CONNECTEDNESS PROPERTIES OF LIMIT SETS

B. H. BOWDITCH

ABSTRACT. We study convergence group actions on continua, and give a criterion which ensures that every global cut point is a parabolic fixed point. We apply this result to the case of boundaries of relatively hyperbolic groups, and consider implications for connectedness properties of such spaces.

0. INTRODUCTION

In this paper, we prove a result concerning the nature of global cut points in continua which admit certain kinds of convergence actions. One of the principal motivations concerns the boundaries of one-ended hyperbolic groups. It was conjectured in [BeM] that such boundaries are locally connected, or equivalently, have no global cut point. This was shown to be the case for strongly rigid groups in [Bo3], [Bo5] and strongly accessible groups in [Bo6]. An ingenious idea of Swarup showed how to adapt these arguments to deal with the general case [Sw]. Inspired by this idea, we reset this result in a broader dynamical context, which will have applications also to limits sets of relatively hyperbolic groups, and in particular to geometrically finite groups. There are close connections between the algebraic structure of such groups and the topology of their boundaries. In this, local and global cut points play a major role. In particular, the main result given here is one of the ingredients in proving that the boundary of a relatively hyperbolic group is locally connected if it is connected (under mild constraints on the parabolic subgroups).

Before stating the main theorem, we need a few definitions. Convergence groups were defined by Gehring and Martin [GeM]. (See also [T] and [Bo7] for some discussion relevant to the present paper.) Let M be a continuum, i.e. a connected compact hausdorff topological space. Suppose that Γ acts as a convergence group on M. (Throughout this paper we use the term "convergence group" for what was called a "discrete convergence group" in [GeM].) Such an action is *minimal* if there is no proper non-empty closed invariant subset. A *parabolic* element of Γ is an infinite order element with precisely one fixed point. If G is a two-ended (i.e. virtually cyclic) subgroup of Γ, we write $\eta_\Gamma(G)$ for the number of ends of the pair (Γ, G). (Note that if Γ splits over G, then $\eta_\Gamma(G) > 1$.) A *loxodromic* subgroup, G, of Γ is one whose limit set, ΛG, consists of precisely two points. Such a group is necessarily two-ended, and $(M \setminus \Lambda G)/G$ is compact hausdorff. We write $\eta_M(G)$ for the number of connected components of $(M \setminus \Lambda G)/G$. If the dynamics of Γ on M in some way reflects the intrinsic geometry of Γ, one might expect the quantities

$\eta_\Gamma(G)$ and $\eta_M(G)$ to be somehow related. All we are really concerned about in this paper is whether these numbers are equal to or greater than 1. (We do not assume them to be finite.)

The statement of the main result is a little technical:

Theorem 0.1. *Let Γ be a one-ended finitely presented group with no infinite torsion subgroup. Let M be a metrisable continuum which admits a minimal convergence action by the group Γ. Suppose that, for any loxodromic subgroup, $G \leq \Gamma$, with $\eta_\Gamma(G) > 1$, we have $\eta_M(G) > 1$. Then, every global cut point of M is a parabolic fixed point.*

In fact, we can get away with a little less. We only really need that if Γ splits over a loxodromic group, G, then $\eta_M(G') > 1$ for some finite index subgroup, G', of G. That some kind of "non-degeneracy" assumption of this nature is necessary is apparent from the well-known examples of surface groups which act as convergence groups on dendrites with no parabolics.

The main application of this result will be the following. Recall that a "relatively hyperbolic group" consists of group, Γ, and a class, $\mathcal{G}$, of preferred "peripheral" subgroups such that Γ is "hyperbolic relative to" $\mathcal{G}$. This notion was defined by Gromov [Gr]. It is discussed further in Section 4. Associated to a relatively hyperbolic group, $(\Gamma, \mathcal{G})$, is a canonical compact metrisable space, $\partial(\Gamma, \mathcal{G})$, called the "boundary" of $(\Gamma, \mathcal{G})$. Now Γ acts as a convergence group on $\partial(\Gamma, \mathcal{G})$ such that the peripheral groups are precisely the maximal parabolic subgroups. In our notation, we shall frequently omit explicit reference to the peripheral structure, $\mathcal{G}$. If the peripheral structure is empty, we are reduced to the standard case of (word) hyperbolic groups.

Theorem 0.2. *Suppose that Γ is a relatively hyperbolic group whose boundary, $\partial\Gamma$, is connected. Suppose that each peripheral subgroup is finitely presented, either one-ended or two-ended, and contains no infinite torsion subgroup. Then every global cut point of $\partial\Gamma$ is a parabolic fixed point.*

This result has two immediate corollaries. The first, in the case where there are no peripheral subgroups, gives us:

Corollary 0.3. *The boundary of a one-ended hyperbolic group has no global cut point.*

As mentioned above, this has been obtained by Swarup [Sw], using results of [Bo3], [Bo6] and [L]. It follows, by [BeM], that such a boundary is locally connected. Thus, for example such a group is semistable at infinity [M], and one can derive the JSJ splitting of Sela [Se] from the local cut point structure of the boundary [Bo4].

The second application is the special case of limit sets of geometrically finite groups in pinched negative curvature [Bo2]. In this case, peripheral subgroups are finitely generated virtually nilpotent [Bo3], and hence automatically satisfy the hypotheses of Theorem 0.2.

Corollary 0.4. *Suppose a group Γ acts as a geometrically finite group on a complete simply connected manifold of pinched negative curvature. If the limit set, $\Lambda\Gamma$, is connected, then every global cut point is a parabolic fixed point.*

Now, given the result of Theorem 0.2, it's not hard to see that, if Γ has a global cut point in its boundary, it must split over a parabolic subgroup relative to the

peripheral subgroups. (This is discussed in [Bo10].) An alternative dynamical approach to this is given in [BoS] in the case where all peripheral groups are one-ended. These results form part of the larger project of showing that the boundary of a relatively hyperbolic group, satisfying the same hypotheses as Theorem 0.2, is locally connected. An overview of this project is described in [Bo9] and [Bo10]. In particular, this shows that the boundary of a geometrically finite kleinian group is locally connected if it is connected. The 3-dimensional case is described in [AnM]. The higher-dimensional case seems to call for a completely different approach, and has been open for some time. The case of kleinian groups is discussed in [BoS].

For the purposes of exposition, we shall first give direct proofs of the two corollaries. The proof of the second corollary can be adapted easily to prove Theorem 0.2, as we discuss in Section 4. Since this calls for additional background material, as set up in [Bo8], we give the main argument in the more familiar context of pinched hadamard manifolds.

Most of this paper was prepared at the University of Melbourne. I am indebted to Walter Neumann and Craig Hodgson for inviting me to take part in the Special Year in geometric group theory, and to the other members of the geometry group for stimulating discussions.

1. Proof of the main theorem

We shall need the following observation, which appears to be folklore. (A proof of the analogous result for the finiteness property FP_n is given in [Bi].)

Lemma 1.1. *Suppose Γ is a finitely presented group which splits as a finite graph of groups all of whose edge groups are finitely presented. Then, all the vertex groups are finitely presented.*

Proof. Let (V, E) be the graph of groups. Given $v \in V$ and $e \in E$, we write $\Gamma(v)$ and $\Gamma(e)$ respectively for the vertex and edge groups. We choose simplicial complexes, $K(v)$ and $K(e)$, with fundamental groups $\Gamma(v)$ and $\Gamma(e)$. We take all the edge complexes, $K(e)$, to be finite. Moreover, if e is incident on v, we can suppose that $K(v)$ contains an embedded homeomorphic copy, $K(v, e)$ of $K(e)$, such that the inclusion induces the given edge-to-vertex group inclusion on the level of fundamental groups. We can assume that the complexes $K(v, e)$ are disjoint, as e ranges over the set, $E(v)$, of edges incident to v, and also that each $K(v, e)$ has an embedded product neighbourhood, $K(v, e) \times [0, 1]$, in $K(v)$. We form a complex, K, by identifying $K(v, e)$ with $K(w, e)$ whenever v and w are the endpoints of some edge e. In this way we get $\pi_1(K) \cong \Gamma$.

Now, let L be a finite complex with $\pi_1(L) \cong \Gamma$. Let $f : L \longrightarrow K$ be a simplicial map inducing an isomorphism of fundamental groups. We can assume that f is injective. (For example, by replacing K by $K \times Q$, where Q is a cone over L, and replacing f by a diagonal map.) Identifying L with its image under f, we can assume that L is a subcomplex of K. In fact, we claim that we can choose the pair K, L so that, for each $e \in E$, the complex $L \cap K(e)$ is connected and carries all of $\Gamma(e) = \pi_1(K(e))$. (The only reason for modifying K in this process is to ensure that L remains embedded.)

The proof of the claim relies on a variation of Stallings's "binding tie" argument. (A similar application of these ideas can be found in [BeF1].) Suppose that $L \cap K(e)$ is disconnected. Let α be an arc in $K(e)$ connecting distinct components of $L \cap K(e)$. Since L carries all of $\pi_1(K)$, we can find an arc, β, in L which is homotopic to α

relative to its endpoints. Thus $\alpha \cup \beta$ spans a disc, D, in K, which we can assume (modifying K if necessary) to be embedded and to satisfy $L \cap D = \beta$. Moreover we can assume that $D \cap \bigcup_{e \in E} K(e)$ consists of a finite set of arcs. (We can easily eliminate closed curves, given that each $K(e)$ is π_1-injective in K.) We now replace L by $L \cup D$. This reduces the total number of components of $L \cap \bigcup_{e \in E} K(e)$. Thus, after a finite number of steps, we arrange that each complex $L \cap K(e)$ is connected. To ensure that each $L \cap K(e)$ carries all of $\Gamma(e)$, we carry out another sequence of operations of this type. This time we choose the arcs α so as to generate what remains of $\Gamma(e)$, using the fact that the edge groups are finitely generated. This proves the claim.

Now it's a simple exercise to check that if $v \in V$, then $L \cap K(v)$ is connected. Let $L(v) = (L \cap K(v)) \cup \bigcup_{e \in E(v)} K(v, e)$. Again, it's easily verified that the inclusion of $L(v)$ in $K(v)$ induces an isomorphism of fundamental groups. In particular, $\Gamma(v) \cong \pi_1(L(v))$ is finitely presented as required.   □

We shall need another lemma about splittings. Let Γ be a group. By a *two-ended splitting* of Γ, we mean a representation of Γ as a finite graph of groups with each edge group two-ended. Also, we shall not allow vertices of degree one with the vertex group equal to the incident edge group. (On the level of trees, this means we are only considering minimal actions.) We say this splitting is *maximal* if none of the vertex groups split further (non-trivially over two-ended subgroups) relative to the incident edge groups. Recall that a splitting is *reduced* in the sense of [BeF1] if it contains no reducible vertices. A *reducible* vertex is a vertex of degree 2, such that one of the incident edge groups is equal to the vertex group. If Γ is finitely presented, then the result of [BeF1] puts a bound on the complexity of a reduced small splitting of Γ, where complexity is measured by the number of edges in the graph of groups. In particular, this deals with two-ended splittings.

We need a slight variation on this. The following result can be deduced from the description of the JSJ splitting of a one-ended finitely presented group (see for example [DS]). For completeness, we summarise the relevant part of the argument.

Lemma 1.2. *Any one-ended finitely presented group admits a maximal two-ended splitting.*

Proof. Let Γ be a finitely presented one-ended group. We start with a reduced two-ended splitting of Γ of maximal complexity, as given by [BeF1]. If it is possible to split one of the vertex groups relative to the incident edge group, then we would obtain a more complex graph of groups, which must therefore contain a reducible vertex. If the theorem fails, then we can iterate this process indefinitely. We can only ever introduce reducible vertices. (For if, at some point, we introduced a non-reducible vertex, we could collapse back all the previously introduced reducible vertices, to obtain a reduced graph of greater complexity than the original.) We thus obtain a sequence of graphs of groups, with arbitrarily long chains of reducible vertices inserted along the edges of the original graph. This subdivides any such edge into "subedges", where the subedge stabilisers each increase or decrease as we move along the edge. Now we cannot have a subedge stabiliser with both adjacent subedge stabilisers strictly smaller (since contracting this subedge to a vertex would introduce a new non-reducible vertex). We are therefore quickly reduced to worrying about the existence of arbitrarily long sequences of subedges with either ascending or descending subedge stabilisers.

Now, infinite ascending chains might indeed occur. However they do not concern us here, since they correspond to splitting the vertex groups trivially relative to the incident edge groups, and hence do not interfere with maximality, as we have defined it. We thus only need to rule out the possibility of infinite descending chains. This is based on Dunwoody's track construction (cf. [D]).

For the purposes of the argument it is convenient to collapse all the irrelevant parts of the graph of groups to vertices, so as to give us an amalgamated free product or HNN extension over a two-ended subgroup. For convenience of exposition, let's suppose we are dealing with an amalgamated free product. The graph consists of a single edge, which we identify with the unit interval, $[0, 1]$. We imagine successively inserting new degree-2 vertices into the graph at the points 2^{-n}, for $n \geq 1$. Taking the corresponding trees, we get a sequence of simplicial trees, T_n, with quotient maps to $[0, 1]$. We now fix a finite 2-dimensional simplicial complex, K, with fundamental group Γ, and define a sequence of Γ-equivariant resolutions, $\phi_n : \tilde{K} \longrightarrow T_n$, where $\tilde{K}$ is the universal cover of K. Passing to the quotient under Γ, these induce a sequence of maps $\psi_m : K \longrightarrow [0, 1]$. If $m \geq n$, we can suppose that $\psi_m^{-1}[2^{-(n+1)}, 1] = \psi_n^{-1}[2^{-(n+1)}, 1]$, and that ψ_m agrees with ψ_n on this preimage.

For each $n \geq 0$, let e_n be an edge of T_n which maps to $[2^{-(n+1)}, 2^{-n}]$ under the quotient map. We choose these edges so that $\Gamma(e_{n+1}) \subseteq \Gamma(e_n)$, and so the edge stabilisers are strictly decreasing. Let $\tilde{t}_n$ be a component of the preimage of the midpoint of e_n which separates $\tilde{K}$ into two unbounded components. Let t_n be the projection of $\tilde{t}_n$ to K. Thus, t_n is a track in K, which represents a splitting of $\Gamma = \pi_1(K)$ over a subgroup of $\Gamma(e_n)$. Since Γ is one-ended, this subgroup must be infinite. Now, we have arranged that all the tracks t_n are disjoint, and so a standard argument (cf. [D]) shows that all but finitely many are parallel, and hence carry the same subgroup of Γ. But, these subgroups are all infinite, and their intersection is finite (since the intersection of the groups $\Gamma(e_n)$ is finite). We thus arrive at a contradiction. □

Before proceeding, we should make a few brief comments on the proof of Lemma 1.2. Strictly speaking, our definition of a reduced graph should have contained a clause relating to the specific case where Γ is a virtual Baumslag-Solitar group. However, this case is easily dealt with separately, and does not concern us here anyway. More details of the argument ruling out descending chains of reducible vertices can be found in [DS]. Examples demonstrating the possibility of infinite ascending chains can be constructed from Abels's example [Ab] of a finitely presented group whose centre is isomorphic to the diadic rationals, as observed by Dunwoody and Sageev.

We now move on to consider actions on $\mathbb{R}$-trees.

Suppose that a group, Γ, acts isometrically on an $\mathbb{R}$-tree, Σ. In this case, we say that a subgroup, G, of Γ, is *parabolic* (on Σ) if it fixes a point of Σ. If Γ is infinite and acts with finite edge stabilisers, then this point is unique. (Note that, in this context, "parabolic actions" are sometimes referred to as "trivial actions" or as "elliptic actions". However, the former term is somewhat misleading, and the latter term is at odds with the terminology of convergence actions. We therefore adopt our slightly non-standard terminology.) We shall be concerned with "stable" actions. This notion can be defined in a number of equivalent ways, for example by demanding that the stabilisers of any decreasing chain of intervals must stabilise.

Suppose that Γ is a finitely presented group with a stable non-parabolic action on an $\mathbb{R}$-tree, Σ, with finite edge-stabilisers. Then, the result of [BeF3] tells us that Γ either is one-ended virtually abelian, or splits non-trivially over a finite or two-ended subgroup. Moreover, if $H_1, \ldots, H_n$ are finitely generated subgroups of Γ which act parabolically on Σ, then, in the latter case, we can choose the splitting of Γ so that each of the subgroups H_i is conjugate into one of the vertex groups.

Lemma 1.3. *Suppose Γ is a finitely presented one-ended group which acts stably on an $\mathbb{R}$-tree Σ with finite edge-stabilisers. Suppose that every two-ended subgroup over which Γ splits is parabolic, and every rank two abelian subgroup of Γ is parabolic. Then Γ is parabolic.*

Proof. Note that it follows that every one-ended virtually abelian subgroup of Γ is parabolic. Suppose, for contradiction, that Γ is not parabolic, and hence not virtually abelian. By Lemma 1.2, Γ admits a maximal two-ended splitting. Let T be the corresponding simplicial tree. Let $V(T)$ be the vertex set of T. Given $v \in V(T)$, let $\Gamma(v)$ denote the vertex stabiliser. By Lemma 2.1, each $\Gamma(v)$ is finitely presented. By hypothesis, each edge stabiliser is parabolic. We claim that each vertex stabiliser is parabolic.

Suppose that $\Gamma(v)$ were not parabolic. From our hypotheses, $\Gamma(v)$ is not one-ended virtually abelian. It also cannot be virtually cyclic (otherwise it would contain an edge group of finite index). Thus, by [BeF3], $\Gamma(v)$ splits over a finite or two-ended subgroup relative to the incident edge groups, contradicting our choice of T, and proving our claim.

From the uniqueness of these fixed points, it now follows that all vertex and edge stabilisers must fix the same point of Σ. Thus Γ is parabolic. □

Note that in the hypotheses we only really need to assume that a certain algebraically predetermined finite set of subgroups of Γ are each parabolic.

Recall that a *dendrite* is a locally connected metrisable continuum which is uniquely arc-connected.

Lemma 1.4. *Suppose that Γ is finitely presented one-ended and contains no infinite torsion subgroup. Suppose that Γ acts as a minimal convergence group on a dendrite, D, in such a way that every two-ended subgroup over which Γ splits is parabolic on D. Then D is trivial (i.e. a point).*

Proof. Suppose D is non-trivial. Then, using either [L] or [Bo5], we construct an $\mathbb{R}$-tree, Σ, on which Γ acts isometrically and non-parabolically with finite edge stabilisers. The non-existence of infinite torsion subgroups ensures that this action is stable. Moreover, if $H_1, \ldots, H_n$ are finitely presented subgroups of Γ which are parabolic on D, then we can assume that they are parabolic on Σ.

Now, any rank-two free abelian subgroup of Γ is necessarily parabolic on D (see, for example, [T]). Let $H_1, \ldots, H_n$ be the edge groups in a maximal splitting of Γ as described in the proof of Lemma 1.3, together with those vertex groups which are virtually abelian. These groups are all parabolic on D, and so can be chosen to be parabolic in Σ. By Lemma 1.3, and the subsequent remark, we arrive at the contradiction that Γ is parabolic on Σ. □

Proof of Theorem 0.1. Define an equivalence relation on M by deeming two points, $x, y \in M$, to be not equivalent if there exists a subset of M consisting of points which individually separate x from y, and which is order isomorphic to the rational

numbers in the natural linear order (arising from the pointwise separation properties of this subset). The quotient of M by this relation is a dendrite, D. If M contains a cut point which is not a parabolic fixed point, then it follows from [Bo3] that D is non-trivial. Moreover, Γ acts as a minimal convergence group on D.

Suppose that Γ splits over a two-ended subgroup, $G \leq \Gamma$. We claim that G is parabolic on D. If G is parabolic on M, then it is certainly parabolic on D, so we can assume that it is loxodromic on M. Thus, its limit set, ΛG, consists of precisely two points, say a and b. Moreover $(M \setminus \Lambda G)/G$ is compact hausdorff. By hypothesis, $(M \setminus \Lambda G)/G$ is disconnected, so we can write it as a disjoint union, $A_1 \sqcup A_2$, of non-empty closed subsets. Now, the preimage, U_i, of A_i in $M \setminus \Lambda G$ is open in $M \setminus \Lambda G$ and hence in M. Thus, $B_i = U_i \cup \Lambda G \subseteq M$ is closed and G-invariant. Moreover, $M = B_1 \cup B_2$ and $\Lambda G = B_1 \cap B_2$.

We claim that B_i is connected. To see this, let K be a connected component of B_i. If $K \cap \Lambda G$ were empty, then we could find a closed and open subset, L, of B_i containing K and which does not meet ΛG. We see that L must be closed and open in M, contradicting the fact that M is connected. This shows that $K \cap \Lambda G \neq \emptyset$. Suppose that $a \in K \cap \Lambda G$. Let $H \leq G$ be the subgroup (of index at most 2) of G which fixes a. Now, K is H-invariant, and so either $\Lambda G \subseteq K$ or $K = \{a\}$. In the former case, we see that $B_i = K$ is connected as required. In the latter case, we deduce, similarly, that $\{b\}$ is a component of B_i, giving the contradiction that $B_i = \Lambda G$.

It now follows that no point of M separates the two points of ΛG. Thus, ΛG collapses to a point in D, and so G is parabolic in D.

The hypotheses of Lemma 1.4 are now satisfied, giving the contradiction that D is trivial. $\square$

It is unclear to what extent the algebraic hypotheses on Γ in Theorem 0.1 are really necessary. It is conceivable it might hold for any finitely generated group, though one might want to interpret a "parabolic fixed point" to mean one whose stabiliser is an infinite (possibly torsion) subgroup. One can drop the assumption that M is metrisable, though to do so would introduce technical complications with little obvious benefit — all likely applications are to metrisable spaces.

2. Proof of the first corollary

Let Γ be a one-ended hyperbolic group (in the sense of Gromov [Gr]). Its boundary, $\partial \Gamma$, is a metrisable continuum, on which Γ acts as a minimal convergence group without parabolics (see [GhH], [F], [T], [Bo7]). Moreover, Γ has no infinite torsion subgroup.

Let X be a Cayley graph of Γ, so that $X \cup \partial \Gamma$ admits a natural compact hausdorff topology. Suppose that $G \leq \Gamma$ is two-ended with $\eta_\Gamma(G) > 1$. Thus, $(X \cup \partial \Gamma \setminus \Lambda G)/G$ is compact hausdorff. There is an open relatively compact subset, N, of X/G such that we can write $(X/G) \setminus N$ as a disjoint union of two closed non-compact subsets A_1 and A_2. Let $\bar{A}_i$ be the closure of A_i in $(X \cup \partial \Gamma \setminus \Lambda G)/G$, and let $B_i = \bar{A}_i \cap (\partial \Gamma \setminus \Lambda G)/G$. We claim that $(\partial \Gamma \setminus \Lambda G)/G = B_1 \sqcup B_2$. To see this, suppose $a \in (\partial \Gamma \setminus \Lambda G)/G$, and consider a path β in X/G tending to a and which is the projection of geodesic ray in X. Since G is two-ended, it quasiconvex, so the lift of N to X is quasiconvex. It follows easily that β must eventually lie in either A_1 or A_2, and so $a \in B_1 \cup B_2$. Moreover any such ray in A_1 must diverge from any such ray in A_2, and so $B_1 \cap B_2 = \emptyset$ as claimed. We see that $\eta_{\partial \Gamma}(G) > 1$.

The hypotheses of Theorem 0.1 are now satisfied. Since Γ has no parabolics on $\partial\Gamma$, there can be no global cut point. This proves Corollary 0.3.

3. Proof of the second corollary

As mentioned in the introduction, the argument we present here will go through, with minor modification, to the case of relatively hyperbolic groups. This will be discussed in Section 4 (where the convex hull of the limit set will be replaced by a locally finite hyperbolic 2-complex). Most of the proof of Corollary 0.4 is given up to dealing with the complication of two-ended parabolic subgroups.

Let Y be a complete simply connected riemannian manifold of pinched negative curvature. We can naturally compactify Y to a closed ball, $Y \cup \partial Y$, by adjoining the ideal sphere, ∂Y. (To avoid any confusion with ideal boundaries, in this section we shall use fr A to denote the topological boundary or "frontier" of a set A.)

Suppose a group, Γ, acts as a geometrically finite group on Y, as discussed in [Bo2]. It was shown there that Γ is finitely generated, and a simple extension of these arguments shows that it must, in fact, be finitely presented. Also, Γ has no infinite torsion subgroup.

Let $\Lambda\Gamma \subseteq \partial Y$ be the limit set of Γ, and let X be the intersection of its closed convex hull with Y. (In fact, it would be more convenient to take a uniform neighbourhood of the convex hull — this ensures that X is a codimension-0 submanifold with C^1 boundary.) Now, $X \cup \Lambda\Gamma$ is a closed subset of $Y \cup \partial Y$. In fact, X is intrinsically a Gromov hyperbolic space with ideal boundary $\partial X = \Lambda\Gamma$. Let $\Pi \subseteq \Lambda\Gamma$ be the set of parabolic fixed points. Given a point $p \in \Pi$, its stabiliser, $\Gamma(p)$, is finitely generated virtually nilpotent [Bo1], and hence, in particular, is either one- or two-ended. The groups, $\Gamma(p)$, are precisely the maximal parabolic subgroups of Γ. Now, each parabolic fixed point, p, is "bounded" in the sense that $(\Lambda\Gamma \setminus \{p\})/\Gamma(p)$ is compact. Moreover, we can choose an open horoball, $B(p)$, about p such that $S(p) = X \cap \text{fr}\, B(p)$ is connected. (To see this, choose any compact set $Q \subseteq \Lambda\Gamma$ such that $\Lambda\Gamma = \bigcup \Gamma(p)Q$, and generators, $\gamma_1, \ldots, \gamma_n$ for $\Gamma(p)$. Now, $Q' = Q \cup \bigcup_{i=1}^n \gamma_i Q$ is a closed subset of $\Lambda\Gamma$ missing the point p, and so we can find a horoball, $B(p)$, whose closure does not meet the convex hull of Q'. Each component of $X \setminus B(p)$ must meet $\Lambda\Gamma$; otherwise we could remove it from the convex hull. It follows that $X \setminus B(p)$ is connected. Since $B(p)$ is connected, and X is simply connected, it follows that $X \cap \text{fr}\, B(p)$ is connected.) Note that $S(p)/\Gamma(p)$ is compact. We refer to $C(p) = X \cap B(p)$ as a *cusp region*. Note that $C(p) \cup \{p\}$ is convex, and hence topologically a closed ball. In particular, $C(p)$ is simply connected and one-ended. We can choose the horoballs $B(p)$ equivariantly with respect to the action of Γ, and with the property that distinct horoballs have disjoint closures. Let $X_0 = X \setminus \bigcup_{p \in \Pi} C(p)$. In this way, X_0/Γ is compact. The connected components of $X \setminus X_0$ are precisely the cusp regions $C(p)$, and the components of the boundary of X_0 in X are precisely the sets $S(p)$.

We write $\Pi = \Pi_1 \sqcup \Pi_2$, where $\Gamma(p)$ (and hence $S(p)$) is one-ended for all $p \in \Pi_1$, and $\Gamma(p)$ (and hence $S(p)$) is two-ended for all $p \in \Pi_2$. Let $X_2 = X \setminus \bigcup_{p \in \Pi_2} C(p)$. We thus have Γ-invariant sets, $X_0 \subseteq X_2 \subseteq X$.

We note that Γ is one-ended if and only if X_0 is one-ended (since X_0/Γ is compact). Also $\Lambda\Gamma$ is connected if and only if X is one-ended (since $\Lambda\Gamma$ is the ideal boundary of X). From the topology of cusp regions described above (specifically the fact that each region $C(p)$ is one-ended, and that the set $S(p)$ is one-ended if

$p \in \Pi_1$), we see that the one-endedness of X is equivalent to that of X_2, and is implied by that of X_0. If it happens that $\Pi_2 = \emptyset$, then these are all equivalent.

We first prove Corollary 0.4 in the case where $\Pi_2 = \emptyset$, i.e. all maximal parabolic subgroups are one-ended. From the previous paragraph, we see that Γ is one-ended. Suppose $G \leq \Gamma$ is loxodromic. Now, the quotient, $(X \cup \Lambda\Gamma \setminus \Lambda G)/G$ is compact. Also G intersects every parabolic subgroup in a finite group, so each cusp region, $C(p)$, projects finite-to-one into X/G. In particular, these projections are connected, and have one-ended boundaries. Removing them we obtain the quotient X_0/G. Now, if $\eta_\Gamma(G) > 1$, then X_0/G has more than one end. Thus, X/G also has more than one end. As in the proof of Corollary 0.3, we see that $(\Lambda\Gamma \setminus \Lambda G)/G$ is disconnected. In other words, $\eta_{\Lambda\Gamma}(G) > 1$. The hypotheses of Theorem 0.1 are now satisfied, proving Corollary 0.4 in this case.

We note that the essential points of the argument were that X is one-ended and Gromov hyperbolic with boundary $\Lambda\Gamma$. The idea to deal with the general case will be to use a "doubling" construction that will preserve this hyperbolicity. We will also want to verify that the resulting space is one-ended. This is perhaps most simply done directly, but to make a cleaner, and more general statement, we translate it into a lemma about groups.

We begin with a definition. Suppose that G is a finitely presented group, and $H_i, \ldots, H_n$ are finitely presented subgroups. We can construct finite complexes, $K(G)$ and $K(H_i)$, with these groups as fundamental groups. We can assume that $K(G)$ contains disjoint embedded copies, $K(G, H_i)$, of each $K(H_i)$. Moreover, we can assume that each $K(G, H_i)$ has a neighbourhood homeomorphic to $K(H_i) \times [0, 1]$, with $K(G, H_i)$ embedded as $K(H_i) \times \{0\}$. We write $\tilde{K}(G)$ etc. for the universal covers of these spaces. We say that G is *one-ended relative to the subgroups* H_i if every compact subset of $\tilde{K}(G)$ whose complement has more than one unbounded component meets a lift of some $K(G, H_i)$. We shall need the following:

Lemma 3.1. *Suppose that a finitely presented group, Γ, splits as a finite graph of groups with all edge-stabilisers two-ended. Suppose each vertex group is one-ended relative to the incident edge groups, and is not itself two-ended. Then Γ is one-ended.*

Proof. For each vertex group, $\Gamma(v)$, and incident edge group, $\Gamma(e)$, we construct finite complexes, $K(\Gamma(v))$, $K(\Gamma(e))$ and $K(\Gamma(v), \Gamma(e))$ as described above, and glue them together, as in the proof of Lemma 1.1, to obtain a finite complex K with fundamental group Γ.

Let (V, E) be the simplicial tree arising from the graph of groups. Given $v \in V$ and $e \in E$, we write $\Gamma(v)$ and $\Gamma(e)$ for the vertex and edge stabilisers in Γ. For each v, we get copy, $\tilde{K}(v)$, of the universal cover of $K(\Gamma(v))$ embedded in $\tilde{K}$. We similarly get a subcomplex, $\tilde{K}(e)$, for each $e \in E$. Note that $\tilde{K}(e)$ is two-ended and separates $\tilde{K}$ into two components.

We first show that if $F \subseteq \tilde{K}(v)$ is a compact subset meeting only one edge complex, $\tilde{K}(e)$, then $\tilde{K}(v) \setminus F$ has precisely one unbounded component. There is no loss in taking F to be a finite subcomplex. Moreover, we can assume that each component of $\tilde{K}(v) \setminus F$ and of $\tilde{K}(e) \setminus F$ is unbounded. Now, each component of $\tilde{K}(v) \setminus F$ meets $\tilde{K}(e)$ (otherwise the boundary of such a component would violate the relative one-endedness hypothesis). If $\tilde{K}(v) \setminus F$ is not connected, we see that it must have precisely two components, say O and U, each meeting $\tilde{K}(e)$ in a

neighbourhood of an end. Let γ be an infinite order element of $\Gamma(e)$. We can suppose that $\tilde{K}(e) \subseteq \bigcup_{n=0}^{\infty} \gamma^{-n}U$. Now we can connect any point, x, in $\tilde{K}(v)$ to $\tilde{K}(e)$ by a compact path, $\alpha \subseteq \tilde{K}(v)$. For sufficiently large n, $\gamma^n\alpha \cap F = \emptyset$, and so $\gamma^n\alpha \subseteq U$, telling us that $\gamma^n x \in U$. We see that $\bigcup_{n=0}^{\infty} \gamma^{-n}U = \tilde{K}(v)$, and so $\bigcap_{n=0}^{\infty} \gamma^{-n}O = \emptyset$. Similarly, we see that $\bigcap_{n=0}^{\infty} \gamma^n U = \emptyset$. We see easily that $\tilde{K}(v) = \bigcup_{n=-\infty}^{\infty} \gamma^n F \cup \bigcup_{n=-\infty}^{\infty} \gamma^n (U \cap \gamma O)$. (Given $x \in \tilde{K}(v)$, consider $\min\{n \mid \gamma^n x \notin O\}$.) But now, $\mathrm{fr}(U \cap \gamma O) \subseteq F \cup \gamma F$ is compact. Also, $(U \cap \gamma O) \cap \tilde{K}(e)$ is relatively compact. It follows that $U \cap \gamma O$ is relatively compact in $\tilde{K}(v)$ (otherwise, pushing it slightly off $\tilde{K}(e)$, we would obtain an unbounded set in $\tilde{K}(v)$ whose boundary is compact and does not meet any edge complex). We conclude that $\tilde{K}(v)/\Gamma(e)$ is compact, and so $\Gamma(e)$ is of finite index in $\Gamma(v)$, giving that $\Gamma(v)$ is two-ended, contrary to the hypotheses.

Now, suppose that $F \subseteq \tilde{K}$ is compact. (We can again take it to be a finite complex.) We claim, by induction on the number of edge complexes which intersect F, that $\tilde{K} \setminus F$ has only one unbounded component. Suppose that $e \in E$. Then $\tilde{K}(e)$ separates $\tilde{K}$ into two pieces, say $\tilde{K}_1$ and $\tilde{K}_2$, whose boundary components, $\mathrm{fr}\,\tilde{K}_1$ and $\mathrm{fr}\,\tilde{K}_2$, are identified with $\tilde{K}(e)$. Now, from the tree structure, we can choose e such that $F \cap \tilde{K}_1$ meets $\mathrm{fr}\,\tilde{K}_1$ but no other edge complex in $\tilde{K}_1$. By the previous paragraph, we can see that the two ends of $\mathrm{fr}\,\tilde{K}_1$ lie in the same component of $\tilde{K}_1 \setminus F$. By the induction assumption (pushing $F \cap \tilde{K}_2$ slightly away from $\mathrm{fr}\,\tilde{K}_2$), we see that each unbounded component of $\tilde{K}_2 \setminus F$ contains at least one end of $\mathrm{fr}\,\tilde{K}_2$. On gluing back together, we see that $\tilde{K} \setminus F$ has only one unbounded component. Thus, by induction, $\tilde{K}$ is one-ended. $\qquad\square$

Another proof of Lemma 3.1 is given in [Bo6]. This uses a different definition of relative one-endedness (namely, G is one-ended relative to $H_1, \ldots, H_n$ if G does not split over any finite group relative to each H_i). That these notions are equivalent follows from a relative version of Stallings's theorem.

We now return to our geometrically finite group, Γ. We observe:

Lemma 3.2. *If Γ is geometrically finite, then the limit set, $\Lambda\Gamma$, is connected if and only if Γ is one-ended relative to its two-ended maximal parabolic subgroups.*

Proof. We have already observed that $\Lambda\Gamma$ is connected if and only if X_2 is one-ended. This is in turn equivalent to saying that X_0 is one-ended relative to the subsets $S(p)$ for $p \in \Pi_2$ (i.e. every compact set, $F \subseteq X_0$, whose complement has more than one unbounded component meets at least one $S(p)$). Since X_0/Γ is compact, it's not hard to see that the last statement is equivalent to saying that Γ is one-ended relative to the groups, $\Gamma(p)$, for $p \in \Pi_2$. $\qquad\square$

We only need the "only if" statement here. We omit the details, since we only need the result to feed into Lemma 3.1, and the argument of Lemma 3.1 can be applied directly to the construction we are about to describe, without translating it into group theoretical terms. A more careful proof of Lemma 3.2 can be found in [Bo9].

Recall that X_2 is obtained by removing from X all sets of the form $C(p)$ for $p \in \Pi_2$. Now $S(p)$ is connected, and $S(p)/\Gamma(p)$ is compact. Although it is not essential to the argument, it will be convenient for the purposes of exposition to assume that each such $S(p)$ is simply connected. If this is not already the case, it can be achieved by equivariantly attaching a number of discs to $S(p)$ in $\mathrm{fr}\,B(p)$,

and then thickening X_2 up slightly on the complement of $B(p)$. It now follows that X_2 is also simply connected. We write fr $X_2 = \bigcup_{p \in \Pi_2} S(p)$ for the topological boundary of X_2 in X. Note that each $S(p)$ is concave in the riemannian sense. It follows that in the induced path metric on X_2, each $S(p)$ is totally geodesic (in the metric space sense) and is quasiisometric to the real line. In particular, $S(p)$ is Gromov hyperbolic. It now follows that X_2 is intrinsically Gromov hyperbolic. This can be seen using the linear isoperimetric inequality. The idea is that if γ is a rectifiable loop in X_2, we can span it by a disc, D, in X, whose area is linearly bounded by the length of γ. If a portion of D enters a region $C(p)$ for $p \in \Pi_2$, we use the fact that $S(p)$ is hyperbolic to push it back into X_2. (This is essentially the same argument as that which shows that a group which is hyperbolic relative to a class of hyperbolic subgroups is intrinsically word hyperbolic; see for example [Bo8].) We write ∂X_2 for its ideal boundary. Now, Γ acts isometrically on X_2, and if $p \in \Pi_2$, then $\Gamma(p)$ is loxodromic on X_2. Write $L(p) \subseteq \partial X_2$ for its fixed point set. Thus, $S(p) \cup L(p)$ is closed in $X_2 \cup \partial X_2$. We obtain, topologically, the space $X \cup \Lambda\Gamma$ by collapsing $L(p)$ to the point p, for each $p \in \Pi_2$.

Now, X_2/Γ is an orbifold with orbifold fundamental group Γ, and with fr X_2/Γ a suborbifold of the orbifold boundary. We construct an orbifold, X_2'/Γ', by doubling X_2/Γ in fr X_2/Γ. Here, X_2' is the universal cover, and Γ' is the orbifold fundamental group of the double. The group, Γ', can be defined abstractly as a graph of groups with two vertex groups each isomorphic to Γ, connected by a set of two-ended edge groups given by (the conjugacy classes in Γ of) two-ended maximal parabolic subgroups. Note that, by Lemmas 3.1 and 3.2, if $\Lambda\Gamma$ is connected, then Γ' is one-ended. Also, Γ' contains no infinite torsion subgroup. (Every finite subgroup of Γ' is conjugate into one of the vertex groups.)

The space, X_2', is obtained by gluing together two copies of X_2 along the components of fr X_2. Since these components are totally geodesic (in the induced path metric), one can verify that X_2' is hyperbolic in the induced path metric. (This follows by the same argument, via the subquadratic isoperimetric inequality, that is given in [BeF2]. Given that, in our case, the sets $S(p)$ are quasiconvex, one can also give a more direct argument to show that every triangle has a centre, i.e. point which lies a bounded distance from each side.) Let $\partial X_2'$ be the ideal boundary of X_2'. We can obtain $X_2' \cup \partial X_2'$ by gluing together copies of $X_2 \cup \partial X_2$ along sets of the form $S(p) \cup L(p)$.

Since X_2' is hyperbolic, Γ' acts as a convergence group on $\partial X_2'$ (see, for example, [Bo7]). Since Γ' is one-ended, $\partial X_2'$ is connected. We are now effectively back in the case in which we have already proved Corollary 0.4 — with Γ' replacing Γ, X_2' replacing X, and $\partial X_2'$ replacing $\Lambda\Gamma$. (We have introduced some geometrically singular codimension-1 submanifolds into the picture but these have no bearing on the argument — they do not affect the topology or the large-scale geometry of the set-up.) We conclude that every global cut-point of $\partial X_2'$ is a parabolic fixed point of Γ'.

Now suppose, for contradiction, that $q \in \partial X \setminus \Pi$ is a global cut point of ∂X. We can write $\partial X \setminus \{q\} = A_0 \sqcup A_1$, where A_0 and A_1 are disjoint nonempty open subsets of ∂X. Let $E_i = A_i \cup \bigcup_{p \in A_i \cap \Pi_2} C(p)$. Thus, E_i is closed in $X \cup \partial X \setminus \{q\}$. Now, $X \cup \partial X \setminus \{q\}$ is metrisable, hence normal, so we can find disjoint open subsets, O_0 and O_1, with $E_i \subseteq O_i$. Let $Q = X \setminus (E_0 \cup E_1) \subseteq X_2$. Thus, $Q \cup \{q\}$ is closed in $X \cup \partial X$.

Now, consider X_2 as a hyperbolic space with boundary ∂X_2. As discussed above, we can view ∂X as a quotient of ∂X_2. Since $q \notin \Pi$, q can be identified as a point of ∂X_2. Moreover, $Q \cup \{q\}$ is closed in $X_2 \cup \partial X_2$. Let $V_i = U_i \cap X_2$. Note that $S(p) \subseteq V_i$ for all $p \in A_i \cap \Pi_2$. Since Π_2 is dense in ∂X, $A_i \cap \Pi_2$ is non-empty.

We now imagine X_2 as isometrically embedded in X_2'. Any point of $X_2' \setminus X_2$ can be connected to X_2 by a path meeting exactly one set $S(p)$ for $p \in \Pi_2$. Moreover, this point p is uniquely determined. Let $W_i \subseteq X_2'$ consist of the set V_i together with all those points of X_2' for which this p lies in $\Pi_2 \cap A_i$. Thus, W_i is open in X_2', and $X_2' \setminus Q = W_0 \sqcup W_1$.

Now, we can identify ∂X_2 as a closed subset of $\partial X_2'$. Let Z_i be the closure of W_i in $\partial X_2' \setminus \{q\}$, and let $F_i = W_i \cap \partial X_2'$. Thus, F_i is non-empty and open in $\partial X_2'$ and $\partial X_2' \setminus \{q\} = F_0 \sqcup F_1$. Thus, q is a global cut point of $\partial X_2'$.

It now follows that q is fixed by an element $\gamma \in \Gamma'$, which is parabolic on X_2'. We see that γ fixes X_2 setwise; in other words, $\gamma \in \Gamma$. Now γ has no periodic points in $\partial X_2' \setminus \{q\} \supseteq \partial X_2 \setminus \{q\}$. It thus has no fixed points in $\partial X \setminus \{q\}$. It follows that γ is parabolic for the action of Γ on X, giving us the contradiction that $q \in \Pi$.

This proves Corollary 0.4 in the general case.

4. Proof of Theorem 0.2

We can adapt the proof of Corollary 0.4 easily to deal with the case of relatively hyperbolic groups. Details of the relevant constructions can be found in [Bo8].

Suppose that Γ is a group and that $\mathcal{G}$ is a set of infinite finitely generated subgroups of $\mathcal{G}$. We shall say that Γ is *hyperbolic relative to* $\mathcal{G}$ if there exists a proper (complete locally compact) hyperbolic path-metric space, X, and a geometrically finite action of Γ on X such that $\mathcal{G}$ is precisely the set of maximal parabolic subgroups. From this, one can deduce that $\mathcal{G}$ is a union of finitely many conjugacy classes of subgroups, that each element of $\mathcal{G}$ is equal to its normaliser, and that the intersection of any two distinct elements of $\mathcal{G}$ is a finite group. We refer to $\mathcal{G}$ as a "peripheral structure" on Γ, and to an element of $\mathcal{G}$ as a "peripheral subgroup". The term "geometrically finite" can be defined by saying that each point of the limit set, $\Lambda\Gamma \subseteq \partial X$, is either a conical limit point or a bounded parabolic fixed point. It turns out that this limit set depends, up to Γ-equivariant homeomorphism, only on Γ and $\mathcal{G}$. It thus gives us a well-defined boundary, $\partial(\Gamma, \mathcal{G})$, of the relatively hyperbolic group, $(\Gamma, \mathcal{G})$. In fact, we can assume that X is "taut" (i.e. that every point of X lies a bounded distance from a biinfinite geodesic) and that the action of Γ on ∂X is minimal, so that the limit set equals ∂X.

It turns out that we can make a number of other simplifying assumptions about the space, X, without any loss of generality. Since we have already dealt with hyperbolic groups, we shall assume that $\mathcal{G} \neq \emptyset$ (though the distinction is artificial — it will be clear that one can give a unified argument). In this case, we can take X to be a 2-complex consisting of a set of ideal hyperbolic triangles glued together so that only finitely many meet along any given edge, and such that the union of two adjacent ideal triangles is an ideal hyperbolic square. Now, Γ acts isometrically on X, and there are finitely many orbits of such triangles. Let $\Pi \subseteq \partial X$ be the set of parabolic points. Thus $\mathcal{G} = \{\Gamma(p) \mid p \in \Pi\}$, where $\Gamma(p)$ is the stabiliser of p. Each triangle is isometrically embedded in X and its three ideal points all lie in Π. We note that we can choose a Γ-invariant system of disjoint horoballs, $(B(p))_{p \in \Pi}$, such that for each $p \in \Pi$, the boundary $S(p) = \text{fr } B(p)$ is a locally finite graph

with $S(p)/\Gamma(p)$ finite, and $B(p)$ is a hyperbolic cone over $S(p)$. Moreover, we can assume that each horoball $B(p)$ is convex in X, and that the distance between distinct horoballs is arbitrarily large in relation to fixed hyperbolicity constants, as well as in relation to the diameters of the graphs $S(p)/\Gamma(p)$.

The proof now proceeds almost exactly as with that of Corollary 0.4, where the complex, X, plays the role of the convex hull of the limit set. It is no longer true that X is simply connected. However, it is "quasi simply connected" in the sense that every loop can be spanned by a disc with holes of a bounded size. In any case, the precise topology of X only served as a convenient language for expressing what were essentially combinatorial or geometric constructions. As before, we can construct the space X_2 by removing all horoballs with two-ended stabilisers. The space X_2' can be constructed explicitly by gluing together copies of X_2' along two-ended subsets which are intrinsically geodesic (in fact, quasigeodesic would be enough). The pieces are glued together in treelike fashion, as dictated by the structure of the doubled group, Γ'. Now, Γ acts cocompactly on X_2'. It is easily verified that X_2' is hyperbolic, and that the action of Γ' on X_2' is geometrically finite. (One can also give a cleaner combinatorial verification of these statements using the ideas of [Bo9].) The remainder of the argument proceeds exactly as with Corollary 0.4.

References

[Ab]   H.Abels, *An example of a finitely presented solvable group*, in "Homological group theory", London Math. Society Lecture Notes Series, No. 36 (ed. C.T.C.Wall), Cambridge University Press (1979) 205–211. MR **82b:**20047

[AnM]   J.W.Anderson, B.Maskit, *On the local connectivity of limit sets of kleinian groups*, Complex Variables **31** (1996) 177–183. MR **98a:**30055

[BeF1] M.Bestvina, M.Feighn, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. **103** (1991) 449–469. MR **92c:**20044

[BeF2] M.Bestvina, M.Feighn, *A combination theorem for negatively curved groups*, J. Differential Geom. **35** (1992) 85–101. MR **93d:**53053

[BeF3] M.Bestvina, M.Feighn, *Stable actions of groups on real trees*, Invent. Math. **121** (1995) 287–321. MR **96h:**20056

[BeM]   M.Bestvina, G.Mess, *The boundary of negatively curved groups*, J. Amer. Math. Soc. **4** (1991) 469–481. MR **93j:**20076

[Bi]   R.Bieri, *Homological dimension of discrete groups*, Queen Mary College Mathematics Notes (1976). MR **57:**6224

[Bo1]   B.H.Bowditch, *Discrete parabolic groups*, J. Differential Geom. **38** (1993) 559–583. MR **94h:**53046

[Bo2]   B.H.Bowditch, *Geometrical finiteness with variable negative curvature*, Duke Math. J. **77** (1995) 229–274. MR **96b:**53056

[Bo3]   B.H.Bowditch, *Treelike structures arising from continua and convergence groups*, Memoirs Amer. Math. Soc. CMP 98:05

[Bo4]   B.H.Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186. CMP 98:17

[Bo5]   B.H.Bowditch, *Group actions on trees and dendrons*, Topology **37** (1998), no. 6, 1275–1298. CMP 98:15

[Bo6]   B.H.Bowditch, *Boundaries of strongly accessible hyperbolic groups*, in "The Epstein Birthday Schrift", Geometry and Topology Monographs, Vol. 1 (ed. I.Rivin, C.Rourke, C.Series) International Press (1998) 59–97.

[Bo7]   B.H.Bowditch, *Convergence groups and configuration spaces*, to appear in "Group Theory Down Under" (ed. J.Cossey, C.F.Miller, W.D.Neumann, M.Shapiro), de Gruyter.

[Bo8]   B.H.Bowditch, *Relatively hyperbolic groups*, preprint, Southampton (1997).

[Bo9]   B.H.Bowditch, *Peripheral splittings of groups*, preprint, Southampton (1997).

[Bo10] B.H.Bowditch, *Boundaries of geometrically finite groups*, to appear in Math. Z.

[BoS] B.H.Bowditch, G.A.Swarup, *Cut points in the boundaries of hyperbolic groups*, in preparation.

[D]   M.J.Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985) 449–457. MR **87d:**20037

[DS]   M.J.Dunwoody, M.E.Sageev, *JSJ splittings for finitely presented groups over slender subgroups*, Invent. Math. **135** (1999) 25–44.

[F]   E.M.Freden, *Negatively curved groups have the convergence property I*, Ann. Acad. Sci. Fenn. Ser. A Math. **20** (1995) 333-348. MR **96g:**20054

[GeM] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I*, Proc. London Math. Soc. **55** (1987) 331–358. MR **88m:**30057

[GhH] E.Ghys, P.de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov*, Progress in Maths. 83, Birkhäuser (1990). MR **94m:**53060

[Gr]   M.Gromov, *Hyperbolic groups*, in "Essays in Group Theory" (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263. MR **89e:**20070

[L]   G.Levitt, *Non-nesting actions on real trees*, Bull. London Math. Soc. **30** (1998) 46–54. MR **99a:**20027

[M]   M.Mihalik, *Semistability at ∞ of finitely generated groups, and solvable groups*, Topology and its Appl. **24** (1986) 259–269. MR **88c:**57005

[Se]   Z.Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II*, Geom. Funct. Anal. **7** (1997) 561–593. MR **98j:**20044

[Sw]   G.A.Swarup, *On the cut point conjecture*, Electron. Res. Announc. Amer. Math. Soc. **2** (1996) 98-100 (Electronic). MR **97f:**20048

[T]   P.Tukia, *Convergence groups and Gromov's hyperbolic metric spaces*, New Zealand J. Math. **23** (1994) 157-187. MR **96c:**30042

Faculty of Mathematical Studies, University of Southampton, Highfield, Southampton SO17 1BJ, Great Britain

E-mail address: bhb@maths.soton.ac.uk