

Boundaries of strongly accessible hyperbolic groups.

B. H. Bowditch

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

bhb@maths.soton.ac.uk

[First draft: March 1996. Revised: November 1997]

Abstract.

We consider splittings of groups over finite and two-ended subgroups. We study the combinatorics of such splittings using generalisations of Whitehead graphs. In the case of hyperbolic groups, we relate this to the topology of the boundary. In particular, we give a proof that the boundary of a one-ended strongly accessible hyperbolic group has no global cut point.

0. Introduction.

In this paper, we shall be interested in splittings of groups over finite and two-ended (i.e. virtually cyclic) subgroups. We shall study combinatorial aspects of such splittings, and in the case of hyperbolic groups, relate this to the topology of the boundary. This fits into a larger programme of analysing such splittings, as described, for example, in [Bo2,Bo4,Bo5].

In the first three sections, we shall study the combinatorics of such splittings in a fairly general context. We consider generalisations of Whitehead graphs as defined in [W]. We use this to give criteria for when a group splits over a finite group relative to a set of two-ended subgroups (Propositions 3.6 and 3.8). In the case of hyperbolic groups we relate this to certain quotients of the boundary. To this end, we generalise ideas of Otal, who studied such questions in the case of a free group [O]. We also study graphs of groups with two-ended edge groups, and in particular, give a criterion for the fundamental group of such a graph to be one-ended. In the case where the vertex groups are free or surface groups, such issues were addressed in [Ma].

In later sections, we give a more detailed analysis of the topology of the boundaries of hyperbolic groups, in relation to splittings over finite and two-ended subgroups. We prove that the boundary of a one-ended strongly accessible hyperbolic group has no global cut point in its boundary (Theorem 9.3). In view of the recent results of Delzant and Potyagailo [DeP], this in fact deals with all hyperbolic groups. Shortly after the original draft of this paper was circulated, Swarup [Swa] gave a direct proof of the general case, using the results of [Bo1,Bo3,L]. This has been elaborated on in [Bo4] which has applications also to relatively hyperbolic groups, as outlined, for example in [Bo5]. It is hoped that some of the ideas of this paper might be applicable to this case.

We shall set out the results of this paper more carefully, beginning with a discussion of strong accessibility.

The notion of accessibility expresses the idea that, under appropriate hypotheses, it should be impossible to split indefinitely a fixed group over a given class of subgroups. Thus, if Γ is a group, and \mathcal{C} is a given class of subgroup (usually constrained by some intrinsic algebraic property) we say that Γ is “accessible over \mathcal{C} ” if we can represent Γ as (the fundamental group of) a finite graph of groups such that none of the vertex groups can be split further over \mathcal{C} relative to incident edge groups. (Such a splitting would give rise to a refined graph of groups. Indeed this is equivalent to the non-existence of a refinement, provided we add a condition banning “reducible” vertices — cf. [BeF1].) Of course, it remains possible that a vertex group may split intrinsically over a subgroup in \mathcal{C} , if we forget how the incident edge groups of our graph were attached. Indeed it is conceivable that one may be able to iterate this process indefinitely. We say that Γ is “strongly accessible” over \mathcal{C} if we can perform such splittings in such a way as to reduce us, in a finite number of steps, to a finite number of conjugacy classes of subgroups, none of which split non-trivially over \mathcal{C} . (Of course, it remains to address the issue of whether there might still exist an alternative sequence of splittings over \mathcal{C} which does not terminate.) Recently Delzant and Potyagailo have shown that any finitely presented group, Γ , is strongly accessible over \mathcal{C} provided that \mathcal{C} is “elementary” (i.e. no element of \mathcal{C} contains a non-cyclic free group, each infinite element of \mathcal{C} is contained in a unique maximal element of \mathcal{C} , and each maximal element of \mathcal{C} is equal to its normaliser in Γ). We shall confine our attention to the case where \mathcal{C} is the class of finite and two-ended (i.e. virtually cyclic) groups. This is particularly natural in the study of hyperbolic groups (in the sense of Gromov [Gr]). Note that this case is included in the result of [DeP].

Stallings’s theorem [St] tells us that an infinite finitely generated group is one-ended if and only if it does not split over any finite subgroup. Dunwoody’s theorem [Du] tells us that any finitely presented group is accessible over finite subgroups. (Over the class of finite groups, the notions of accessibility and strong accessibility coincide.) Thus, we may represent a finitely presented group as a finite graph of groups where all the edge groups are finite, and all the vertex groups are finite or one-ended. Moreover, the vertex groups are finitely presented. The infinite vertex groups are precisely the maximal one-ended subgroups, and hence canonically determined.

Now, a one-ended finitely presented group admits a canonical splitting over two-ended subgroups, namely the JSJ splitting. (See [RS,DuS,FuP], or in the context of hyperbolic groups, [Se,Bo2].) The vertex groups are again finitely presented. (It is a general fact that if a finitely presented group splits over a finitely presented subgroups, then all the vertex groups are again finitely presented — see, for example, [Bo4].) The vertex groups might be one, two or infinite-ended, so we can start all over again, splitting all the vertex groups into finite or one-ended components over finite groups. In this way we get a canonical splitting process. In this paper, we shall not make any explicit use of the JSJ splitting. We will work with arbitrary splittings over two-ended subgroups. Of course, this introduces some ambiguity into the splitting process. The extent to which such choices are significant is explored in Section 9.

This leads naturally on the question of when a finite graph of groups with two-ended edge groups gives rise to a one-ended group. This turns out to be the case if and only if none of the vertex groups splits over a finite group relative to the incident edge groups (Theorem

2.3). This criterion is only non-vacuous for those vertex group which are infinite-ended. The case where all vertex groups are free is considered in [Ma]. We aim to generalise some of these ideas to more general groups.

This, in turn, leads us to criteria for recognising when a given group, G , does not split over a finite group relative to a given finite set of two-ended subgroups of G . In the case where G is accessible over finite groups, we may as well assume that all the two-ended subgroups are “hyperbolic” relative to a maximal splitting of G over finite subgroups. One condition therefore relates to the combinatorics of the axes of these subgroups in the corresponding G -tree. The case of free groups is fairly classical. It relates to work of Whitehead and was elaborated on by Otal.

To give the idea, let’s begin by considering the case of a free group, F , and a single non-trivial element, $\gamma \in F$. We want to recognise when γ is “indecomposable” i.e. cannot be conjugated into any proper free factor of F . This can be interpreted topologically. Note that the boundary, ∂F , of F is a cantor set. We define an equivalence relation, \approx , on ∂F , by deeming that $x \approx y$ if and only if either $x = y$ or x and y are the fixed points of some conjugate of γ . Now, it’s easily verified that this relation is closed, and so the (equivariant) quotient, $\partial F/\approx$ is compact hausdorff. It was shown in [O] that γ is indecomposable if and only if $\partial F/\approx$ is connected. (Otal goes on to show that, in this case, $\partial F/\approx$ is locally connected and has no global cut point.)

A combinatorial criterion for indecomposability is formulated in [W] (though not explicitly stated in quite these terms). Let a_1, a_2, \dots, a_n be a system of free generators for F . Let w be a reduced cyclic word in the a_i ’s and their inverses representing (the conjugacy class of) γ . Let \mathcal{G} be the graph with vertex set $a_1, \dots, a_n, a_1^{-1}, \dots, a_n^{-1}$, and with $a_i^{\epsilon_i}$ deemed to be adjacent to $a_j^{\epsilon_j}$ if and only if the string $a_i^{\epsilon_i} a_j^{-\epsilon_j}$ occurs somewhere in w (where $\epsilon_i, \epsilon_j \in \{-1, 1\}$). The graph \mathcal{G} is commonly referred to as the “Whitehead graph”. Suppose we choose the generating set so as to minimise the length of the word w . Then (a simple consequence of) Whitehead’s lemma tells us that γ is indecomposable if and only if \mathcal{G} is connected. Moreover in such a case, \mathcal{G} has no cut vertex. (In fact, one can give an algorithm to obtain such a minimising set of generators from any given one.)

This can be reinterpreted in terms of what we shall call “arc systems”. Let T be the Cayley graph of F with respect to free generators $a_1 \dots a_n$. Thus, T is a simplicial tree, whose boundary, ∂T , may be naturally identified with ∂F . The element γ determines a biinfinite arc, β , in T , namely the axis of γ . Let \mathcal{B} be the set of images of β under Γ . We refer to \mathcal{B} as a (Γ -invariant) “arc system”. We can reconstruct the Whitehead graph from this arc system in a simple combinatorial fashion, as described in Section 3. Note that each arc of \mathcal{B} determines a pair of points in ∂T . We refer to the set of unordered pairs arising in this way as a “formal arc system”. Note that this is essentially the same structure as the equivalence relation, \approx on ∂F , described above.

Note that all the above discussion applies equally well if we replace γ by a finite set, $\{\gamma_1, \dots, \gamma_p\}$, of non-trivial elements. This set is said to be indecomposable if F cannot be represented as a free product in such a way that each element is conjugate into a vertex group. This can again be interpreted in terms of arc systems.

One can generalise these notions to arbitrary accessible groups (i.e. accessible over finite subgroups), in particular, to finitely presented groups. This is most conveniently

described in the particular case of a hyperbolic group. Suppose Γ is an infinite-ended hyperbolic group. We can decompose its boundary, $\partial\Gamma$, as a disjoint union of two Γ -invariant sets, $\partial_0\Gamma$ and $\partial_\infty\Gamma$, where $\partial_\infty\Gamma$ is the set of singleton components of $\partial\Gamma$. Algebraically this corresponds to the action of Γ on a simplicial tree, T , without edge inversions, with finite quotient, and with finite edge stabilisers and finite or one-ended vertex stabilisers. Such an action is given by the accessibility theorem [Du]. Each of the vertex groups is quasiconvex, and hence intrinsically hyperbolic. Now, $\partial_\infty\Gamma$ can be canonically identified with ∂T , and the connected components of $\partial_0\Gamma$ are precisely the boundaries of the infinite vertex stabilisers. The infinite vertex stabilisers are, in fact, precisely the maximal one-ended subgroups of Γ . Note that Γ is virtually free if and only if $\partial_0\Gamma = \emptyset$.

Suppose now that $H \subseteq \Gamma$ is a two-ended subgroup of Γ . Its limit set is an unordered pair of points of $\partial\Gamma$. We say that H is “hyperbolic” if these two endpoints lie in $\partial_\infty\Gamma$. Otherwise, they both lie in a component of $\partial_0\Gamma$, and we refer to H as “elliptic”. Suppose that \mathcal{H} is a finite set of two-ended subgroups. We obtain an equivalence relation, \approx , on $\partial\Gamma$, by identifying the limit points of each conjugate of each element of \mathcal{H} , exactly as in the case of a free group. As before, $\partial\Gamma/\approx$, is compact hausdorff (see Section 5). We shall see that Γ splits over a finite group relative to \mathcal{H} if and only if $\partial\Gamma/\approx$ is connected.

We shall also give a combinatorial means of recognising when Γ splits in this way. Only those elements of \mathcal{H} which are hyperbolic play any part in this, so we may as well assume that they are all hyperbolic. Thus \mathcal{H} determines a formal arc system on $\partial_\infty\Gamma$ which we are identifying with ∂T . We thus get an arc system, \mathcal{B} , on T , where the arcs are precisely the axes of conjugates of elements of \mathcal{H} . From this arc system we may construct analogues of Whitehead graphs, which give us a combinatorial mean of describing indecomposability. In fact, this discussion applies equally well to any accessible group Γ , as discussed in Section 3.

Put together with Theorem 2.3 discussed above, this gives a combinatorial means of recognising the one-endedness of a group represented as a graph of groups over two-ended subgroups, as for example, in the JSJ decomposition. This therefore generalises some of the results of Martinez [Ma]. It is also natural to ask when a group represented in this way is (word) hyperbolic, assuming that all the vertex groups are intrinsically (word) hyperbolic. One can give a fairly complete answer to this question from the criterion given in [BeF2]. In this particular case, it is sufficient to ban Baumslag-Solitar subgroups (including rank-2 free abelian subgroups), which might arise from a cycle in the graph which conjugates an infinite order edge-stabilising element to itself, or to a power of itself.

The original motivation for this study was to give a proof of the cut-point conjecture for strongly accessible hyperbolic groups, based on the results of [Bo1] and [Bo3]. Shortly afterwards, Swarup [Swa] showed how to adapt to these ideas to deal with the general case, without assuming strong accessibility. This has been elaborated on in [Bo4]. Other approaches to some of these constructions have been suggested by the work of Levitt [L] and Swenson [Swe]. Some of these results have applications also to relatively hyperbolic groups, as outlined, for example, in [Bo5]. In this context, global cut points can indeed occur. With the possibility that these ideas may have some relevance to this broader setting, we include our original proof, but relegate it to Section 8. As remarked earlier, we see, in retrospect, from the work of [DeP], that this does indeed take care of all one-ended

hyperbolic groups.

The cut point conjecture, formulated in [BeM], asserts that the boundary of a one-ended hyperbolic group has no global cut point. This is equivalent to local connectedness of the boundary [BeM,Bo2], which has a number of algebraic implications (see for example the construction of the JSJ splitting in [Bo2]). It is also closely related to semistability at infinity (see [Mi]), as observed by Geoghegan, and reported in [BeM]. The semistability of an accessible group is equivalent to the semistability of each of its maximal one-ended subgroups. Suppose, then, that Γ is a one-ended hyperbolic group. It was shown in [BeM] that $\partial\Gamma$ naturally compactifies the Rips complex, so as to give a contractible ANR, with $\partial\Gamma$ embedded as a Z -set. It follows that semistability at infinity for Γ is equivalent to $\partial\Gamma$ being pointed 1-movable, the latter property being intrinsic to $\partial\Gamma$. Moreover, it was shown in [Kr] that a metrisable continuum is pointed 1-movable if and only if it has the shape of a Peano continuum. (Generalisations to higher dimensional connectivity are given in [Fe].) It follows that if Γ is one-ended hyperbolic, then $\partial\Gamma$ is semistable at infinity if and only if $\partial\Gamma$ has the shape of a Peano continuum. (I'm indebted to Ross Geoghegan for explaining to me how this works.) The cut point conjecture thus implies that all hyperbolic groups are semistable at infinity. It has been conjectured that all finitely presented groups are semistable at infinity.

As mentioned above, the cut point conjecture is now known [Bo1,Bo3,L,Swa,Bo4]. In Section 8, we show that if a one-ended hyperbolic group, Γ , splits as a finite graph of groups over two-ended subgroups, then any global cut point of $\partial\Gamma$ must arise a global cut point in the boundary of a maximal one-ended subgroup of one of the vertex groups of this splitting. Thus together with the results of [Bo1,Bo3], this implies the cut point conjecture for strongly accessible groups. We remark that, with the main result of [MiT] in place of Theorem 8.1 of this paper, one can deduce semistability at infinity for strongly accessible hyperbolic groups directly from the results of [Bo1,Bo3,MiT].

This paper is a reworking of an earlier preprint of the same title (March 1996), prepared at the university of Melbourne. The main focus of the original was the proof of Theorem 9.3. In the present paper (revised November 1997), we have elaborated considerably on the question of splittings relative to two-ended subgroups, and how this is reflected in the topology of the boundary. The main discussion of cut points thus appears in Section 8, though some of the results of Section 7 lead up to this.

The structure of this paper is roughly as follows. In Section 1, we explore some general facts about groups accessible over finite groups. In Section 2, we give a criterion (Theorem 2.3) for a finite graph of groups with two-ended edge groups to be one-ended. In Section 3, we study arc systems on trees and their connections to Whitehead graphs. In Section 4, we give an overview of some general facts about quasiconvex splittings. In Section 5, we look at certain quotients of the boundaries of hyperbolic groups, and relate this to some of the combinatorial results of Section 3. In Section 6, we set up some of the general machinery for analysing the topology of the boundaries of hyperbolic groups which split over two-ended subgroups. In Section 7, we look at some implications concerning connectedness properties of boundaries. In Section 8, we apply this specifically to global cut points. Finally, in Section 9, we discuss further the question of strong accessibility of groups over finite and two-ended subgroups.

Much of the material of the original version of this paper was worked out while visiting the University of Auckland. The first draft was written at the University of Melbourne. I would like to thank Gaven Martin as well as Craig Hodgson and Walter Neumann for their respective invitations. The paper was substantially revised in Southampton, with most of the material of Sections 1, 2, 3 and 5 added. In particular, I would like to thank Martin Dunwoody for helpful conversations in relation to such matters.

1. Trees and splittings.

In this section, we introduce some terminology and notation relating to simplicial trees and group splittings.

Let T be a simplicial tree, which we regard a 1-dimensional CW-complex. We write $V(T)$ and $E(T)$ respectively for the vertex set and edge set. Given $v, w \in V(T)$, we write $\text{dist}(v, w)$ for the distance between v and w , in other words, the number of edges in the arc connecting v to w . If $\vec{e} \in \vec{E}(T)$ and $v \in V(T)$, we say that \vec{e} “points towards” v if $\text{dist}(v, \text{tail}(\vec{e})) = \text{dist}(v, \text{head}(\vec{e})) + 1$.

If $S \subseteq T$ is a subgraph, we write $V(S) \subseteq V(T)$ and $E(S) \subseteq E(T)$ for the corresponding vertex and edge sets. A subtree of T is a connected subgraph. Of particular interest are “rays” and “biinfinite arcs” (properly embedded subsets homeomorphic to $[0, \infty)$ and \mathbf{R} respectively.)

We may define the ideal boundary, ∂T , of T , as the set of cofinality classes of rays in Σ . We shall only be interested in ∂T as a set. (In fact, $T \cup \partial T$ can be given a natural compact topology as a dendron, as discussed in [Bo1]. It can also be given a finer topology by viewing T has a Gromov hyperbolic space, and ∂T as its Gromov boundary.) If $S \subseteq T$ is a subgraph, we write $\partial S \subseteq \partial T$ for the subset arising from those rays which lie in S . Note that if β is a biinfinite arc, then $\partial\beta$ contains precisely two points, $x, y \in \partial T$. We say that β connects x to y .

Further discussion of general simplicial trees will be given in Sections 2 and 3. We now move on to consider group actions on trees.

Let G be a group. A G -tree is a simplicial tree, T , admitting a simplicial action of G without edge inversions. If $v \in V(T)$ and $e \in E(T)$, we write $G_T(v)$ and $G_T(e)$ for the corresponding vertex and edge stabilisers respectively. Where there can be no confusion, we shall abbreviate these to $G(v)$ and $G(e)$. Such a tree gives rise to a splitting of G as a graph of groups, G/T . We shall say that T is *cofinite* if T/G is finite. We shall usually assume that T is *minimal*, i.e. that there is no proper G -invariant subtree. This is the same as saying that T has no terminal vertex, or, on the level of the splitting, that no vertex group of degree one is equal to the incident edge groups. Such a vertex will be referred to as a *trivial vertex*. A subset (usually a subgroup) H , of G is *elliptic* with respect to T , if it lies inside some vertex stabiliser. If \mathcal{H} is a set of subsets of G , we say that the splitting is *relative to* \mathcal{H} , if every element of \mathcal{H} is an elliptic subset. We note that any finite subgroup of a group is elliptic with respect to every splitting. Thus any splitting of any group is necessarily relative to the set of all finite subgroups.

Suppose that F is a G -invariant subgraph of T , we can obtain a new G -tree, Σ , by

collapsing each component of F to a point. We speak of the splitting T/G as being a *refinement* of the splitting Σ/G . Note that one may obtain a refinement of a given graph of groups, if one of the vertex groups splits relative to its incident edge groups.

We say that a G -tree, T' , is a *subdivision* of T , if it is obtained by inserting degree-2 vertices into the edges of T in a G -equivariant fashion. Suppose that Σ is another G -tree. A *folding* of T onto Σ is a G -equivariant map of T onto Σ such that each edge of T gets mapped homeomorphically onto an edge of Σ . A *morphism* of T onto Σ is a folding of some subdivision of T . Such maps are necessarily surjective provided that Σ is minimal. Clearly a composition of morphisms is a morphism.

We say that T *dominates* Σ (or that the splitting T/G *dominates* Σ/G) if there exists a morphism from T to Σ . It's not hard to see that this is equivalent to saying that every vertex stabiliser in T is elliptic with respect to Σ . We say that T and Σ are *equivalent* if each dominates the other. This is equivalent to saying that a subset of G is elliptic with respect to T if and only if it is elliptic with respect to Σ .

Suppose that T is cofinite. If T dominates Σ , then Σ is also cofinite. In this case, any morphism from T to Σ expands combinatorial distances by at most a bounded factor (namely the maximum number of edges into which we need to subdivide a given edge of T to get a folding.) Also, any two morphisms remain a bounded distance apart. In particular, any self-morphism of a cofinite tree is a bounded distance from the identity map, and is thus a quasiisometry. Suppose that T and Σ are equivalent, and that $\phi : T \rightarrow \Sigma$ is a morphism. Let $\psi : \Sigma \rightarrow T$ be any morphism. Now, since ψ expands distances by a bounded factor, and $\psi \circ \phi$ is a quasiisometry, it follows that ϕ is itself a quasiisometry. In summary, we have shown:

Lemma 1.1 : *If T and Σ are equivalent cofinite G -trees, then any morphism from T to Σ is quasiisometry.* \diamond

We see from the above discussion that there is a natural bijective correspondence between the boundaries, ∂T and $\partial \Sigma$, of T and Σ .

This is of particular interest in the case of accessible groups. We shall discuss these in more detail shortly. For the moment, we just note that if G is accessible (over finite groups), then there exists a cofinite G -tree with finite edge stabilisers such that all vertex stabilisers are either finite or one-ended. In this case, the infinite vertex groups are precisely the maximal one-ended subgroups, and hence canonically determined. We have also observed that finite groups are always elliptic in any splitting. It follows that if T' is another such G -tree, then T and T' are equivalent. In particular ∂T and $\partial T'$ can be canonically (and hence G -equivariantly) identified. We can thus associate to any accessible group, G , a canonical G -set, $\partial_\infty G$, which we may identify with the boundary of any such G -tree.

Clearly in the case of a free group, we just recover the usual boundary. More generally, if G is (word) hyperbolic (and hence accessible) then we may identify $\partial_\infty G$ with the set of singleton components of the boundary, ∂G . In fact, as discussed in the introduction, we can write ∂G as a disjoint union $\partial_0 G \sqcup \partial_\infty G$, where each component of $\partial_0 G$ is the boundary of a maximal one-ended subgroup of G .

We shall need a further observation regarding morphisms. Suppose that T and Σ are equivalent cofinite G -trees with finite edge stabilisers, and that $\phi : T \rightarrow \Sigma$ is any

morphism. Now, if ϵ is an edge of the subdivided tree, T' , and $\gamma \in \Gamma$ is such that ϵ and $\gamma\epsilon$ get mapped to the same edge of Σ , then $\gamma \in \Gamma_\Sigma(\phi\epsilon)$. Now, $\Gamma_\Sigma(\phi\epsilon)$ is finite, and T/Γ is finite. We therefore conclude that $\phi^{-1}e$ must be finite for each edge, e , of Σ . We deduce:

Lemma 1.2 : *Suppose that T and Σ are equivalent cofinite G -trees with finite edge-stabilisers. If $\phi : T \rightarrow \Sigma$ is a morphism, and $S \subseteq \Sigma$ is a finite subtree, then $\phi^{-1}S$ is contained in a finite subtree of T . \diamond*

We shall need to elaborate a little on the notion of accessibility over finite groups. For the remainder of this section, all splittings will be assumed to be over finite groups, and the term “accessible” is assumed to mean “accessible over finite groups”.

We shall say that a graph of groups is *reduced* if no vertex group of degree one or two is equal to an incident edge group. (Every graph of groups is a refinement of a reduced graph.) We say that a group G is “accessible” if there is a bound on the the complexity (as measured by the number of edges) of a splitting of G as a reduced graph of groups (with finite edge groups). Among graphs of maximal complexity, one for which the sum of the orders of the edge stabilisers is minimal will be referred to as a “complete splitting”. By Dunwoody’s theorem [Du], any finitely presented group is accessible. (This has been generalised to splittings over small subgroups by Bestvina and Feighn [BeF1].)

This can be rephrased in terms of one-ended subgroups. For this purpose, we define a group to be *one-ended* if it is infinite and does not split non-trivially (over any finite subgroup). Thus, by Stallings’s theorem, this coincides with the usual topological notion for finitely generated groups. Suppose that G is accessible, and we take a complete splitting of G . Now any splitting of a vertex group is necessarily relative to the incident edge groups, and so would give rise to a refined splitting. It is possible that this refined splitting may no longer be reduced, but in such a case, we can coalesce two vertex groups, to produce a reduced graph with one smaller edge stabiliser than the original, thereby contradicting completeness. In summary, we see that all the vertex groups of a complete splitting are either finite or one-ended. In fact, we see that the infinite vertex groups are precisely the maximal one-ended subgroups. It turns out that there is a converse to this statement: any group which can be represented as a finite graph of groups with finite edge groups and with all vertex groups finite or one-ended is necessarily accessible (see [DiD]).

Finally, suppose that G is accessible, and we represent it as a finite graph of groups over finite subgroups. Now each vertex group must be accessible. Taking complete splittings of each of the vertex groups, we can see that we can refine the original splitting in such a way that all the vertex groups are finite or one-ended. (It is possible that this refinement might not be reduced.)

We shall make some further observations about accessible groups in connection with strong accessibility in Section 9.

2. Splittings over two-ended subgroups.

The main aim of this section will be to give a proof of Theorem 2.3. We first introduce some terminology regarding “arc systems” which will be relevant to later sections.

Let T be a simplicial tree.

Definition : An *arc system*, \mathcal{B} , on T consists of a set of biinfinite arcs in T .

We say that \mathcal{B} is *edge-finite* if at most finitely many elements of \mathcal{B} contain any given edge of T .

If G is a group, and T is a G -tree, then we shall assume that an arc system on T is G -invariant.

Recall that a subgroup, H , of G is “elliptic” if it fixes a vertex of T . If H is two-ended (i.e. virtually cyclic) then either H is elliptic, or else there is a biinfinite β in T which is H -invariant. In the latter case, we say that H is *hyperbolic* and that β is the *axis* of H . Clearly, the H -stabiliser of any edge of \mathcal{B} is finite.

Suppose now that all edge stabilisers of T are finite. Then every hyperbolic two-ended subgroup of G lies in a unique maximal two-ended subgroup of G , namely the setwise stabiliser of the axis. Note also that there are only finitely many two-ended subgroups, H , with a given axis, \mathcal{B} , and with the number of edges of β/H bounded. In particular, we see that only finitely many G -conjugates of a given hyperbolic two-ended subgroup, H , can share the same axis.

Suppose, now, that \mathcal{H} is a finite union of conjugacy classes of two-ended subgroups of G , and that \mathcal{B} is the set of all axes of all hyperbolic elements of \mathcal{H} . (In other words, \mathcal{B} is an arc-system with \mathcal{B}/Γ finite, and such that the setwise stabiliser of each element of \mathcal{B} is infinite, and hence two-ended.) We note:

Lemma 2.1 : *The arc system \mathcal{B} is edge-finite.*

Proof : We want to show that any given edge lies in a finite number of elements of \mathcal{B} . Without loss of generality, we can suppose that \mathcal{B} consists of the orbit of a single arc, β . Let H be the setwise stabiliser of β . Choose any edge $e \in T$. Let $K \leq G$ be the stabiliser of e . Without loss of generality, we may as well suppose that $e \in E(\beta)$. Note that $E(\beta)/H$ is finite. Now, the G -orbit, Ge , of e meets $E(\beta)$ in an H -invariant set consisting of finitely many H -orbits, say $Ge \cap E(\beta) = Hg_1e \cup Hg_2e \cup \dots \cup Hg_ne$, where $g_i \in G$.

Suppose that $e \subseteq g\beta$, for some $g \in G$. Now $g^{-1}e \in E(\beta)$, so $g^{-1}e = hg_ie$ for some $h \in H$, and $i \in \{1, \dots, n\}$. Thus $ghg_i \in K$, so $gH = kg_i^{-1}H$ for some $k \in K$. Since K is finite, there are finitely many possibilities for the right coset gH , and hence for the arc $g\beta$. \diamond

Now, let \mathcal{H} be any finite union of conjugacy classes of two ended subgroups of G , as above. Recall that to say that G splits over a finite subgroup relative to H means that there is a non-trivial G -tree with finite edge stabilisers, and with each element of \mathcal{H} elliptic with respect to T . We can always take such a G -tree to be cofinite, and indeed to have only one orbit of edges. We say that \mathcal{H} is *irreducible* if G does not split over any finite group relative to H .

In Section 3, we shall give a general criterion for irreducibility in terms of arc systems. For the moment, we note:

Lemma 2.2 : *Suppose that G is a group and that T is a G -tree with finite edge stabilisers. Suppose that \mathcal{H} is a finite union of conjugacy classes of two-ended subgroups of G . Let \mathcal{B} be the arc system consisting of the set of axes of hyperbolic elements of G . If H is indecomposable, then each edge of T lies in at least two elements of \mathcal{B} .*

Proof : Suppose that $T \neq \bigcup \mathcal{B}$. Then, collapsing each component of $\bigcup \mathcal{B}$ to a point, we obtain another G -tree, Σ , with finite edge stabilisers. Moreover, each element of \mathcal{H} is elliptic with respect to Σ , contradicting indecomposability.

We thus have $T = \bigcup \mathcal{B}$. Suppose, for contradiction, that there is an edge of T which lies in precisely one element of \mathcal{B} . We may as well suppose that this is true of all edges of T . (For if not, let F be the union of all edges of T which lie in at least two elements of \mathcal{B} . Collapsing each component of F to a point, we obtain a new G -tree. We replace \mathcal{B} by the set of axis of those elements of \mathcal{H} which remain hyperbolic. Thus each element of the new arc system is the result of collapsing an element of the old arc system along a collection of disjoint compact subarcs.)

We now construct a bipartite graph, Σ , with vertex set an abstract disjoint union of $V(T)$ and \mathcal{B} , by deeming $x \in V(T)$ and $\beta \in \mathcal{B}$ to be adjacent in Σ if $x \in \beta$ in T . Now, it's easily verified that Σ is a simplicial tree, and that the stabiliser of each pair (x, \mathcal{B}) is finite. In other words, Σ is a G -tree with finite edge stabilisers. Finally, we note that each element of \mathcal{H} is elliptic in Σ . This again contradicts the indecomposability of \mathcal{H} . \diamond

We now move on to considering splittings over two-ended subgroups. Suppose that Γ is a group, and that Σ is a cofinite Γ -tree (with no terminal vertex) and with two-ended edge-stabilisers. We can write $V(\Sigma)$ as a disjoint union, $V(\Sigma) = V_1(\Sigma) \sqcup V_2(\Sigma) \sqcup V_\infty(\Sigma)$, depending on whether the corresponding vertex stabiliser is one, two or infinite-ended. Note that $V_2(\Sigma)$ is precisely the set of vertices of finite degree.

We remark that if there is a bound on the order of finite subgroups of Γ , and there are no infinitely divisible elements, then each two-ended subgroup lies in a unique maximal two-ended subgroup. In this case, we can refine our splitting so that for each vertex $v \in V_1(\Sigma) \cup V_\infty(\Sigma)$, the incident edge groups are all maximal two-ended subgroups of $\Gamma(v)$. This is automatically true of the JSJ splitting of hyperbolic groups (as described in [Bo2]), for example, though we shall have no need to assume this in this section.

It is fairly easy to see that the one-endedness or otherwise of Γ depends only on the infinite-ended vertex groups, $\Gamma(v)$ for $v \in V_\infty(\Sigma)$. In one direction, it easy to see that if one of these groups splits over a finite group relative to incident edge groups, then we can refine our splitting so that one of the new edge groups is finite. Hence Γ is not one-ended. In fact, we also have the converse. Recall that a ‘‘trivial vertex’’ of a splitting is a vertex of degree 1 such that the vertex group equals the adjacent edge group (i.e. it corresponds to a terminal vertex of the corresponding tree).

Theorem 2.3 : *Suppose we represent a group, Γ , as finite graph of groups with two-ended vertex groups and no trivial vertices. Then, Γ is one-ended if and only if none of the infinite-ended vertex groups split intrinsically over a finite subgroup relative to the incident edge groups.*

Proof : Let Σ be the Γ -tree corresponding to the splitting, and write $V(\Sigma) = V_1(\Sigma) \sqcup V_2(\Sigma) \sqcup V_\infty(\Sigma)$ as above. Given $v \in V(\Sigma)$ let $\Delta(v) \subseteq E(\Sigma)$ be the set of incident edges. We are supposing that for each $v \in V_\infty(\Sigma)$, the set of incident edge stabilisers, $\{\Gamma_\Sigma(e) \mid e \in \Delta(v)\}$, is indecomposable in the group $\Gamma_\Sigma(v)$. This is therefore true for all $v \in V(\Sigma)$. We aim to show that Γ is one-ended.

Suppose, for contradiction, that there exists a non-trivial minimal G -tree, T , with finite edge stabilisers. Let \mathcal{B} be the arc system on T consisting of the axes of those Σ -edge stabilisers, $\Gamma_\Sigma(e)$, which are hyperbolic with respect to T . By Lemma 2.1, \mathcal{B} is edge-finite.

Suppose, first, that $\mathcal{B} = \emptyset$, i.e. each group $\Gamma_\Sigma(e)$ for $e \in E(\Sigma)$ is elliptic in T . Suppose $v \in V(\Sigma)$. Since $\{\Gamma_\Sigma(e) \mid e \in \Delta(v)\}$ is indecomposable in $\Gamma_\Sigma(v)$, it follows that $\Gamma_\Sigma(v)$ must be elliptic in T . It therefore fixes a unique vertex of T . Suppose $w \in V(\Sigma)$ is adjacent to v . Since $\Gamma_\Sigma(v) \cap \Gamma_\Sigma(w)$ is infinite, it follows that $\Gamma_\Sigma(w)$ must also fix the same vertex of T . Continuing in this way, we conclude that this must be true of all Σ -vertex stabilisers. We therefore arrive at the contradiction that Γ fixes a vertex of T .

We deduce that $\mathcal{B} \neq \emptyset$. Now, choose any $\beta \in \mathcal{B}$ and any edge $\epsilon \in E(\beta)$. By construction, β is the axis of some edge stabiliser $\Gamma_\Sigma(e_0)$ for $e_0 \in E(\Sigma)$. Let $v \in V(\Sigma)$ be an endpoint of e_0 . Now, $\Gamma_\Sigma(e_0) \subseteq \Gamma_\Sigma(v)$, so $\Gamma_\Sigma(v)$ is not elliptic in T . It follows that $v \notin V_1(\Sigma)$. If $v \in V_2(\Sigma)$, then β is the axis in T of $\Gamma_\Sigma(v)$, and hence of any edge $e_1 \in E(\Sigma)$ adjacent to e_0 . In particular, ϵ lies in the axis of $\Gamma_\Sigma(e_1)$. If $v \in V_\infty(\Sigma)$, let $T(v)$ be the unique minimal $\Gamma_\Sigma(v)$ -invariant subtree of T . Let $\mathcal{B}(v)$ be the set of axis of hyperbolic elements of $\{\Gamma_\Sigma(e) \mid e \in \Delta(v)\}$. Thus, $\mathcal{B}(v) \subseteq \mathcal{B}$ is an arc system on $T(v)$, and $\beta \in \mathcal{B}(v)$. By Lemma 2.2, there is some $\beta' \in \mathcal{B}(v) \setminus \{\beta\}$ with $\epsilon \in E(\beta')$. Now, β' is the axis of $\Gamma_\Sigma(e_1)$ for some edge $e_1 \in E(\Sigma)$ adjacent to e_0 , as in the case where $v \in V_2(\Sigma)$. Now, in the same way, we can find some edge e_2 incident on the other endpoint of e_1 , so that $\Gamma_\Sigma(e_2)$ is hyperbolic in T and contains ϵ in its axis. Continuing, we get an infinite sequence of edges, $(e_n)_{n \in \mathbf{N}}$, which form a ray in Σ , and which all have this property.

Now, since \mathcal{B} is edge-finite, we can pass to a subsequence so that the axes of the groups $\Gamma_\Sigma(e_n)$ are constant. Since Σ is cofinite, we can find an edge $e \in E(\Sigma)$ and an element $\gamma \in \Gamma$ which is hyperbolic in Σ , and such that the axes of $\Gamma_\Sigma(e)$ and $\Gamma_\Sigma(\gamma e) = \gamma \Gamma_\Sigma(e) \gamma^{-1}$ in T are equal to α , say. In particular, $\gamma \alpha = \alpha$. Now, $\Gamma_\Sigma(e)$ has finite index in the setwise stabiliser of α , and so some power of γ lies in $\Gamma_\Sigma(e)$, contradicting the fact that γ is hyperbolic in Σ .

This finally contradicts the existence of the Γ -tree T . ◇

We note that Theorem 2.3 gives a means of describing the indecomposability of a set of two-ended subgroups in terms of the ‘‘doubled’’ group, as follows.

Suppose that G is a group, and that \mathcal{H} is a union of conjugacy classes of subgroups. We form a graph of groups with two vertices as follows. We take two copies of G as vertices, and connect them by a set of edges, one for each conjugacy class of subgroup in \mathcal{H} . We associate to each edge the corresponding group. We refer to the fundamental group of this graph of groups as the *double* of G in \mathcal{H} , and write it as $D(G, \mathcal{H})$. For example, if H is any subgroup of G and \mathcal{H} is its conjugacy class, then we just get the amalgamated free product, $D(G, \mathcal{H}) \cong G *_H G$.

From Theorem 2.3, we deduce immediately:

Corollary 2.4 : *Suppose that G is a group, and that \mathcal{H} is a union of finitely many conjugacy classes of two-ended subgroups. Then, \mathcal{H} is indecomposable in G if and only if the double, $D(G, \mathcal{H})$, is one-ended. \diamond*

We note that Theorem 2.3 can be extended to allow for one-ended edge groups. The hypotheses remain unaltered. We simply demand that no vertex group splits over a finite group relative to the set of two-ended incident edge groups. The argument remains essentially unchanged. If, however, we allow for infinite-ended edge groups, then Theorem 2.3 and Corollary 2.4 may fail.

Consider, for example, a one-ended group, K , with an infinite order element $a \in K$. Let G be the free product $K * \mathbf{Z}$, and write $b \in G$ for the generator of the \mathbf{Z} factor. Let $H \leq G$ be the subgroup generated by a and b . Thus, H is free of rank 2. Now, the conjugacy class of H is indecomposable in G . (For suppose that T is a G -tree with finite edge stabilisers and with H elliptic. Now, since K is one-ended, it is also elliptic. Since $K \cap H$ is infinite, and since $K \cup H$ generates G , we arrive at the contradiction that G is elliptic.) However, $G *_H G$ is not one-ended. In fact, $G *_H G \cong (K *_{\langle a \rangle} K) * \mathbf{Z}$. We remark that by taking $\langle a \rangle$ to be malnormal in K (for example taking K to be any torsion-free one-ended word hyperbolic group, and taking a to be any infinite order element) we can arrange that H is malnormal in G .

3. Indecomposable arc systems

In this section, we look further at arc systems and give a combinatorial characterisation of indecomposability. First, we introduce some additional notation concerning trees.

Suppose $S \subseteq T$ is a subtree. We write $\pi_S : T \cup \partial T \rightarrow S \cup \partial S$ for the natural retraction. Thus, $\pi_S((T \cup \partial T) \setminus (S \cup \partial S)) \subseteq V(S) \subseteq S$. If $R \subseteq S$ is another subtree, then $\pi_R \circ \pi_S = \pi_R$. Moreover, $\pi_R|_{(S \cup \partial S)}$ is defined intrinsically to S .

If $v \in V(S)$, then $T \cap \pi_S^{-1}(v)$ is a subtree of T , which we denote by $F(S, v)$. Note that $F(S, v) \cap S = \{v\}$, and that $\partial F(S, v) = \partial T \cap \pi_S^{-1}(v)$. Also, $T = S \cup \bigcup_{v \in V(S)} F(S, v)$.

We begin by describing generalisations of Whitehead graphs. For the moment, we do not need to introduce group actions.

Let T be a simplicial tree. We write $\mathcal{S}(T)$ for the set of finite subtrees of T . We can think of $\mathcal{S}(T)$ as a directed set under inclusion. Given $S \in \mathcal{S}(T)$, we define an equivalence relation, \approx_S , on ∂T by writing $x \approx_S y$ if $\pi_S x = \pi_S y$. In other words, $x \approx_S y$ if and only if the arc connecting x to y meets S in at most one point. Clearly, if $S \subseteq R \in \mathcal{S}(T)$, then \approx_R is finer than \approx_S . We therefore get a direct limit system of equivalence relations indexed by $\mathcal{S}(T)$. The direct limit (i.e. intersection) of these relations is just the equality relation on ∂T .

Suppose now that \mathcal{B} is an arc system on T . We have another equivalence relation, $\approx_{\mathcal{B}}$, on ∂T defined as follows. We write $x \approx_{\mathcal{B}} y$ if $x = y$ or if there exists some $\beta \in \mathcal{B}$ such that $\partial\beta = \{x, y\}$. If the intersection of any two arcs of \mathcal{B} is compact (as in most of the cases in which we shall be interested) then this is already an equivalence relation. If not, we take $\approx_{\mathcal{B}}$ to be the transitive closure of this relation.

Given $S \in \mathcal{S}(T)$, let $\sim_{S, \mathcal{B}}$ be the transitive closure of the union of the relations \approx_S

and $\approx_{\mathcal{B}}$. Thus, the relations $\sim_{S,\mathcal{B}}$ again form a direct limit system indexed by $\mathcal{S}(T)$. We write $\sim_{\mathcal{B}}$ for the direct limit.

Definition : We say that the arc system \mathcal{B} is *indecomposable* if there is just one equivalence class of $\sim_{\mathcal{B}}$ in ∂T

We can give a more intuitive description of this construction which ties in with Whitehead graphs as follows. We fix our arc system \mathcal{B} . If $S \in \mathcal{S}(T)$, we abbreviate $\sim_{S,\mathcal{B}}$ to \sim_S . Note that, if $Q \subseteq \partial T$ is a \sim_S -equivalence class, then $Q = \partial T \cap \pi_S^{-1} \pi_S Q$. Let $\mathcal{W}(S)$ be the collection of all sets of the form $\pi_S Q$, as Q runs over the set, $\partial T / \sim_S$, of \sim_S -classes. Thus, $\mathcal{W}(S)$ gives a partition of the subset $\bigcup \mathcal{W}(S)$ of $V(S)$. We refer to $\mathcal{W}(S)$ as a “subpartition” of $V(S)$ (i.e. a collection of disjoint subsets). There is a natural bijection between $\mathcal{W}(S)$ and the set $\partial T / \sim_S$.

Let $\mathcal{B}(S) \subseteq \mathcal{B}$ be the set of arcs which meet S in a non-trivial interval (i.e. non-empty and not a point). If $\beta \in \mathcal{B}(S)$, we write $I(\beta)$ for the interval $\beta \cap S$, thought of abstractly, and write $\text{fr } I(\beta)$ for the set consisting of its two endpoints. Let $Z(S)$ be the disjoint union $Z(S) = \bigsqcup_{\beta \in \mathcal{B}(S)} I(\beta)$, and let $\text{fr } Z(S) = \bigsqcup_{\beta \in \mathcal{B}(S)} \text{fr } I(\beta)$. There is a natural projection $p : Z(S) \rightarrow S$ with $p(\text{fr } Z(S)) \subseteq V(S)$. Now let $\mathcal{G}(S)$ be the quotient space $Z(S) / \cong$, where \cong is the equivalence relation on $Z(S)$ defined by $x \cong y$ if and only if $x = y$ or $x, y \in \text{fr } Z(S)$ and $px = py$. We see that $\mathcal{G}(S)$ is a 1-complex, with vertex set, $V(\mathcal{G}(S))$, arising from $\text{fr } Z(S)$. The map p induces a natural map from $\mathcal{G}(S)$ to S , also denoted by p . Now, $p|_{V(\mathcal{G}(S))}$ is injective, and $p(V(\mathcal{G}(S))) = \bigcup \mathcal{W}(S)$, where $\mathcal{W}(S)$ is the subpartition of $V(S)$ described earlier. Moreover, an element of $\mathcal{W}(S)$ is precisely the vertex set of connected component of $\mathcal{G}(S)$. If \mathcal{B} is edge-finite, then $\mathcal{G}(S)$ will be a finite graph.

To relate this to the theory of Whitehead graphs, the following observation will be useful.

Lemma 3.1 : *Suppose that $S_1, S_2 \in \mathcal{S}(T)$ are such that $S_1 \cap S_2$ consists of a single edge $e \in E(S_1) \cap E(S_2)$. If $\mathcal{G}(S_1)$ and $\mathcal{G}(S_2)$ are 2-vertex connected, then so is $\mathcal{G}(S)$.*

Proof : Let $S = S_1 \cup S_2 \in \mathcal{S}(T)$. Let $V_1 = V(S_2) \setminus V(S_1)$ and $V_2 = V(S_1) \setminus V(S_2)$. Write $W_i = p^{-1}(V_i) \subseteq V(\mathcal{G}(S))$ so that $V(\mathcal{G}(S)) = W_1 \sqcup W_2$. Let \mathcal{G}_i be the full subgraph spanned by W_i . Then $\mathcal{G}(S_i)$ is obtained by collapsing \mathcal{G}_i to a single vertex. The result therefore follows from the following observation, of which we omit the proof. \diamond

Lemma 3.2 : *Suppose that \mathcal{G} is a connected graph and that \mathcal{G}_1 and \mathcal{G}_2 are disjoint connected subgraphs. Write \mathcal{G}'_i for the result of collapsing \mathcal{G}_i to a single point in the graph \mathcal{G} . If \mathcal{G}'_1 and \mathcal{G}'_2 are both 2-vertex connected, then so is \mathcal{G} .* \diamond

Suppose $v \in V(T)$. Write $S(v)$ for the subtree consisting of the union of all edges incident on v . If T is a locally finite, then $S(v) \in \mathcal{S}(T)$. Applying Lemma 3.1 inductively we conclude:

Lemma 3.3 : *Suppose that \mathcal{B} is an arc system on the locally finite tree, T . If $\mathcal{G}(S(v))$ is 2-vertex connected for all $v \in V(T)$, then \mathcal{B} is indecomposable. \diamond*

The classical example of this, as discussed in the introduction, is that of Whitehead graphs. Suppose that G is a free group with free generators a_1, \dots, a_n . Let T be the Cayley graph of G with respect to these generators. Thus, T is locally finite cofinite G -tree.

Let $\{\gamma_1, \dots, \gamma_p\}$ be a finite set of non-trivial elements of G . It's easy to see that the indecomposability of the set of cyclic subgroups $\{\langle \gamma_1 \rangle, \dots, \langle \gamma_p \rangle\}$ (as defined in Section 2) is equivalent to that of $\{H_1, \dots, H_p\}$ where H_k is the maximal cyclic subgroup containing $\langle \gamma_k \rangle$. For this reason, we don't lose any generality by taking the elements γ_k to be indivisible, though this is not essential for what we are going to say.

Now, let \mathcal{B} be the arc system consisting of the set of axes of all conjugates of the elements γ_i . Now, the graph $\mathcal{G}(S(v))$ is independent of the choice of vertex $v \in V(T)$, so we may write it simply as \mathcal{G} . We can construct \mathcal{G} abstractly as the graph with vertex set $\{a_1, \dots, a_n, a_1^{-1}, \dots, a_n^{-1}\}$ where the number of edges connecting $a_i^{\epsilon_i}$ to $a_j^{\epsilon_j}$ equals the total number of times the subword $a_i^{\epsilon_i} a_j^{-\epsilon_j}$ occurs in the (disjoint union of the) reduced cyclic words representing elements γ_k (where $\epsilon_i, \epsilon_j \in \{-1, 1\}$). Thus, the total number of edges in \mathcal{G} equals the sum of the cyclically reduced word lengths of the elements γ_k . The fact that we are taking reduced cyclic words tells us immediately that there are no loops in \mathcal{G} . We call \mathcal{G} the *Whitehead graph*. This agrees with the description in the introduction, except that we are now allowing for multiple edges. (To recover the description of the introduction, and that of the original paper [W], we can simply replace each multiple edge by a single edge. This has no consequence for what we are going to say.)

By Lemma 3.3, we see immediately that:

Proposition 3.4 : *If \mathcal{G} is 2-vertex connected, then \mathcal{B} is indecomposable. \diamond*

We shall see later, in a more general context, that the indecomposability of \mathcal{B} is equivalent to the indecomposability of the set of subgroups $\{\langle \gamma_1 \rangle, \dots, \langle \gamma_p \rangle\}$.

By a "cut vertex" of \mathcal{G} we mean a vertex of \mathcal{G} which separates the component in which it lies. Now, if \mathcal{G} contains a cut vertex, one can change the generators (in an explicit algorithmic fashion) so as to reduce the total length of \mathcal{G} (allowing multiple edges) — cf. [W]. Thus, after a linearly bounded number of steps, we arrive at a Whitehead graph with no cut vertex. (It follows that if we choose generators so as to minimise the sum of the cyclically reduced word lengths of the γ_k , then the Whitehead graph will have this property.) In this case, the Whitehead graph is either disconnected or 2-vertex connected. In the former case, \mathcal{B} is clearly not indecomposable, whereas in the latter case it is (by Proposition 3.4). There is therefore a linear algorithm to decide indecomposability for a finite set of elements in a free group.

We remark that we can also recognise a free generating set by the same process. If $p = n$, then $\{\gamma_1, \dots, \gamma_n\}$ forms a free generating set if and only if a minimal Whitehead graph (or any Whitehead graph without cut vertices) is a disjoint union of n bigons. (If the elements γ_i are all indivisible, then any component with 2 vertices must be a bigon.) The algorithm arising out of this procedure was one of the main motivations of the original paper [W].

We want to generalise some of this discussion of indecomposability to the context of groups accessible over finite groups, as alluded to in Section 2.

For the moment, suppose that G is any group, and that T and Σ are equivalent cofinite G -trees with finite edge stabilisers. There are morphisms $\phi : T \rightarrow \Sigma$ and $\psi : \Sigma \rightarrow T$. These morphisms are quasiisometries, and hence induce a canonical bijection between ∂T and $\partial \Sigma$. In this case, it is appropriate to deal with formal arc systems, i.e. (G -invariant) sets of unordered pairs of elements of $\partial T \equiv \partial \Sigma$. Such a formal arc system determines an arc system, \mathcal{B} , on T and one, \mathcal{A} , on Σ . There is a bijection between \mathcal{B} and \mathcal{A} such that corresponding arcs have the same ideal endpoints. Thus, if $\beta \in \mathcal{B}$, then $\phi(\beta)$ is a subtree of Σ , with $\partial\phi(\beta) \equiv \partial\beta$. We see that the corresponding arc, $\alpha \in \mathcal{A}$ is the unique biinfinite arc contained in $\phi(\beta)$. Note that we get relations $\sim_{\mathcal{B}}$ and $\sim_{\mathcal{A}}$ on $\partial T \equiv \partial \Sigma$, from the direct limit construction described earlier. Our first objective will be to check that these are equal. It follows that the indecomposability of \mathcal{A} and \mathcal{B} are equivalent (Lemma 3.5). We thus get a well-defined notion of indecomposability of formal arc systems for such trees.

Suppose then, that T and Σ are G -trees as above, and that $\phi : T \rightarrow \Sigma$ and $\psi : \Sigma \rightarrow T$ are morphisms. For the moment, we can forget about the group G . All we need is the fact (Lemma 1.2) that the preimage of any finite tree under one of these morphisms is contained in another finite tree.

Suppose that $S \in \mathcal{S}(T)$. For clarity, we write $\approx_{S,T}$ for the relation on ∂T abbreviated to \approx_S in the previous discussion (i.e. $x \approx_{S,T} y$ if $\pi_S x = \pi_S y$). We thus have a direct limit system $(\approx_{S,T})_{S \in \mathcal{S}(T)}$. We similarly get another direct limit system $(\approx_{R,\Sigma})_{R \in \mathcal{S}(\Sigma)}$. We claim that these are cofinal. In other words, for each $S \in \mathcal{S}(T)$, there is some $R \in \mathcal{S}(\Sigma)$ such that the relation $\approx_{R,\Sigma}$ is finer than $\approx_{S,T}$, and conversely, swapping the roles of T and Σ .

Suppose, then, that $S \in \mathcal{S}(T)$. By Lemma 1.2, there is some $R \in \mathcal{S}(\Sigma)$ such that $\psi^{-1}S \subseteq R$. Suppose that $x, y \in \partial T \equiv \partial \Sigma$, and let β and α be the arcs in T and Σ respectively connecting x to y . Thus, $\beta \subseteq \psi(\alpha)$. Suppose that $x \approx_{R,\Sigma} y$. In other words, $|\alpha \cap R| \leq 1$. Now, $\beta \cap S \subseteq \psi(\alpha) \cap S = \psi(\alpha \cap \psi^{-1}S) \subseteq \psi(\alpha \cap R)$, and so $|\beta \cap S| \leq 1$. Thus $x \approx_{S,T} y$ as required. Swapping the roles of T and Σ , we deduce the cofinality of the direct limit systems as claimed.

Now, suppose that \mathcal{B} and \mathcal{A} are arc systems on T and Σ respectively, giving rise to the same formal arc system. We get identical relations $\approx_{\mathcal{B}} = \approx_{\mathcal{A}}$ on $\partial T = \partial \Sigma$, as defined earlier. Now, it follows that the direct limit systems $(\approx_{S,\mathcal{B}})_{S \in \mathcal{S}(T)}$ and $(\approx_{R,\mathcal{A}})_{R \in \mathcal{S}(\Sigma)}$ are cofinal, and so give rise to the same direct limit, namely $\sim_{\mathcal{B}} = \sim_{\mathcal{A}}$, as claimed earlier.

In particular, we see that \mathcal{B} is indecomposable if and only if \mathcal{A} is. In summary, reintroducing the group action, we have shown:

Lemma 3.5 : *Suppose that T and Σ are equivalent cofinite G -trees with finite edge stabilisers. Suppose that \mathcal{B} and \mathcal{A} are arc systems on T and Σ respectively, corresponding to the same formal arc system on $\partial T \equiv \partial \Sigma$. Then, \mathcal{B} is indecomposable if and only if \mathcal{A} is indecomposable. \diamond*

Suppose, now, that G is accessible over finite groups. As discussed in Section 1, we can associate to G a set $\partial_{\infty}G$, which we can identify with the boundary of any cofinite G -tree with finite edge stabilisers and finite and one-ended vertex stabilisers. We refer to

such trees as *complete* G -trees. Any two complete G -trees are equivalent, so by Lemma 3.5, it makes sense to speak about a formal arc system on $\partial_\infty G$ as being indecomposable.

Suppose, now that $H \leq G$ is a two-ended subgroup. We say that H is *elliptic* if it lies inside some one-ended subgroup of G . Thus H is elliptic if and only if it is elliptic with respect to some (and hence any) complete G -tree. Otherwise, we say that H is *hyperbolic*. In this case, there is a unique H -invariant unordered pair of points in $\partial_\infty G$ which we denote by ΛH . Thus, ΛH is the pair of endpoints of the axis of H in any complete G -tree. We refer to ΛH as the *limit set* of H . We note that if H' is another hyperbolic two-ended subgroup, and $\Lambda H \cap \Lambda H' \neq \emptyset$, then H and H' are commensurable, and hence lie in the same maximal two-ended subgroup.

Let \mathcal{H} be a finite union of conjugacy classes of hyperbolic two-ended subgroups of G . Recall that \mathcal{H} is “indecomposable” if we cannot write G as a non-trivial amalgamated free product or HNN-extension with each element of H conjugate into a vertex group. It is easy to see that this property depends only on the commensurability classes of the elements of \mathcal{H} , so we may, if we wish, take all the elements of H to be maximal two-ended subgroups, in which case their limit sets are all disjoint. Note that we get a formal arc system, $\{\Lambda H \mid H \in \mathcal{H}\}$, on $\partial_\infty G$. We claim:

Proposition 3.6 : *If the formal arc system $\{\Lambda H \mid H \in \mathcal{H}\}$ is indecomposable, then \mathcal{H} is indecomposable.*

Proof : Suppose not. Then there is a non-trivial cofinite G -tree, T , with finite edge stabilisers and with each element of \mathcal{H} elliptic with respect to T . Now, as discussed in Section 1, we can refine the splitting T/G to a complete splitting, giving us a complete G -tree, Σ . We can recover T by collapsing T along a disjoint union of subtrees. Each element of H fixes setwise one of these subtrees.

Now, let \mathcal{B} be the arc system on Σ given by the formal arc system, in other words, the set of axes of elements of \mathcal{H} . Thus each axis lies inside one of the collapsing subtrees. In particular, $\Sigma \neq \bigcup \mathcal{B}$, and so \mathcal{B} is decomposable. \diamond

We shall prove a converse to Proposition 3.6 in the case where G is finitely generated. For this we shall need a relative version of Stallings’s theorem.

Let G be a finitely generated group, and let X be a Cayley graph of X (or, indeed, any graph on which G acts with finite vertex stabilisers and finite quotient). Given a subset $A \subseteq V(X)$ we write $E_A \subseteq E(X)$ for the set of edges with precisely one endpoint in A . Thus, to say that X has “more than one end” means that we can find an infinite subset, $A \subseteq V(X)$ such that its complement $B = V(X) \setminus A$ is also infinite, and such that $E_A = E_B$ is finite. Thus, Stallings’s theorem [St] tells us that in such a case, G splits over a finite group.

Suppose, now that $H \leq G$ is a two ended subgroup, and that $C \subseteq V(X)$ is an H -orbit of vertices (or any H -invariant subset with C/H finite). Now, for all but finitely many G -images, gC , of C , we have either $gC \subseteq A$ or $gC \subseteq B$. For the remainder, we have three possibilities: either $gC \cap A$ is finite or $gC \cap B$ is finite, or else both of these subsets give us a neighbourhood of an end of H . We shall not say more about the last case, since it precisely the case we wish to rule out. Note that this classification of does not depend

on the choice of H -orbit, C . A specific relative version of Stallings's theorem says the following:

Lemma 3.7 : *Suppose G is a finitely group and \mathcal{H} is a finite union of conjugacy classes of two-ended subgroups. Let X be a Cayley graph of G . Suppose we can find an infinite set, $A \subseteq V(X)$, such that E_A is finite and $B = V(X) \setminus A$ is infinite. Suppose that for any $H \in \mathcal{H}$ either $A \cap C$ or $B \cap C$ is finite for some (hence every) H -orbit of vertices, C . Then, \mathcal{H} is decomposable (i.e. G splits over a finite group relative to \mathcal{H}). \diamond*

In fact, a much stronger result follows immediately from the results of [DiD]. It may be stated as follows. Suppose G is any finitely generated group, and $A \subseteq G$ is an infinite subset, whose complement $B = G \setminus A$ is also infinite. Suppose that the symmetric difference of A and Ag is finite for all $g \in G$. Suppose that H_1, \dots, H_n are subgroups such that for all $g \in G$ and all $i \in \{1, \dots, n\}$ either $gH_i \cap A$ or $gH_i \cap B$ is finite. Then G splits over a finite group relative to $\{H_1, \dots, H_n\}$. (In fact, it's sufficient to rule out G being a non-finitely generated countable torsion group.)

Alternatively, one can deduce Lemma 3.7, as we have stated it, by applying Stallings's theorem to the double, $D(G, \mathcal{H})$, and using Corollary 2.4. We briefly sketch the argument. We may construct a Cayley graph, Y , for $D(G, \mathcal{H})$ by taking lots of copies of X , and stringing them together in a treelike fashion. Let's focus on a particular copy of X , which we take to be acted upon by G . Now each adjacent copy of X corresponds to an element $H \in \mathcal{H}$, and is connected ours by an H -orbit of edges. We refer to such edges as "amalgamating edges". The amalgamating edges corresponding to H are attached to X by an H -orbit, C_H , of vertices of X . By hypothesis, either $C_H \cap A$ is finite, in which case, we write E_H for the set of amalgamating edges which have an endpoint in $C_H \cap A$, or else, $C_H \cap B$ is finite, in which case, we write E_H for the set of amalgamating edges which have an endpoint in $C_H \cap B$. Now, for all but finitely many H , the set E_H is empty. Thus, the set $E_{\mathcal{H}} = \bigcup_{H \in \mathcal{H}} E_H$ is finite, and so $E_0 = E_A \cup E_{\mathcal{H}} \subseteq E(Y)$ is finite. Now, E_0 separates Y into two infinite components. Thus, by Stallings's theorem, $D(G, \mathcal{H})$ splits over a finite group, and so by Corollary 2.4, \mathcal{H} is decomposable. With the details filled in, this gives another proof of Lemma 3.7.

We are now ready to prove a converse to Proposition 3.6:

Proposition 3.8 : *Suppose that G is a finitely generated accessible group. Suppose that \mathcal{H} is a finite union of conjugacy classes of hyperbolic two-ended subgroups. If \mathcal{H} is indecomposable, then the formal arc system, $\{\Lambda H \mid H \in \mathcal{H}\}$, on $\partial_{\infty} G$, is indecomposable.*

Proof : Let T be a complete G -tree, and let \mathcal{B} be the corresponding arc system on T , i.e. the set of axes of elements of \mathcal{H} . Suppose, for contradiction, that \mathcal{B} is decomposable. In other words, we can find $S \in \mathcal{T}$ such that there is more than one \sim_S -class. By taking projections of \sim_S -classes as discussed in Section 1, we can write $V(S)$ as a disjoint union of non-empty subsets, $V(S) = W_1 \sqcup W_2$ with the property that if $\beta \in \mathcal{B}$, then β meets S , if at all, in compact interval (or point) with either both endpoints in W_1 or both endpoints in W_2 . Let $F_i = \pi_S^{-1}W_i$. Thus, $T = S \cup F_1 \cup F_2$, and each component of each F_i is a subtree meeting S in a single point.

Now, let X be a Cayley graph of G . Let $f : V(X) \rightarrow V(T)$ be any G -equivariant map. Let $A_i = f^{-1}F_i \subseteq V(X)$. Thus, $V(X) = A_1 \sqcup A_2$. Moreover, it is easily seen that $E_{A_1} = E_{A_2}$ is finite. (For example, extend f equivariantly to a map $f : X \rightarrow T$ so that each edge of X gets mapped to a compact interval of T . Only finitely many G -orbits of such an interval can contain a given edge of T . Now, the image of an edge of E_{A_1} connects a vertex of F_1 to a vertex of F_2 , and hence contains an edge of S . There are only finitely many such edges.)

Finally, suppose that $H \in \mathcal{H}$. Let $\beta \in \mathcal{B}$ be the axis of H . Without loss of generality, we can suppose that both ends of β are contained in F_1 . Now suppose that C is any H -orbit of vertices of X . Then $f(C)$ remains within a bounded distance of \mathcal{B} , from which we see easily that $f(C) \cap F_2$ is finite. Thus, $C \cap A_2$ is finite.

We have verified the hypotheses of Lemma 3.7, and so \mathcal{H} is decomposable, contrary to our hypotheses. \diamond

Note that Propositions 3.6 and 3.8 apply, in particular, to any finitely presented group, and even more specifically, to any hyperbolic group, G . In the latter case, $\partial_\infty G$ can be identified as a subset of the Gromov boundary, ∂G , as discussed in Section 2. If $H \leq G$ is a hyperbolic two-ended subgroup, then $\Lambda H \subseteq \partial G$ is the limit set of H by the standard definition. This ties in with the discussion of equivalence relations on ∂G in the introduction, and will be elaborated on in Section 5.

4. Quasiconvex splittings of hyperbolic groups.

For most of the rest of this paper, we shall be confining our attention to hyperbolic groups. We shall consider how some of the general constructions of Sections 1–3 relate to the topology of the boundary in this case. Before we embark on this, we review some general facts about quasiconvex splittings of hyperbolic groups (i.e. splittings over quasiconvex subgroups). This elaborates on the account given in [Bo2].

Throughout the rest of this paper, we shall use the notation $\text{fr } A$ to denote the topological boundary (or “frontier”) of a subset, A , of a larger topological space. We reserve the symbol “ ∂ ” for ideal boundaries.

Let Γ be any hyperbolic group. Let X be any locally finite connected graph on which Γ acts freely and cocompactly (for example a Cayley graph of Γ). We put a path metric, d , on X by assigning a positive length to each edge in a Γ -invariant fashion. Let $\partial\Gamma \equiv \partial X$ be the boundary of Γ . We may put a metric on $\partial\Gamma$ as follows. Choose any $a \in V(X)$. Given $x, y \in \partial\Gamma$, let $\delta = \min\{d(a, \alpha)\}$, as α ranges over all biinfinite geodesics in X connecting x to y . Let $\rho_0(x, y) = \lambda^\delta$, where $\lambda \in (0, 1)$ is constant. Provided we choose λ small enough in relation to the constant of hyperbolicity, ρ_0 will be a quasiultrametric on $\partial\Gamma$, as described in [GhH]. In general, ρ_0 need not be continuous on $\partial\Gamma$. However, it is bilipschitz equivalent to a canonical Γ -invariant metric, ρ , which induces the standard compact topology on $\partial\Gamma$. Although all the arguments of this paper can be expressed in purely topological terms, it will be convenient to have recourse to this metric.

Note that if $G \leq \Gamma$ is quasiconvex, then it is intrinsically hyperbolic, and we may identify its boundary, ∂G , with its limit set $\Lambda G \subseteq \partial\Gamma$. Note that G acts properly discon-

tinuously on $\partial\Gamma \setminus \Lambda G$. The setwise stabiliser of ΛG in Γ is precisely the commensurator, $\text{Comm}(G)$, of G in Γ (i.e. the set of all $g \in \Gamma$ such that $G \cap gGg^{-1}$ has finite index in G). In this case, G has finite index in $\text{Comm}(G)$. In fact, $\text{Comm}(G)$ is the unique maximal subgroup of Γ which contains G as finite index subgroup. We say that G is *full* if $G = \text{Comm}(G)$.

We shall use the following notation. If $f : Z \rightarrow [0, \infty)$ is a function from some set Z to the nonnegative reals, we write “ $f(z) \rightarrow 0$ for $z \in Z$ ” to mean that $\{z \in Z \mid f(z) \geq \epsilon\}$ is finite for all $\epsilon > 0$. We similarly define “ $f(z) \rightarrow \infty$ for $z \in Z$ ”.

Lemma 4.1 : *If $G \leq \Gamma$ is quasiconvex and $x \in \partial\Gamma$, then $\rho(gx, \Lambda G) \rightarrow 0$ for $g \in G$.*

Proof : Since G acts properly discontinuously on $\partial\Gamma \setminus \Lambda G$, there can be no accumulation point of the G -orbit of x in this set. \diamond

We want to go on to consider splittings of Γ . For this, we shall want to introduce some further notation regarding trees.

By a “directed edge” we mean an edge together with an orientation. We write $\vec{E}(T)$ for the set of directed edges. We shall always use the convention that $e \in E(T)$ represents the undirected edge underlying the directed edge $\vec{e} \in \vec{E}(T)$. We write $\text{head}(\vec{e})$ and $\text{tail}(\vec{e})$ respectively for the head and tail of \vec{e} . We use $-\vec{e}$ for the same edge oriented in the opposite direction, i.e. $\text{head}(-\vec{e}) = \text{tail}(\vec{e})$ and $\text{tail}(-\vec{e}) = \text{head}(\vec{e})$. If $\vec{e} \in \vec{E}(T)$ and $v \in V(T)$, we say that \vec{e} “points towards” v if $\text{dist}(v, \text{tail}(\vec{e})) = \text{dist}(v, \text{head}(\vec{e})) + 1$.

If $v \in V(T)$, let $\Delta(v) \subseteq E(T)$ be the set of edges incident on v , and let $\vec{\Delta}(v) = \{\vec{e} \in \vec{E}(T) \mid \text{head}(\vec{e}) = v\}$. Thus, the degree of v is $\text{card}(\Delta(v)) = \text{card}(\vec{\Delta}(v))$.

Given $\vec{e} \in \vec{E}(T)$, we write $\Phi(\vec{e}) = \Phi_T(\vec{e})$ for the connected component of T minus the interior of e which contains $\text{tail}(\vec{e})$. Thus, $V(\Phi(\vec{e}))$ is the set of vertices, v , of T such that \vec{e} points away from v .

Given $v \in V(T)$, we shall write $\vec{\Omega}(v) \subseteq \vec{E}(T)$ for the set of directed edges which point towards v . Thus, for each edge $e \in E(T)$, precisely one of the pair $\{\vec{e}, -\vec{e}\}$ lies in $\vec{\Omega}(v)$. Note that $\vec{e} \in \vec{\Omega}(v)$ if and only if $v \notin \Phi(\vec{e})$. Clearly $\vec{\Delta}(v) \subseteq \vec{\Omega}(v)$.

We now return to our hyperbolic group, Γ . Suppose that Γ acts without edge inversions on a simplicial tree, Σ , with Σ/Γ finite. We suppose that this action is minimal. Given $v \in V(\Sigma)$ and $e \in E(\Sigma)$, write $\Gamma(v)$ and $\Gamma(e)$ respectively for the corresponding vertex and edge stabilisers. Note that $\Gamma(v)$ is finite if and only if v has finite degree in Σ . If $v, w \in V(\Sigma)$ are the endpoints of an edge $e \in E(\Sigma)$, then $\Gamma(e) = \Gamma(v) \cap \Gamma(w)$.

As in [Bo2], we may construct a Γ -equivariant map $\phi : X \rightarrow \Sigma$ such that each edge of X either gets collapsed onto a vertex of Σ or mapped homeomorphically onto a closed arc in Σ . (Note that, after subdividing X if necessary, we can assume that, in the latter case, this closed arc is an edge of Σ .) Since the action of Γ is minimal, ϕ is surjective.

A proof of the following result can be found in [Bo2], though it appears to be “folklore”.

Proposition 4.2 : *If $\Gamma(e)$ is quasiconvex for each $e \in E(\Sigma)$, then $\Gamma(v)$ is quasiconvex for each $v \in V(\Sigma)$.* \diamond

We refer to such a splitting as a *quasiconvex splitting*.

We note that if a vertex group, $\Gamma(v)$, of a quasiconvex splitting has the property that all incident edge groups are of infinite index in $\Gamma(v)$, then $\Gamma(v)$ must be full in the sense described above. In other words, $\Gamma(v)$ is the setwise stabiliser of $\Lambda\Gamma(v)$. This will be the case in most situations of interest (in particular where all edge groups are finite or two-ended, but $\Gamma(v)$ is not).

Note that, if $v, w \in V(\Sigma)$, then $\Gamma(v) \cap \Gamma(w)$ is quasiconvex (since the intersection of any two quasiconvex subgroups is quasiconvex [Sh]). We see that $\Lambda\Gamma(v) \cap \Lambda\Gamma(w) = \Lambda(\Gamma(v) \cap \Gamma(w))$. In particular, if v, w are the endpoints of an edge $e \in E(\Sigma)$, then $\Lambda\Gamma(v) \cap \Lambda\Gamma(w) = \Lambda\Gamma(e)$.

As described in [Bo2], there is a natural Γ -invariant partition of $\partial\Gamma$ as $\partial\Gamma = \partial_0\Gamma \sqcup \partial_\infty\Gamma$, where $\partial_0\Gamma = \bigcup_{v \in V(\Sigma)} \Lambda\Gamma(v)$, and $\partial_\infty\Gamma$ is naturally identified with $\partial\Sigma$. Note that $\partial_\infty\Gamma$ is dense in $\partial\Gamma$, provided that Σ is non-trivial. (In the case where the edge stabilisers are all finite, this agrees with the notion introduced for accessible groups in Section 2.)

Given $\vec{e} \in \vec{E}(\Sigma)$, we write

$$\Psi(\vec{e}) = \partial\Phi(\vec{e}) \cup \bigcup_{v \in V(\Phi(\vec{e}))} \Lambda\Gamma(v).$$

It's not hard to see that $\Psi(\vec{e})$ is a closed $\Gamma(e)$ -invariant subset of $\partial\Gamma$. Moreover, $\Psi(\vec{e}) \cup \Psi(-\vec{e}) = \partial\Gamma$ and $\Psi(\vec{e}) \cap \Psi(-\vec{e}) = \text{fr } \Psi(\vec{e}) = \Lambda\Gamma(e)$.

Recall that, if $v \in V(\Sigma)$, then $\vec{\Delta}(v)$ is defined to be the set of all directed edges of Σ with head at v . Now, $V(\Sigma) = \{v\} \sqcup \bigsqcup_{\vec{e} \in \vec{\Delta}(v)} V(\Phi(\vec{e}))$ and $\partial\Sigma = \bigsqcup_{\vec{e} \in \vec{\Delta}(v)} \partial\Phi(\vec{e})$. It follows that:

Lemma 4.3 : $\partial\Gamma = \Lambda\Gamma(v) \cup \bigcup_{\vec{e} \in \vec{\Delta}(v)} \Psi(\vec{e})$. ◇

Moreover, for each $\vec{e} \in \vec{\Delta}(v)$, we have $\Lambda\Gamma(v) \cap \Psi(\vec{e}) = \Lambda\Gamma(e)$.

The above assertions become more transparent, given the following alternative description of $\Psi(\vec{e})$.

Let $m(e)$ be the midpoint of the edge e , and let $I(\vec{e})$ be the closed interval in Σ consisting of the segment of e lying between $m(e)$ and $\text{tail}(\vec{e})$. Let $Q(e) = \phi^{-1}(m(e)) \subseteq X$ and $R(\vec{e}) = \phi^{-1}(\Phi(\vec{e}) \cup I(\vec{e})) \subseteq X$, where $\phi : X \rightarrow \Sigma$ is the map described above. Note that $Q(e) = \text{fr } R(\vec{e}) = R(\vec{e}) \cap R(-\vec{e})$. By the arguments given in [Bo2], we see easily that $Q(e)$ and $R(\vec{e})$ are quasiconvex subsets of X . Moreover, $\Psi(\vec{e}) = \partial R(\vec{e})$.

Note that the collection $\{Q(e) \mid e \in E(\Sigma)\}$ is locally finite in X . It follows that, for any fixed $a \in X$, we have $d(a, Q(e)) \rightarrow \infty$ for $e \in E(\Sigma)$.

Now, fix some vertex, $v \in V(\Sigma)$. Recall that $\vec{\Omega}(v)$ is defined to be the set of all directed edges pointing towards v . Choose any $b \in \phi^{-1}(v) \subseteq X$. Now, if $\vec{e} \in \vec{\Omega}(v)$, we have $v \notin \Phi(\vec{e}) \cup I(\vec{e})$, and so $b \notin R(\vec{e})$. Since $Q(e) = \text{fr } R(\vec{e})$, we have $d(b, R(\vec{e})) = d(b, Q(e))$. It follows that $d(b, R(\vec{e})) \rightarrow \infty$ for $\vec{e} \in \vec{\Omega}(v)$. In fact, we see that $d(a, R(\vec{e})) \rightarrow \infty$ given any fixed basepoint, $a \in X$. Now, there are only finitely many Γ -orbits of directed edges, and so the sets $R(\vec{e})$ are uniformly quasiconvex. From the definition of the metric ρ on $\partial\Gamma$, it follows easily that $\text{diam}(\Psi(\vec{e})) \rightarrow 0$, where diam denotes diameter with respect to ρ . In summary, we have shown:

Lemma 4.4 : For any $v \in V(\Sigma)$, $\text{diam}(\Psi(\vec{e})) \rightarrow 0$ for $\vec{e} \in \vec{\Omega}(v)$. \diamond

We now add the hypothesis that $\Gamma(e)$ is infinite for all $e \in E(\Sigma)$.

Suppose $v \in V(\Sigma)$ and suppose K is any closed subset of $\Lambda\Gamma(v)$. Let $\vec{\Delta}_K(v) = \{\vec{e} \in \vec{\Delta}(v) \mid \Lambda\Gamma(e) \subseteq K\}$, and let $\Upsilon(v, K) = K \cup \bigcup_{\vec{e} \in \vec{\Delta}_K(v)} \Psi(\vec{e}) \subseteq \partial\Gamma$.

Lemma 4.5 : The set $\Upsilon(v, K)$ is closed in $\partial\Gamma$.

Proof : Suppose $x \notin \Upsilon(v, K)$. In particular, $x \notin K$, so $\epsilon = \rho(x, K) > 0$. Now, if $\vec{e} \in \vec{\Delta}_K(v)$ and $\rho(x, \Psi(\vec{e})) < \epsilon/2$, then $\text{diam}(\Psi(\vec{e})) > \epsilon/2$ (since $K \cap \Psi(\vec{e}) \supseteq \Lambda\Gamma(e)$, which, by the hypothesis on edge stabilisers, is non-empty). By Lemma 4.4, this occurs for only finitely many such \vec{e} . Since each $\Psi(\vec{e})$ is closed, it follows that $\rho(x, \Upsilon(v, K))$ is attained, and hence positive. In other words, $x \notin \Upsilon(v, K)$ implies $\rho(x, \Upsilon(v, K)) > 0$. This shows that $\Upsilon(v, K)$ is closed. \diamond

5. Quotients.

In this section, we aim to consider quotients of boundaries of hyperbolic groups, and to relate this to indecomposability, thereby generalising some of the results of [O].

First, we recall a few elementary facts from point-set topology [Ke, HoY]. Let M be a hausdorff topological space. A subset of M is *clopen* if it is both open and closed. We may define an equivalence relation on M by deeming two points to be related if every clopen set containing one must also contain the other. The equivalence classes are called *quasicomponents*. A *component* of M is a maximal connected subset. Components and quasicomponents are always closed. Every component is contained in a quasicomponent, but not conversely in general. However, if M is compact, these notions coincide. Thus, if K and K' are distinct components of a compact hausdorff space, M , then there is a clopen subset of M containing K , but not meeting K' .

Suppose that M is a compact hausdorff space, and that \approx is an equivalence relation on M . If the relation \approx is closed (as a subset of $M \times M$), then the quotient space, M/\approx is hausdorff.

The compact spaces of interest to us here will be the boundaries of hyperbolic groups. Suppose that G is a hyperbolic group, and that ∂G is its boundary. Now, any two ended subgroup, H , of G is necessarily quasiconvex, so its limit set, $\Lambda H \subseteq \partial G$, consists of pair of points. If H' is another two-ended subgroup, and $\Lambda H \cap \Lambda H' \neq \emptyset$, then H and H' are commensurable, and so lie in a common maximal two-ended subgroup. In particular, $\Lambda H = \Lambda H'$ (cf. the discussion of accessible groups in Section 3).

Suppose that \mathcal{H} is a union of finitely many conjugacy classes of two-ended subgroups of G . Let $\approx_{\mathcal{H}}$ be the equivalence relation defined on ∂G defined by $x \approx_{\mathcal{H}} y$ if and only if either $x = y$ or there exists $H \in \mathcal{H}$ such that $\Lambda H = \{x, y\}$. Now, it's a simple consequence of Lemma 4.1 that the relation $\approx_{\mathcal{H}}$ is closed. We write $M(G, \mathcal{H})$ for the quotient space $\partial G/\approx_{\mathcal{H}}$. Thus:

Lemma 5.1 : $M(G, \mathcal{H})$ is compact hausdorff. \diamond

We aim to describe when $M(G, \mathcal{H})$ is connected. Clearly, if G is one-ended so that ∂G is connected, this is necessarily the case. We can thus restrict attention to the case when G is infinite-ended.

Let T be a complete G -tree. As in Section 3, we can define $\partial_\infty G$ as ∂T . This also agrees with the notation introduced in Section 4, thinking of T as a quasiconvex splitting of G . In particular, we can identify $\partial_\infty G$ as a subset of ∂G . This set $\partial_0 G = \partial G \setminus \partial_\infty G$ is a disjoint union of the boundaries of the infinite vertex stabilisers of T , i.e. the maximal one-ended subgroups. In other words, the components of $\partial_0 G$ are precisely the boundaries of the maximal one-ended subgroups of G .

Let \mathcal{H} be a set of one-ended subgroups as above. The subset, \mathcal{H}_0 , of \mathcal{H} consisting of those subgroups in \mathcal{H} which are hyperbolic (i.e. with both limit points in $\partial_\infty G$), defines a formal arc system on $\partial_\infty G$. We aim to show that $M(G, \mathcal{H})$ is connected if and only if this arc system is indecomposable. This, in turn, we know to be equivalent to asserting that \mathcal{H}_0 is irreducible.

In fact, it's easy to see that the elliptic elements of \mathcal{H} have no bearing on the connectivity or otherwise of $M(G, \mathcal{H})$. For this reason, we may as well suppose, for simplicity, that \mathcal{H} consists entirely of hyperbolic two-ended subgroups. We therefore aim to show:

Theorem 5.2 : *Let G be an infinite-ended hyperbolic group, and let \mathcal{H} be a union of finitely many conjugacy classes of hyperbolic two-ended subgroups. Then, the quotient space $M(G, \mathcal{H})$ is connected if and only if \mathcal{H} is indecomposable.*

First, we set about proving the “only if” bit. Let T be a complete G -tree. Thus, $\partial_\infty G$ is identified with ∂T , and \mathcal{H} determines an arc system, \mathcal{B} , on T . We know (Propositions 3.6 and 3.8) that the indecomposability of \mathcal{H} is equivalent to the indecomposability of \mathcal{B} .

We shall say that a subgraph, F , of T is *finitely separated* if there are only finitely many edges of T with precisely one endpoint in F . Now, it's not hard to see that F is finitely separated if and only if it's a finite union of finite intersections of subtrees of the form $\Phi(\vec{e})$ for $\vec{e} \in \vec{E}(T)$ (recalling the notation of Section 4).

Now, given a subgraph, $F \subseteq T$, we write

$$A(F) = \partial F \cup \bigcup_{v \in V(F)} \Lambda G(v)$$

(so that $A(T) = \partial G$). If F is finitely separated, then $A(F)$ is a finite union of finite intersections of sets of the form $\Psi(\vec{e})$, which are each closed by the remarks of Section 4. We conclude:

Lemma 5.3 : *If $F \subseteq T$ is a finitely separated subgraph, then $A(F)$ is closed in ∂G .* \diamond

We can now prove:

Lemma 5.4 : *If $M(G, \mathcal{H})$ is connected, then the arc system \mathcal{B} is indecomposable.*

Proof : Suppose, to the contrary, that \mathcal{B} is decomposable. Then, exactly as in the proof of Proposition 3.8, we can find two disjoint finitely separated subgraphs, F_1 and F_2 of T with $V(T) = V(F_1) \sqcup V(F_2)$ and $\partial T = \partial F_1 \sqcup \partial F_2$, and such that for each $\beta \in \mathcal{B}$, either $\partial\beta \subseteq \partial F_1$ or $\partial\beta \subseteq \partial F_2$. We see that $\partial G = A(F_1) \sqcup A(F_2)$.

Let $q : \partial G \rightarrow \partial G / \approx_{\mathcal{H}} = M(G, \mathcal{H})$ be the quotient map. Now, from the construction, we see that if $x \approx_{\mathcal{H}} y$ then either $x, y \in \partial F_1 \subseteq A(F_1)$ or $x, y \in \partial F_2 \subseteq A(F_2)$. We therefore get that $M(G, \mathcal{H}) = q(A(F_1)) \sqcup q(A(F_2))$. But applying Lemma 5.3, the sets $q(A(F_i))$ are both closed in $M(G, \mathcal{H})$, contrary to the assumption that $M(G, \mathcal{H})$ is connected. \diamond

Of course, we could slightly streamline this argument, if we make use of Proposition 3.8, and choose an appropriate complete G -tree, though it seems more satisfying to give a direct proof.

We now proof the converse. This will rely on the relative version of Stallings's theorem discussed in Section 3. One can also give direct, though more complicated, argument using results of Section 7 and ideas of Section 3, though we shall not pursue this here.

Lemma 5.5 : *If \mathcal{H} is indecomposable, then $M(G, \mathcal{H})$ is connected.*

Proof : Suppose, for contradiction, that we can write $M(G, \mathcal{H})$ as the disjoint union of two non-empty closed sets, K_1 and K_2 . Let $L_i \subseteq \partial G$ be the preimage of K_i under the quotient map $\partial G \rightarrow M(G, \mathcal{H})$. Thus, $\partial G = L_1 \sqcup L_2$. Let X be a Cayley graph of G . Now, we can give $X \cup \partial G$ a natural G -invariant topology as a compact metrisable space. Since $X \cup \partial G$ is normal, we can find disjoint open subsets, $U_i \subseteq X \cup \partial G$ with $L_i \subseteq U_i$. Now, $(X \cup \partial G) \setminus (U_1 \cup U_2) \subseteq X$ is compact, and so lies inside a finite subgraph, Y , of X . Let $A = U_1 \cap V(X)$ and let $B = V(X) \setminus A$. We need to verify that A satisfies the hypotheses of Lemma 3.7.

Note that $A \cup L_1$ and $B \cup L_2$ are both closed in $X \cup \partial G$. We see that A and B are both infinite. Recall that $E_A = E_B$ is the set of edges of X which have one endpoint in A and the other in B . Now, $E_A \subseteq E(Y)$, and so E_A is finite.

Finally, suppose that $H \in \mathcal{H}$ and that $C \subseteq V(X)$ is an H -orbit of vertices of X . Now, $C \cup \partial H$ is closed in $X \cup \partial G$. Without loss of generality we can suppose that $\Lambda H \subseteq L_1$. Since $B \cup L_2 \subseteq X \cup \partial G$ is closed, we see that $C \cap B$ is finite.

We have verified the hypotheses of Lemma 3.7, and so we arrive at the contradiction that \mathcal{H} is decomposable. \diamond

This concludes the proof of Theorem 5.2.

6. Splittings of hyperbolic groups over finite and two-ended subgroups.

Suppose that a hyperbolic group splits over a collection of two-ended subgroups. We may in turn try to split each of the vertex groups over finite groups, thus giving us a two-step series of splittings. We want to study how the combinatorics of such splittings are reflected in the topology of the boundary. The combinatorics can be described in terms of the trees associated to each step of the splitting, together with arc systems on the trees of the second step which arise from the incident edge groups of the first step. Thus, for

example, combining the results of Sections 2 and 3 we see that the boundary of the original group is connected if and only if each of the arc systems is indecomposable. It is possible to give a direct topological proof of this fact along using the methods of the next section, though we shall not do this explicitly. In this section, we shall set up a few main lemmas in a general context. In the Section 7, we shall introduce connectedness hypotheses into the picture. We explore applications to global cut points in Section 8.

Suppose that Γ is a hyperbolic group, and that Σ is a cofinite Γ -tree with two-ended edge stabilisers. Note that this is necessarily a quasiconvex splitting (as described in Section 4), since a two-ended subgroup of a hyperbolic group is necessarily quasiconvex (see, for example, [GhH]). We shall fix some vertex, $\omega \in V(\Sigma)$, and write $G = \Gamma(\omega)$. We suppose that G is not two-ended. By Proposition 4.2, G is quasiconvex, and hence intrinsically hyperbolic. We shall, in turn, want to consider splittings of G over finite groups, so to avoid any confusion later on, we shall alter our notation, so that it is specific to this situation.

Let Ξ be an indexing set which is in bijective correspondence with the set, $\vec{\Delta}(\omega)$, of directed edges of Σ with heads at ω . Thus, G permutes the elements of Ξ . There are finitely many G -orbits (since $\vec{\Delta}(\omega)/\Gamma(\omega)$ is finite). Given $\xi \in \Xi$, we write $H(\xi)$ for the stabiliser, in G , of ξ . Thus, if $\vec{e} \in \vec{\Delta}(\omega)$ is the edge corresponding to ξ , then $H(\xi) = \Gamma(\vec{e})$. In particular, $H(\xi)$ is two-ended. Let $J(\xi) = \Psi(\vec{e})$. Thus, $J(\xi)$ is a closed $H(\xi)$ -invariant subset of ΛG . Moreover, $\text{fr } J(\xi) = J(\xi) \cap \Lambda G = \Lambda H(\xi)$ consists of a pair of distinct points.

In this notation, we have:

Lemma 6.1 : $\partial\Gamma = \Lambda G \cup \bigcup_{\xi \in \Xi} J(\xi)$. ◇

Lemma 6.2 : $\text{diam } J(\xi) \rightarrow 0$ for $\xi \in \Xi$. ◇

Here, Lemma 6.1 is a rewriting of Lemma 4.3, and Lemma 6.2 is a restriction of Lemma 4.4.

If $K \subseteq \Lambda G$ is closed, we write $\Xi(K) = \{\xi \in \Xi \mid \text{fr } J(\xi) \subseteq K\}$, and write $\Upsilon(K) = K \cup \bigcup_{\xi \in \Xi(K)} J(\xi)$. Thus, Lemma 4.5 says that:

Lemma 6.3 : $\Upsilon(K)$ is a closed subset of $\partial\Gamma$. ◇

These observations tell us all we need to know about the groups $H(\xi)$ and sets $J(\xi)$ for the rest of this section. Thus, for the moment, we can forget how they were constructed.

Now, G is intrinsically hyperbolic, with ∂G identified with ΛG . We write $\Lambda G = \Lambda_0 G \sqcup \Lambda_\infty G$, corresponding to the partition $\partial G = \partial_0 G \sqcup \partial_\infty G$, as described in Section 5. Let T be a complete G -tree, so that $\partial T \equiv \Lambda_\infty G$. We write $V_{\text{fin}}(T)$ and $V_{\text{inf}}(T)$ respectively, for the sets of vertices of T of finite and infinite degree. Thus, $\Lambda_0 G = \bigsqcup_{v \in V(T)} \Lambda G(v)$. We note that if T is non-trivial (i.e. not a point), then $\Lambda_\infty G$ is dense in ΛG .

Given $\xi \in \Xi$, the subgroup $H(\xi)$ is two-ended. It is either elliptic or hyperbolic with respect to the G -tree T . We write Ξ_{ell} and Ξ_{hyp} , respectively, for the sets of $\xi \in \Xi$ such that $H(\xi)$ is elliptic or hyperbolic.

If $\xi \in \Xi_{\text{ell}}$, then $H(\xi)$ fixes a unique vertex $v(\xi) \in V_{\text{inf}}(T)$, so that $H(\xi) \subseteq G(v(\xi))$ and $\text{fr } J(\xi) \subseteq \Lambda G(v(\xi))$. Given $v \in V(T)$, we write $\Xi_{\text{ell}}(v) = \{\xi \in \Xi \mid H(\xi) \subseteq G(v)\}$. Thus

$\Xi_{\text{ell}}(v) \subseteq \Xi_{\text{ell}}$, and $\Xi_{\text{ell}}(v) = \emptyset$ for all $v \in V_{\text{fin}}(T)$. In fact, $\Xi_{\text{ell}} = \bigsqcup_{v \in V(T)} \Xi_{\text{ell}}(v)$.

Given $\xi \in \Xi_{\text{hyp}}$, we write $\beta(\xi) \subseteq T$ for the unique biinfinite arc in T preserved setwise by $H(\xi)$. Note that, under the identification of ∂T and $\Lambda_0 G$, we have $\partial\beta(\xi) = \Lambda H(\xi)$.

Suppose that $F \subseteq T$ is a finitely separated subgraph. Recall from Section 5 that $A(F)$ is defined as $A(F) = \partial F \cup \bigcup_{v \in V(F)} \Lambda G(v)$. Thus, by Lemma 5.3, $A(F)$ is closed in ΛG and hence in $\partial\Gamma$. We abbreviate $A(\Phi(\vec{e}))$ to $A(\vec{e})$. (So that $A(\vec{e})$ has the form $\Psi(\vec{e})$ in the notation of Section 4.)

If $F \subseteq T$ is finitely separated, we write $\Xi(F) = \Xi(A(F)) = \{\xi \in \Xi \mid \text{fr } J(\xi) \subseteq A(F)\}$. Thus, $\xi \in \Xi_{\text{ell}} \cap \Xi(F)$ if and only if $v(\xi) \in V(F)$. Also, $\xi \in \Xi_{\text{hyp}} \cap \Xi(F)$ if and only if $\partial\beta(\xi) \subseteq \partial F$.

If $\vec{e} \in \vec{E}(T)$, we shall abbreviate $\Xi(\vec{e}) = \Xi(\Phi(\vec{e}))$. Thus, $\xi \in \Xi(\vec{e})$ if and only if \vec{e} points away from $v(\xi)$ or $\beta(\xi)$. Suppose $v_0 \in V(T)$. Let $\alpha \subseteq T$ be the arc joining v_0 to $v(\xi)$ or to the nearest point of $\beta(\xi)$. Then, $\{\vec{e} \in \vec{\Omega}(v_0) \mid \xi \in \Xi(\vec{e})\}$ consists of the directed edges in α which point towards v_0 . In particular, this set is finite. Indeed, if $\Xi_0 \subseteq \Xi$ is finite, we see that $\{\vec{e} \in \vec{\Omega}(v_0) \mid \Xi_0 \cap \Xi(\vec{e}) \neq \emptyset\}$ is finite.

If $F \subseteq T$ is a finitely separated subgraph, we write

$$B(F) = A(F) \cup \bigcup_{\xi \in \Xi(F)} J(\xi).$$

In other words, $B(F) = \Upsilon(A(F))$, as defined earlier in this section. Thus, by Lemma 6.3, we have:

Lemma 6.4 : *The set $B(F) \subseteq \partial\Gamma$ is closed, for any finitely separated subgraph, F , of T .* \diamond

If $\vec{e} \in \vec{E}(T)$, we abbreviate $B(\vec{e}) = B(\Phi(\vec{e}))$.

Lemma 6.5 : *If $v_0 \in V(T)$, then $\text{diam } B(\vec{e}) \rightarrow 0$ for $\vec{e} \in \vec{\Omega}(v_0)$.*

Proof : Suppose $\delta > 0$. By Lemma 6.2, there is a finite subset $\Xi_0 \subseteq \Xi$ such that if $\xi \in \Xi \setminus \Xi_0$ then $\text{diam } J(\xi) \leq \delta/3$. Let $\vec{\Omega}_0 = \{\vec{e} \in \vec{\Omega}(v_0) \mid \Xi_0 \cap \Xi(\vec{e}) \neq \emptyset\}$. As observed above, $\vec{\Omega}_0$ is finite. Let $\vec{\Omega}_1 = \{\vec{e} \in \vec{\Omega}(v_0) \mid \text{diam } A(\vec{e}) \geq \delta/3\}$. By Lemma 4.4, $\vec{\Omega}_1$ is also finite.

Suppose $\vec{e} \in \vec{\Omega}(v_0) \setminus (\vec{\Omega}_0 \cup \vec{\Omega}_1)$. Suppose $x \in B(\vec{e})$. If $x \notin A(\vec{e})$, then $x \in J(\xi)$ for some $\xi \in \Xi(\vec{e})$. Since $\vec{e} \notin \vec{\Omega}_0$, $\Xi_0 \cap \Xi(\vec{e}) = \emptyset$, so $\xi \notin \Xi_0$. Therefore, $\text{diam } J(\xi) \leq \delta/3$. Now, $\text{fr } J(\xi) \subseteq A(\vec{e})$, and so $\rho(x, A(\vec{e})) \leq \delta/3$. This shows that $B(\vec{e})$ lies in a $(\delta/3)$ -neighbourhood of $A(\vec{e})$. Now, since $\vec{e} \notin \vec{\Omega}_1$, $\text{diam } A(\vec{e}) < \delta/3$ and so $\text{diam } B(\vec{e}) < \delta$. \diamond

Recall, from Section 3, that if $S \subseteq T$ is a subtree, then there is a natural projection $\pi_S : T \cup \partial T \rightarrow S \cup \partial S$. If $v \in V(S)$, we write $F(S, v)$ for the subtree $T \cap \pi_S^{-1}v$. If $R \subseteq S$ is a subtree, then we see that $F(S, v) \subseteq F(R, \pi_R v)$. Recall that $\vec{\Delta}(S) = \{\vec{e} \in \vec{E}(T) \mid \text{head}(\vec{e}) \in S, \text{tail}(\vec{e}) \notin S\}$. If $v \in V(S)$, set $\vec{\Delta}(S, v) = \vec{\Delta}(S) \cap \vec{\Delta}(v)$. We write $\vec{\Omega}(S)$ for the set of all directed edges pointing towards S , i.e. $\vec{\Omega}(S) = \bigcap_{v \in V(S)} \vec{\Omega}(v)$. Clearly,

$\vec{\Delta}(S) \subseteq \vec{\Omega}(S)$. Also if $R \subseteq S$ is a subtree, then $\vec{\Omega}(S) \subseteq \vec{\Omega}(R)$. If $v \in V(T) \setminus V(R)$, let $\vec{e}(R, v)$ be the directed edge with head at $\pi_R v$ which lies in the arc joining v to $\pi_R v$. In other words, $\vec{e}(R, v)$ is the unique edge in $\vec{\Delta}(R)$ such that $v \in \Phi(\vec{e}(R, v))$. Note that, if $v \in V(S) \setminus V(R)$, then $F(S, v) \subseteq \Phi(\vec{e}(R, v))$.

Let \mathcal{T} be the set of all finite subtrees of T . Given $\delta > 0$, let

$$\begin{aligned} \mathcal{T}_1(\delta) &= \{S \in \mathcal{T} \mid (\forall \vec{e} \in \vec{\Delta}(S))(\text{diam } B(\vec{e}) < \delta)\} \\ \mathcal{T}_2(\delta) &= \{S \in \mathcal{T} \mid (\forall v \in V(S) \cap V_{\text{fin}}(T))(\text{diam } B(F(S, v)) < \delta)\} \\ \mathcal{T}_3(\delta) &= \{S \in \mathcal{T} \mid (\forall v \in V(S) \cap V_{\text{inf}}(T))(\forall \vec{e} \in \vec{\Delta}(S, v))(\rho(\Lambda G(v), B(\vec{e})) < \delta)\}. \end{aligned}$$

Let $\mathcal{T}(\delta) = \mathcal{T}_1(\delta) \cap \mathcal{T}_2(\delta) \cap \mathcal{T}_3(\delta)$.

It is really the collection $\mathcal{T}(\delta)$ in which we shall ultimately be interested. It can be described a little more directly as follows. A finite tree, S , lies in $\mathcal{T}(\delta)$ if and only if for each $v \in V(S)$, we have either $v \in V_{\text{fin}}(T)$ and $\text{diam } B(F(S, v)) < \delta$ or else $v \in V_{\text{inf}}(T)$ and for all $\vec{e} \in \vec{\Delta}(S, v)$ we have $\text{diam } B(\vec{e}) < \delta$ and $\rho(\Lambda G(v), B(\vec{e})) < \delta$. It is this formulation we shall use in applications.

Note that if $R \in \mathcal{T}_1(\delta)$, then, in fact, $\text{diam } B(\vec{e}) < \delta$ for all $\vec{e} \in \vec{\Omega}(R)$. We see that if $R \in \mathcal{T}_1(\delta)$, $S \in \mathcal{T}$ and $R \subseteq S$, then $S \in \mathcal{T}_1(\delta)$. More to the point, we have:

Lemma 6.6 : *If $R \in \mathcal{T}(\delta)$, $S \in \mathcal{T}$ and $R \subseteq S$, then $S \in \mathcal{T}(\delta)$.*

Proof : As observed above, $S \in \mathcal{T}_1(\delta)$.

Suppose that $v \in V(S) \cap V_{\text{fin}}(T)$. If $v \in V(R)$, then $F(S, v) \subseteq F(R, v)$, and so $B(F(S, v)) \subseteq B(F(R, v))$. Therefore, $\text{diam } B(F(S, v)) \leq \text{diam } B(F(R, v)) < \delta$, since $R \in \mathcal{T}_2(\delta)$. On the other hand, if $v \notin V(R)$, then $F(S, v) \subseteq \Phi(\vec{e}(R, v))$, so $\text{diam } B(F(S, v)) \leq \text{diam } B(\vec{e}(R, v)) < \delta$, since $R \in \mathcal{T}_1(\delta)$. This shows that $S \in \mathcal{T}_2(\delta)$.

Finally, suppose $v \in V(S) \cap V_{\text{inf}}(T)$ and $\vec{e} \in \vec{\Delta}(S, v)$. If $v \in V(R)$, then $\vec{e} \in \vec{\Delta}(R, v)$, so $\rho(\Lambda G(v), B(\vec{e}))$, since $R \in \mathcal{T}_3(\delta)$. On the other hand, if $v \notin V(R)$, then $\{v\} \cup \Phi(\vec{e}) \subseteq F(R, \vec{e}(R, v))$, and so $\Lambda G(v) \cup B(\vec{e}) \subseteq B(F(R, \vec{e}(R, v)))$. But $\text{diam } B(F(R, \vec{e}(R, v))) < \delta$, since $R \in \mathcal{T}_1(\delta)$. In particular, $\rho(\Lambda G(v), B(\vec{e})) < \delta$. This shows that $S \in \mathcal{T}_3(\delta)$. \diamond

Lemma 6.7 : $\mathcal{T}(\delta) \neq \emptyset$.

Proof : Using Lemma 6.5, we can certainly find some $R \in \mathcal{T}_1(\delta)$. We form another finite tree, $S \supseteq R$, by adjoining a finite number of adjacent edges as follows. If $v \in V(R) \cap V_{\text{fin}}(T)$, we add all edges which are incident on v . If $v \in V(R) \cap V_{\text{inf}}(T)$, we add all those incident edges, e , which correspond to $\vec{e} \in \vec{\Delta}(R, v)$ for which $\rho(\Lambda G(v), B(\vec{e})) \geq \delta$. By Lemma 4.1, and the fact that $\vec{\Delta}(v)/G(v)$ is finite, there are only finitely many such \vec{e} . We thus see that S is finite. The fact that $S \in \mathcal{T}(\delta)$ follows by essentially the same arguments as were used in the proof of Lemma 6.6. \diamond

7. Connectedness properties of boundaries of hyperbolic groups

In this section, we continue the analysis of Section 6, bringing connectedness assumptions into play.

Suppose, as before, that Γ is a hyperbolic group, and that Σ is a cofinite Γ -tree with two-ended edge stabilisers. We now add the assumption that Γ is one ended, so that $\partial\Gamma$ is a continuum. In this case, we note:

Lemma 7.1 : *For each $\vec{e} \in \vec{E}(\Sigma)$, the set $\Psi(\vec{e})$ is connected.*

Proof : Since $\Gamma(e)$ is two-ended, we have $\text{fr } \Psi(\vec{e}) = \Lambda\Gamma(e) = \{a, b\}$, where $a, b \in \Psi(\vec{e})$ are distinct. Moreover, $\Psi(\vec{e})$ is closed and $\Gamma(e)$ -invariant. Also $\Psi(\vec{e}) \neq \{a, b\}$, since it must, for example, contain all points of $\partial\Phi(\vec{e})$.

Let K be a connected component of $\Psi(\vec{e})$. We claim that $K \cap \{a, b\} \neq \emptyset$. To see this, suppose $a, b \notin K$. There are subsets $K_1, K_2 \subseteq \Psi(\vec{e})$, containing K , with $a \notin K_1, b \notin K_2$, and which are clopen in $\Psi(\vec{e})$. Let $L = K_1 \cap K_2$. Thus, $K \subseteq L \subseteq \Psi(\vec{e}) \setminus \text{fr } \Psi(\vec{e})$. Since $\Psi(\vec{e})$ is closed in $\partial\Gamma$, so is L , and since $\Psi(\vec{e}) \setminus \partial\Psi(\vec{e})$ is open in $\partial\Gamma$, so also is L . In other words, L is clopen in M , contradicting the hypothesis that $\partial\Gamma$ is connected.

Suppose, then, that $a \in K$. Let $H \leq \Gamma(e)$ be the subgroup (of index at most 2) fixing a (and hence b). We see that K is H -invariant. Now $\Lambda H = \{a, b\}$ so either $b \in K$, or $K = \{a\}$. In the former case, we see that $K = \Psi(\vec{e})$, showing that $\Psi(\vec{e})$ is connected. In the latter case, we see, by a similar argument, that the component of K containing b equals $\{b\}$, giving the contradiction that $\Psi(\vec{e}) = \{a, b\}$. \diamond

Now, as in Section 6, we focus on one vertex $\omega \in V(\Sigma)$, and write $G = \Gamma(\omega)$. Let T be a complete G -tree. Now, $\Lambda G = \Lambda_0 G \sqcup \Lambda_\infty G$, where $\Lambda_0 G = \bigsqcup_{v \in V(T)} \Lambda G(v)$ and $\Lambda_\infty G$ is identified with ∂T . It is possible that T may be trivial, but most of the following discussion will be vacuous in that case. If not, then $\Lambda_\infty G$ is dense in ΛG .

We now reintroduce the notation used in Section 6, namely $\Xi, J(\xi), H(\xi), B(\vec{e})$, etc. Note that if $\xi \in \Xi$ corresponds to the directed edge \vec{e} of Σ , then $J(\xi)$ equals $\Psi(\vec{e})$ and the closure of $\partial\Gamma \setminus J(\xi)$ in $\partial\Gamma$ equals $\Psi(-\vec{e})$ (in the notation of Section 4). Thus, rephrasing Lemma 7.1, we get:

Lemma 7.2 : *For each $\xi \in \Xi$, the set $J(\xi)$ is connected. Moreover, the closure of $\partial\Gamma \setminus J(\xi)$ in $\partial\Gamma$ is also connected.* \diamond

Let $\mathcal{B} = \{\beta(\xi) \mid \xi \in \Xi_{\text{hyp}}\}$. Now, Ξ_{hyp}/G is finite, so Lemma 2.1 tells us that:

Lemma 7.3 : *The arc system \mathcal{B} is edge-finite.* \diamond

Now, since Γ is one-ended, the set of two-ended subgroups $\mathcal{H} = \{H(\xi) \mid \xi \in \Xi_{\text{hyp}}\}$ is indecomposable. Since \mathcal{B} is the set of axes of elements of \mathcal{H} , we see by Proposition 3.8 that:

Lemma 7.4 : \mathcal{B} is indecomposable. \diamond

Alternatively, one can give a direct proof of Lemma 7.4 along the lines of Lemma 5.4. Thus, if \mathcal{B} is decomposable, we can find two finitely separated subgraphs, F_1 and F_2 , of T , so that $\partial G = A(F_1) \sqcup A(F_2)$, and such that for all $\xi \in \Xi_{\text{hyp}}$, either $\partial\beta(\xi) \in \partial F_1$, or $\partial\beta(\xi) \in \partial F_2$. It follows that $\partial\Gamma = B(F_1) \sqcup B(F_2)$ are closed in $\partial\Gamma$, contradicting the assumption that $\partial\Gamma$ is connected.

To go further, we shall want some more general observations and notation regarding simplicial trees. For the moment, T can be any simplicial tree, and \mathcal{B} any arc system on T .

In Section 3, we associated to any finite subtree, $S \subseteq T$, an equivalence relation, $\sim_S = \sim_{S,\mathcal{B}}$, on ∂T . This, in turn, gives us a subpartition, $\mathcal{W}(S)$, of the set $V(S)$ of vertices of S . The elements of $\mathcal{W}(S)$ are the vertex sets of the connected components of the Whitehead graph, $\mathcal{G}(S)$.

More generally, we shall say that a subtree, S , of T is *bounded* if it has finite diameter in the combinatorial metric on T . In particular, every arc of \mathcal{B} meets S , if at all, in a compact interval (or point). We define the equivalence relation, $\sim_S = \sim_{S,\mathcal{B}}$ on ∂T in exactly the same way as for finite trees. We also get a graph $\mathcal{G}(S)$, and a subpartition, $\mathcal{W}(S)$ of $V(S)$ as before. Note that if \mathcal{B} is edge-finite, then $\mathcal{G}(S)$ is locally finite.

We have already observed that if $R \subseteq S$ is a subtree of S , then the relation \sim_R is coarser than the relation \sim_S (i.e. $x \sim_S y$ implies $x \sim_R y$). Moreover, the subpartition, $\mathcal{W}(R)$ of $V(R)$ can be described explicitly in terms of the subpartition $\mathcal{W}(S)$ and the map $\pi_R|_{V(S)} : V(S) \rightarrow V(R)$. To do this, define \cong to be the equivalence relation on $\mathcal{W}(S)$ generated by relations of the form $W \cong W'$ whenever $\pi_R W \cap \pi_R W' \neq \emptyset$. An element of $\mathcal{W}(R)$ is then a union of sets of the form $\pi_R W$ as W ranges over some \cong -class in $\mathcal{W}(S)$. For future reference, we note:

Lemma 7.5 : Suppose $R \subseteq S$ are bounded subtrees of T . If $W \in \mathcal{W}(S)$, $W \subseteq V(R)$, and $W \cap \pi_R(V(S) \setminus V(R)) = \emptyset$, then $W \in \mathcal{W}(R)$.

Proof : If $W' \in \mathcal{W}(S)$ and $W \cap \pi_R W' \neq \emptyset$, then $W \cap W' \neq \emptyset$. (To see this, choose $v \in W'$ with $\pi_R v \in W \subseteq V(R)$. Since $W \cap \pi_R(V(S) \setminus V(R)) = \emptyset$, it follows that $v \in V(R)$, so $\pi_R v = v$. Thus $v \in W \cap W'$.) Since $W, W' \in \mathcal{W}(S)$ we thus have $W = W'$, so $W' = \pi_R W'$. This shows that any set of the form $\pi_R W'$ for $W' \in \mathcal{W}(S)$ which meets W must, in fact, be equal to W . From the description of $\mathcal{W}(R)$ given above, we see that $W \in \mathcal{W}(R)$. \diamond

Given a directed edge $\vec{e} \in \vec{E}(T)$, let $\mathcal{S}(\vec{e})$ be the set of finite subtrees, S , of T with the property that $\vec{\Delta}(\text{head}(\vec{e})) \cap \vec{E}(S) = \{\vec{e}\}$ (i.e. $e \subseteq S$, and $\text{head}(\vec{e})$ is a terminal vertex of S). Given $S \in \mathcal{S}(\vec{e})$, we define the equivalence relation \simeq_S on $\partial\Phi(\vec{e})$ to be the transitive closure of relations of the form $x \simeq_S y$ whenever $\pi_S x = \pi_S y$ or $\partial\beta = \{x, y\}$ for some $\beta \in \mathcal{B}$, with $\beta \subseteq \Phi(\vec{e})$. Clearly, if $x \simeq_S y$ then $x \sim_S y$. Also, if $R, S \in \mathcal{S}(\vec{e})$ with $R \subseteq S$, then $x \simeq_S y$ implies $x \simeq_R y$. We can also define a subpartition, $\mathcal{W}(S, \vec{e})$, of $V(S) \setminus \{\text{head}(\vec{e})\}$, in a similar manner to $\mathcal{W}(S)$, as described in Section 3.

Suppose now that \mathcal{B} is edge-finite and indecomposable, and suppose $S \in \mathcal{S}(\vec{e})$. Sup-

pose $Q \subseteq \partial\Phi(\vec{e})$ is a \simeq_S -class. Since there is only one \sim_S -class, there must be some $\beta \in \mathcal{B}$ with one endpoint in Q and one endpoint in $\partial\Phi(-\vec{e})$. Thus, $e \subseteq \beta$. It follows that the number of \simeq_S -classes is bounded by the number of arcs in \mathcal{B} containing the edge e . By the edge-finiteness assumption, this number is finite. It follows that, as the trees $S \in \mathcal{S}(\vec{e})$ get bigger, the relations \simeq_S must stabilise. More precisely, there is a (unique) equivalence relation, \simeq , on $\partial\Phi(\vec{e})$ such that the set $\mathcal{S}_0(\vec{e}) = \{S \in \mathcal{S}(\vec{e}) \mid \simeq_S = \simeq\}$ contains all but finitely many elements of $\mathcal{S}(\vec{e})$. Note that if $R \in \mathcal{S}_0(\vec{e})$, $S \in \mathcal{S}(\vec{e})$, and $R \subseteq S$, then $S \in \mathcal{S}_0(\vec{e})$. Note also that there are finitely many \simeq -classes.

We now return to the set-up described earlier, with T a complete G -tree, and with $\mathcal{B} = \{\beta(\xi) \mid \xi \in \Xi_{\text{hyp}}\}$. We have seen that \mathcal{B} is edge-finite and indecomposable. We note:

Lemma 7.6 : *Suppose $\vec{e} \in \vec{E}(T)$ and $x, y \in \partial\Phi(\vec{e})$. If $x \simeq y$, then x and y lie in the same connected component of $B(\vec{e})$.*

Proof : Suppose, for contradiction that x and y lie in different components of $B(\vec{e})$. We can partition $B(\vec{e})$ into two closed subsets, $B(\vec{e}) = K \sqcup L$, with $x \in K$ and $y \in L$.

Let $\delta = \frac{1}{2}\rho(K, L) > 0$. By Lemma 6.7, we can find some $R \in \mathcal{T}(\delta)$. By Lemma 6.6, we can suppose that $S = R \cap (e \cup \Phi(\vec{e})) \in \mathcal{S}_0(\vec{e})$. (For example, take R to be the smallest tree containing a given element of $\mathcal{T}(S)$ and a given element of $\mathcal{S}_0(\vec{e})$.) Thus, $\simeq_S = \simeq$, so in particular, $x \simeq_S y$. Note that, if $v \in V(S) \setminus \{\text{head}(\vec{e})\}$, then $F(R, v) = F(S, v)$ (in the notation of Section 2).

Now, from the definition of the relation \simeq_S , we have a finite sequence, $x = x_0, x_1, \dots, x_n = y$ of points of $\partial\Phi(\vec{e})$, such that for each i , either $\pi_S x_i = \pi_S x_{i+1}$, or there is some $\xi \in \Xi_{\text{hyp}}$, with $\partial\beta(\xi) = \{x_i, x_{i+1}\}$. Now, $\partial\Phi(\vec{e}) \subseteq B(\vec{e}) = K \sqcup L$, so for each i , either $x_i \in K$ or $x_i \in L$. We claim, by induction on i , that $x_i \in K$ for all i .

Suppose, then, that $x_i \in K$. Suppose first, that $\{x_i, x_{i+1}\} = \partial\beta(\xi)$ for some $\xi \in \Xi_{\text{hyp}}$. We have that $x_i, x_{i+1} \in J(\xi) \subseteq B(\vec{e})$. Moreover, by Lemma 6.1, $J(\xi)$ is connected. It follows that $x_{i+1} \in K$.

We can thus suppose that $\pi_S x_i = \pi_S x_{i+1} = v \in V(S) \setminus \{\text{head}(\vec{e})\}$. Thus, $x_i, x_{i+1} \in \partial F(S, v) = \partial F(R, v) \subseteq B(F(R, v))$. Now, if $v \in V_{\text{fin}}(T)$, then, since $R \in \mathcal{T}(\delta)$, we have $\text{diam } B(F(R, v)) < \delta$. Therefore, $\rho(x_i, x_{i+1}) < \delta$ and so $x_{i+1} \in K$. Thus, we can assume that $v \in V_{\text{inf}}(T)$. Since $x_i \in \partial F(R, v)$, we have $x_i \in \partial\Phi(\vec{e})$ for some $\vec{e} \in \vec{\Delta}(R, v)$. Again, since $R \in \mathcal{T}(\delta)$, we have $\text{diam } B(\vec{e}) < \delta$ and $\rho(B(\vec{e}), \Lambda G(v)) < \delta$. Thus, $\rho(x_i, \Lambda G(v)) < 2\delta$. Similarly, $\rho(x_{i+1}, \Lambda G(v)) < 2\delta$. Now, $\Lambda G(v)$ is connected, and so it again follows that $x_{i+1} \in K$.

Thus, by induction on i , we arrive at the contradiction that $y = x_n \in K$. This shows that x and y lie in the same component of $B(\vec{e})$ as required. \diamond

Now, fix some $v \in V_{\text{inf}}(T)$, so that $G(v)$ is one-ended, and $\Lambda G(v)$ is a subcontinuum of $\partial\Gamma$.

We say that a $G(v)$ -invariant subtree, S , of T is *stable about* v if $S \cap \Phi(\vec{e}) \in \mathcal{S}_0(\vec{e})$ for all $\vec{e} \in \vec{\Delta}(v)$. Note that, since $\vec{\Delta}(v)/G(v)$ is finite, $S/G(v)$ is finite. In particular, we see that S is bounded (i.e. has finite diameter). Note that, since S contains every edge of T incident on v , we have $\pi_S \partial T \subseteq V(S) \setminus \{v\}$. Let $\sim_S = \sim_{S, \mathcal{B}}$ be the equivalence relation on

∂T as defined in Section 3 (in the case of finite trees). We remark that \sim_S is independent of the choice of stable tree, S , since it is easily seen to be definable purely in terms of the arc system \mathcal{B} , and the relations, \simeq for $\vec{e} \in \vec{\Delta}(v)$. We shall thus write \sim_S simply as \sim . Clearly, \sim is $G(v)$ -invariant. (It need not be trivial, since we are only assuming that S is bounded.)

We can certainly construct a stable tree about v by taking $S = \bigcup_{\vec{e} \in \vec{\Delta}(v)} S(\vec{e})$. In this case, $S \cap \Phi(\vec{e}) = S(\vec{e}) \in \mathcal{S}_0(\vec{e})$.

Note that we get a subpartition, $\mathcal{W}(S)$, of $V(S)$, as described in Section 3. Note that $\bigcup \mathcal{W}(S) \subseteq \pi_S \partial T$. In particular, $v \notin \bigcup \mathcal{W}(S)$.

Lemma 7.7 : *The setwise stabiliser, in $G(v)$, of every \sim -class is infinite.*

Proof : As described in Section 3, each \sim -class corresponds to an element of $\mathcal{W}(S)$. Moreover, $(\bigcup \mathcal{W}(S))/G(v) \subseteq V(S)/G(v)$ is finite. Thus, the lemma is equivalent to asserting that each element of $\mathcal{W}(S)$ is infinite.

Suppose, to the contrary, that $W \in \mathcal{W}(S)$ is finite. Let $\vec{\Delta}_0 = \{\vec{e} \in \vec{\Delta}(v) \mid W \cap S(\vec{e}) \neq \emptyset\}$, and let $R = \bigcup_{\vec{e} \in \vec{\Delta}_0} S(\vec{e})$. Thus, R is a finite subtree of S , and $W \subseteq V(R)$. Moreover, $\pi_R(V(S) \setminus V(R)) = \{v\}$, so, in particular, $W \cap \pi_R(V(S) \setminus V(R)) = \emptyset$. Thus, by Lemma 7.5, $W \in \mathcal{W}(R)$. But $v \in \bigcup \mathcal{W}(R)$ (since any element of $\partial \Phi(\vec{e})$ for $\vec{e} \in \vec{\Delta}(v) \setminus \vec{\Delta}_0$ projects to v under π_R). Thus, $\mathcal{W}(R) \neq \{W\}$. This shows that there is more than one \sim_R -class, contradicting the fact that \mathcal{B} is indecomposable. \diamond

Finally, we note:

Lemma 7.8 : *If $x, y \in \partial T$ with $x \sim y$, then x and y lie in the same quasicomponent of $\partial \Gamma \setminus \Lambda G(v)$.*

Proof : In fact, we shall show that x and y both lie in a compact connected subset, K , of $\partial \Gamma \setminus \Lambda G(v)$.

By the definition of the relation $\sim = \sim_S$, we can assume that either $\pi_S x = \pi_S y$ or there is some $\xi \in \Xi_{\text{hyp}}$ with $\partial \beta(\xi) = \{x, y\}$.

In the former case, let $w = \pi_S x = \pi_S y$. Thus, $w \in V(S(\vec{e}))$ for some $\vec{e} \in \vec{\Delta}(v)$. Since $S(\vec{e}) \in \mathcal{S}_0(\vec{e})$, we have $x \simeq y$, and so, by Lemma 7.6, x and y lie in the same component of $B(\vec{e})$. Call this component K . Thus, K is closed in $B(\vec{e})$ and hence in $\partial \Gamma$. Note that, from the definition of $B(\vec{e})$, we have $B(\vec{e}) \cap \Lambda G(v) = \emptyset$ and so $K \cap \Lambda G(v) = \emptyset$.

In the latter case, set $K = J(\xi)$. Thus, by Lemma 6.1, K is connected. Also $K \cap \Lambda G = \{x, y\} \subseteq \partial T$, and so, again, $K \cap \Lambda G(v) = \emptyset$. \diamond

8. Global cut points.

In this section, we set out the ‘‘inductive step’’ of the proof that a strongly accessible hyperbolic group has no global cut points in its boundary. In the light of the result announced in [DeP], we see that this, in fact, applies to all one-ended hyperbolic groups.

A more direct proof of the general case was given in [Swa] using the results of [Bo1,Bo3,L]. (See also [Bo4].)

Specifically, we shall show:

Theorem 8.1 : *Suppose that Γ is a one-ended hyperbolic group. Suppose that we represent Γ as a finite graph of groups over two-ended subgroups. Suppose that each maximal one-ended subgroup of each vertex group has no global cut point in its boundary (as an intrinsic hyperbolic group). Then, $\partial\Gamma$ has no global cut point.*

Before we start on the proof, we give a few general definitions and observations relating to global cut points.

Suppose that M is any continuum, i.e. a compact connected hausdorff space. (For the moment, the compactness assumption is irrelevant.) If $p \in M$, and $O, U \subseteq M$, we write OpU to mean that O and U are non-empty open subsets and that M is (set theoretically) a disjoint union $M = O \sqcup \{p\} \sqcup U$. Note that $\text{fr}O = \text{fr}U = \{p\}$. Also, it's not hard to see that $O \cup \{p\}$ and $U \cup \{p\}$ are connected. (More discussion of this is given in [Bo1].) We say that a point $p \in M$ is a *global cut point* if there exist $O, U \subseteq M$ with OpU .

Definition : If $Q \subseteq M$ is any subset, and $p \in M$, we say that Q is *indivisible in M at p* if whenever we have $O, U \subseteq M$ with OpU , then either $Q \cap O = \emptyset$ or $Q \cap U = \emptyset$.

If $R \subseteq M$ is another subset, we say that Q is *indivisible in M over R* , if it is indivisible in M at every point of R .

We say that Q is (*globally*) *indivisible* in M if it is indivisible at every point of M .

We say that M is (*intrinsically*) *indivisible* if it is indivisible in itself.

In other words, M is indivisible if and only if it does not contain a global cut point.

Obviously, if $P \subseteq Q \subseteq M$ and Q is indivisible in M , then so is P . Also any intrinsically indivisible subcontinuum of M is indivisible in M . We shall need the following simple observations:

Lemma 8.2 : *If $P, Q \subseteq M$ are indivisible in M , and $\text{card}(P \cap Q) \geq 2$, then $P \cup Q$ is indivisible in M .*

Proof : Suppose OpU . Choose any $x \in P \cap Q \setminus \{p\}$. We can assume that $x \in O$, so that $P \cap U = Q \cap U = \emptyset$. Thus $(P \cup Q) \cap U = \emptyset$. ◇

Lemma 8.3 : *Suppose that \mathcal{Q} is a chain of indivisible subsets of M (i.e. if $P, Q \in \mathcal{Q}$, then $P \subseteq Q$ or $Q \subseteq P$). Then $\bigcup \mathcal{Q}$ is indivisible.*

Proof : Suppose OpU , and $x \in O \cap (\bigcup \mathcal{Q})$ and $y \in U \cap (\bigcup \mathcal{Q})$. Then $x, y \in Q$ for some $Q \in \mathcal{Q}$, contradicting the indivisibility of Q . ◇

Lemma 8.4 : *If Q is indivisible in M , then so is its closure, \bar{Q} .*

Proof : If OpU , then we can assume that $O \cap Q = \emptyset$, so $O \cap \bar{Q} = \emptyset$. \diamond

Now, let Γ be a one-ended hyperbolic group, and let Σ be a cofinite Γ -tree with two-ended edge stabilisers. We begin with the following observation:

Lemma 8.5 : *If $\Lambda\Gamma(v)$ is indivisible in $\partial\Gamma$ for all $v \in V(\Sigma)$, then $\partial\Gamma$ is indivisible.*

Proof : Note that if $v, w \in V(\Sigma)$ are adjacent, then $\Gamma(v) \cap \Gamma(w)$ is two-ended, so $\Lambda\Gamma(v) \cap \Lambda\Gamma(w) = \Lambda(\Gamma(v) \cap \Gamma(w))$ consists of a pair of points. Thus, by Lemma 8.2, $\Lambda\Gamma(v) \cap \Lambda\Gamma(w)$ is indivisible in $\partial\Gamma$. By an induction argument, we see that $\bigcup_{v \in V(S)} \Lambda\Gamma(v)$ is indivisible for any finite subtree, $S \subseteq \Sigma$. Taking an exhaustion of Σ by an increasing sequence of finite subtrees, and applying Lemma 8.3, we see that $\bigcup_{v \in V(\Sigma)} \Lambda\Gamma(v)$ is indivisible. But this set is dense in $\partial\Gamma$ (since it is non-empty and Γ -invariant). The result follows by Lemma 8.4. \diamond

In fact, it's enough to verify the hypotheses of Lemma 8.5 for those $v \in V(\Sigma)$ for which $\Gamma(v)$ is not two-ended. To see this, first note that if α is a finite arc connecting two points $v_0, v_1 \in V(\Sigma)$ such that $\Gamma(v)$ is two ended for all $v \in V(\alpha) \setminus \{v_0, v_1\}$, then the groups $\Gamma(e)$ and $\Gamma(v)$ are all commensurable for all $e \in E(\alpha)$ and $v \in V(\alpha) \setminus \{v_0, v_1\}$. Now, since Γ is hyperbolic and not two-ended, there must be some $v_0 \in V(\Sigma)$ such that $\Gamma(v_0)$ is not two-ended. Suppose that $v \in V(\Sigma)$ is some other vertex. Connect v to v_0 by an arc in Σ , and let w be the first vertex of this arc for which $\Gamma(w)$ is not two-ended. Thus, $\Gamma(v) \cap \Gamma(w)$ has finite index $\Gamma(v)$, and so $\Lambda\Gamma(v) \subseteq \Lambda\Gamma(w)$. Clearly, if $\Lambda\Gamma(w)$ is indivisible in $\partial\Gamma$, then so is $\Lambda\Gamma(v)$.

As in Section 7, we now fix $\omega \in V_{\text{inf}}(\Sigma)$ and set $G = \Gamma(\omega)$. We are interested in the indivisibility properties of ΛG as a subset of $\partial\Gamma$. We aim to show that if ΛG is indivisible in $\partial\Gamma$ at each point of $\Lambda_0 G$, then it is (globally) indivisible in $\partial\Gamma$ (Corollary 8.8). Moreover, if $\Lambda G(v)$ is indivisible in $\partial\Gamma$ at some point $p \in \Lambda G(v)$, then ΛG is also indivisible in $\partial\Gamma$ at p (Proposition 8.9). As a corollary, we deduce (Corollary 8.10) that if $\Lambda G(v)$ is indivisible in $\partial\Gamma$ for all $v \in V(T)$, then ΛG is indivisible in $\partial\Gamma$. (Note that this is the essential ingredient in showing that $\partial\Gamma$ is intrinsically indivisible, as in Lemma 8.5.)

Recall the notation Ξ , $J(\xi)$, $H(\xi)$, $B(\vec{e})$ etc. from Section 6. We begin with the following observation:

Lemma 8.6 : *ΛG is indivisible in $\partial\Gamma$ over $\partial\Gamma \setminus \Lambda G$.*

Proof : Suppose $p \in \partial\Gamma \setminus \Lambda G$. Then, by Lemma 6.1, $p \in J(\xi) \setminus \text{fr } J(\xi)$ for some $\xi \in \Xi$. Let K be the closure of $\partial\Gamma \setminus J(\xi)$ in $\partial\Gamma$. By Lemma 7.2, K is connected. Moreover $\Lambda G \subseteq K$. Suppose $O, U \subseteq M$ with OpU . Without loss of generality, we can suppose that $K \cap U = \emptyset$. (Otherwise $O \cap K$ and $U \cap K$ would partition K .) But $\Lambda G \subseteq K$, and so $\Lambda G \cap U = \emptyset$. \diamond

Recall the notation $\mathcal{S}_0(\vec{e})$, \simeq_S etc. from Section 7.

For each $\vec{e} \in \vec{E}(T)$, we shall choose $S(\vec{e}) \in \mathcal{S}_0(\vec{e})$. We do this equivariantly with respect to the action of G . Thus, $N = \max\{\text{diam } S(\vec{e}) \mid \vec{e} \in \vec{E}(T)\} < \infty$ (where diam

denotes diameter with respect to combinatorial distance in T).

Lemma 8.7 : ΛG is indivisible in $\partial\Gamma$ over $\Lambda_\infty G$.

Proof : Clearly, we can assume that $\Lambda_\infty G$ is non-empty, and hence dense in ΛG . Suppose that $p \in \Lambda_\infty G$, and $O, U \subseteq \partial\Gamma$ with OpU . If $O \cap \Lambda G \neq \emptyset$, then $O \cap \Lambda_\infty G \neq \emptyset$, and similarly for U . Thus, suppose, for contradiction, that there exist $x \in O \cap \Lambda_\infty G$ and $y \in U \cap \Lambda_\infty G$. Clearly x, y and p are all distinct.

Now, let $v \in V(T)$ be the median of the points $x, y, p \in \partial T$. In other words, v is the unique intersection point of the three arcs connecting the points x, y and p pairwise. Let α be the ray from v to p , and let $w \in V(T)$ be that vertex at distance $N + 1$ from v along α . Let \vec{e} be the directed edge of α pointing towards p with $\text{head}(\vec{e}) = w$ (so that $\text{dist}(v, \text{tail}(\vec{e})) = N$). Thus $x, y \in \partial\Phi(\vec{e})$ and $p \in \partial\Phi(-\vec{e})$.

Write $S = S(\vec{e})$, so that $\text{diam } S \leq N < \text{dist}(v, w)$. Now v is the nearest point to w in the biinfinite arc connecting x to y . We see that this arc does not meet S , and so $\pi_S x = \pi_S y$. In particular, $x \simeq_S y$, and so, since $S \in \mathcal{S}_0(\vec{e})$, we have $x \simeq y$. By Lemma 7.6, x and y lie in the same component of $B(\vec{e})$. But, $\partial\Phi(-\vec{e}) \cap B(\vec{e}) = \emptyset$, and so $p \notin B(\vec{e})$. But this contradicts the fact that p separates x from y . (More formally, $O \cap B(\vec{e})$ and $U \cap B(\vec{e})$ partition $B(\vec{e})$ into two non-empty open sets.) \diamond

Putting Lemma 8.7 together with Lemma 8.6, we obtain:

Corollary 8.8 : If ΛG is indivisible in $\partial\Gamma$ over $\Lambda_0 G$, then ΛG is (globally) indivisible in $\partial\Gamma$. \diamond

Next, we show:

Proposition 8.9 : If $\Lambda G(v)$ is indivisible in $\partial\Gamma$ at the point $p \in \Lambda G(v)$, then ΛG is indivisible in $\partial\Gamma$ at p .

Proof : First, note that if T is trivial, then $G = G(v)$, so there is nothing to prove. We can thus assume that T is non-trivial.

Suppose that $O, U \subseteq \partial\Gamma$ with OpU . Since $\Lambda G(v)$ is indivisible in $\partial\Gamma$ at p , we can assume that $U \cap \Lambda G(v) = \emptyset$. We claim that $U \cap \Lambda G = \emptyset$. Since $\Lambda_\infty G$ is dense in ΛG , it's enough to show that $U \cap \Lambda_\infty G = \emptyset$.

Suppose, to the contrary, that there is some $x \in U \cap \Lambda_\infty G$. Let $G_0 \subseteq G(v)$ be the setwise stabiliser of the \sim -class of x . By Lemma 7.7, G_0 is infinite. Now a hyperbolic group cannot contain an infinite torsion subgroup (see for example [GhH]) and so we can find some $g \in G_0$ of infinite order.

Now, for each $i \in \mathbf{Z}$, $g^i x \sim x$, so, by Lemma 7.8, there is a connected subset (in fact a subcontinuum), K , containing x and $g^i x$, with $K \cap \Lambda G(v) = \emptyset$. Since $p \in \Lambda G(v)$, we have $K \subseteq \partial\Gamma \setminus \{p\}$. Thus, $K \subseteq U$. (Otherwise $O \cap K$ and $U \cap K$ would partition K .) In particular, $g^i x \in U$. Now, as $i \rightarrow \infty$, the sequences $g^i x$ and $g^{-i} x$ converge on distinct points, $a, b \in \Lambda G_0 \subseteq \Lambda G(v)$. Since $U \cup \{p\}$ is closed, we have $a, b \in U \cup \{p\}$, and so, without loss of generality, $a \in U$. But now, $a \in U \cap \Lambda G(v)$, contradicting the assumption that $U \cap \Lambda G(v) = \emptyset$. \diamond

Putting Proposition 8.9 together with Corollary 8.8, we get:

Corollary 8.10 : *Suppose that, for all $v \in V_{\text{inf}}(T)$, the continuum $\Lambda G(v)$ is indivisible in $\partial\Gamma$ over $\Lambda G(v)$. Then, ΛG is (globally) indivisible in $\partial\Gamma$. \diamond*

Of course, it's enough to suppose that each continuum $\Lambda G(v)$ is intrinsically indivisible.

Finally, putting Corollary 8.10 together with Lemma 8.5, we get the main result of this section, namely Theorem 8.1.

9. Strongly accessible groups.

In this final section, we look once more at the property of strong accessibility over finite and two-ended subgroups. We begin with general groups, and specialise to finitely presented groups. We finish by showing how Theorem 8.1, together with the results of [Bo1,Bo3] imply that the boundary of a one-ended strongly accessible hyperbolic group has no global cut point (Theorem 9.3).

As discussed in the introduction, the issue of strong accessibility is concerned with sequences of splittings of over a class of subgroups (in particular, the class of finite and two-ended subgroups), and when such sequences must terminate. In general, this may depend on the choices of splittings that we make at each stage of the process. We first describe a few general results which imply, at least for finitely presented groups, that we can assume that at any given stage, we can split over finite groups whenever this is possible.

Suppose, for the moment, that Γ is any group, and that G_1 and G_2 are one-ended subgroups with $G_1 \cap G_2$ infinite. Then the group, $\langle G_1 \cup G_2 \rangle$, generated by G_1 and G_2 is also one-ended. (For if not, there is a non-trivial action of $\langle G_1 \cup G_2 \rangle$ on a tree, T , with finite edge stabilisers. Now, since the groups, G_i are one-ended, they each fix a unique vertex of T . Since $G_1 \cap G_2$ is infinite, this must be the same vertex, contradicting the non-triviality of the action.) Note that essentially the same argument works if G_1 is one-ended and G_2 is two-ended.

Similarly, suppose that $G \leq \Gamma$ is one-ended, and $g \in \Gamma$ with $G \cap gGg^{-1}$ infinite. Then $\langle G, g \rangle$ is one-ended. (Since if $\langle G, g \rangle$ acts on a tree, T , with finite edge stabilisers, then G and gGg^{-1} must fix the same unique vertex of T . Thus, g must also fix this vertex, again showing that the action is trivial.) Recall that the commensurator, $\text{Comm}(G)$, of G is the set of elements $g \in \Gamma$ such that $G \cap gGg^{-1}$ has finite index in G . Thus, $\text{Comm}(G)$ is a subgroup of Γ containing G . We see that if G is one-ended, then so is $\text{Comm}(G)$.

Now, suppose that Γ is accessible over finite groups. Then every one-ended subgroup of Γ is contained in a unique maximal one-ended subgroup of Γ . Each maximal one-ended subgroup is equal to its commensurator, and there are only finitely many conjugacy classes of such subgroups. If G is a maximal one-ended subgroup, and $H \leq G$ is two-ended, then either $H \leq G$ or else $H \cap G$ is finite. Moreover, H can lie in at most one maximal one-ended subgroup. These observations follow from the remarks of the previous two paragraphs. They can also be deduced by considering the action of H on a complete Γ -tree.

Now, suppose that Γ splits as an amalgamated free product or HNN-extension over a two-ended subgroup. This corresponds to a Γ -tree, Σ , with just one orbit of edges, and with two-ended edge stabiliser. We consider two cases, depending on whether or not the edge group is elliptic or hyperbolic, i.e. whether or not it lies in a one-ended subgroup of Γ .

Consider, first, the case where the edge stabiliser of Σ does not lie in a one-ended subgroup, and hence intersects every one-ended subgroup in a finite group. In this case, we have:

Lemma 9.1 : *Suppose $v \in V(\Sigma)$. Then, each maximal one-ended subgroup of $\Gamma(v) = \Gamma_\Sigma(v)$ is a maximal one-ended subgroup of Γ . Moreover, every maximal one-ended subgroup of Γ arises in this way (for some $v \in V(\Sigma)$).*

Proof : Suppose, first, that G is any one-ended subgroup of Γ . Then, G must lie inside some (unique) vertex stabiliser $\Gamma(v)$. (Otherwise, G would split over a group of the form $G \cap H$, where H is an edge-stabiliser. But $G \cap H$ is finite, contradicting the fact that G is one-ended.) If G is maximal in Γ , then clearly it is also maximal in $\Gamma(v)$.

Conversely, suppose that G is a maximal one-ended subgroup of a vertex stabiliser, $\Gamma(v)$. Let G' be the unique maximal one-ended subgroup of Γ containing G . By the first paragraph, G' lies inside some vertex group, which must, in this case, be $\Gamma(v)$. By maximality in $\Gamma(v)$, we must therefore have $G = G'$. \diamond

The second case is when an edge group lies inside some one-ended subgroup. To consider this case, fix an edge e of Σ , with endpoints $v, w \in V(\Sigma)$. Now, $\Gamma(e)$ lies inside a unique maximal one-ended subgroup, Γ_0 , of Γ . Any other maximal one-ended subgroup of Γ must intersect $\Gamma(e)$ in a finite subgroup. In this case, we have:

Lemma 9.2 : *Γ_0 splits as an amalgamated free product or HNN extension over $\Gamma(e)$, with incident vertex groups equal to $\Gamma_0 \cap \Gamma(v)$ and $\Gamma_0 \cap \Gamma(w)$. Each maximal one-ended subgroup of $\Gamma(v)$ is a maximal one-ended subgroup of $\Gamma_0 \cap \Gamma(v)$ or of Γ (and similarly for w). Every maximal one-ended subgroup of $\Gamma_0 \cap \Gamma(v)$ arises in this way. Each maximal one-ended subgroup of Γ is conjugate, in Γ , to Γ_0 or to a maximal one-ended subgroup of $\Gamma(v)$ or $\Gamma(w)$.*

Proof : Suppose G is a maximal one-ended subgroup of Γ . Either G contains some edge-stabiliser, so that some conjugate of G contains $\Gamma(e)$ and hence equals Γ_0 , or else G meets each edge stabiliser in a finite group. In the latter case, we see, as in Lemma 9.1, that G is a maximal one-ended subgroup of a vertex group.

Now suppose that G is a maximal one-ended subgroup of $\Gamma(v)$. Let G' be the maximal one-ended subgroup of Γ containing G . From the first paragraph, we see that either $G' = \Gamma_0$, or G' is a maximal one-ended subgroup of $\Gamma(v)$. In the former case, we see that $G \subseteq \Gamma_0 \cap \Gamma(v)$, and must therefore be maximal one-ended in $\Gamma_0 \cap \Gamma(v)$. The latter case, we obtain $G = G'$.

Finally suppose that G is a maximal one-ended subgroup of $\Gamma_0 \cap \Gamma(v)$. Let G' be the maximal one-ended subgroup of $\Gamma(v)$ containing G . From the previous paragraph, we see

that $G' \subseteq \Gamma_0 \cap \Gamma(v)$, so $G = G'$.

It remains to show that Γ_0 splits over $\Gamma(e)$ in the manner described. This amounts to showing that if H is an edge stabiliser and a subgroup of $\Gamma_0 \cap \Gamma(v)$, then H is conjugate in $\Gamma_0 \cap \Gamma(v)$ to $\Gamma(e)$, (and similarly for w).

We know that there must be some $g \in \Gamma(v)$ such that $H = g\Gamma(e)g^{-1}$ (since Γ itself is assumed to be an amalgamated free product or HNN extension over $\Gamma(e)$). Now, $H \subseteq \Gamma_0 \cap g\Gamma_0g^{-1}$. Since H is infinite, it follows that the group generated by Γ_0 and $g\Gamma_0g^{-1}$ must be one-ended, and so, by maximality, must equal Γ_0 . Hence, $g\Gamma_0g^{-1} = \Gamma_0$. In particular, $g \in \text{Comm}(\Gamma_0)$. But, from the earlier discussion, $\text{Comm}(\Gamma_0) = \Gamma_0$, and so $g \in \Gamma_0 \cap \Gamma(v)$ as required. \diamond

We now go on to describe the notion of strong accessibility. To set up the notation, let Γ be any group, and let \mathcal{C} be any conjugacy-invariant set of subgroups of Γ . (In the case of interest, \mathcal{C} will be the set of all finite and two-ended subgroups of Γ .) We want to look at sequences of splittings of Γ over \mathcal{C} , where the only information retained at each stage will be the vertex groups of the previous splittings. In other words, we get a sequence of conjugacy invariant sets of subgroups of Γ . (In fact, if \mathcal{C} is closed under isomorphism, we can just view these as isomorphism classes of groups.) Note that finite groups can never split non-trivially, and so for our purposes, we can throw away finite subgroups whenever they arise.

To be more formal, suppose that \mathcal{J} and \mathcal{J}' are both conjugacy invariant sets of subgroups of Γ . We say that \mathcal{J}' is obtained by splitting \mathcal{J} over \mathcal{C} if it has the form $\mathcal{J}' = \bigcup_J \mathcal{J}(J)$, where $\mathcal{J}(J)$ is the set of (Γ -conjugacy classes of) infinite vertex groups of some splitting of J as a finite graph of groups over \mathcal{C} , and where J ranges over a conjugacy transversal in \mathcal{J} . Thus, a sequence of splittings of Γ over \mathcal{C} consists of a sequence, $\mathcal{J}_0, \mathcal{J}_1, \mathcal{J}_2, \dots$, where $\mathcal{J}_0 = \{\Gamma\}$, and each \mathcal{J}_{i+1} is obtained as a splitting of \mathcal{J}_i over \mathcal{C} in the manner just described. Note that, by induction, each of the sets \mathcal{J}_i is a finite union of conjugacy classes in Γ . Note also that we can assume, if we wish, by introducing some intermediate steps, that each \mathcal{J}_{i+1} is obtained from \mathcal{J}_i by splitting one of the conjugacy classes of \mathcal{J}_i as an amalgamated free product or HNN extension, while leaving the remaining groups unchanged. We say that the sequence terminates, if for some n , none of the elements of \mathcal{J}_n split non-trivially over \mathcal{C} . We say that Γ is *strongly accessible* over \mathcal{C} if there exists such a sequence which terminates.

Suppose that \mathcal{J} is a union of conjugacy classes of subgroups of Γ , each accessible over finite groups. Let $\mathcal{F}(\mathcal{J}) = \bigcup_{J \in \mathcal{J}} \mathcal{F}(J)$, where $\mathcal{F}(J)$ is the set of maximal one-ended subgroups of J . Thus $\mathcal{F}(\mathcal{J})$ is obtained by \mathcal{J} by splitting over the class of finite subgroups of Γ , in the sense defined above.

Let us now suppose that Γ is finitely presented, and that \mathcal{C} is the set of all finite and one-ended subgroups of Γ . Suppose that $(\mathcal{J}_i)_i$ is a sequence of splitting of Γ over \mathcal{C} . By induction, each element of each \mathcal{J}_i is finitely presented and hence accessible over finite groups. We can thus form a sequence $(\mathcal{F}_i)_i$ where $\mathcal{F}_i = \mathcal{F}(\mathcal{J}_i)$. Now, we can assume that \mathcal{J}_{i+1} is obtained from \mathcal{J}_i by splitting an element of \mathcal{J}_i as an amalgamated free product or HNN extension either over a finite group or over a two-ended group. In the former case, we see that $\mathcal{F}_{i+1} = \mathcal{F}_i$. In the latter case, we see, from Lemmas 9.1 and 9.2, that \mathcal{F}_{i+1} is obtained from \mathcal{F}_i by first splitting some element over a two-ended subgroup, and

then, if necessary splitting over some finite subgroups to reduce ourselves again to one-ended groups. Thus, after inserting some intermediate steps if necessary, we can suppose that the sequence $(\mathcal{F}_i)_i$ is also a sequence of splittings of Γ over \mathcal{C} . If the sequence $(\mathcal{J}_i)_i$ terminates at \mathcal{J}_n , then $\mathcal{F}_n = \mathcal{F}(\mathcal{J}_n) = \mathcal{J}_n$, so $(\mathcal{F}_i)_i$ also terminates (and in the same set of subgroups).

In summary, we see that if Γ is finitely presented, and strongly accessible over \mathcal{C} , then we can find a terminating sequence of splittings over \mathcal{C} where we split over finite groups wherever possible (in priority to splitting over two-ended subgroups). In other words, we only ever need to split one-ended groups over two-ended subgroups and to split infinite-ended and two-ended groups over finite subgroups.

Finally, suppose that Γ is a strongly accessible one-ended hyperbolic group, and that $\mathcal{J}_0, \mathcal{J}_1, \dots, \mathcal{J}_n$ is a sequence of splitting of Γ over finite and one-ended subgroups, which terminates in \mathcal{J}_n . In this case, each elements of each \mathcal{J}_i is quasiconvex, and hence intrinsically hyperbolic. Moreover, we can suppose, as above, that the only groups we ever split over two-ended groups are one-ended.

Now, each element of \mathcal{J}_n is one-ended and does not split over any two-ended subgroup. From the results of [Bo1,Bo3], we see that each element of \mathcal{J}_n has no global cut point in its boundary. Now, applying Theorem 8.1 inductively, we conclude that this is also true of Γ .

We have shown:

Theorem 9.3 : *Suppose that Γ is a one-ended hyperbolic group which is strongly accessible over finite and two-ended subgroups. Then, $\partial\Gamma$ has no global cut point. \diamond*

As mentioned in the introduction, Delzant and Potyagailo have shown that every finitely presented group, Γ , is strongly accessible over any “elementary” class of subgroups, \mathcal{C} . In particular, this deals with the case where Γ is hyperbolic, and where \mathcal{C} is the set of finite and two-ended subgroups of Γ . We thus conclude that the boundary of any one-ended hyperbolic group has no global cut point, and is thus locally connected by the result of [BeM].

References.

- [BeF1] M.Bestvina, M.Feighn, *Bounding the complexity of simplicial actions on trees* : Invent. Math. **103** (1991) 449–469.
- [BeF2] M.Bestvina, M.Feighn, *A combination theorem for negatively curved groups* : J. Differential Geom. **35** (1992) 85–101.
- [BeM] M.Bestvina, G.Mess, *The boundary of negatively curved groups* : J. Amer. Math. Soc. **4** (1991) 469–481.
- [Bo1] B.H.Bowditch, *Treelike structures arising from continua and convergence groups* : to appear in Memoirs Amer. Math. Soc.

- [Bo2] B.H.Bowditch, *Cut points and canonical splittings of hyperbolic groups* : to appear in Acta Math.
- [Bo3] B.H.Bowditch, *Group actions on trees and dendrons* : to appear in Topology.
- [Bo4] B.H.Bowditch, *Connectedness properties of limit sets* : preprint, Melbourne (1996).
- [Bo5] B.H.Bowditch, *Peripheral splittings of groups* : preprint, Southampton (1997).
- [DeP] T.Delzant, L.Potyagailo : paper in preparation.
- [DiD] W.Dicks, M.J.Dunwoody, *Groups acting on graphs* : Cambridge Studies in Advanced Mathematics No. 17, Cambridge University Press (1989).
- [Du] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449–457.
- [DuS] M.J.Dunwoody, M.E.Sageev, *JSJ-splittings for finitely presented groups over slender subgroups* : preprint, Southampton (1996).
- [Fe] S.Ferry, *A stable converse to the Vietoris-Smale theorem with applications to shape theory* : Trans. Amer. Math. Soc. **261** (1980) 369–386.
- [FuP] K.Fujiwara, P.Papasoglu, *JSJ decompositions of finitely presented groups and complexes of groups* : preprint (1997).
- [GhH] E.Ghys, P.de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Maths. 83, Birkhäuser (1990).
- [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.
- [HoY] J.G.Hocking, G.S.Young, *Topology* : Addison-Wesley (1961).
- [Ke] J.L.Kelley, *General topology* : Graduate Texts in Mathematics, 21, Springer-Verlag (reprint of Van Nostrand edition 1955).
- [Kr] J.Krasinkiewicz, *Local connectedness and pointed 1-movability* : Bull. Acad. Polon. Sci. Sér. Sci. Math. Atronom. Phys. **25** (1977) 1265–1269.
- [L] G.Levitt, *Non-nesting actions on real trees* : to appear in Bull. London Math. Soc.
- [Ma] I.Martinez, *Bord d'un produit amalgamé sur \mathbf{Z} de deux groupes libres ou de surface* : preprint, Orsay (1992).
- [Mi] M.Mihalik, *Semistability at ∞ of finitely generated groups, and solvable groups* : Topology and its Appl. **24** (1986) 259–264.
- [MiT] M.L.Mihalik, S.T.Tschantz, *Semistability of amalgamated products and HNN-extensions* : Memoirs Amer. Math. Soc. No. 471, American Mathematical Society, Providence, Rhode Island (1992).
- [O] J.-P.Otal, *Certaines relations d'équivalence sur l'ensemble des bouts d'un groupe libre* : J. London Math. Soc. **46** (1992) 123–139.

- [RS] E.Rips, Z.Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition* : Ann. of Math. **146** (1997) 53–109.
- [Se] Z.Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II* : to appear in GAFA.
- [Sh] H.Short, *Quasiconvexity and a theorem of Howson's* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 168–176.
- [St] J.R.Stallings, *Group theory and three-dimensional manifolds* : Yale Math. Monographs No. 4, Yale University Press, New Haven (1971).
- [Swa] G.A.Swarup, *On the cut point conjecture* : Electron. Res. Announc. Amer. Math. Soc. **2** (1996) 98–100 (Electronic).
- [Swe] E.L.Swenson, *A cut point tree for a continuum* : preprint (1997).
- [W] J.H.C.Whitehead, *On certain sets of elements in a free group* : Proc. London Math. soc. **41** (1936) 48–56.