the boundary and relates it to the local geometry of the Cayley graph. Let $C = 3\delta$ as in Lemma 1.16, and let $M = 6C + 2\delta + 3$. We then have the following.

Definition 1.18. We say that Γ satisfies $\ddagger_n$ for $n \geq 1$ if for every $R \geq 0$ and every $x, y \in \Gamma$ such that $d(x, 1) = d(y, 1) = R$ and $d(x, y) \leq M$ we have that there is a path through $\Gamma - B_1(R)$ connecting x to y , of length at most n .

Theorem 1.19 ([5, Props. 3.1 and 3.2]). *Let G be a hyperbolic group, with δ -hyperbolic Cayley graph Γ . Then ∂G is connected if and only if there exists $n \geq 1$ such that Γ satisfies $\ddagger_n$.*

Note that through the algorithm presented in [11], one can decide if a given hyperbolic group G is one-ended, and if so this algorithm will output n such that $\ddagger_n$ holds in the Cayley graph associated with the given generators of G . Their algorithm also applies to relatively hyperbolic groups, but we will not use that here.

1.3. Quasiconvex subgroups and splittings

Recall that a subset of a geodesic space is termed *convex* if it contains any geodesic between any two of its points. This notion is too precise for the setting of groups, so we must “quasify” it.

Let $Q \geq 0$, and let G be a hyperbolic group with Cayley graph Γ . Then a subgroup $H \leq G$ is called *Q -quasiconvex* if for every $h, h' \in H$, any geodesic path between h, h' in Γ is contained in the closed Q -neighbourhood of H . The quasiconvexity of H does not depend on the choice of Cayley graph for G , given that G is hyperbolic, and quasiconvex subgroups of hyperbolic groups are also hyperbolic. The following characterisation of quasiconvexity will be important.

Lemma 1.20. *Let G be a hyperbolic group, $H \leq G$ a finitely generated subgroup, and fix word metrics on these groups. Then H is quasiconvex if and only if the inclusion map $H \hookrightarrow G$ is a (λ, ε) -quasi-isometric embedding for some $\lambda \geq 1$, $\varepsilon \geq 0$.*

Moreover, given a presentation of G and generators of H , if H is quasiconvex then these constants can be found algorithmically.

The above follows immediately from Kapovich's algorithm [21] for computing the quasiconvexity constant Q , together with [9, Ch. III.Γ, Lem. 3.5].

Quasiconvex subgroups interact with the Gromov boundary in a very controlled way. Informally, if $H \leq G$ is a quasiconvex subgroup, then this induces an inclusion $\partial H \hookrightarrow \partial G$. More precisely, we need the idea of “limit sets” in the boundary.

Definition 1.21. Let G be a hyperbolic group with Cayley graph Γ , and H some subgroup of G . Then define the *limit set* of H as

$$\Lambda H = \{p \in \partial \Gamma : p = \lim_{n \rightarrow \infty} h_n \text{ for some } h_n \in H\}.$$

It is a standard fact that if H is quasiconvex, then $\Lambda H \cong \partial H$, and ΛH is a closed H -invariant subset of ∂G (see, e.g., [22]). Combining Lemmas 1.16 and 1.17 we immediately arrive at the following, which will play a big role in the next section.

Lemma 1.22. *Let Γ be a δ -hyperbolic Cayley graph of a hyperbolic group G , and let H be a Q -quasiconvex subgroup. There exists some computable constant $\eta = \eta(\delta, Q) > 0$ with the following property. If γ is a geodesic ray in Γ based at 1 such that $\gamma(\infty) \in \Lambda H$, then γ is contained in the closed η -neighbourhood of H .*

Another important fact about quasiconvex subgroups of hyperbolic groups is that they are “almost” self-commensurating. More formally we have the following classical result due to Arzhantseva [2] and Kapovich–Short [24].

Theorem 1.23 ([2, Thm. 2]). *Let G be a hyperbolic group and H a quasiconvex subgroup. Then H has finite index in $\text{Comm}_G(H)$.*

Corollary 1.24. *Let H be a quasiconvex subgroup of a hyperbolic group G . Then the commensurator $\text{Comm}_G(H)$ is the unique maximal finite index overgroup of H .*

Proof. Let $K \leq G$ be some finite index overgroup of H , and let $k \in K$. Clearly $H^k \subset K$, so $H^k \cap H \leq H \cap K = H$ and H^k has finite index in K . Then,

$$|K : H \cap H^k| = |K : H| \cdot |H : H \cap H^k|,$$

so $|H : H \cap H^k| < \infty$. Similarly, $|H^k : H \cap H^k| < \infty$, and so $k \in \text{Comm}_G(H)$. ■

We conclude this section with a description of the vertex groups of a splitting in which the edge groups are quasiconvex. We say that a subgroup is *full* if it does not have any finite index overgroups.

Proposition 1.25 ([7, Prop. 1.2]). *Let G be a one-ended hyperbolic group, and suppose that G splits over a quasiconvex subgroup as an amalgam or an HNN extension. Then the vertex groups of this splitting are full quasiconvex.*

Proof. The fact that they are quasiconvex is proven in [7], so we will just show here that they are necessarily also full.

Let G acts simplicially, minimally, and without inversions on a tree T with a single-edge orbit, such that the edge stabilisers are quasiconvex. Let v be a vertex of T , and let G_v denote the stabiliser of v . Suppose that G_v is not full, so there is finite index overgroup $K > G_v$. Then the K -orbit of v is finite, so there is some finite subtree $\Sigma \subset T$ stabilised by K . Let u be the geometric centre of Σ . If u is a midpoint of an edge e , then it must be the case that the action of K inverts this edge. We assumed that G acts without inversions so this cannot happen, and so u is a vertex of T . Since u is distinct from v (else $K = G_v$), we see that there is a vertex w adjacent to v such that

$$G_v \leq G_w \leq K.$$

In particular, G_v has finite index in G_w .

Let e be the edge connecting v to w . We split into two cases. Firstly, suppose that this splitting is an amalgam. Then $G_e = G_v \cap G_w = G_v$, but then the given splitting is of the form $G = G_w *_{G_v} G_v$, and so $G_w = G$ and G fixes a point on T .

Suppose instead that G acts on T with a single vertex orbit, that is, that this splitting is an HNN extension. Then there is some $g \in G$ such that $G_w = G_v^g$. Suppose that $G_w > G_v$, then we find the following sequence

$$G_v < G_v^g < G_v^{g^2} < G_v^{g^3} < G_v^{g^4} < \dots,$$

where each has finite index in the last. But then G_v has no maximal finite index overgroup, which contradicts Corollary 1.24 since G_v is quasiconvex. Thus, $G_w = G_v = G_e$. In particular, by translating this picture around we see that the stabiliser of every edge in T is equal to the stabiliser of either of its endpoints. Since T is connected, it follows that G_v actually fixes T pointwise and is in fact the kernel of the action of G on T . Thus, G_v is a normal subgroup of G . Since G_v is quasiconvex, this can only happen if G_v is finite or has finite index in G [24]. By assumption, G is one-ended and so G_v has finite index in G . But then G must act on T with a global fixed point. With this contradiction, we conclude that G_v must be full. ■

The following definition will be helpful.

Definition 1.26 (Lonely subgroups). Let G be a group, then a subgroup $H \leq G$ is said to be *lonely* if there does not exist another subgroup H' distinct from H such that H is commensurable with H' .

We then have the following dichotomy, which will play a central role in our algorithmic results.

Proposition 1.27. *Let G be a hyperbolic group and H a quasiconvex subgroup such that either*

- (1) $H \neq \text{Comm}_G(H)$, or
- (2) H is lonely.

Then H is associated with a splitting if and only if there is some semi-nested H -almost invariant set X .

Proof. Applying Theorem 1.9, it is clear that if we have a semi-nested H -almost invariant set, then H is associated with a splitting.

Conversely, suppose that G splits over a subgroup H' commensurable with H . If H is lonely, then the result is clear as we necessarily split over H , thus there is a nested H -almost invariant set as in Example 1.6. So assume instead that $H \neq \text{Comm}_G(H)$. We now split into two cases.

Firstly, suppose that H' has infinite index in the vertex group(s) of its splitting. We have by Theorem 1.23 that $\text{Comm}_G(H)$ is commensurable with H' , so it will be strictly contained in a vertex group of this splitting, up to conjugacy. We can then transform our splitting over H' easily into a splitting over $C := \text{Comm}_G(H)$ via, for example,

$$G = A *_H B = A *_H (C *_C B) = (A *_H C) *_C B.$$

The HNN case is similar. Since we have a non-trivial splitting over C , there is an almost nested C -almost invariant set by Example 1.6. Clearly such a set is also H -almost invariant, and so we are done.

Suppose instead that H' has finite index in one of the vertex stabilisers. By Proposition 1.25, we know that the vertex groups of the splitting over H' are full, so by Corollary 1.24 we deduce that one vertex group of the splitting over H' is precisely $\text{Comm}_G(H)$. Let T be the Bass–Serre tree of this splitting. Then H fixes a vertex of T , namely the vertex stabilised by $\text{Comm}_G(H)$. Let $e \in ET$ be the edge stabilised by H' with endpoints $v, u \in VT$, and suppose without loss of generality that H fixes v . Then the H -orbit of e is a finite collection of edges abutting v . Since $H \neq \text{Comm}_G(H)$, there are multiple H -orbits of edges abutting v . Let $F = \{he : h \in H\}$. Then deleting the interior of the edges in F separates T into a disjoint collection of subtrees. Let T_v be the subtree which contains v , then set $X = \phi^{-1}(VT_v)$, where $\phi : G \rightarrow T$ is the map constructed in Example 1.6. It is clear that X is a non-trivial H -almost invariant set, and

$$\{g \in G : gX \text{ crosses } X\} \subset G_v = \text{Comm}_G(H),$$

where $G_v \leq G$ is the stabiliser of v . Thus, X is semi-nested by definition. ■

Remark 1.28. It might seem strange that in the hypothesis of Proposition 1.27 we ask that H be distinct from its commensurator. This rules out, for example, maximal cyclic subgroups (which are often associated with splittings). This exclusion can be explained by an example.

Suppose our group G splits as an amalgam $G = A *_C B$ where $|B : C| < \infty$. Then the vertex group B is associated with a splitting, by definition. However, it is certainly not

clear if G contains some semi-nested B -invariant subset (it is helpful to check where the construction given in the above proof fails in this case). This technicality can be overcome by replacing our given quasiconvex subgroup with one of its finite index subgroups (if one exists), which explains why it is helpful to assume that our given subgroup is residually finite.

2. Limit set complements

The key aim of this section is to understand the connectivity of $\partial G - \Lambda H$ via local geometry. In particular, we will generalise a toolbox introduced by Barrett in [3].

2.1. Annular neighbourhoods

The Gromov boundary encodes a wide variety of data about our group, but we cannot access it directly. Thus, we must characterise the presence of the topological features we care about via some kind of equivalent local geometric feature in the Cayley graph.

For what follows, we will need to fix some notation. Throughout this section G will be a one-ended hyperbolic group with δ -hyperbolic Cayley graph Γ , and H will be a Q -quasiconvex subgroup with a fixed finite generating set Y . Note that given a presentation of a hyperbolic group, one can effectively compute the hyperbolicity constant δ of the corresponding Cayley graph via [30]. Given $x, y \in \Gamma$, we may denote by $[x, y]$ a choice of geodesic path between x and y . Let $C = 3\delta$ as in Lemma 1.16, $D = D(1, 0)$ as in Lemma 1.17. Fix a visual metric ρ on ∂G , with parameter a and multiplicative constants k_1, k_2 . We will make use of some methods from [5], so recall Theorem 1.19 and fix $n \geq 0$ such that $\ddagger_n$ holds in Γ . Finally, choose $\lambda \geq 1, \varepsilon \geq 0$ such that the inclusion map $H \hookrightarrow G$ is a (λ, ε) -quasi-isometric embedding as in Lemma 1.20, and $\eta = \eta(\delta, Q)$ as in Lemma 1.22.

Lemma 2.1. *Let $p \in \partial G - \Lambda H$. Then for $\gamma \in p$, we have that*

$$\lim_{t \rightarrow \infty} d(\gamma(t), H) = \infty.$$

Proof. Suppose not, so $\liminf_{t \rightarrow \infty} d(\gamma(t), H) < \infty$. Then there exists a sequence $t_i \rightarrow \infty$, and some $A \geq 0$ such that $d_G(\gamma(t_i), H) < A$ for all i . Let $h_i \in H$ be such that $d_G(\gamma(t_i), h_i) < A$.

Recall that we fixed a generating set Y of H , and thus a word metric on H . Let δ_H be such that H is δ_H -hyperbolic. Choose geodesic rays μ_i in H based at 1 such that μ_i passes within $3\delta_H$ of h_i . Let $\iota : H \hookrightarrow G$ be the inclusion map, which is a (λ, ε) -quasi-isometric embedding. Then $\iota(\mu_i)$ is a (λ, ε) -quasi-geodesic ray at 1. Let v_i be a geodesic ray in Γ with the same endpoints as $\iota(\mu_i)$. Then,

$$d(v_i, \gamma(t_i)) < A + 3\lambda\delta_H + \varepsilon + D(\lambda, \varepsilon),$$

for every i . It is now routine to see that $v_i(\infty) \rightarrow p$ in ∂G . But ΛH is a closed subset of ∂G , so $p \in \Lambda H$. This is a contradiction, so the lemma follows. $\blacksquare$

We now introduce some notation borrowed from [3]. Let $0 \leq r \leq K \leq R$ be constants, where R is possibly infinite. Denote by $N_R(H)$ the closed R -neighbourhood of H in Γ , and by $N_{r,R}(H)$ the annular region

$$N_{r,R}(H) = \{x \in \Gamma : r \leq d(x, H) \leq R\}.$$

Let $C_K(H) = \{x \in \Gamma : d(x, H) = K\}$, and finally let $A_{r,R,K}(H)$ be the union of the components of $N_{r,R}(H)$ which intersect $C_K(H)$. The geometric significance of $A_{r,R,K}(H)$ is that we take an annular region about H , but discard the components which are “too close” to H . This is illustrated in Figure 1. Following Barrett, we will demonstrate a correspondence between the components of $\partial G - \Lambda H$ and the components of $A_{r,R,K}(H)$, provided r , R , and K are sufficiently large.

Next, we introduce “shadows.” Given a component U of $A_{r,\infty,K}(H)$, define its *shadow* $\mathcal{S}U$ as the set of points $p \in \partial G$ such that for any ray $\gamma \in p$ based at 1, we have that $\gamma(t)$ is in U for all sufficiently large t . We also extend the definition of a shadow to a union of components $V = \bigcup_i U_i$ by defining $\mathcal{S}V = \bigcup_i \mathcal{S}U_i$.

We now prove five lemmas, generalised from [3]. Originally, these results related to bi-infinite geodesics, but careful inspection of their proofs reveals that we only really need the fact that a geodesic path in a hyperbolic space is a quasiconvex set. This allows us to generalise some of the machinery in this paper from geodesics to quasiconvex subgroups. Apart from some small changes, the following proofs are practically reproduced verbatim from [3].

Lemma 2.2 ([3, Lem. 1.17]). *For $r > \eta$, $K \geq r$, we have that*

$$\partial G - \Lambda H = \bigcup_U \mathcal{S}U,$$




Figure 1. The region $A_{r,R,K}(H)$ surrounding H . In this cartoon, H is depicted as a line for illustrative purposes.

where U ranges over the connected components of $A_{r,\infty,K}(H)$. Moreover, $\mathcal{S}U \cap \mathcal{S}V = \emptyset$ for distinct components U, V .

Proof. It is clear from the definition of $\mathcal{S}U$ that distinct components have disjoint shadows. Since $r > \eta$, we have by Lemma 1.22 that $\mathcal{S}U$ and ΛH are disjoint. Now let $p \in \partial G - \Lambda H$. We need to check that $p \in \mathcal{S}U$ for some component U of $A_{r,\infty,K}(H)$. Let $\gamma \in p$, then there is some t_0 such that for $t \geq t_0$ we have that $d(\gamma(t), H) \geq r + D$.

Let U be the component of $A_{r,\infty,K}(H)$ containing $\gamma(t)$ for $t \geq t_0$. We claim that if γ' is another ray in p , then $\gamma'(t')$ lies in U for sufficiently large t' . We know that the Hausdorff distance between γ and γ' is at most D . In particular, for every t' there is some t such that $d(\gamma(t), \gamma'(t')) \leq D$, and by the triangle inequality t and t' satisfy $|t - t'| \leq D$. Then if $t' \geq t_0 + D$, we must have that $t \geq t_0$, and so the segment between $\gamma'(t')$ and $\gamma(t)$ lies in $A_{r,\infty,K}(H)$, and thus $\gamma'(t') \in U$. It follows $p \in \mathcal{S}U$. ■

Lemma 2.3 ([3, Lem. 1.18]). *Let $r > \eta$ and $K \geq r + Q + \delta + C$. Then for every component U of $A_{r,\infty,K}(H)$, we have that $\mathcal{S}U$ is non-empty.*

Proof. Let $x \in C_K(H) \cap U$, and using Lemma 1.16 choose a geodesic ray γ based at 1 which passes within C of x , say $d(\gamma(t), x) \leq C$. Since $K > r + C$, we have that $\gamma(t) \in U$. In particular,

$$d(\gamma(t), H) > r + Q + \delta.$$

Suppose now that for some $t' \geq t$, we have that $d(\gamma(t'), H) \leq r$. Let $y \in H$ be a nearest point projection of $\gamma(t')$. Then consider the geodesic triangle $[1, y, \gamma(t')]$, where the segment $[1, \gamma(t')]$ is precisely an initial segment of γ . Then $\gamma(t)$ is either δ -close to $[y, \gamma(t')]$ or $[1, y]$. In any case,

$$d(\gamma(t), H) \leq r + Q + \delta,$$

which is a contradiction. Thus, $\gamma(t') \in U$ for every $t' \geq t$. It now follows by an argument similar to that in the proof of Lemma 2.2 that $\gamma(\infty) \in \mathcal{S}U$. ■

Lemma 2.4 ([3, Lem. 1.19]). *Let $r > \eta$, $K > r$, then for every component U of $A_{r,\infty,K}(H)$, we have that $\mathcal{S}U$ is open and closed in $\partial G - \Lambda H$.*

Proof. Let $\gamma \in p \in \mathcal{S}U$, and for $t \geq 0$ let

$$V_t(\gamma) = \{\beta(\infty) : \beta \text{ is a ray based at } 1, d(\gamma(t), \beta(t)) < 2\delta + 1\}.$$

As t varies, this forms a basis of neighbourhoods about p . By Lemma 2.1, there exists a t_0 such that for all $t \geq t_0$, we have that

$$d(\gamma(t), H) > r + Q + 7\delta + 1.$$

We claim that $V_{t_0}(\gamma) \subseteq \mathcal{S}U$, which implies that $\mathcal{S}U$ is open.

Let $\beta \in q \in V_{t_0}(\gamma)$, so by definition we have $d(\beta(t_0), \gamma(t_0)) < 2\delta + 1$. Then consider another ray $\beta' \in q$ also, so $d(\beta(t_0), \beta'(t_0)) < 4\delta$. We need to show that $\beta'(t) \in U$ for all $t \geq t_0$. Suppose this is not the case. We have that $\beta'(t_0)$ is in U , so for this not to hold we must have that β' “leaves” U at some point. Thus for some $t' \geq t_0$, we have that $d(\beta'(t'), H) \leq r$, say for some $y \in H$ we have $d(\beta'(t'), y) \leq r$.

Consider a geodesic path $[1, y]$. Since H is Q -quasiconvex, this path is contained in the Q -neighbourhood of H . Combine this with the fact that the triangle $[1, \beta'(t'), y]$ is δ -slim, then inspection of Figure 2 reveals that

$$d(\gamma(t_0), H) \leq r + Q + 7\delta + 1.$$

This is a contradiction, so $\mathcal{S}U$ is open.

To conclude, note that the shadows $\mathcal{S}U$ form a disjoint open cover of $\partial G - \Lambda H$. It follows that they must also be closed. $\blacksquare$

For the next lemma, recall that a_1, a_2 , and k denote the parameters of our fixed visual metric on ∂G . Also recall that n has been fixed so that $\frac{1}{3}n$ holds in Γ .

Lemma 2.5 ([3, Lem. 1.20]). *Suppose that r satisfies*

$$r > 2 \log_a \left(\frac{k_2(n-1)}{k_1(1-a^{-1})} \right) + M + 8\delta + \eta + C,$$

and let $K > r$. Then for every component U of $A_{r,\infty,K}(H)$, $\mathcal{S}U$ is contained within exactly one connected component of $\partial G - \Lambda H$. Moreover, every component of $\partial G - \Lambda H$ is path connected.

Proof. Let $p, q \in \mathcal{S}U$, and let $\alpha_1 \in p, \alpha_2 \in q$. Then there is some t_1, t_2 such that $\alpha_1(t_1), \alpha_2(t_2) \in U$. Let $\phi : [0, \ell] \rightarrow U$ be a path connecting $\alpha_1(t_1)$ to $\alpha_2(t_2)$, parameterised by arc length.




Figure 2. Bounding the distance between $\gamma(t_0)$ and H .

We follow the methodology of Bestvina–Mess [5] and “project” this path to the boundary. For every $i \in \mathbb{Z} \cap [0, \ell]$, let β_i be a ray from 1, passing through a point $\beta_i(m_i) = z_i$ such that $d(z_i, \phi(i)) \leq C$. We will show that $\beta_i(\infty)$ and $\beta_{i+1}(\infty)$ can be connected by a path in ∂G avoiding ΛH . This implies the result. Note that each $\beta_i(\infty)$ is not in ΛH , since $r \geq \eta + C$.

For notational convenience, assume that $i = 0$. Let n be such that $\ddagger_n$ holds in Γ . For every n -adic $t \in [0, 1]$, by induction on the power k of the denominator, we define β_t , satisfying

$$d(\beta_{j/n^k}(m_i + k), \beta_{j+1/n^k}(m_i + k)) \leq M,$$

where $M = 6C + 2\delta + 3$, for each $0 \leq j < n^k$. Note that the base case of this induction holds because $M \geq 2C + 1$. Secondly, the triangle inequality then gives the following lower bound:

$$\begin{aligned} (\beta_{j/n^k}(\infty) \cdot \beta_{j+1/n^k}(\infty))_1 &\geq \liminf_{n_1, n_2} (\beta_{j/n^k}(n_1) \cdot \beta_{j+1/n^k}(n_2))_1 \\ &\geq (\beta_{j/n^k}(m_0 + k) \cdot \beta_{j+1/n^k}(m_0 + k))_1 \\ &= m_0 + k - M/2. \end{aligned}$$

We therefore deduce that

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

Inductively applying the triangle inequality and summing the geometric series, we thus obtain the bound that for every n -adic rational $t \in [0, 1]$, we have

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

We then define a path $\psi : [0, 1] \rightarrow \partial G$ by $\psi(t) = \beta_t(\infty)$, for every n -adic t , and extending to $[0, 1]$ continuously. We have shown that this path is contained within a ball of radius

$$\frac{k_2(n-1)a^{-m_0 + \frac{M}{2}}}{1 - a^{-1}} \quad (1)$$

around $\beta_0(\infty)$.

We now seek to place a lower bound on $\rho(\beta_0(\infty), \Lambda H)$ in terms of m_0 . Let γ be a ray from 1 with limit in ΛH . So γ is contained in the η -neighbourhood of H . By [9, Ch. III.H, 3.17], we have that

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

So let $n_1, n_2 \geq m_0$. We will show that $(\beta_0(n_1) \cdot \gamma(n_2))_1$ can be approximated by $(\beta_0(m_0) \cdot \gamma(m_0))_1$. Since $d(\beta_0(m_0), H) \geq r > \delta + \eta$, we have that $d(\beta_0(m_0), \gamma) > \delta$. So there exists some $a \in [\beta_0(n_1), \gamma(n_2)]$ such that $d(\beta_0(m_0), a) \leq \delta$. Similarly, we have

that $d(\gamma(m_0), \beta_0) > \delta$, as $r > 2\delta + \eta$. Thus, again there exists a point $b \in [\beta_0(n_1), \gamma(n_2)]$ such that $d(\gamma(m_0), b) \leq \delta$.

We now check that a and b are “in order.” Suppose that b is closer to $\beta_0(n_1)$ than to a . Then by considering the geodesic triangle $[a, \beta_0(m_0), \beta_0(n_1)]$, we see that $d(b, \beta_0) \leq 2\delta$. But this implies that $d(\beta_0(m_0), \gamma(m_0)) \leq 6\delta$, which is a contradiction since $r > 6\delta + \eta$. Therefore, the picture we have looks something like Figure 3.

We can thus conclude that

$$d(\beta_0(n_1), \gamma(n_2)) = d(\beta_0(n_1), a) + d(a, b) + d(b, \gamma'(n_2)).$$

We can then compute the following inequality:

$$\begin{aligned} & (\beta_0(n_1) \cdot \gamma(n_2))_1 - (\beta_0(m_0) \cdot \gamma(m_0))_1 \\ &= \frac{1}{2}[n_1 + n_2 - d(\beta_0(n_1), \gamma(n_2))] - \frac{1}{2}[2m_0 - d(\beta_0(m_0), \gamma(m_0))] \\ &= \frac{1}{2}[(n_1 - m_0) + (n_2 - m_0) + d(\beta_0(m_0), \gamma(m_0)) - d(\beta_0(n_1), \gamma(n_2))] \\ &= \frac{1}{2}[d(\beta_0(n_1), \beta_0(m_0)) - d(\beta_0(n_1), a) + d(\gamma(n_2), \gamma(m_0)) - d(\gamma(n_2), b) \\ &\quad + d(\beta_0(m_0), \gamma(m_0)) - d(a, b)] \\ &\leq \frac{1}{2}[\delta + \delta + 2\delta] = 2\delta. \end{aligned}$$

Applying the bound (2), it follows that

$$\begin{aligned} (\beta_0(\infty) \cdot \gamma(\infty))_1 &\leq (\beta_0(m_0) \cdot \gamma(m_0))_1 + 4\delta \\ &= \frac{1}{2}[2m_0 - d(\beta_0(m_0), \gamma(m_0))] + 4\delta \\ &\leq m_0 - \frac{1}{2}(r - \eta) + 4\delta. \end{aligned}$$

This then implies the following lower bound on the distance between $\gamma(\infty)$ and $\beta_0(\infty)$:

$$\rho(\beta_0(\infty), \gamma(\infty)) \geq k_1 a^{-m_0 - 4\delta + \frac{r-\eta}{2}}.$$




Figure 3. The point a is closer to $\beta_0(n_1)$ than b .

We combine this inequality with (1), and a simple calculation reveals that our choice of r ensures that the path constructed between p and q avoids ΛH . The result follows. ■

Lemma 2.6 ([3, Lem. 1.21]). *If $R > 4\delta + Q + \max\{r + 4\delta + 1, K\}$, then the inclusion*

$$A_{r,R,K}(H) \hookrightarrow A_{r,\infty,K}(H)$$

induces a bijection between connected components.

Proof. Surjectivity is obvious since $R \geq K$, so we need only show injectivity. Let $x, y \in C_K(H)$ lie in the same component of $A_{r,\infty,K}(H)$. We claim that the shortest path between x and y in $N_{r,\infty}(H)$ actually lies inside $N_{r,R}(H)$, which implies the result.

Let $\phi : [0, \ell] \rightarrow \Gamma$ be the shortest such path, parameterised by length. Then suppose that for some $s \in [0, \ell]$, we have that $d(\phi(s), H) > R$. Let $[t_0, t_1] \subset [0, \ell]$ be the maximal subinterval containing s , such that for all $t \in [t_0, t_1]$, we have

$$d(\phi(t), H) \geq r + 4\delta + 1.$$

We claim that $\phi|_{[t_0, t_1]}$ is an $(8\delta + 2)$ -local geodesic. Indeed, for any $t \in [t_0, t_1]$, we have that the segment

$$\phi|_{[t-4\delta-1, t+4\delta+1] \cap [t_0, t_1]}$$

is contained in $N_{r+4\delta+1,\infty}(H)$. Therefore, any segment from $\phi(\max\{t_0, t - 4\delta - 1\})$ to $\phi(\min\{t_1, t + 4\delta + 1\})$ lies in $N_{r,\infty}(H)$. By the minimality of ϕ , it follows that $\phi|_{[t_0, t_1]}$ is an $(8\delta + 2)$ -local geodesic. Apply [9, Ch. III.H, 1.13], we then have that any geodesic path between $\phi(t_0)$ and $\phi(t_1)$ is contained in the 2δ -neighbourhood of $\phi|_{[t_0, t_1]}$. Maximality of $[t_0, t_1]$ means that either $t_0 = 0$ or $d(\phi(t_0), H) = r + 4\delta + 1$. In both cases, it follows that $d(\phi(t_0), H) \leq \max\{r + 4\delta + 1, K\}$. Identical reasoning also shows that $d(\phi(t_1), H) \leq \max\{r + 4\delta + 1, K\}$.

Let s_0, s_1 be closest-point projections of $\phi(t_0), \phi(t_1)$, respectively. Now, consider the quadrilateral $[s_0, s_1, \phi(t_1), \phi(t_0)]$. Any quadrilateral in Γ is 2δ -thin, so by inspecting Figure 4 we see that in any case, there is a path from $\phi(s)$ to H of length at most $4\delta + Q + \max\{r + 4\delta + 1, K\} = R$. This contradicts our choice of s , and so the result follows. ■

In summary, the above lemmas give us the following.

Proposition 2.7. *Let G be a one-ended hyperbolic group with δ -hyperbolic Cayley graph Γ , and let H be a Q -quasiconvex subgroup. Then there exist computable values $r \leq K \leq R$, such that the map*

$$\pi_0(A_{r,R,K}(H)) \rightarrow \pi_0(A_{r,\infty,K}(H)) \xrightarrow{\delta} \pi_0(\partial G - \Lambda H)$$

is a well-defined bijection.




Figure 4. The point $\phi(s)$ cannot be too far from H .

For the remainder of this paper, we fix r, R, K such that the above is satisfied. We conclude this section with a final remark that H acts on $\pi_0(A_{r,R,K}(H))$ and $\pi_0(\partial G - \Lambda H)$ by permutations, and the above bijection is easily seen to be equivariant with respect to this action.

2.2. Ends of pairs, revisited

We will conclude this section by using the above machinery to relate the number of (filtered) ends of a quasiconvex subgroup to its limit set complement. We begin with the following characterisation of filtered ends.

Theorem 2.8. *Let G be a one-ended hyperbolic group and H a quasiconvex subgroup. Then $\tilde{e}(G, H)$ is equal to the number of components of $\partial G - \Lambda H$.*

Proof. Let $N \in \mathbb{N} \cup \{\infty\}$ denote the number of components of $\partial G - \Lambda H$. Choose a ray γ_i based in each component C_i of $\partial G - \Lambda H$. If $i \neq j$, it is easy to see that there is no filtered homotopy between γ_i and γ_j . This proves that $\tilde{e}(G, H) \geq N$.

Conversely, fix a component C of $\partial G - \Lambda H$. It is clear if γ is a ray such that $\gamma(\infty) \in C$, then γ is a filtered ray. Now, if γ' is another ray such that $\gamma'(\infty) \in C$, then by Lemma 2.5 we have that C is path connected. So consider a path through C between $\gamma(\infty)$ and $\gamma'(\infty)$, then it is easy to see that this induces a filtered homotopy between γ and γ' in the Cayley complex. Thus, $\tilde{e}(G, H) \leq N$. ■

Note that H acts on $\partial G - \Lambda H$ by homeomorphisms. In particular, H permutes the connected components of this set. This gives us the following.

Theorem 2.9. *Let G be a one-ended hyperbolic group and H a quasiconvex subgroup. Then $e(G, H)$ is equal to the number of H -orbits of components of $\partial G - \Lambda H$.*

Proof. Recall by Proposition 1.12 that $e(G, H)$ is equal to the number of ends of the coset graph $H \backslash \Gamma$. Note that each right coset Hg satisfies the following: $d(x, H)$ is constant across all $x \in Hg$. We call this value the *height* of a coset. It is easy to see that at any given height there is only finitely many cosets.

Now, we will form the coset graph of H by a sequence of identifications, in a way which will make the conclusion clear. First, identify all cosets which lie in $N_r(H)$ to points. Next, identify all points in $N_{r,R}(H) - A_{r,R,K}(H)$ which lie in the same coset. Now, in each component of $A_{r,\infty,K}(H)$, identify vertices which lie in the same coset of H . At this stage, it is clear we have a graph with $\tilde{e}(G, H)$ ends. The only remaining identifications to be made are identifying some of these “ends.” We have that the number of ends of the final graph is then equal to the number of H -orbits of components of $A_{r,\infty,K}(H)$, which is equal to the number of H -orbits of components of $\partial G - \Lambda H$. But the number of ends of this graph is equal to $e(G, H)$ by Proposition 1.12, and so the result follows. ■

The results of this section lead us to a straightforward proof of the following corollary.

Corollary 2.10. *Let G be a one-ended hyperbolic group, and H a quasiconvex subgroup. Then $e(G, H)$ is finite. In particular, $e(G, H) \leq |B_{2l+R}(1)|$, where l is the longest length of a given generator of H .*

Proof. The action of H on the components of $\partial G - \Lambda H$ is induced by the action of H on the components of $A_{r,R,K}(H)$. The latter is a locally finite graph on which H acts cocompactly, and so this must have finitely many H -orbits of connected components, since the quotient graph is finite. It follows that there can only be finitely many H -orbits of components of $\partial G - \Lambda H$, and the result follows by Theorem 2.9. The precise bound follows from just bounding above the size of a fundamental domain for the action of H on $A_{r,R,K}(H)$. ■

The above results thus give us a full description of all possible H -almost invariant subsets of G , up to equivalence: Any such set is equivalent to some union of H -orbits of components of $A_{r,\infty,K}(H)$. To ease notation, if U is a union of connected components of $A_{r,\infty,K}(H)$, then let

$$U^* := A_{r,\infty,K}(H) - U.$$

It is easy to see that the symmetric difference $U^* \Delta (G - U)$ is H -finite, so this overloading of notation is not a problem for our purposes.

Now, note that if U is a connected component of $A_{r,\infty,K}(H)$, then gU is a connected component of $A_{r,\infty,K}(gH)$. The connected components of the latter are in one-to-one correspondence with the components of $\partial G - \Lambda gH$, and we can define $\mathcal{S}(gU)$ in the obvious way to realise this correspondence. Note that one has to be careful with the choice of basepoint, but this does not really matter as if we increase r accordingly for each g , then it is easy to see that the choice of basepoint does not affect the properties of $\mathcal{S}$. We conclude

this section by characterising via the boundary what it means for an H -almost invariant set X to cross itself, in the sense of Definition 1.5.

Proposition 2.11. *Let $X \subset G$ be an H -almost invariant set, and let U be the unique union of connected components of $A_{r,\infty,K}(H)$ which is equivalent to X . Then X crosses itself if and only if there exists $g \in G$ such that*

$$\mathcal{S}U \cap \mathcal{S}(gU^*), \quad \mathcal{S}U \cap \mathcal{S}(gU), \quad \mathcal{S}(U^*) \cap \mathcal{S}(gU), \quad \mathcal{S}(U^*) \cap \mathcal{S}(gU^*)$$

are all non-empty.

Proof. It can be seen that $U \cap gU$ is H -infinite if and only if it contains arbitrarily large balls [27, Remark 1.13]. This clearly implies that $\mathcal{S}U \cap \mathcal{S}(gU) \neq \emptyset$.

Conversely, $\mathcal{S}U \cap \mathcal{S}(gU)$ is an open subset of ∂G , so if it is non-empty, then it contains some open ball $B \subset \mathcal{S}U \cap \mathcal{S}(gU)$. It follows quickly from this that $U \cap gU$ must contain balls of arbitrary diameter. By symmetry in the other four cases, the result follows. ■

3. Detection via digraphs

So far, we have related the connectivity of the limit set complement to the connectivity of a certain subgraph of the Cayley graph, but this subgraph is still infinite. In order to obtain any sort of algorithm, we will need some way of understanding the global connectivity properties by looking only at a finite piece.

In this section, we will make headway towards this goal by constructing a certain *labelled digraph*, whose connectivity and language encode a great deal of information about the properties of the subgraph $A_{r,R,K}(H)$.

3.1. Digraphs and their languages

We begin by briefly introducing labelled digraphs. This area has many applications to the study of subgroups of free groups, and a good survey of this rich theory can be found in [23].

Definition 3.1 (Labelled digraphs). Let S be a finite set of symbols, closed under taking formal inverses, that is, $S = S^{-1}$. An S -digraph Δ is a finite graph where every edge is oriented and is labelled by some $s \in S$.

We also require that if there is a directed edge e labelled by $s \in S$ from v to v' , then there is an edge labelled by s^{-1} from v' to v , which we denote by e^{-1} .

Note that in this definition we allow the possibility of single-edge loops, and we do not require our graphs to be connected. We now set up some notation.

We write $V\Delta$ for the vertex set of Δ . Denote by $o(e)$ and $t(e)$ the initial and terminal vertices of the oriented edge e , respectively. Denote by $\text{Lab}(e)$ the label of the oriented edge e , so $\text{Lab}(e^{-1}) = \text{Lab}(e)^{-1}$. A *path* through Δ is a sequence of oriented edges $e_1 e_2 \cdots e_n$ such that $o(e_{i+1}) = t(e_i)$. We call this path a *loop* if $o(e_1) = t(e_n)$. We define the *inverse* of the path p as the path $p^{-1} := e_n^{-1} \cdots e_1^{-1}$. Given p as above, define $o(p) = o(e_1)$, $t(p) = t(e_n)$. We say that two paths p, q are *freely equal* if $o(p) = o(q)$, $t(p) = t(q)$, and the words $\text{Lab}(p)$ and $\text{Lab}(q)$ are equal in the free group $F(S)$.

Given a set of symbols S , denote by S^* the set of finite strings in S . If $p = e_1 \cdots e_n$ is a path through Δ , define its label $\text{Lab}(p) \in S^*$ as the formal string $\text{Lab}(p) = \text{Lab}(e_1) \cdots \text{Lab}(e_n)$. We have the following definition.

Definition 3.2. Let Δ be an S -digraph, and let v, v' be vertices of Δ . Define the *language of Δ from v to v'* as the set

$$\mathcal{L}(\Delta, v, v') = \{\text{Lab}(p) : p \text{ is a path in } \Delta \text{ from } v \text{ to } v'\} \subset S^*.$$

Let H be a group and $S \subset H$ a finite symmetric subset. Denote by $\pi : S^* \rightarrow H$ the obvious projection. Then we have the following key lemma.

Lemma 3.3. *Let Δ be an S -digraph, where S is a finite subset of a group H . Then for every pair v_0, v_1 of vertices of Δ , there is a finitely generated subgroup $K_{v_0} \leq H$ and a finite subset $T_{v_0, v_1} \subset H$ such that*

$$\pi(\mathcal{L}(\Delta, v_0, v_1)) = \bigcup_{t \in T_{v_0, v_1}} K_{v_0} t.$$

Moreover, T_{v_0, v_1} and generators of K_{v_0} can be computed effectively.

Proof. To ease notation, let $\mathcal{L} = \mathcal{L}(\Delta, v_0, v_1)$. We first construct K_{v_0} and T_{v_0, v_1} . Consider the following sets of paths through Δ :

- $\Pi(v, v') = \{p : p \text{ is a simple path through } \Delta \text{ with } o(p) = v, t(p) = v'\}.$
- $\Lambda(v) = \{\ell : \ell \text{ is a simple loop through } \Delta \text{ with } o(\ell) = t(\ell) = v\}.$
- $\Lambda'(v) = \{p\ell p^{-1} : v' \in V\Delta, \ell \in \Lambda(v'), p \in \Pi(v, v')\}.$

Intuitively, $\Lambda'(v)$ is the set of “lollipop loops” based at v , which are formed by attaching a simple loop to a simple path and its inverse. Let $Y = \{\pi(\text{Lab}(\gamma)) : \gamma \in \Lambda'(v_0)\}$, and set $K_{v_0} = \langle Y \rangle$. Also define

$$T_{v_0, v_1} = \{\pi(\text{Lab}(p)) : p \in \Pi(v_0, v_1)\}.$$

The rest of this proof will be dedicated to showing the following statement. We claim that every path γ through Δ with $o(\gamma) = v_0$ is freely equal to a path of the form $\alpha_1 \cdots \alpha_r q$, where each $\alpha_i \in \Lambda(v_0)$ and $q \in \Pi(v_0, t(\gamma))$. From this the lemma follows immediately.

With the above claim in mind, let γ be a path in Δ from v_0 to v_1 and assume without loss of generality that $\text{Lab}(\gamma)$ is freely reduced. If γ is a simple path or simple loop, then

the statement is trivial, so assume γ self-intersects somewhere away from $o(\gamma) = v_0$. We then proceed by induction on the length of γ . We may decompose γ into the form

$$\gamma = p_0 \ell_1 p_1 \ell_2 p_2 \cdots \ell_k p_k,$$

where the p_j are (possibly trivial) simple paths such that $t(p_j) = o(p_{j+1})$ and each ℓ_j is a non-trivial simple loop in $\Lambda(o(p_j))$. For each $1 \leq j \leq k$, let

$$\ell'_j := p_0 p_1 \cdots p_{j-1} \ell_j p_{j-1}^{-1} \cdots p_1^{-1} p_0^{-1}.$$

Inspection reveals then that γ is freely equal to the path $\gamma' := \ell'_1 \ell'_2 \cdots \ell'_k p_0 p_1 \cdots p_k$. Consider the subpath $\gamma_j = p_0 \cdots p_j$. Each γ_j is a path of strictly shorter length than γ . By the inductive hypothesis we have that γ_j is of the form $s_1 \cdots s_k q$ where each $s_i \in \Lambda'(v_0)$ and $q \in \Pi(v_0, t(p_j))$. Now we have that

$$\ell'_{j+1} = \gamma_{j-1} \ell_j \gamma_{j-1}^{-1} = s_1 \cdots s_i q \ell_j q^{-1} s_i^{-1} \cdots s_1^{-1}.$$

Since $q \ell_j q^{-1} \in \Lambda'(v_0)$, we have that ℓ'_j is thus freely equal to a path of the form $\alpha_1 \cdots \alpha_r$ where each $\alpha_i \in \Lambda'(v_0)$. If we apply the inductive hypothesis once more to γ_k , then the claim follows. ■

3.2. Adjacency digraphs

Let G be a finitely generated group and Γ a Cayley graph of G . Let H be a finitely generated subgroup of G . Let $F \subset \Gamma$ be a finite subgraph, and consider $X = HF$. Then X is a subgraph of Γ on which H acts cocompactly. We say that F is a *finite H -cover* of X . In principle, X will have many finite H -covers.

Given such a finite set of vertices F , we form its *H -adjacency set*, or just its *adjacency set* S_F as follows. Let

$$S_F = \{s \in H : s \neq 1, sF \cap F \neq \emptyset\}.$$

Informally, S_F is the finite set containing $s \in H$ such that sF intersects F . Note that S_F is symmetric, that is, $S_F = S_F^{-1}$. The fact that S_F is finite follows from F being finite and Γ being locally finite.

We form the S_F -digraph Δ_F , called the *F -adjacency graph*, as follows. The vertex set $V\Delta_F$ is precisely the set $\pi_0(F)$ of connected components of F , and there is a directed edge labelled by $s \in S_F$ from v to v' if $v \cap sv' \neq \emptyset$. In particular, this implies that v and sv' are contained in the same connected component of X . It is easy to check that Δ_F is indeed a well-defined digraph. See Figure 5 for an example of this construction.

The relevance of the adjacency digraph Δ_F to our problem is given by the following lemma.

Lemma 3.4. *Let $v, v' \in \pi_0(F)$, and $h \in H$. Then v and hv' are contained in the same connected component of $X = HF$ if and only if $h \in \pi(\mathcal{L}(\Delta_F, v, v'))$.*




Figure 5. An example of the construction of an adjacency digraph. Here, $\pi_0(F) = \{A, B, C, D\}$ and $S_F = \{a, b, a^{-1}, b^{-1}\}$. Due to symmetry we have not drawn the translates of F by a^{-1} or b^{-1} , nor the corresponding edges of Δ_F .

Proof. Firstly, suppose that $h = \pi(w)$, where $w = s_1 \cdots s_n \in \mathcal{L}(\Delta_F, v, v')$. Let $\gamma = e_1 \cdots e_n$ be a path through Δ_F between v and v' labelled by w , so e_i is labelled by s_i . Let $v = v_0$ and $v_i = t(e_i)$ for each $1 \leq i \leq n$. Observe that by definition we have that v_i intersects $s_{i+1}v_{i+1}$ in Γ , and so $s_{i+1}v_{i+1}$ and v_i must be contained within the same component of X . Since $w \in \mathcal{L}(\Delta_F, v, v')$, we have that $v_n = v'$. Note also that $s_{i+1} \cdots s_n \in \mathcal{L}(\Delta_F, v_i, v_n)$. We proceed by induction on n . Clearly if $n = 1$, then the result is true by the definition of Δ_F .

Suppose then that $n > 1$. Then $w' = s_2 \cdots s_n \in \mathcal{L}(\Delta_F, v_2, v_n)$. Let $h' = \pi(w')$, then by the inductive hypothesis we have that v_1 and $h'v_n$ are contained within the same component of X . Then v_0 is contained within the same component as s_1v_1 by definition. But this is contained in the same component as $s_1h'v_n = hv'$.

Conversely, suppose that v and hv' lie in the same component of X for some $h \in H$. Let γ be a path through X between v and hv' . The path γ will pass through a sequence of translates of components of F , say

$$v = h_0v_0, h_1v_1, h_2v_2, \dots, h_lv_l = hv',$$

where $h_i \in H$, $h_0 = 1$, $v_i \in \pi_0(F)$, and $h_{i-1}v_{i-1} \cap h_iv_i \neq \emptyset$ for every $1 \leq i \leq n$. By definition, we then have that $s_i := h_{i-1}^{-1}h_i \in S_F$, and there is an edge in Δ_F from v_{i-1} to v_i labelled by s_i . We thus deduce that $s_1 \cdots s_n \in \mathcal{L}(\Delta_F, v, v')$. Since $h = s_1 \cdots s_n$, the result follows. $\blacksquare$

The next result is a characterisation of when the graph $X = HF$ has only finitely many components. Given $v, v' \in \pi_0(F)$, recall the definition of K_v and $T_{v,v'}$ from Lemma 3.3.

Lemma 3.5. *Let F be a finite H -cover of $X = HF$. Then X has finitely many connected components if and only if for every $v \in \pi_0(F)$ we have that K_v has finite index in H .*

Proof. By Lemmas 3.3 and 3.4, note that $h \in K_v$ if and only if there is a path through X from v to hv , and thus the same also holds for $h'v$ and $h'hv$ for any $h' \in H$. So suppose first that some K_v has finite index in H and let T be a finite left-transversal of K_v . Then $TK_vv = Hv$, and this intersects at most $|T|$ components of X . Repeat this for every $v \in \pi_0(F)$ and see that X has at most

$$\sum_{v \in \pi_0(F)} |H : K_v| < \infty$$

components.

Conversely, fix $v \in \pi_0(F)$ such that $|H : K_v| = \infty$ and let $T = \{h_i : i \in \mathbb{N}\}$ be an infinite left-transversal in H of K_v . We claim that every $h_i v$ lies in a distinct component of X . Indeed, suppose that $h_i v$ and $h_j v$ were contained in the same component of X for some $i \neq j$. Then by Lemmas 3.3 and 3.4 we have that $h_i^{-1}h_j \in K_v$. This contradicts our choice of T , so the lemma follows. ■

It is clear that H permutes the connected components of X . Note that if X is connected, then Δ_F is certainly connected, though the converse is not necessarily true. Instead, the number of connected components of Δ_F actually encodes the following.

Lemma 3.6. *The number of connected components of Δ_F is equal to the number of H -orbits of components of $X = HF$.*

Proof. By Lemma 3.4, we see that $v, v' \in \pi_0(F)$ are connected by a path in Δ_F if and only if v is joined by a path through X to some H -translate of v' . The lemma then follows immediately. ■

3.3. First algorithms

The first application of our digraph machinery is the following. Throughout this subsection we fix constants $0 \leq r \leq K \leq R < \infty$ and a Cayley graph Γ of our one-ended hyperbolic group G .

Proposition 3.7. *Let H be a quasiconvex subgroup of a one-ended hyperbolic group G . Then there is an algorithm which, given $x \in G$, will decide if $x \in A_{r,R,K}(H)$.*

Proof. First, we show that one can decide membership of $N_{r,R}(H)$. Given $x \in G$, compute the finite balls $U_1 = B_r(x)$, $U_2 = B_R(x)$ in Γ . Then $x \in N_{r,R}(H)$ if and only if $U_1 \cap H = \emptyset$ and $U_2 \cap H \neq \emptyset$. This is decidable, since membership of H is decidable. Similarly, we can decide membership of $C_K(H)$. If x is not in $N_{r,R}(H)$, then terminate and return ‘no.’

Now, note that H acts cocompactly on $N_{r,R}(H)$, so compute a finite H -cover F containing x . This can be achieved by, for example, letting l be the length of the longest given generator of H , and choosing $F = N_{r,R}(H) \cap B_{2l+R}(x)$. Let $v_x \in \pi_0(F)$ be the component of F containing x . We now form the adjacency digraph Δ_F and mark the vertices

$v \in \pi_0(F)$ which intersect $C_K(H)$ in Γ . One can then check that $x \in A_{r,R,K}(H)$ if and only if v_x lies in a connected component of Δ_F containing a marked vertex. ■

We now have the following algorithms, which will allow us to distinguish (filtered) ends from one another.

Proposition 3.8. *There is an algorithm which, given $x, y \in A_{r,R,K}(H)$, will decide if x and y lie in distinct H -orbits of connected components of $A_{r,R,K}(H)$.*

Proof. Use Proposition 3.7 to compute a finite H -cover F of $A_{r,R,K}(H)$, and let $v_x, v_y \in \pi_0(F)$ be such that $x \in v_x, y \in v_y$. Then, as remarked in the proof of Lemma 3.6, we need only form the adjacency digraph Δ_F and check whether v_x and v_y lie in the same connected component of Δ_F . ■

Proposition 3.9. *There is an algorithm which, given $x, y \in A_{r,R,K}(H)$ and a solution to the generalised word problem for H , will decide if x and y lie in distinct connected components of $A_{r,R,K}(H)$.*

Proof. Find a finite H -cover F for $A_{r,R,K}(H)$ containing both x and y . Let $v_x, v_y \in \pi_0(F)$ be the components of F containing x and y , respectively. By Lemma 3.4, we have that x and y are contained in the same component if and only if $1 \in \pi(\mathcal{L}(\Delta_F, v_x, v_y))$. Form the subgroup K_{v_x} and set of words T_{v_x, v_y} as in Lemma 3.3, then x and y lie in different components if and only if $T_{v_x, v_y} \cap K_{v_x}$ is empty. Using our given solution to the generalised word problem in H , this is decidable. ■

We now turn to reproving Proposition 3.9, but we drop the hypothesis that H has a solvable generalised word problem and replace it with the condition that H has finitely many filtered ends in G . The key observation that makes this problem tractable in this case is the following.

Proposition 3.10. *If $A_{r,R,K}(H)$ has finitely many components, then there is an algorithm which, given $x, y \in A_{r,R,K}(H)$, will decide if x and y lie in distinct connected components of $A_{r,R,K}(H)$.*

Proof. It is known by Lemma 3.5 that, for every $v \in \pi_0(F)$, the subgroup $K_v \leq H$ constructed in Lemma 3.3 has finite index in H . In particular, K_v is quasiconvex in H , and so we can decide membership of K_v . The algorithm then proceeds exactly as in Proposition 3.9. ■

3.4. Counting (filtered) ends of pairs

We conclude this section by giving algorithms to compute $e(G, H)$ and $\tilde{e}(G, H)$. Fix constants $0 \leq r \leq K \leq R < \infty$ such that Proposition 2.7 is satisfied. Let F be a finite H -cover for $A_{r,R,K}(H)$, and form the adjacency digraph Δ_F . Note that the construction

of this digraph is completely effective. We immediately have the following new algorithm for computing $e(G, H)$, which was first shown to be computable by Vonseel [43].

Theorem 3.11. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasiconvex subgroup H , will output $e(G, H)$.*

Proof. One simply counts connected components of Δ_F . The construction of Δ_F is completely effective, so the result follows immediately from Lemma 3.6 and Theorem 2.9. ■

Computing $\tilde{e}(G, H)$ poses more problems. In particular, we need to somehow be able to decide if K_v has finite index in H , which can be seen to be undecidable for an arbitrary choice of H via the Rips construction (see, e.g., [4]). Thus, in light of Lemma 3.5, we cannot expect to be able to decide if $\tilde{e}(G, H) = \infty$ without either adding further hypotheses to H (e.g., this is decidable if H is free [23]), or somehow further controlling the structure of $A_{r,R,K}(H)$. It is not clear whether the latter of these is even possible, which suggests this problem may be undecidable for arbitrary choices of quasiconvex $H \leq G$. We can, however, at least give the following two algorithms.

Theorem 3.12. *There is an algorithm which, upon input of a one-ended hyperbolic group G , generators of a quasiconvex subgroup H , a solution to the generalised word problem for H , and an integer $N \geq 0$, will decide whether $\tilde{e}(G, H) \geq N$. In particular, we can decide if $\partial G - \Lambda H$ is connected.*

Proof. Compute a finite H -cover F_0 of $A_{r,R,K}(H)$, then inductively define $F_{i+1} := YF_i \cup F_i$, where Y is a symmetric generating set for H . Thus, we have an increasing sequence of H -covers (F_i) , where each strictly contains the last. For each i , let N_i denote the number of components of $A_{r,R,K}(H)$ which intersect F_i . This number is computable by Proposition 3.9, and we can conclude that $e(G, H) \geq N_i$. If there is some i such that $N_i = N_{i+1}$, then since Y is a generating set, it follows that $N_j = N_i$ for all $j > i$. Given N as input, our algorithm will run until $N_i \geq N$ for some i or terminate if the sequence $(N_i)_i$ stabilises. By the above, this will always halt. ■

Theorem 3.13. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasiconvex subgroup H , will terminate if and only if $\tilde{e}(G, H)$ is finite. Moreover, upon termination it will output the value of $\tilde{e}(G, H)$.*

Proof. We proceed as before and compute a finite H -cover F_0 of $A_{r,R,K}(H)$, then inductively define $F_{i+1} := YF_i \cup F_i$, where Y is a symmetric generating set for H . Thus, we have an increasing sequence of H -covers

$$F_0 \subset F_1 \subset F_2 \subset \cdots,$$

where each strictly contains the last. Moreover, $\bigcup_i F_i = A_{r,R,K}(H)$. For each $i \geq 0$, we now run the following search. For each component of YF_i , search for a path through

$A_{r,R,K}(H)$ into F_i . If this process terminates for a given $i \geq 0$, then it follows from an easy induction argument that there is a path from any point in $A_{r,R,K}(H) = HF_i$ back to F_i travelling through $A_{r,R,K}(H)$. In particular, this means that every connected component of $A_{r,R,K}(H)$ intersects F_i .

Clearly such an F_i exists if and only if $A_{r,R,K}(H)$ has finitely many components, which is equivalent to the condition that $\tilde{e}(G, H) < \infty$ by Proposition 2.7 and Theorem 2.8. If we do find such an F_i , then to compute the exact value of $\tilde{e}(G, H)$ we may use Proposition 3.10 to decide how many distinct components of $A_{r,R,K}(H)$ intersect F_i . By our choice of F_i , this will then be precisely the total number of components of $A_{r,R,K}(H)$. ■

4. Searching for splittings and crossings

In this section, we apply the above tools to the problem of deciding if a given quasiconvex subgroup is associated with a splitting. In short, we run two searches in parallel – one search for a splitting and another search for obstructions to splittings.

4.1. An algorithm to search for splittings

The first step to searching for splittings over subgroups commensurable with H is to be able to recognise such subgroups. We will achieve this by deciding membership of $\text{Comm}_G(H)$.

Proposition 4.1. *Let G be a hyperbolic group, then there is an algorithm which, on input of generators of a quasiconvex subgroup H , will decide if $|G : H| < \infty$.*

Proof. We have that $|G : H|$ is finite if and only if $\partial G = \Lambda H$. By Proposition 2.7, this is true if and only if $A_{r,R,K}(H) = \emptyset$ for suitably chosen r, R, K . This can be decided by computing a finite H -cover F of $N_{r,R}(H)$ as in the proof of Proposition 3.7 and checking if F intersects $C_K(H)$. ■

Proposition 4.2. *Let G be a hyperbolic group. Given generators of a quasiconvex subgroup $H \leq G$, then membership of the commensurator $\text{Comm}_G(H)$ is decidable.*

Proof. Let $g \in G$, then since G is hyperbolic we have that H^g is quasiconvex. Moreover, $H \cap H^g$ is quasiconvex and we can compute an explicit generating set for this group via [16]. We then use Proposition 4.1 to decide if $|H : H \cap H^g|$ and $|H^g : H \cap H^g|$ are finite. This decides whether $g \in \text{Comm}_G(H)$. ■

Note that $\text{Comm}_G(H)$ is itself quasiconvex, and so given a generating set of this subgroup we would have that the membership problem would be decidable via Kapovich's algorithm. However, we are not given generators of $\text{Comm}_G(H)$, but of H . So, what

the above proposition tells us is that we can decide membership of the commensurator in spite of this problem. Applying this observation, we produce the following algorithm which searches for splittings where the edge group is commensurable with H .

Proposition 4.3. *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasiconvex subgroup H , and terminates if and only if H is associated with a splitting.*

Proof. Enumerate presentations of G via Tietze transformations. If a given presentation has the general form of an amalgam or HNN extension, run Kapovich's algorithm [21] on the generators of the edge group, which terminates if and only if this subgroup is quasiconvex and outputs a quasiconvexity constant Q if it does terminate. This procedure enumerates splittings of G over quasiconvex subgroups.

Given a particular quasiconvex splitting of G , say over H' , we can decide if H' is commensurable with H as follows. Using Proposition 4.2 we decide if $H' \leq \text{Comm}_G(H)$ and $H \leq \text{Comm}_G(H')$. It is easy to check these two relations hold if and only if H is commensurable with H' . This completes the algorithm. ■

4.2. An algorithm to search for crossings

Recall Proposition 2.11, which characterised crossings via intersections of shadows in ∂G . We now characterise these intersections via local geometry and present an algorithm which terminates if and only if such a crossing exists. We first need the following technical lemma.

Lemma 4.4. *Let U_1, U_2 be unions of connected components of $A_{r,\infty,K}(H)$. The intersection $\mathcal{S}U_1 \cap \mathcal{S}(gU_2)$ is non-empty if and only if there exists some $x \in U_1 \cap gU_2$ such that*

$$d(x, H) > K \quad \text{and} \quad d(x, gH) > K + 5\delta + |g|.$$

Proof. Firstly, suppose that such an x exists, then let γ be a ray based at 1 passing within $C = 3\delta$ of x , as in Lemma 1.16. As in the proof of Lemma 2.3, we see that $\gamma(\infty) \in \mathcal{S}U_1$ since γ passes within C of $C_K(H)$. Secondly, let γ' be a geodesic ray based at g such that $\gamma'(\infty) = \gamma(\infty)$. Then the Hausdorff distance between γ and γ' is at most $5\delta + |g|$ (apply, e.g., [14, Exer. 11.86]). Again, as in the proof of Lemma 2.3, we see that $\gamma'(\infty) \in \mathcal{S}(gU_2)$, and we are done.

Conversely, let $\gamma \in p \in \mathcal{S}U_1 \cap \mathcal{S}(gU_2)$. By the definition of $\mathcal{S}$ we have that there is some t_0 such that for all $t \geq t_0$, $\gamma(t) \in U_1 \cap gU_2$. Moreover, by Lemma 2.1, we have for $i = 1, 2$ that $d(\gamma(t), H_i) \rightarrow \infty$ as $t \rightarrow \infty$. Thus, by setting $x = \gamma(t)$ for some sufficiently large t , we are done. ■

We are now ready to search for crossings. This algorithm will check every choice of H -almost invariant set and search for any crossings. It will terminate if and only if it finds

a crossing for every such choice. Note that since $e(G, H) < \infty$ by Corollary 2.10, there is only finitely many possible H -almost invariant sets to check, up to equivalence.

Proposition 4.5. *There exists an algorithm which, on input of a one-ended hyperbolic group G and generators of a quasiconvex subgroup H , will terminate if and only if for every choice $X \subset G$ of H -almost invariant set, we have that X is not semi-nested.*

Proof. We begin by picking representative choices for every equivalence class of non-trivial H -almost invariant subsets $X_1, \dots, X_n$. In particular, each X_i is a union of H -orbits of connected components of $A_{r,\infty,K}(H)$. For notational convenience, we identify X_i^* with its representative in this list.

Enumerate elements $g \in G - \text{Comm}_G(H)$ via Proposition 4.2. For each $i = 1, \dots, n$, search for some $x_{i,1} \in X_i \cap gX_i$, $x_{i,2} \in X_i \cap g(X_i^*)$, $x_{i,3} \in X_i^* \cap gX_i$, and $x_{i,4} \in X_i^* \cap g(X_i^*)$ such that the conditions in Lemma 4.4 are met for each $x_{i,j}$. We terminate our search if and only if we find such an $x_{i,j}$ for every i, j . By Proposition 2.11 and Lemma 4.4, this will terminate if and only if every H -almost invariant subset is not semi-nested. ■

4.3. Splitting detection

We now present the final result of this paper, an algorithm to detect splittings over quasiconvex subgroups. We split this algorithm into two cases, and firstly we consider the situation that we know *a priori* that our subgroup has finitely many filtered ends.

Theorem 4.6. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasiconvex subgroup H such that $\tilde{e}(G, H) < \infty$, will decide whether H is associated with a splitting. Furthermore, if such a splitting exists, then the algorithm will output this splitting.*

Proof. Let $U_1, \dots, U_n$ be the components of $A_{r,\infty,K}(H)$, and let $H' \leq H$ be a finite index subgroup which fixes each individual component, so each U_i is an H' -almost invariant subset. Any other subgroup which is commensurable with H will have the same set of filtered ends, so this is the “finest” set of H'' -almost invariant subsets for any subgroup H'' commensurable with H . It therefore follows that if H' does not admit a semi-nested H' -almost invariant set, then neither does any other subgroup which is commensurable with H .

With the above in mind we run two algorithms in parallel. We search for a splitting over a subgroup commensurable with H via Proposition 4.3, and concurrently run the algorithm in Proposition 4.5 on H' . By the above discussion, exactly one of these will terminate, and if the former algorithm terminates, then it will output a presentation of a splitting over a subgroup commensurable with H . ■

Indeed, if $\tilde{e}(G, H)$ is not finite, then we cannot rely on the machinery used above, as the stabiliser of some component of $\partial G - \Lambda H$ may have infinite index in H . Moreover,

referring back to the discussion in Section 3.4, we are unlikely to be able to decide if $\tilde{e}(G, H)$ is finite for arbitrary quasiconvex H , at least with just the current tools presented in this paper.

Recall that a subgroup H in G is said to be lonely if there is no subgroup $H' \neq H$ such that H is commensurable to H' . For a quasiconvex subgroup H of a hyperbolic group G , this condition is equivalent to saying that $H = \text{Comm}_G(H)$ and H has no finite quotients.

Theorem 4.7. *There is an algorithm which, upon input of a one-ended hyperbolic group G , generators of a quasiconvex subgroup H , and knowledge of whether H is lonely in G , will decide whether H is associated with a splitting. Furthermore, if such a splitting exists, then the algorithm will output this splitting.*

Proof. Firstly, begin running the algorithm from Proposition 4.3, which will terminate if and only if H is associated with a splitting.

Concurrently we run the following. If H is not lonely, then simultaneously search for an element $g \in \text{Comm}_G(H) - H$ and a finite index subgroup H' of H . If we find the former, then continue, and if we find the latter, then replace H with H' and then continue. At least one of these will terminate, and this ensures that $H \neq \text{Comm}_G(H)$. If H is lonely, then just continue. We now run the algorithm presented in Proposition 4.5. By Proposition 1.27, exactly one of these two procedures will terminate, and if the first algorithm terminates, then it will output a presentation of a splitting over a subgroup commensurable with H . ■

Corollary 4.8. *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasiconvex, residually finite subgroup H . This algorithm will then decide if H is associated with a splitting and will output such a splitting if one exists.*

It is conjectured that the problem of deciding whether a given hyperbolic group has a finite quotient is undecidable, and in fact this problem is known to be equivalent to the well-known conjecture that there exists a hyperbolic group which is not residually finite [10]. Assuming this conjecture, it would be undecidable whether a given quasiconvex subgroup is lonely. This means that the hypothesis in Theorem 4.7 is likely necessary unless we place further restrictions on H , such as requiring H be residually finite.

Acknowledgements. My thanks go to Panos Papazoglou for suggesting this problem and for many helpful discussions. I am also grateful to Sam Hughes and Ric Wade for their detailed feedback, to Michah Sageev and Henry Wilton for fruitful exchanges, and to Thomas Delzant for pointing me towards Vonseel's work. Finally, I thank the referees for their helpful suggestions.

References

- [1] S. I. Adian, Unsolvability of some algorithmic problems in the theory of groups (in Russian). *Tr. Moskov. Mat. Obshch.* **6** (1957), 231–298 Zbl 0080.24101 MR 0095872
- [2] G. N. Arzhantseva, On quasiconvex subgroups of word hyperbolic groups. *Geom. Dedicata* **87** (2001), no. 1–3, 191–208 Zbl 0994.20036 MR 1866849
- [3] B. Barrett, Computing JSJ decompositions of hyperbolic groups. *J. Topol.* **11** (2018), no. 2, 527–558 Zbl 1498.20110 MR 3828057
- [4] G. Baumslag, C. F. Miller III, and H. Short, Unsolvable problems about small cancellation and word hyperbolic groups. *Bull. London Math. Soc.* **26** (1994), no. 1, 97–101 Zbl 0810.20025 MR 1246477
- [5] M. Bestvina and G. Mess, The boundary of negatively curved groups. *J. Amer. Math. Soc.* **4** (1991), no. 3, 469–481 Zbl 0767.20014 MR 1096169
- [6] W. W. Boone, The word problem. *Proc. Nat. Acad. Sci. USA* **44** (1958), no. 10, 1061–1065 Zbl 0086.24701 MR 0101267
- [7] B. H. Bowditch, Cut points and canonical splittings of hyperbolic groups. *Acta Math.* **180** (1998), no. 2, 145–186 Zbl 0911.57001 MR 1638764
- [8] B. H. Bowditch, Splittings of finitely generated groups over two-ended subgroups. *Trans. Amer. Math. Soc.* **354** (2002), no. 3, 1049–1078 Zbl 0988.20027 MR 1867372
- [9] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999 Zbl 0988.53001 MR 1744486
- [10] M. R. Bridson and H. Wilton, The triviality problem for profinite completions. *Invent. Math.* **202** (2015), no. 2, 839–874 Zbl 1360.20020 MR 3418245
- [11] F. Dahmani and D. Groves, Detecting free splittings in relatively hyperbolic groups. *Trans. Amer. Math. Soc.* **360** (2008), no. 12, 6303–6318 Zbl 1196.20050 MR 2434288
- [12] F. Dahmani and V. Guirardel, The isomorphism problem for all hyperbolic groups. *Geom. Funct. Anal.* **21** (2011), no. 2, 223–300 Zbl 1258.20034 MR 2795509
- [13] G.-A. Diao and M. Feighn, The Grushko decomposition of a finite graph of finite rank free groups: an algorithm. *Geom. Topol.* **9** (2005), no. 4, 1835–1880 Zbl 1093.20022 MR 2175158
- [14] C. Druţu and M. Kapovich, *Geometric group theory*. Colloq. Publ. 63, Amer. Math. Soc., Providence, RI, 2018 Zbl 1447.20001 MR 3753580
- [15] R. Geoghegan, *Topological methods in group theory*. Grad. Texts in Math. 243, Springer, New York, 2008 Zbl 1141.57001 MR 2365352
- [16] R. Gitik, On intersections of conjugate subgroups. *Internat. J. Algebra Comput.* **27** (2017), no. 4, 403–419 Zbl 1393.20013 MR 3668102
- [17] M. Gromov, Hyperbolic groups. In *Essays in group theory*, Math. Sci. Res. Inst. Publ. 8, pp. 75–263, Springer, New York, 1987 Zbl 0634.20015 MR 0919829
- [18] C. H. Houghton, Ends of locally compact groups and their coset spaces. *J. Austral. Math. Soc.* **17** (1974), no. 3, 274–284 Zbl 0289.22005 MR 0357679
- [19] W. Jaco, D. Letscher, and J. H. Rubinstein, Algorithms for essential surfaces in 3-manifolds. In *Topology and geometry: commemorating SISTAG*, Contemp. Math. 314, pp. 107–124, Amer. Math. Soc., Providence, RI, 2002 Zbl 1012.57029 MR 1941626
- [20] W. Jaco and U. Oertel, An algorithm to decide if a 3-manifold is a Haken manifold. *Topology* **23** (1984), no. 2, 195–209 Zbl 0545.57003 MR 0744850

- [21] I. Kapovich, Detecting quasiconvexity: algorithmic aspects. In *Geometric and computational perspectives on infinite groups*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 25, pp. 91–99, Amer. Math. Soc., Providence, RI, 1996 Zbl [0859.20026](#) MR [1364182](#)
- [22] I. Kapovich and N. Benakli, Boundaries of hyperbolic groups. In *Combinatorial and geometric group theory*, Contemp. Math. 296, pp. 39–93, Amer. Math. Soc., Providence, RI, 2002 Zbl [1044.20028](#) MR [1921706](#)
- [23] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups. *J. Algebra* **248** (2002), no. 2, 608–668 Zbl [1001.20015](#) MR [1882114](#)
- [24] I. Kapovich and H. Short, Greenberg’s theorem for quasiconvex subgroups of word hyperbolic groups. *Canad. J. Math.* **48** (1996), no. 6, 1224–1244 Zbl [0873.20025](#) MR [1426902](#)
- [25] P. H. Kropholler and M. A. Roller, Relative ends and duality groups. *J. Pure Appl. Algebra* **61** (1989), no. 2, 197–210 Zbl [0691.20036](#) MR [1025923](#)
- [26] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*. Reprint of the 1977 Edition (Ergebnisse der Mathematik und ihrer Grenzgebiete), Classics in Math., 89, Springer, Berlin, 1977 Zbl [0368.20023](#) MR [0577064](#)
- [27] G. Niblo, M. Sageev, P. Scott, and G. A. Swarup, Minimal cubings. *Internat. J. Algebra Comput.* **15** (2005), no. 2, 343–366 Zbl [1085.20026](#) MR [2142089](#)
- [28] G. A. Niblo, The singularity obstruction for group splittings. *Topology Appl.* **119** (2002), no. 1, 17–31 Zbl [0991.20024](#) MR [1881707](#)
- [29] P. S. Novikov, On the algorithmic unsolvability of the word problem in group theory (in Russian). *Trudy Mat. Inst. Steklov.*, vol. 44, Izdat. Akad. Nauk SSSR, Moscow, 1955 Zbl [0068.01301](#) MR [0075197](#)
- [30] P. Papasoglu, An algorithm detecting hyperbolicity. In *Geometric and computational perspectives on infinite groups*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 25, pp. 193–200, Amer. Math. Soc., Providence, RI, 1996 Zbl [0857.20017](#) MR [1364185](#)
- [31] M. O. Rabin, Recursive unsolvability of group theoretic problems. *Ann. of Math. (2)* **67** (1958), no. 1, 172–194 Zbl [0079.24802](#) MR [0110743](#)
- [32] P. Scott, Ends of pairs of groups. *J. Pure Appl. Algebra* **11** (1977), no. 1–3, 179–198 Zbl [0368.20021](#) MR [0487104](#)
- [33] P. Scott, The symmetry of intersection numbers in group theory. *Geom. Topol.* **2** (1998), no. 1, 11–29 Zbl [0897.20029](#) MR [1608688](#)
- [34] P. Scott and G. A. Swarup, Splittings of groups and intersection numbers. *Geom. Topol.* **4** (2000), no. 1, 179–218 Zbl [0983.20024](#) MR [1772808](#)
- [35] P. Scott and G. A. Swarup, Regular neighbourhoods and canonical decompositions for groups. *Astérisque* **289** (2003), 239 Zbl [1036.20028](#) MR [2032389](#)
- [36] P. Scott and T. Wall, Topological methods in group theory. In *Homological group theory*, London Math. Soc. Lecture Note Ser. 36, pp. 137–203, Cambridge Univ. Press, Cambridge-New York, 1979 Zbl [0423.20023](#) MR [0564422](#)
- [37] Z. Sela, The isomorphism problem for hyperbolic groups. I. *Ann. of Math. (2)* **141** (1995), no. 2, 217–283 Zbl [0868.57005](#) MR [1324134](#)
- [38] J.-P. Serre, *Trees*. Springer Monogr. Math., Springer, Berlin, 2003 Zbl [1013.20001](#) MR [1954121](#)
- [39] J. Stallings, *Group theory and three-dimensional manifolds*. Yale Math. Monogr. 4, Yale University Press, New Haven, Conn.-London, 1971 Zbl [0241.57001](#) MR [0415622](#)
- [40] J. R. Stallings, On torsion-free groups with infinitely many ends. *Ann. of Math. (2)* **88** (1968), no. 2, 312–334 Zbl [0238.20036](#) MR [0228573](#)

- [41] G. A. Swarup, [On the cut point conjecture](#). *Electron. Res. Announc. Amer. Math. Soc.* **2** (1996), no. 2, 98–100 Zbl [0868.20032](#) MR [1412948](#)
- [42] N. W. M. Touikan, [Detecting geometric splittings in finitely presented groups](#). *Trans. Amer. Math. Soc.* **370** (2018), no. 8, 5635–5704 Zbl [1494.20065](#) MR [3803145](#)
- [43] A. Vorseel, [Ends of Schreier graphs of hyperbolic groups](#). *Algebr. Geom. Topol.* **18** (2018), no. 5, 3089–3118 Zbl [1499.20109](#) MR [3848409](#)

Received 22 March 2023.

Joseph Paul MacManus

Mathematical Institute, University of Oxford, Andrew Wiles Building, Observatory Quarter,
Oxford OX2 6GG, UK; macmanus@maths.ox.ac.uk