Cut points and canonical splittings of hyperbolic groups.

B. H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
bhb@maths.soton.ac.uk

0. Introduction.

In this paper, we give a construction of the JSJ splitting of a one-ended hyperbolic group (in the sense of Gromov [Gr]), using the local cut point structure of the boundary. In particular, this gives the quasiisometry invariance of the splitting, as well the annulus theorem for hyperbolic groups. The canonical nature of the splitting is also immediate from this approach.

The notion of a JSJ splitting, in this context, was introduced by Sela [Se], who constructed such splittings for all (torsion free) hyperbolic groups. They take their name from the analogy with the characteristic submanifold construction for irreducible 3-manifolds described by Jaco and Shalen [JaS] and Johannson [Jo] (developing a theory outlined earlier by Waldhausen). The JSJ splitting gives a description of the set of all possible splittings of the group over two-ended subgroups, and thus tells us about the structure of the outer automorphism group.

We shall take as hypothesis here, the fact that the boundary is locally connected, i.e. a "Peano continuum". This is now known to be the case for all one-ended hyperbolic groups, from the results of [Bo1,Bo2,L,Sw,Bo5], as we shall discuss shortly. This uses the fact that local connectedness is implied by the non-existence of a global cut point [BeM].

A generalisation of the JSJ splitting to finitely presented groups has been given by Rips and Sela [RS]. The methods of [Se] and [RS] are founded on the theory of actions on **R**-trees. They consider only splittings over infinite cycic groups. It seems that their methods run into problems if one wants to consider, for example, splittings over infinite dihedral groups (see [MNS]).

A more general approach to this has recently been described by Dunwoody and Sageev [DuSa] using tracks on 2-complexes. Fujiwara and Papasoglu have obtained similar results using actions on products of trees [FuP]. These methods work in a more general context than those of this paper. (They deal with splittings of finitely presented groups over "slender" subgroups.) However, one looses some information about the splitting. For example, it is not known if the splitting is quasiisometry invariant in this generality. The annulus theorem would appear to generalise, though this does not follow immediately. A proof of the latter has recently been claimed in [DuSw] for finitely generated groups. (We shall return to this point later.) We shall see that for hyperbolic groups, all these results can be unified in one approach.

As we have suggested, deriving the spitting from an analysis of the boundary enables us to conclude that certain topological properties of the boundary are reflected in the structure of the group. For example, we see that the splitting is non-trivial if and only if the boundary has a local cut point (see Theorem 6.2). In fact, much information about the splitting can be read off immediately, without any knowledge of how the group acts on the boundary. In the course of the analysis, we shall also give an elementary proof that the boundary has no global cut point in the case where it is assumed to be locally connected (Proposition 5.4), without reference to the papers cited earlier. This is a converse to the result of Bestvina and Mess [BeM].

Our formulation of the JSJ splitting differs slightly from the Sela's original, though in the case of torsion free groups, it amounts to the same thing. Of course, for our result to apply to one-ended hyperbolic groups in general, we need the somewhat non-trivial fact that the boundary is necessarily locally connected. We suspect that many of the ideas of this paper can also be applied to relatively hyperbolic groups, a matter we aim to pursue in the future. Indeed, most of the analysis proceeds in a general dynamical context, without any specific geometric input.

In general, the boundary, $\partial\Gamma$, of a hyperbolic group, Γ, is a compact metrisable topological space, on which Γ acts as a discrete convergence group without parabolics, in the sense described in [GeM1]. Now, $\partial\Gamma$ is connected (i.e. a continuum) if and only if Γ is one-ended. By Stallings's theorem [St], this, in turn, happens if and only if Γ is not finite or virtually cyclic (i.e. two-ended) and does not split over a finite subgroup. In this paper we shall also assume that $\partial\Gamma$ is a Peano continuum. As noted above, this is equivalent to saying that $\partial\Gamma$ has no global cut point. In [Bo1,Bo2] this was shown to be the case if Γ does not split over any two-ended subgroup. (An alternative argument can be given via Levitt's generalisation of [Bo2] in [L].) This was generalised to strongly accessible groups in [Bo3]. Swarup showed how to adapt these arguments to deal with the general case [Sw]. This can be placed in a more general dynamical context [Bo5]. (It remains open as to whether every hyperbolic group must be strongly accessible.)

A "local cut point" can be defined as a point $x \in \partial\Gamma$ such that $\partial\Gamma \setminus \{x\}$ has more than one end. Putting the results of this paper together with those of [Bo1,Bo2] mentioned above, we deduce that a hyperbolic group has a local or global cut point in its boundary if and only if it either splits over a two-ended subgroup or is a "virtual semitriangle group" as defined below. In particular, we see that this property is quasiisometry invariant (provided we rule out cocompact fuchsian groups). Another consequence is the "annulus theorem" of Scott and Swarup [ScS] (stated in the torsion-free case), namely that a (non-fuchsian) hyperbolic group, Γ, splits over a two-ended subgroup if and only if there is some two-ended subgroup, $G \leq \Gamma$, such that the pair (Γ, G) has more than one end. As mentioned earlier, this has recently been generalised to finitely generated groups by Dunwoody and Swenson [DuSw]. (We also remark that a different algebraic adaptation of such 3-manifold results, in the form on a torus theorem for 3-dimensional Poincaré duality groups, had been obtained earlier by Kropholler — see [K] and the references therein.)

Note that fuchsian groups play a special role in this theory. By a *fuchsian group* we mean a non-elementary finitely generated group which acts properly discontinuously on the hyperbolic plane. We do not assume that the action is faithful, only that its kernel is finite. This kernel is canonically determined as the unique maximal finite normal subgroup, and so the quotient 2-orbifold is also canonically determined. The group action is cocompact if and only if the 2-orbifold is closed, or, equivalently, if the group is a virtual (closed) surface group. In this case, the ideal boundary of the group is a circle. Such a group splits over a two-ended group if and only if the 2-orbifold is neither a sphere with three cone points nor a triangle with mirrors. This is what we mean by a "virtual semitriangle group". (A semitriangle group is the orbifold fundamental group of a sphere with three cone points.)

By a *bounded fuchsian group*, we mean a non-elementary fuchsian group which is convex cocompact but not cocompact. Thus, the convex core of the quotient is a compact orbifold with non-empty boundary consisting of a disjoint union of compact 1-orbifolds (circles or intervals with mirrors). In this case, the ideal boundary of the group is cantor set with a natural cyclic order. The *peripheral subgroups* are the maximal two-ended subgroups which project to the fundamental groups of the boundary 1-orbifolds. The conjugacy classes of peripheral subgroups are thus in bijective correspondence with the boundary components.

The essential features of the JSJ splitting can be summarised as follows. Our formulation differs slightly from that given in [Se], though the result is more or less equivalent (at least in the torsion-free case). These matters will be elaborated on in Section 6.

Theorem 0.1 : *Suppose that Γ is a one-ended hyperbolic group, which is not a cocompact fuchsian group. Suppose that $\partial\Gamma$ is locally connected (or has no global cut point). Then there is a canonical splitting of Γ as a finite graph of groups such that each edge group is two-ended, and each vertex group is of one of the following three types:*

(1) a two-ended subgroup,

(2) a maximal "hanging fuchsian" subgroup, or

(3) a non-elementary quasiconvex subgroup not of type (2).

These types are mutually exclusive, and no two vertices of the same type are adjacent. Every vertex group is a full quasiconvex subgroup. Moreover, the edge groups that connect to any given vertex group of type (2) are precisely the peripheral subgroups of that group.

Finally, if $G \leq \Gamma$ is a two-ended subgroup such that (Γ, G) has more than one end, then G can be conjugated into one of the edge groups, or one of the vertex groups of type (1) or (2).

A "hanging fuchsian" subgroup, G, of Γ, is a virtually free quasiconvex subgroup together with a collection of "peripheral" two-ended subgroups, which arise from an isomorphism of G with a bounded fuchsian group. For a more careful description, see Section 6. (Thus a hanging fuchsian group coincides with what Sela calls a "quadratically hanging" subgroup in the case where Γ is torsion-free.) By a "full" (quasiconvex) group, we mean one which is not a finite index subgroup of any strictly larger subgroup of Γ.

Since the splitting is canonical, it is respected by any automorphism of Γ. ¿From this once can deduce that a finite index subgroup of the outer automorphism group is generated by "Dehn twists" about the edges, and 2-orbifold mapping classes arising form the hanging fuchsian subgroups. Thus it is virtually a direct product of a free abelian group and finitely many 2-orbifold mapping class groups [Se].

In the course of the analysis, we also get topological information about the boundary. It turns out that every local cut point plays a role in the JSJ splitting. More precisely, each local cut point lies in the limit set of (a conjugate of) a vertex group of type (1) or (2). It follows that the valency of any point (i.e. the number of ends of $\partial\Gamma \setminus \{x\}$) is always finite. In fact the maximum value of the valency of any point equals the maximal number of ends of a pair (Γ, G), as G ranges over all two-ended subgroups. This maximum is always finite.

The proof we give here proceeds by studying the topology of the boundary. In our analysis, the circle arises as a special case. This case has been analysed by Tukia [T1], except for certain exceptional cases which were subsequently dealt with independently by Gabai [Ga] and Casson and Jungreis [CJ]. In such a case, the group is a cocompact fuchsian group. In fact, Tukia's result is more general that this. For example it applies to the case of cyclically ordered cantor sets — a fact we shall use in order to describe the hanging fuchsian subgroups.

Most of the results given above can be arrived at without explicit reference to hyperbolic groups. Instead, we work with a uniform convergence group action on a metrisable Peano continuum. Such groups are necessarily hyperbolic [Bo6], but for most of the analysis we shall not need any geometric input. Apart from the reference to Tukia's work mentioned above, all the arguments are fairly elementary topology. There is one final point at which we need to refer back to hyperbolic groups — in order to verify that hanging fuchsian subgroups are indeed finitely generated (as required by the definition). In some ways, it would be nice to produce an argument which avoided this.

We have already observed that our splitting is canonical — it arises naturally out of the action of Γ on $\partial\Gamma$. Indeed, correctly formulated, it is unique. This uniqueness is best described in terms of the action of Γ on a simplicial tree Σ. This is done formally in Section 6. We note that the tree Σ arises purely out of the topology of $\partial\Gamma$. The action of Γ on $\partial\Gamma$ then induces an action on Σ.

The structure of the paper is roughly as follows. In Section 1, we give a brief account of how splittings of groups are reflected in the topology of the boundary. In Section 2, we describe certain abstract "order" structures which are meant to capture something the arrangement of local cut points in a (Peano) continuum. In Section 3, we carry out this analysis for a Peano continuum. In Section 4, we give a summary of some general results about convergence groups. In Section 5, we derive the JSJ splitting for a "uniform" convergence group acting on a metrisable Peano continuum (i.e. an action which is properly discontinuous and cocompact on the space of distinct triples). Finally, in Section 6, we give a summary of the results applied specifically to the case of a hyperbolic group.

1. Quasiconvex splittings of hyperbolic groups.

In this section, we recall some basic facts about hyperbolic groups in order to establish some terminology and notation. We also describe how splittings over quasiconvex subgroups are reflected in the topology of the boundary. It will be the aim of the rest of the paper to examine the converse implications in the case of splittings over two-ended subgroups.

Many of the statements given here concerning quasiconvex splittings appear to be "folklore", though I have found no explicit reference. These results are also needed for [Bo3].

Let Γ be a hyperbolic group, and X its Cayley graph (or any graph with a discrete cocompact Γ-action). Let $V(X)$ be the vertex set of X. We give X a Γ-invariant path metric, d, by assigning to each edge a length of 1. Thus $\partial\Gamma \equiv \partial X$. Now, Γ acts properly discontinuously and cocompactly on the space of distinct triples in ∂X (see Section 4). In particular, it acts as a convergence group without parabolics in the sense of Gehring and Martin [GeM1].

Any infinite order element, γ, in Γ is "loxodromic". In particular, $\mathrm{fix}(\gamma) = \{x \in \partial\Gamma \mid \gamma x = x\}$ consists of precisely two points. Moreover, $\langle\gamma\rangle$ acts properly discontinuously and cocompactly on $\partial\Gamma \setminus \mathrm{fix}(\gamma)$. Let $G(\gamma) = \{g \in \Gamma \mid g\,\mathrm{fix}(\gamma) = \mathrm{fix}(\gamma)\}$. Then $\langle\gamma\rangle$ has finite index in $G(\gamma)$. In fact, $G(\gamma)$ is the unique maximal virtually cyclic subgroup of Γ containing γ. Two loxodromics have a common fixed point if and only if they lie in a common virtually cyclic subgroup (and hence have both fixed points in common).

For brevity, we shall use the term "loxodromic" in preference to "infinite order element", and (in view of Stallings's theorem) we use "two-ended group" for "virtually cyclic group".

Given a loxodromic $\gamma \in \Gamma$, we write $l(\gamma) = \min\{d(x, \gamma x) \mid x \in X\}$ for the *translation length* of γ. Thus, $l(\gamma)$ is a conjugacy invariant. There are only finitely many conjugacy classes of elements of translation length (at most) any given number. (A more natural and interesting conjugacy invariant is the "stable length" as defined by Gromov [Gr], but the simpler definition above will serve for our purposes.)

A subgroup $G \leq \Gamma$ is *quasiconvex* if the G-orbit of some (and hence any) point of X is quasiconvex in X. Equivalently, G is quasiconvex if and only if there is some G-invariant quasiconvex subset $Q \subseteq X$ such that Q/G is compact (i.e. a finite graph). Thus, G is itself a hyperbolic group, and the "limit set" $\Lambda G \subseteq \partial\Gamma$ may be naturally (and hence G-equivariantly) identified with ∂G. Every two-ended subgroup is quasiconvex, with limit set consisting of two points.

If $G \leq \Gamma$ is quasiconvex, then the setwise stabiliser of ΛG is precisely the commensurator, $\mathrm{Comm}(G)$, of G in Γ. In this case, G has finite index in $\mathrm{Comm}(G)$. In fact, $\mathrm{Comm}(G)$ is the unique maximal subgroup of Γ which contains G as a subgroup of finite index. We say that G is *full* if $G = \mathrm{Comm}(G)$. (Note that, for any quasiconvex group, G, the group $\mathrm{Comm}(G)$ is full quasiconvex.)

If $G, H \leq \Gamma$ are both quasiconvex, then so is $G \cap H$ (see [Sho]). It follows that $\Lambda(G \cap H) = \Lambda G \cap \Lambda H$. In particular, if $\Lambda G \cap \Lambda H \neq \emptyset$, then $G \cap H$ contains a loxodromic.

Suppose $G \leq \Gamma$ is quasiconvex, and $\gamma \in G$ is loxodromic. All Γ-conjugates of γ in G have the same translation length, and hence fall into finitely many G-conjugacy classes. Put another way, only finitely many conjugates of G in Γ can contain a given loxodromic — or a given two-ended group. Putting this together with the observation of the previous paragraph, and the fact that distinct maximal two-ended subgroups cannot have a common fixed point, we obtain:

Lemma 1.1 : *Suppose $G \subseteq \Gamma$ is quasiconvex, and $x \in \partial\Gamma$. Then there are at most finitely many conjugates, G' of G in Γ such that $x \in \Lambda G'$.* ◇

We now go on to describe splittings of Γ over quasiconvex subgroups.

Suppose that Σ is a simplicial tree with vertex set $V(\Sigma)$ and edge set $E(\Sigma)$. Suppose that Γ acts simplicially on Σ. We can suppose that Γ acts minimally (i.e. there is no proper Γ-invariant subtree). Also, by subdividing the edges if necessary, we can suppose that there are no edge inversions. Given $v \in V(\Sigma)$ and $e \in E(\Sigma)$, we write $\Gamma(e)$ and $\Gamma(v)$ respectively for the vertex and edge stabilisers. If $e \in E(\Sigma)$ has endpoints $v, w \in V(\Sigma)$, then $\Gamma(e) = \Gamma(v) \cap \Gamma(w)$.

We shall assume that the quotient graph, Σ/Γ is finite. Thus, Γ is the fundamental group of a finite graph of groups (see for example [DiD]). We shall be interested in the case where Γ is hyperbolic, and each edge stabiliser is quasiconvex.

In such a case, we can construct a Γ-equivariant map $\phi : X \longrightarrow \Sigma$ by choosing arbitrarily the image, $\phi(v)$, for each v in a (finite) Γ-transversal of $V(X)$. This determines $\phi|V(X)$. Each edge of X is then mapped linearly onto a geodesic segment in Σ. Since the action on Σ is minimal, ϕ in necessarily surjective. On the level of quotients, this map is finite to one, except where it might collapse a finite subgraph of X/Γ to a point.

Given $e \in E(\Sigma)$, let $m(e)$ be the midpoint of e. Let $Q(e) = \phi^{-1}(m(e))$. Thus, $Q(e)/\Gamma$ is compact and non-empty, and so $Q(e)$ is quasiconvex.

Given $v \in V(\Sigma)$, let $E(v) \subseteq E(\Sigma)$ be the set of edges incident on v. Let $S(v) \subseteq \Sigma$ be the connected subset of Σ consisting of v together with the segments of each $e \in E(v)$ lying between v and $m(e)$. Thus $\partial S(v) = \{m(e) \mid e \in E(v)\}$. Let $Q(v) = \phi^{-1}S(v)$. Thus, $Q(v)$ is $\Gamma(v)$-invariant, and $Q(v)/\Gamma(v)$ is compact. Note that if $e \in E(v)$, then $Q(e) \subseteq Q(v)$. Suppose $v, w \in V(\Sigma)$. If v, w are not adjacent, then $Q(v) \cap Q(w) = \emptyset$, whereas if v and w are the endpoints of some edge $e \in E(\Sigma)$, the $Q(v) \cap Q(w) = Q(e)$. Since ϕ is surjective, the collection $\{Q(v)|v \in V(\Sigma)\}$ gives a locally finite cover of X.

Suppose β is a geodesic segment in X with both endpoints in $Q(e)$ for some $e \in E(\Sigma)$. Since $Q(e)$ is quasiconvex, β remains a bounded distance from $Q(e)$. In particular, its projection to Σ under ϕ has bounded diameter. Since there are finitely many conjugacy classes of edge stabilisers, this bound can be taken to be uniform.

Now suppose that $\alpha \subseteq X$ is a geodesic segment connecting two points of $Q(v)$ for some $v \in V(\Sigma)$. Since Σ is a tree, we see that each component of α which lies outside $Q(v)$ is of the type described in the previous paragraph, i.e. it connects two points of $Q(e)$ for some $e \in E(v)$, and thus remains within a bounded distance of $Q(e)$. Now, $Q(e) \subseteq Q(v)$ and so we deduce that α remains a bounded distance from $Q(v)$. If follows that $Q(v)$ is quasiconvex. Since $Q(v)/\Gamma(v)$ is compact, we see that $\Gamma(v)$ is quasiconvex. We have shown:

Proposition 1.2 : *Suppose that a hyperbolic group Γ splits as a finite graph of groups with each edge group quasiconvex. Then each vertex group is quasiconvex.* $\diamondsuit$

In certain circumstances, we can also say that the vertex stabilisers are full quasiconvex subgroups. Suppose $v \in V(\Sigma)$, and $\Gamma(v) \leq H \leq \Gamma$, with $[H : \Gamma(v)] < \infty$. Then, the H-orbit, Hv of v is finite. Thus, Hv has a well defined geometrical centre, $a \in \Sigma$. Now, a is H-invariant, and is either a vertex or the midpoint of an edge. By connecting v by and arc to a, we see that either $[H : \Gamma(v)] = 2$ and H is the stabiliser of an edge incident to v, so that there is an edge inversion (a possibility we have been disallowing), or there is a vertex, w, adjacent to v with $\Gamma(v) \leq \Gamma(w)$ and $[\Gamma(w), \Gamma(v)] < \infty$. Thus, if we can also rule out the latter possibility in any particular situation, we can conclude that all vertex stabilisers are full.

We can associate to Σ and "ideal boundary", $\partial\Sigma$, which we can think of as cofinality classes of geodesic rays in Σ. (We are only interested in $\partial\Sigma$ as a set, though it turns out that $\Sigma \cup \partial\Sigma$ can be given a natural topology as a dendrite — see [Bo1] for some discussion of this.)

Now suppose that α is a geodesic ray in X. We have observed that $\phi \circ \alpha$ can only double back on itself over bounded distances. Thus, either $\phi \circ \alpha$ converges in some ideal point in $\partial\Sigma$, or else it eventually remains within a bounded distance of some $v \in V(\Sigma)$. In the latter case, we see that α must converge on some limit point in $\Lambda\Gamma(v)$.

Now, suppose α and β are geodesic rays in X with $\phi \circ \alpha$ and $\phi \circ \beta$ both unbounded in Σ. Then $\phi \circ \alpha$ and $\phi \circ \beta$ converge on the same ideal point in $\partial\Sigma$ if and only if α and β remain a bounded distance apart. The "if" bit is fairly trivial. For the "only if" bit, note that α and β both pass in turn through a sequence of quasiconvex sets, $Q(v_1), Q(v_2), Q(v_3), \ldots$, where $(v_i)_{i \in \mathbf{N}}$ is a sequence of vertices of Σ converging on the ideal point of $\partial\Sigma$. By the local finiteness of $\{Q(v) \mid v \in V(\Sigma)\}$, we see that $d(a, Q(v_i)) \to \infty$, where a is any fixed point of X. Since the $Q(v_i)$ are uniformly quasiconvex, it's a simple geometric exercise to show that α and β remain a bounded distance apart — and thus converge to the same point of $\partial\Gamma$. We have essentially shown:

Proposition 1.3 : *The set $\partial\Gamma \setminus \bigcup_{v \in V(\Sigma)} \Lambda\Gamma(v)$ can be naturally identified with $\partial\Sigma$.* $\diamondsuit$

Note that if $e \in E(\Sigma)$ has endpoints $v, w \in V(\Sigma)$, then $\Gamma(e) = \Gamma(v) \cap \Gamma(w)$ so $\Lambda\Gamma(e) = \Lambda\Gamma(v) \cap \Lambda\Gamma(w)$. Also, by Lemma 1.1, we know that only finitely many of the sets $\Lambda\Gamma(v)$ for $v \in V(\Sigma)$ can meet any given point of $\partial\Gamma$.

Of particular interest to us is the case where Γ has one end, so that $\partial\Gamma$ is connected, and where each of the edge stabilisers is two-ended.

Let $e \in E(\Sigma)$, so that $\Lambda\Gamma(e)$ consists of two points. Now, $m(e)$ splits $\Sigma \cup \partial\Sigma$ into two components $\Sigma_1 \cup \partial\Sigma_1$ and $\Sigma_2 \cup \partial\Sigma_2$. If $v \in V(\Sigma_1)$ and $w \in V(\Sigma_2)$, then either $\Lambda\Gamma(v) \cap \Lambda\Gamma(w) = \emptyset$, or there is some loxodromic $\gamma \in \Gamma(v) \cap \Gamma(w) = \Gamma(e)$. It follows that in this case, $\Lambda\Gamma(v) \cap \Lambda\Gamma(w) = \Lambda\Gamma(e)$. Note that there are only finitely many $v \in V(\Sigma)$ for which $\Lambda\Gamma(e) \subseteq \Lambda\Gamma(v)$. We see that we get a natural partition of $\partial\Gamma \setminus \Lambda\Gamma(e)$ as $U_1 \sqcup U_2$, where

$$U_i = \partial\Sigma \cup \bigcup_{v \in V(\Sigma_i)} \Lambda\Gamma(v) \setminus \Lambda\Gamma(e).$$

By similar arguments to those already given, it's not hard to see that that the sets U_1 and U_2 are open in $\partial\Gamma$. (This is most easily seen by observing that $\phi^{-1}\Sigma_i$ is a quasiconvex subset of X, by a similar argument to Proposition 1.2. Moreover, $U_i \cup \Lambda\Gamma(e)$ is precisely the ideal boundary of this set, and hence closed in ∂X.) In particular, $\partial\Gamma \setminus \Lambda\Gamma(e)$ is disconnected. This in turn implies that $X/\Gamma(e)$ has more than one end, in other words, the pair $(\Gamma, \Gamma(e))$ has more than one end. This is an instance of a much more general fact that if a finitely generated group splits over a finitely generated subgroup, then the pair has more than one end.

So far, we have not made any assumption of local connectedness, though in that case, one can say more.

Suppose that $\partial\Gamma$ has no global cut point, so that $\partial\Gamma$ is a Peano continuum [BeM]. If $x, y \in \partial\Gamma$, then $\partial\Gamma \setminus \{x, y\}$ has finitely many components, the closure of each containing both x and y. If $G \leq \Gamma$ is two-ended, set $e(G)$ to be the number of components of $\partial\Gamma \setminus \Lambda G$. If $\gamma \in \Gamma$ is loxodromic, set $e(\gamma) = e(\langle\gamma\rangle)$. Note that some power, γ^n, of γ fixes setwise each component of $\partial\Gamma \setminus \mathrm{fix}(\gamma)$ (where $n \leq e(\gamma)!$). Thus, $(\partial\Gamma \setminus \mathrm{fix}(\gamma))/\langle\gamma^n\rangle$ has precisely $e(\gamma) = e(\gamma^n)$ components. But these components are in bijective correspondence with the the ends of $X/\langle\gamma^n\rangle$ and so, by definition, the pair $(\Gamma, \langle\gamma^n\rangle)$ has precisely $e(\gamma)$ ends. It follows that we could alternatively define $e(G)$ as the maximum number of ends of a pair (Γ, G') where G' ranges over the finite index subgroups of G.

Now, a consequence of our construction will be that if $G \leq \Gamma$ is a two-ended subgroup with $e(G) \geq 2$, then G can be conjugated into a vertex group of type (1) or (2) in the JSJ splitting. In particular, if there exists such a subgroup, then the splitting will be non-trivial. This gives the result of Scott and Swarup [ScS] referred to in the introduction. Another consequence is that there are only finitely many conjugacy classes of maximal two-ended subgroups G of Γ such that $e(G) \geq 3$.

2. Order structures.

In this section, we summarise two kinds of structure we use to describe the arrangement of local cut points in a (Peano) continuum.

The first structure is what we call a *pretree*. This consists of a set, T, together with a ternary "betweenness" relation, denoted xyz for $x, y, z \in T$. We speak of y as lying (strictly) "between" x and z. The relation should satisfy the following axioms:

(T0) If xyz, then $x \neq z$.

(T1) The relations xyz and xzy can never hold simultaneously.

(T2) The relation xyz holds if and only if zyx holds.

(T3) If xyz holds and $w \neq y$, then either xyw or wyz holds.

These axioms appear in a paper by Ward [W]. They are studied in some detail in [AN] and [Bo1]. They describe very general structures, in that most other treelike structures can be viewed as special cases of pretrees.

One can show that any finite subset, F, of a pretree can be embedded in a finite simplicial tree, τ, in such a way that xyz holds in F if and only if y separates x from z in τ. In fact, one could take this as an alternative definition of a pretree. It also gives a simple means of verifying many statements about pretrees.

Given $x, y \in T$, we write $(x, y) = \{z \in T \mid xzy\}$ and $[x, y] = \{x, y\} \cup (x, y)$. We refer to such sets as *intervals*. We see that such an interval carries a natural linear order (given an order on the pair $\{x, y\}$). If $x, y, z \in T$, then $[x, y] \cap [y, z] \cap [z, x]$ contains at most one point. If such a point exists, we refer to it as the *median* of (x, y, z), and denote it by $\mathrm{med}(x, y, z)$. If T is such that every three points have a median, then we refer to it as a *median pretree*.

We say that a pretree is *discrete* if $[x, y]$ is finite for all $x, y \in T$. A discrete median pretree is the same as a $\mathbf{Z}$-tree (except that the axioms of a $\mathbf{Z}$-tree are usually given in terms of the closed intervals, $[x, y]$, rather than the betweenness relation — see, for example

[Sha]). It can thus be realised as the vertex set of a simplicial tree, where "betweenness" is interpreted in the obvious way.

Every pretree can be embedded in a median pretree (see [AN] or [Bo1]). Of more significance here, is the fact that every discrete pretree, T, can be embedded in discrete median pretree, Θ. A simple way to describe such an embedding is as follows. A "star", F, in T is a maximal "null and full" subset of T. In other words, F satisfies the following. Firstly, there is no relation of the form xyz where $y \in T$ and $x, z \in F$ (F is "null and full"), and secondly, if $x \in T \setminus F$, there exist $y, z \in F$ with xyz (maximality). We can take Θ to be the disjoint union of T, together with the set of all stars in T containing at least 3 elements. One can show that Θ admits a natural structure of a median pretree, which induced the original pretree structure on T. (For more details, see [Bo1]).

An example of a pretree is given by any connected hausdorff topological space, M, where the relation xyz is interpreted to mean that y separates x from z in M (i.e. we can write $M \setminus \{y\} = U \sqcup V$, where $U \ni x$ and $V \ni z$ are open subsets of M). Proofs are given in [W], [Bo1] and [Bo4]. In particular, we derive the well-known fact that if $x, y \in M$, then the set, $[x, y]$, of points separating x from y, together with $\{x, y\}$, has a natural linear order.

More generally, we can speak of a "linearly ordered subset" of M; in other words, a subset, $L \subseteq M$ with a linear order, $<$, such that if $x, y, z \in L$, then xyz if and only if $x < y < z$ or $z < y < x$. Thus, any interval is such a set. For future reference, we note that if $x < y < z < w$, then $\{x, z\}$ separates y from w in M. Also, $M \setminus \{y, z\}$ has at least 3 components.

So far, we have not made use of compactness or local connectedness. In this special case, we can say more [HY]:

Lemma 2.1 : *Suppose M is a Peano continuum, and $x, y \in M$. Then $[x, y]$ is a closed subset of M. Moreover, the subspace topology on $[x, y]$ agrees with the order topology.* $\diamondsuit$

As a corollary, we note that the closure of a linearly ordered set of cut points in a Peano continuum is also linearly ordered.

(We remark that another advantage of working with Peano continua is that it allows us to deal with components of open subsets, whereas, for an arbitrary continuum, it is more appropriate to work with quasicomponents — for a locally connected space, these notions coincide [HY].)

The other type of "order structure" in which we are interested is, in some sense, at the opposite extreme to that of a pretree, namely a cyclic order. This is a familiar notion. It can be defined as a set, Δ, together with a 4-ary relation, denoted $\delta(x, y, z, w)$ for $x, y, z, w \in \Delta$ with the following property. If $F \subseteq \Delta$ is any finite subset, then we can embed F in the circle S^1, such that if $x, y, z, w \in F$, then $\delta(x, y, z, w)$ holds if and only if x and z lie in different components of $S^1 \setminus \{y, w\}$. It's not hard to write down explicit axioms for δ, though we won't bother here. (Note that it would be enough to assume our finite set, F, has at most 5 elements.) When dealing with cyclic orders we implicitly assume that $\text{card}(\Delta) \geq 4$.

Every cyclic order carries a natural order topology. A base for the open sets is given by the collection of sets of the form $\{x \in \Delta \mid \delta(a,b,c,x)\}$ where a,b,c are any three distinct points of Δ. We say that Δ is *separable* if it has a countable dense subset. Any compact separable cyclically ordered set can be embedded as a closed subset of the circle.

We say that two points $x,y \in \Delta$ are *adjacent* if there do not exist points $z,w \in \Delta$ such that $\delta(x,z,y,w)$. A *jump* in Δ is an unordered pair of adjacent points. Let $J(\Delta)$ be the set of all jumps. Note that two jumps can only intersect in an isolated point. Note also that Δ is a cantor set, if and only if it is compact, separable, contains no isolated points and $\bigcup J(\Delta)$ is dense in Δ.

Suppose now that M is a continuum, and that $\sigma \subseteq M$. Given $x,y,z,w \in \sigma$, write $\delta(x,y,z,w)$ if y and w are separated by $\{x,z\}$. We say that σ is a *cyclically separating set* if the relation δ is a cyclic order on σ.

Lemma 2.2 : *Suppose that M is a Peano continuum, and $\sigma \subseteq M$ is a cyclically separating set. Then so is its topological closure, $\bar{\sigma}$. Moreover, the subspace topology on $\bar{\sigma}$ (or on σ) agrees with the cyclic order topology.*

Proof : We can effectively reduce this to the case of linear orders. Choose any $a \in \sigma$, and let $M(a)$ be the space obtained by adjoining the space of ends to $M \setminus \{a\}$. Thus $M(a)$ is a Peano continuum, and $\sigma \setminus \{a\}$ is a linearly ordered set of cut points in $M(a)$. We can now apply Lemma 2.1, and the subsequent observation to this set. The result can be deduced by first splitting σ into two subintervals, and using the above to deal separately with the closures of each.

(One could alternatively go back to first principles and adapt the arguments of [HY] to this situation.) $\diamondsuit$

Suppose now that M is a Peano continuum with no global cut point. Given distinct points $x,y \in M$, let $\mathcal{U}(x,y)$ be the set of components of $M \setminus \{x,y\}$. Thus, $\mathcal{U}(x,y)$ is finite, and each element $U \in \mathcal{U}(x,y)$ is open and connected with $\partial U = \{x,y\}$.

Suppose $\Delta \subseteq M$ is a closed cyclically separating set. Let $\mathcal{U}(\Delta)$ be the set of components of $M \setminus \Delta$. Suppose $\theta = \{x,y\} \in J(\Delta)$ is a jump. We have $\operatorname{card} \mathcal{U}(x,y) \geq 2$. Moreover, there is some $U \in \mathcal{U}(x,y)$ such that $\Delta \subseteq \bar{U} = U \cup \{x,y\}$. Let $\mathcal{U}_\Delta(\theta) = \mathcal{U}(x,y) \setminus \{U\}$.

Lemma 2.3 : $\mathcal{U}(\Delta) = \bigsqcup_{\theta \in J(\Delta)} \mathcal{U}_\Delta(\theta)$.

Proof : Suppose $U \in \mathcal{U}(\Delta)$. Since M has no global cut point, ∂U has at least two elements. Suppose $x,y \in \partial U \subseteq \Delta$. Now, we must have $\{x,y\} \in J(\Delta)$ (for if $z,w \in \Delta$ with $\delta(x,z,y,w)$, we would have $U \cup \{x,y\} \subseteq M \setminus \{z,w\}$ giving the contradiction that x and y lie in the same component of $M \setminus \{z,w\}$). Let $\theta = \{x,y\}$. Since $x,y \in \partial U$ were arbitrary, and it is impossible for three distinct elements of Δ to be mutually adjacent, we see that $\partial U = \{x,y\}$, and so $U \in \mathcal{U}(x,y)$. Now $U \cap \Delta = \emptyset$, and so $U \in \mathcal{U}_\Delta(\theta)$. We have shown that $\mathcal{U}(\Delta) \subseteq \bigcup_{\theta \in J(\Delta)} \mathcal{U}_\Delta(\theta)$.

Conversely, suppose $\theta \in J(\Delta)$ and $U \in \mathcal{U}_\Delta(\theta)$. Since $U \cap \Delta = \emptyset$ and U is connected, we must have $U \subseteq V$ for some $V \in \mathcal{U}(\Delta)$. Thus, $V \in \mathcal{U}_\Delta(\theta')$ for some $\theta' \in J(\Delta)$. Since

$U \cap \Delta = \emptyset$, we have $\partial U \subseteq \partial V$. Thus $\theta = \partial U = \partial V = \theta'$, and so $U = V \in \mathcal{U}(\Delta)$ as required.

Finally, to the see that we have a disjoint union, note that in $U \in \mathcal{U}_\Delta(\theta)$, then $\partial U = \theta$. $\diamondsuit$

We define Π be the space of distinct unordered pairs in M. Thus, Π takes its topology by identifying it as $M \times M$ minus the diagonal, quotiented out by the involution which swaps the entries in a pair.

Note that if Δ is a cyclically separating set, then $J(\Delta) \subseteq \Pi$. A simple consequence of Lemma 2.2 is:

Lemma 2.4 : *$J(\Delta)$ is a discrete subset of Π.* $\diamondsuit$

3. The structure of local cut points.

Throughout this section, (except where otherwise stated), M will be a Peano continuum with no global cut points. We discuss the structure of local cut points in M. We also introduce the notion of an "annulus" which will be used throughout the rest of the paper.

Suppose $x \in M$. We define the *valency* of x, denoted by $\mathrm{val}(x) \in \mathbf{N} \cup \{\infty\}$, to be the number of ends of the locally compact space $M \setminus \{x\}$. We say that x is a *local cut point* if $\mathrm{val}(x) \geq 2$. Given $n \in \mathbf{N} \cup \{\infty\}$, we write $M(n) = \{x \in M \mid \mathrm{val}(x) = n\}$ and $M(n+) = \{x \in M \mid \mathrm{val}(x) \geq n\}$. We shall be particularly interested in the subsets $M(2)$ and $M(3+)$. (In the case in which we are really interested, it turns out that $M(\infty) = \emptyset$, so it doesn't much matter whether or not we include this in $M(3+)$.)

Given $x, y \in M$, we defined, in Section 2, $\mathcal{U}(x, y)$ to be the set of components of $M \setminus \{x, y\}$. We write $N(x, y) = \mathrm{card}\, \mathcal{U}(x, y) \in \mathbf{N}$. Note that $N(x, y) \leq \min\{\mathrm{val}(x), \mathrm{val}(y)\}$. We define a relation $\sim$ on $M(2)$, by $x \sim y$ if and only if either $x = y$ or $N(x, y) = 2$. Note that if $x \sim y$ and $x \neq y$, then x and y are both local cut points.

The following construction will be useful in deducing a few basic properties of this relation. Suppose that $F \subseteq M$ is a finite set. Let $\mathcal{C}(F)$ be the set of components of $M \setminus F$. Thus $\mathcal{C}(F)$ is finite. If $U \in \mathcal{C}(F)$, we say that U is *adjacent* to $x \in F$ if $x \in \bar{U}$. We can thus define the bipartite graph $\mathcal{G} = \mathcal{G}(F)$ with vertex set $V(\mathcal{G}) = V_0(\mathcal{G}) \sqcup V_1(\mathcal{G})$, where $V_0(\mathcal{G}) = F$ and $V_1(\mathcal{G}) = \mathcal{C}(F)$. We define edge set $E(\mathcal{G})$ by joining $x \in V_0(\mathcal{G})$ to $u \in V_1(\mathcal{G})$ if they are adjacent in the sense already defined.

We make the following observations. Firstly, $\mathcal{G}$ is a connected bipartite graph. If $x \in F = V_0(\mathcal{G})$, then $\deg(x) \leq \mathrm{val}(x)$, where $\deg(x)$ is the degree of x in $\mathcal{G}$. No point of $V_0(\mathcal{G})$ can separate $\mathcal{G}$ (otherwise it would be a global cut point for M). It follows that each vertex of $V_1(\mathcal{G})$ has degree at least 2.

Lemma 3.1 : *The relation $\sim$ is an equivalence relation on $M(2)$.*

Proof : It is clear that $\sim$ is reflexive and symmetric. So, suppose that $x \sim y$ and $x \sim z$. We claim that $y \sim z$. We can assume that x, y, z are all distinct. Let $F = \{x, y, z\}$, and let $\mathcal{G} = \mathcal{G}(F)$ be the graph described above. Thus, $\deg(x) \leq 2$, and the sets $\{x, y\}$ and $\{x, z\}$ both separate $\mathcal{G}$.

Suppose there is some vertex $u \in V_1(\mathcal{G})$ connected to each element of F. Since $\{x, y\}$ separates, there must be some component, C_1, of $\mathcal{G} \setminus \{x, y\}$ with $\partial C_1 = \{x, y\}$. Similarly, there is another component, C_2, with $\partial C_2 = \{x, z\}$. Also the component, C_3, containing u has $\partial C_3 = \{x, y, z\}$. These components are all distinct, so we get the contradiction that $\deg(x) \geq 3$.

We conclude that every vertex in $V_1(\mathcal{G})$ has degree 2. Now since x is not a global cut point, there must be a vertex $u \in V_1(\mathcal{G})$ adjacent to both y and z. Similarly, there is a vertex v adjacent to both z and x, and a vertex w adjacent to both x and y. Thus, $xvzuywx$ is a circuit in $\mathcal{G}$. Since x cannot be connected to u, we see that $\{y, z\}$ separates $\mathcal{G}$, and so $y \sim z$ as required. $\diamondsuit$

(We remark that we have only really used that fact that $x \in M(2)$.)

Lemma 3.2 : *If $\sigma \subseteq M(2)$ is a $\sim$-equivalence class, then σ is a cyclically separating set.*

Proof : Suppose $F \subseteq \sigma$ is finite. Let $\mathcal{G} = \mathcal{G}(F)$. Thus, each vertex of $V_0(\mathcal{G})$ has degree at most 2.

Suppose, for contradiction, that some $u \in V_1(\mathcal{G})$ has degree at least 3. Thus, u is adjacent to the a subset of vertices $W \subseteq V_0(\mathcal{G})$ with $\operatorname{card}(W) \geq 3$. Since any pair, x, y, of distinct elements of W separates $\mathcal{G}$, we see that there must be at least one component, C, of $\mathcal{G} \setminus W$, with $\partial C = \{x, y\}$. But, since $\operatorname{card}(W) \geq 3$, this contradicts the fact that the degree of each vertex of W is at most 2.

We deduce that every vertex of $\mathcal{G}$ has degree at most 2. Since $\mathcal{G}$ connected, it must be either an arc or a circle. But since no point of $V_0(\mathcal{G})$ separates, it cannot be an arc, and is thus a circle. Now the separation properties of $F = V_0(\mathcal{G})$ are the same in $\mathcal{G}$ as in M. We deduce that σ is cyclically separating. $\diamondsuit$

Lemma 3.3 : *Suppose that $x, y, z, w \in M(2)$. Suppose that $z \sim w$ and that z and w lie in distinct components of $M \setminus \{x, y\}$. Then, $x \sim y \sim z \sim w$.*

Proof : Let $F = \{x, y, z, w\}$, and $\mathcal{G} = \mathcal{G}(F)$. Since $\{x, y\}$ separates z and w, there can be no vertex of $V_1(\mathcal{G})$ adjacent to both z and w. Suppose that there is some $u \in V_1(\mathcal{G})$ of degree at least 3. Then $\deg(u) = 3$, and, without loss of generality, u is adjacent to the points x, y and z. Since $\mathcal{G}$ is connected, there must be another vertex $v \in V_1(\mathcal{G})$ adjacent to w, and either x or y, say x. Since v is not adjacent to z, it must be adjacent to y. Since $\deg(y) \leq 2$, and z and w are non-separating, this must account for all of $\mathcal{G}$. But now, $\mathcal{G} \setminus \{z, w\}$ is connected, contradicting the hypothesis that $z \sim w$.

We thus conclude that every vertex of $\mathcal{G}$ has degree at most 2. But now, as in the proof of Lemma 3.2, we deduce that $\mathcal{G}$ is a circle. In particular, $x \sim y \sim z \sim w$. $\diamondsuit$

Lemma 3.4 : *Suppose that U_1, U_2, U_3 are disjoint open connected subsets of M. Suppose that $x_i, y_i \in U_i \cap M(2)$, for $i = 1, 2, 3$. Suppose that $x_1 \sim x_2 \sim x_3$ and $y_1 \sim y_2 \sim y_3$. Then, $x_1 \sim y_1 \sim x_2 \sim y_2 \sim x_3 \sim y_3$.*

Proof : We can suppose that $x_i \neq y_i$ for each i. Since $x_1, x_2, x_3 \in M(2)$, we see that $M \setminus \{x_1, x_2, x_3\}$ has precisely three components, each adjacent to two of the x_i. Let W_1, W_2, W_3 be these components, with $x_i \notin \bar{W}_i$. Now $\{x_1\} \cup W_2 \cup W_3$ is a connected component of $M \setminus \{x_2, x_3\}$, and U_1 is a connected subset of $M \setminus \{x_2, x_3\}$. Thus, $U_1 \subseteq \{x_1\} \cup W_2 \cup W_3$. Since $y_1 \in U_1 \setminus \{x_1\}$, we see that $y_1 \in W_2 \cup W_3$. Similarly, $y_2 \in W_3 \cup W_1$ and $y_3 \in W_1 \cup W_2$. Now, without loss of generality, we can suppose that $y_1 \in W_2$. Since $y_2 \in W_3 \cup W_1$, we see that y_1 and y_2 lie in different components of $M \setminus \{x_1, x_3\}$. Also, by hypothesis, $y_1 \sim y_2$. Thus, applying Lemma 3.4 to the set $\{x_1, x_3, y_1, y_2\}$, we see that $x_1 \sim x_3 \sim y_1 \sim y_2$. The result now follows from Lemma 3.1. ♢

As an immediate corollary, we get:

Lemma 3.5 : *If $\sigma, \tau \subseteq M(2)$ are $\sim$-classes, with* $\mathrm{card}(\bar{\sigma} \cap \bar{\tau}) \geq 3$*, then $\sigma = \tau$.* ♢

Corollary 3.6 : *If σ and τ are $\sim$-classes with $\bar{\sigma} = \bar{\tau}$, then $\sigma = \tau$.*

Proof : If σ is finite, then $\sigma = \bar{\sigma} = \bar{\tau} = \tau$. Otherwise, apply Lemma 3.5. ♢

Lemma 3.7 : *If σ is a $\sim$-class, then $\bar{\sigma} \setminus \sigma \subseteq M(3+)$.*

Proof : If σ were finite, then $\sigma = \bar{\sigma}$, so we can assume that σ is infinite. Choose any distinct $a, b, c \in \sigma$. Suppose $x \in \bar{\sigma} \setminus \sigma$. By Lemma 2.2, σ is cyclically separating, so, without loss of generality, $\{a, x\}$ separates b from c. In particular, $N(a, x) \geq 2$. But if $\mathrm{val}(x) = 2$, then $a \sim x$, and so we get $x \in \sigma$. We must therefore have $\mathrm{val}(x) \geq 3$. ♢

Another relation on M we shall be considering is defined as follows. Given $x, y \in M$, we write $x \approx y$ if $x \neq y$ and $N(x, y) = \mathrm{val}(x) = \mathrm{val}(y) \geq 3$. (Recall that $N(x, y)$ is the number of components of $M \setminus \{x, y\}$. The relation can thus be interpreted as asserting that this number is as large as possible given the valency of x or of y.)

Lemma 3.8 : *If $x \approx y$ and $x \approx z$, then $y = z$.*

Proof : Let $n = \mathrm{val}(x)$. Suppose that $y \neq z$, so that x, y and z are all distinct. Now there are n components of $M \setminus \{x, y\}$, only one of which contains the point z. The rest are all components of $M \setminus \{x, y, z\}$. Thus there are (at least) $n - 1$ components of $M \setminus \{x, y, z\}$ adjacent to x and y but not to z. Similarly, there are at least $n - 1$ such components adjacent to x and z but not to y. Thus, $n = \mathrm{val}(x) \geq 2(n - 1) > n$ since $n \geq 3$. This contradiction shows that $y = z$. ♢

In other words, some subset of $M(3+)$ is partitioned into pairs of the form $\{x, y\}$ where $x \approx y$. We refer to such a pair as a $\approx$-*pair*. For a general Peano continuum, the relation

$\approx$ would appear unnaturally restrictive. However, in the special case which interests us in Section 5, we shall see that this partition accounts for all of $M(3+)$.

Before going on to discuss these relations further, we introduce a notion that will be used throughout the rest of this paper. For this definition, and the immediately ensuing discussion, we do not require that M has no global cut point.

Definition : An *annulus*, A, consists of an ordered pair, (A^-, A^+), of disjoint closed connected subsets of M, such that if U is a component of $M \setminus (A^- \cup A^+)$, then $\bar{U} \cap A^- \neq \emptyset$ and $\bar{U} \cap A^+ \neq \emptyset$.

We write $R(A) = M \setminus (A^- \cup A^+)$. We write $\mathcal{U}(A)$ for the set of components of $R(A)$. Thus, each element of $\mathcal{U}(A)$ is an open connected subset of M, and $\mathcal{U}(A)$ is finite. We write $N(A) = \operatorname{card} \mathcal{U}(A)$.

The clause about the closure of each component of $R(A)$ meeting both A^- and A^+ is largely for convenience. It is easy to arrange this. Suppose that $B = (B^-, B^+)$ is any ordered pair of disjoint closed connected subsets of M. Let $\mathcal{U}(B)$ be the set of all components of $M \setminus (B^- \cup B^+)$. Let $\mathcal{U}^\pm(B) = \{U \in \mathcal{U}(B) \mid \bar{U} \cap B^\pm \neq \emptyset \text{ and } \bar{U} \cap B^\mp = \emptyset\}$, and let $\mathcal{U}^0(B) = \{U \in \mathcal{U}(B) \mid \bar{U} \cap B^- \neq \emptyset \text{ and } \bar{U} \cap B^+ \neq \emptyset\}$. Thus $\mathcal{U}(B) = \mathcal{U}^0(B) \sqcup \mathcal{U}^-(B) \sqcup \mathcal{U}^+(B)$. Let $A^\pm = B^\pm \cup \bigcup \mathcal{U}^\pm(B)$. Then, it's not hard to see that $A = (A^-, A^+)$ is an annulus with $\mathcal{U}(A) = \mathcal{U}^0(B)$. We shall write $A = B'$ for this construction.

Note that if A is an annulus, then $M \setminus A^-$ and $M \setminus A^+$ are both connected open subsets of M.

Given two annuli A and B, we write $A < B$ to mean that $M = \operatorname{int} A^+ \cup \operatorname{int} B^-$. Thus, if $A < B$, we have $A^- \subseteq \operatorname{int} B^-$ and $B^+ \subseteq \operatorname{int} A^+$. It's easily seen that the relation $<$ is a partial order on the set of all annuli in M.

Given any closed set $K \subseteq M$ and annulus A, we write $K < A$ to mean that $K \subseteq \operatorname{int} A^+$, and $A < K$ to mean that $K \subseteq \operatorname{int} A^+$. For $x \in M$ we define $x < A$ and $A < x$ respectively to mean $\{x\} < A$ and $A < \{x\}$. Note that if A and B are annuli, then $A < B$ is equivalent to $(M \setminus \operatorname{int} A^+) < B$.

We can think of the relation $<$ as describing "nesting" of annuli. There is another relation of interest, namely $\ll$, which can be thought of as describing inclusion of annuli.

Given annuli, A and B, we write $A \ll B$ to mean that $B^- \subseteq \operatorname{int} A^-$ and $B^+ \subseteq \operatorname{int} A^+$, or in other words, $B^- < A < B^+$. Note that $R(A) \subseteq R(B)$. Clearly $\ll$ is also a partial order.

Lemma 3.9 : *Suppose A, B are annuli with $A \ll B$. Then there is a natural surjective map $f : \mathcal{U}(A) \longrightarrow \mathcal{U}(B)$ such that $U \subseteq f(U)$ for all $U \in \mathcal{U}(A)$.*

Proof : We have $R(A) \subseteq R(B)$. If $U \in \mathcal{U}(A)$, then U is a connected subset of $R(A)$ and thus of $R(B)$. It thus lies in some element of $\mathcal{U}(B)$ which we set to be $f(U)$.

To see that f is surjective, suppose $V \in \mathcal{U}(B)$. Now $\bar{V}$ meets both B^- an B^+ and so V meets both A^- and A^+. Thus $V \cap R(A) \neq \emptyset$ (otherwise $V \cap A^+$ and $V \cap A^-$ would partition V into two closed subsets). It follows that V meets some component, U, of $R(A)$.

Thus $U \cup V$ is a connected subset of $R(B)$ and so $U \cup V = V$. In other words $U \subseteq V$ so $f(U) = V$. $\diamondsuit$

Corollary 3.10 : *If $A \ll B$, then $N(A) \geq N(B)$.* $\diamondsuit$

So far, none of our discussion of annuli has assumed that there is no global cut point. For the rest of this section, however, we shall reinstate this hypothesis.

Note that, in this case, any pair, (x, y), of distinct points of M is an annulus. (To be more precise, we should write $(\{x\}, \{y\})$, though we shall not bother about this distinction). In this case, the notation $\mathcal{U}(x, y)$ and $N(x, y)$ agrees with that previously defined.

Lemma 3.11 : *Suppose $x \in M(n)$, where $n < \infty$, and $F \subseteq M \setminus \{x\}$ is compact. Then, there is an annulus, (K, x), (or more properly $(K, \{x\})$) with $F < (K, x)$ and $N(K, x) = n$.*

Proof : By definition, $M \setminus \{x\}$ has n ends. So, given any compact set $F \subseteq M \setminus \{x\}$, we can find another compact set, $F' \subseteq M$, with $F \subseteq \mathrm{int}\, F'$ and such that $M \setminus (F' \cup \{x\})$ has precisely n unbounded components. Let K be the union of F' together with all the bounded components of $M \setminus F'$ (i.e. $(K, x) = (F', x)'$). $\diamondsuit$

By a similar argument, we obtain:

Lemma 3.12 : *If $x \in M(\infty)$, and $F \subseteq M \setminus \{x\}$ is compact, then there is some annulus (K, x) with $F < (K, x)$, and $N(K, x)$ arbitrarily large.* $\diamondsuit$

Lemma 3.13 : *Suppose A is an annulus. Suppose (K_i, x_i) for $i = 1, \ldots, p$ are annuli with $A \ll (K_i, x_i)$ and $N(K_i, x_i) \geq 3$ for all i. If the points x_i are all distinct, then $p < n^n$, where $n = N(A)$.*

Proof : For each i, we get a surjective map $f_i : \mathcal{U}(A) \longrightarrow \mathcal{U}(K_i, x_i)$ as described by Lemma 3.7. Such a map gives rise to a partition of $\mathcal{U}(A)$ into disjoint non-empty subsets — given by the preimages of elements of $\mathcal{U}(K_i, x_i)$.

Suppose that for a given i, the sets $U_1, U_2, U_3 \in \mathcal{U}(A)$ lie in distinct elements of this partition, so that $f(U_1)$, $f(U_2)$ and $f(U_3)$ are all distinct. Choose any $y_1 \in U_1$, $y_2 \in U_2$ and $y_3 \in U_3$. Thus, x_i separates each distinct pair from $\{y_1, y_2, y_3\}$ in $M \setminus F_i$ and hence in $M \setminus A^-$. (Recall that $M \setminus F_i$ and $M \setminus A^-$ are connected.) In other words, $x_i = \mathrm{med}(y_1, y_2, y_3)$ in the pretree structure on $M \setminus A^-$ as described in Section 2.

Suppose that for some $i \neq j$, the maps f_i and f_j give rise to same partition of $\mathcal{U}(A)$. Then, we could choose the same sets U_1, U_2, U_3 and points y_1, y_2, y_3 for each. Thus, in $M \setminus A^-$, we get that $x_i = \mathrm{med}(y_1, y_2, y_3) = x_j$, contradicting the hypothesis that the x_i are all distinct.

It follows that each of the partitions are different. Since there are less than n^n partitions of an n-set, the result follows. $\diamondsuit$

Of course, one could do a lot better than n^n. We only really care that this number is finite. A similar argument yields:

Lemma 3.14 : *Suppose that A is an annulus. Suppose (K_i, x_i) for $i \in 1, \ldots, p$ are annuli with $A \ll (K_i, x_i)$ and $N(K_i, x_i) \geq 2$ for all i. If $x_i \not\sim x_j$ for all i, then $p \leq n^n$.*

Proof : As with Lemma 3.13. This time, if $p > n^n$, we obtain $y_1, y_2 \in M$ and $i \neq j$ such that both x_i and x_j separate y_1 from y_2 in $M \setminus A^-$. It follows that $x_i \sim x_j$. ♢

We can now draw some conclusions from these observations. Let Π be the space of unordered pairs of distinct points in M, as described at the end of Section 2. Let $\Pi(3+) \subseteq \Pi$ be the subset of $\approx$-pairs, i.e. pairs $\{x, y\}$ such that $x \approx y$.

Lemma 3.15 : $\Pi(3+)$ *is a discrete subset of* Π.

Proof : Suppose $\{x, y\} \in \Pi$. Let U, V be disjoint connected open neighbourhoods of x and y. Let $A = (\bar{U}, \bar{V})'$ be the corresponding annulus. Now if $(z, w) \in (U \times V) \cap \Pi(3+)$, then we can think of (z, w) as an annulus with $A \ll (z, w)$ and $N(z, w) \geq 3$. By Lemma 3.13 and Lemma 3.8, there can only be many such pairs (z, w). In other words, $(U \times V) \cap \Pi(3+)$ is finite. Now, any compact subset, $K \subseteq \Pi$ can be covered by finitely many sets of the form $U \times V$, and so $K \cap \Pi(3+)$ is finite. In other words, $\Pi(3+)$ is discrete. ♢

Let $\Pi(2) \subseteq \Pi$ be the set of pairs, $\{x, y\}$ such that $x, y \in M(2)$ and $x \sim y$ and $x \neq y$. We can define the equivalence relation $\sim$ on $\Pi(2)$ by $\{x, y\} \sim \{x', y'\}$ if $x \sim x' \sim y \sim y'$.

As with Lemma 3.15, using Lemma 3.14 in place of Lemma 3.13, we obtain:

Lemma 3.16 : *Any compact subset of* Π *can meet only finitely many* $\sim$*-classes in* $\Pi(2)$. ♢

Put another way, it's impossible for an infinite set of distinct $\sim$-classes to accumulate at two distinct points of M. Clearly, this also applies to closures of $\sim$-classes.

Given $x, y \in M$, and $\zeta \in \Pi$, we write $x\zeta y$ to mean that x and y lie in different components of $M \setminus \zeta$. We say that a subset P of Π *accumulates* at some point $a \in M$ if for every neighbourhood, U, of a, there are infinitely many $\zeta \in P$ with $\zeta \subseteq U$.

Lemma 3.17 : *Suppose $a, b \in M$ are distinct. Suppose that $P \subseteq \Pi(2) \cup \Pi(3+)$ is such that $a\zeta b$ for a all $\zeta \in P$, and that if $\zeta, \eta \in P \cap \Pi(2)$ with $\zeta \sim \eta$, then $\zeta = \eta$. Then either P is finite, or it accumulates at a or at b (or both).*

Proof : By Lemmas 3.15 and 3.16, if the conclusion fails, then P must accumulate at some point $c \in M \setminus \{a, b\}$. Now, $M \setminus \{c\}$ is connected, locally connected and compact and so admits an exhaustion by compact connected sets. Thus, there is an open neighbourhood $U \ni c$ such that $M \setminus U$ is connected and contains both a and b. Now if $\zeta \in P$, with $\zeta \subseteq U$, we would have $a\zeta b$. But $a, b \in M \setminus U \subseteq M \setminus \zeta$, giving a contradiction. ♢

The next objective is to construct a discrete pretree, T, based on the structure of local cut points. In the case where there is a convergence group action (Section 5), this construction will give rise to a simplicial tree Σ which describes the JSJ splitting.

Let $T_1 = \Pi(3+)$ be the set of $\approx$-pairs in $M(3+)$, and let $T_2 = M(2)/\!\sim$ be the set of $\sim$-classes in $M(2)$. Let $T = T_1 \sqcup T_2$. Note that if $\zeta, \eta \in T$ are distinct, then $\zeta \cap \eta = \emptyset$.

Lemma 3.18 : *Suppose $\zeta, \eta \in T$ are distinct. Suppose that $a, b \in \zeta$ are distinct, and that $U \in \mathcal{U}(a,b)$. If $\eta \cap U \neq \emptyset$, then $\eta \subseteq U$.*

Proof : If the result fails, then there are points $x, y \in \eta$ separated in M by $\{a, b\}$. It follows that $\eta \in T_2$, so that $x, y \in M(2)$. (For if $\eta = \{x, y\} \in T_1$, then there would be some $V \in \mathcal{U}(x,y)$ with $a, b \notin U$, so $\bar{V} = V \cup \{x, y\}$ would connect x to y in $M \setminus \{a, b\}$.) If $\zeta \in T_2$, then $a, b, x, y \in M(2)$, and so, by Lemma 3.3, we get that $a \sim b \sim x \sim y$. This gives the contradiction that $\zeta = \eta$. We are thus reduced to the case where $\zeta = \{a, b\} \in T_1$. Now, $N(x, y) = 2$, so we can let $\mathcal{U}(x,y) = \{W_1, W_2\}$. Since $\{a, b\}$ separates x and y, the points a and b cannot lie in the same component of $M \setminus \{x, y\}$. Thus, we can assume that $a \in W_1$ and $b \in W_2$. But $N(a, b) \geq 3$, so there is some component, $C \in \mathcal{U}(a,b)$, with $x, y \notin C$. Thus, $\bar{C} = C \cup \{a, b\} \subseteq M \setminus \{x, y\}$. But now $\bar{C} \cup W_1 \cup W_2$ is a connected subset of $M \setminus \{x, y\}$. We thus arrive at a contradiction. ♢

Given $\eta \in T$ and $x, y \in M$, we write $x\eta y$ if there are distinct points $a, b \in \eta$ such that x and y are separated in M by $\{a, b\}$. Given $\zeta, \eta, \theta \in T$, we write $\zeta\eta\theta$ if $(\exists x \in \zeta)(\exists y \in \theta)(x\eta y)$. In view of Lemma 3.18, this is the same as saying that there are distinct $a, b \in \eta$ and distinct $U, V \in \mathcal{U}(a,b)$ such that $\zeta \subseteq U$ and $\theta \subseteq V$.

Lemma 3.19 : *With the ternary relation thus defined, T is a pretree.*

Proof : Properties (T0) and (T1) are immediate. Property (T3) is also elementary. Suppose $\zeta\eta\theta$ and $\xi \neq \eta$. We can find $a, b \in \eta$ and $U, V \in \mathcal{U}(a,b)$ such that $\zeta \subseteq U$ and $\theta \subseteq V$. Choose any $x \in \xi$. If $x \notin U$ then $\zeta\eta\theta$, and if $x \notin V$ then $\theta\eta\xi$.

To deduce (T2), suppose, for contradiction, that $\zeta\eta\theta$ and $\zeta\theta\eta$. We can find $a, b \in \eta$ such that ζ and θ are subsets of distinct elements of $\mathcal{U}(a,b)$, and we can find $c, d \in \theta$ such that ζ and η are subsets of distinct elements of $\mathcal{U}(c,d)$. Let U be the element of $\mathcal{U}(a,b)$ containing ζ, and let V be the element of $\mathcal{U}(c,d)$ containing η. In particular, $a, b \in V$. Now, $\bar{U} = U \cup \{a, b\}$ is connected. Also $c, d \notin U$, so $\bar{U} \subseteq M \setminus \{c, d\}$. We see that $\bar{U} \subseteq V$. In particular, $\zeta \subseteq V$, contradicting the fact that $\{c, d\}$ separates ζ from $\eta \supseteq \{a, b\}$. ♢

Proposition 3.20 : *T is a discrete pretree.*

Proof : Given Lemma 3.19, it remains to verify that if $\zeta, \eta \in T$ are distinct, then the pretree interval $[\zeta, \eta]$ is finite.

Choose distinct $a, b \in \zeta$ and $c, d \in \eta$. Let U_a, U_b, U_c and U_d be disjoint open neighbourhoods of a, b, c and d respectively. Suppose $\zeta\theta\eta$, so that there is a pair of distinct points $x, y \in \theta$ such that $\{x, y\}$ separates ζ from η. In particular, $\{x, y\}$ separates both a and b from both c and d. Now, clearly $\{x, y\}$ cannot be a subset of both U_a and U_b, nor a subset of both U_c and U_d. Applying Lemma 3.17, we see that there are only finitely many possibilities for $\{x, y\}$ up the the relation $\sim$. Thus, there are only finitely many possibilities for θ. ♢

¿From the discussion in Section 2, we see that T can be embedded as a subset of the set of vertices in a simplicial tree. (Note that any singleton $\sim$-class will be a terminal vertex in this tree. In the next section we rule out the possibility of such singletons in the case of discrete group actions — Corollary 5.15)

In summary, we have shown that the two kinds of subsets of M, namely $\sim$-classes in $M(2)$ and $\approx$-pairs in $M(3+)$, together can be embedded in a natural way in a simplicial tree. In the case that interests us (Section 5), we shall see that these subsets account for all local cut points in M.

4. Convergence groups.

In this section, we describe some general properties of convergence groups which we shall be using in the next section. Some general references are [GeM1,T1,T2,Fr,Bo4].

The notion of a convergence group was defined by Gehring and Martin [GeM1] in the context of groups acting on topological spheres. Most of the basic theory generalises without problems to compact hausdorff spaces (or at least to compact metrisable spaces) — see, for example [T2]. Here, we shall be principally interested in what we shall call "uniform convergence groups" — see, for example, [Bo4].

It is shown in [Bo6] that, in fact, uniform convergence groups are precisely hyperbolic groups acting on their boundaries. We shall not refer to that result here — the extra geometric information does not seem to help much in this context.

Let M be a compact hausdorff space. Let $\Phi(M)$ be the space of distinct (ordered) triples in M, i.e. $M \times M \times M$ minus the large diagonal. Note that $\Phi(M)$ is locally compact hausdorff. Suppose that a group, Γ, acts by homeomorphism on M. We get an induced action on $\Phi(M)$. The group, Γ, is said to be a *convergence group* on M if the action on $\Phi(M)$ is properly discontinuous.

If M is metrisable, this is equivalent to the following hypothesis, which was the original definition in [GeM1], and is the most frequently used formulation of convergence group. Suppose that $(\gamma_n)_{n\in\mathbf{N}}$ is a sequence of distinct elements of Γ. Then, there is a subsequence $(\gamma_i)_i$ and points $\lambda, \mu \in M$ such that the maps $\gamma_i|(M \setminus \{\lambda\})$ converge locally uniformly to μ. (Note that it follows that $\gamma_i^{-1}|(M \setminus \{\mu\})$ converge locally uniformly to λ.) In this hypothesis, we allow for the possibility that $\lambda = \mu$. The equivalence of these definition for actions on spheres is proven in [GeM2]. Their argument would seem to generalise unchanged to Peano continua. A general proof is given in [Bo4]. (In fact we don't need to assume that M is metrisable, provided we reformulate the Gehring-Martin definition in terms of nets rather than sequences.)

One can classify the elements of Γ into *elliptic* (finite order), *parabolic* (infinite order with a single fixed point) and *loxodromic* (infinite order with two fixed points). If $\gamma \in \Gamma$ is loxodromic, we can write its fixed point set, $\mathrm{fix}(\gamma) = \{\mathrm{fix}^-(\gamma), \mathrm{fix}^+(\gamma)\}$, where $\mathrm{fix}^-(\gamma)$ and $\mathrm{fix}^+(\gamma)$ are, respectively, the repelling and attracting fixed points. In this case, the cyclic group, $\langle\gamma\rangle$ acts properly discontinuously and cocompactly on $M \setminus \mathrm{fix}(\gamma)$. The discussion of loxodromics in Section 1 applies equally well to this more general situation.

Note that if $K \subseteq M$ is closed with $\mathrm{card}(K) \geq 2$ and fixed setwise by a loxodromic γ, then $\mathrm{fix}(\gamma) \subseteq K$. We also have [T2]:

Lemma 4.1 : *Suppose $K \subseteq M$ is closed and $\gamma \in \Gamma$ is such that $\gamma K \subseteq \operatorname{int} K$, then γ is loxodromic, with $\operatorname{fix}^+(\gamma) \in \operatorname{int} K$ and $\operatorname{fix}^-(\gamma) \in M \setminus K$.* $\diamondsuit$

We say that a subgroup of Γ is *elementary* if it is finite or two-ended.

It is common to allow, in the definition of an elementary group, any group which fixes a point of M. Such a group must either be a torsion group, or else contain a parabolic element. However, neither of these possibilities can arise in the case of interest to us. (Parabolic elements are ruled out by Lemma 4.4. It is also well known that torsion subgroups cannot occur in hyperbolic groups — see, for example, [GhH]. In any case, all groups which we claim to be non-elementary will be seen to contain a loxodromic element, and so cannot be torsion groups.)

The following is a trivial observation:

Lemma 4.2 : *Suppose $G \leq \Gamma$ is non-elementary, and $K \subseteq M$ is a non-empty closed G-invariant subset. Then G acts as a convergence group on K.* $\diamondsuit$

We shall be mainly dealing with a restricted class of convergence groups. For this, it's convenient to assume that M is a perfect compact hausdorff space. (Recall that "perfect" means "having no isolated points.)

Definition : We say that Γ acts as a *uniform convergence group* on M if it acts properly discontinuous and cocompactly on the space of distinct triples, $\Phi(M)$.

It follows easily that there is compact subset, $\Phi_0 \subseteq \Phi(M)$, such that $\Phi(M) = \bigcup \Gamma \Phi_0$.

Direct dynamical proofs of the following three lemmas can be found in [Bo4].

Lemma 4.3 : *The action of a uniform convergence group is minimal (i.e. there is no proper nonempty closed invariant subset).* $\diamondsuit$

Lemma 4.4 : *A uniform convergence group has no parabolic elements.* $\diamondsuit$

Lemma 4.5 : *If Γ acts as a uniform convergence group on a Peano continuum, then Γ is finitely generated and one-ended.* $\diamondsuit$

In particular Γ does not split over any finite subgroup in this case.

It's well known that a hyperbolic group acts as a uniform convergence group on its boundary (a proof is given in [Bo4]). The converse is given in [Bo6]. For a non-elementary hyperbolic group Γ, the results stated above can be deduced by direct geometric arguments. Note that $\partial\Gamma$ is perfect and metrisable. The fact that Γ acts as a convergence group on $\partial\Gamma$ by the Gehring-Martin definition is shown directly in [T2] and [F]. There are several arguments to show that Γ has no parabolics (see for example [GhH]). Also, it's easy to see that $\partial\Gamma$ is connected if and only if Γ is one-ended (again, see for example [GhH]).

Returning to our set-up of a group Γ, acting on a perfect compact hausdorff space, M, we formulate a notion of quasiconvexity for subgroups in dynamical terms.

Given a closed subset $\Lambda \subseteq M$, we define $\Phi_M(\Lambda) = \{(x, y, z) \in \Phi(M) \mid x, y \in \Lambda\}$. Thus, $\Phi_M(\Lambda)$ is a closed subset of $\Phi(M)$.

Elementary (i.e. finite or two-ended) subgroups of Γ are always deemed to be quasiconvex. Otherwise:

Definition : A non-elementary subgroup, $G \leq \Gamma$, is *quasiconvex* if there is a non-empty closed perfect G-invariant set $\Lambda \subseteq M$ such that $\Phi_M(\Lambda)/G$ is compact.

(In fact, the assumption that Λ is perfect is redundant — it is a consequence of the definition given that M is perfect. However, we shall only apply it in cases where we already know Λ to be perfect.)

For most purposes in this section, we only explicitly use an apparently weaker property, namely that $\Phi(\Lambda)/G$ is compact, where $\Phi(\Lambda) = \{(x, y, z) \in \Phi(M) \mid x, y, z \in \Lambda\}$ is the space of distinct ordered triples in Λ. In fact, this turns out to be equivalent to quasiconvexity (see [Bo4]).

Note that if G is quasiconvex, then G acts as a uniform convergence group on Λ. In particular, from Lemma 4.3, we see:

Lemma 4.6 : *The set Λ is the unique minimal non-empty closed G-invariant subset of M.* $\diamondsuit$

In particular, Λ is uniquely determined by G. We refer to it as the *limit set* of G, and write $\Lambda = \Lambda G$. In the case where G is two-ended, we set ΛG to be the fixed point set of any loxodromic in G. (This definition agrees with the standard one for convergence groups.)

Note that G has finite index in the setwise stabiliser of Λ (since the setwise stabiliser acts properly discontinuously on $\Phi(\Lambda)$ and $\Phi(\Lambda)/G$ is compact).

Before continuing, we make few remarks about two-ended subgroups which will be relevant to the subsequent discussion.

Note that if G is any two-ended group, then the end-preserving subgroup is a normal subgroup of index at most 2 in G. It can be defined purely group theoretically in terms of the action of G on its Cayley graph. Also, any infinite order element $g \in G$ determines an ordering on the pair of ends of G, according to which end any given forward orbit tends. We can refer to this as the "direction of translation" of g. It's well known that a 2-ended groups acts properly discontinuously and isometrically on the real line (see for example [DiD]). It's not hard to see that the end-preserving subgroup of G and the direction of translations of elements are respected by such a representation. Note that a subgroup, G, of a convergence group Γ, is two-ended if and only if it preserves setwise an unordered pair of distinct points, $\{x, y\} \subseteq M$. In this case, the pointwise stabiliser of x and of y is precisely the end-preserving subgroup of G. Moreover, if $\gamma \in G$ is loxodromic, then $y = \mathrm{fix}^+(\gamma)$ if and only of γ translates in the direction of y in the intrinsic group-theoretical sense.

As mentioned in the introduction, we have a special interest in convergence groups acting on circles. Such actions have been completely analysed by Tukia [T1] and the subsequent papers of Gabai [Ga] and Casson and Jungreis [CJ]. It turns out that, up to topological conjugacy, such a group can be represented by a group acting isometrically and properly discontinuously on the hyperbolic plane, $\mathbf{H}^2$, where we are considering the induced action on the ideal circle, $\partial\mathbf{H}^2$. We state explicitly two special cases of this. Firstly:

Theorem 4.7 : [T1,Ga,CJ] *Suppose that Γ is a acts as a uniform convergence group on circle, S^1. Then, there is a properly discontinuous cocompact isometric action of Γ on $\mathbf{H}^2$, and a Γ-equivariant homeomorphism of S^1 to $\partial\mathbf{H}^2$.* ♢

We refer to such a group as a cocompact "fuchsian group". Group theoretically, this is the same as saying that Γ is a virtual surface group. There is a slight distinction between our formulation and that given in the references cited, in that we are not assuming the action to be effective. However the distinction is essentially trivial, since the kernel of such an action can be simply characterised as the unique maximal finite normal subgroup. (Note that the fact that Γ is finitely generated is easy to see in this case, since its quotient by the kernel of the action is the orbifold fundamental group of the compact 3-orbifold $\Phi(S^1)/\Gamma$.)

The following result follows from Tukia's original paper [T1]:

Theorem 4.8 : *Suppose that Δ is a cyclically ordered cantor set, and that Γ acts as a minimal discrete convergence group (without parabolics) on Δ, preserving the cyclic order. Then, there is a properly discontinuous action of Γ on $\mathbf{H}^2$ (without parabolics), and a cyclic-order-preserving Γ-equivariant homeomorphism from Δ onto the limit set of the Γ-action on $\partial\mathbf{H}^2$.*

Proof : In the terminology of [T1], the action of Γ is of the "second kind". Thus, in view of Theorem 6B(f) of that paper, it suffices to show that that the action of Γ on Δ extends to a convergence action on the circle, S^1.

As in Section 3, let $J(\Delta)$ the set of jumps in Δ, which in this case are all disjoint. Suppose $\theta \in J(\Delta)$. Note that if $\gamma \in \Gamma$ were loxodromic with $\mathrm{fix}(\gamma) \cap \theta \neq \emptyset$, then in fact $\mathrm{fix}(\gamma) = \theta$. (Otherwise there would be $\langle\gamma\rangle$-orbits of points converging on each element of θ from either side, showing that θ could not be a jump.) We also see that the setwise stabiliser of a jump is either finite or two-ended.

Now, we can certainly find a cyclic order-preserving embedding of Δ in S^1. The jumps of Δ are then in bijective correspondence with the complementary open intervals of $S^1 \setminus \Delta$.

Let $\theta = \{x, y\} \in J(\Delta)$, and let G be the setwise stabiliser of θ. Let $I \subseteq S^1$ be the complementary open interval with $\partial I = \{x, y\}$. We want to define an action of G on I. Suppose, first, that G is two-ended. Now as discussed earlier in this section, G admits a properly discontinuous isometric action on the real line. Now there is a natural correspondence between the ends of the real line and the ends of G, and thus, in turn, with the pair $\{x, y\}$, which are the ends of the interval I. We now take any homeomorphism of I with the real line respecting this identification of ends. We thus conjugate the action of G under this homeomorphism, to give and action if G on I. The case where G is finite is simpler. In this case, we just take any homeomorphism of I with the real line. We define the action of $\Gamma \in G$ on the real line by setting g to be the identity if g fixes x and y, and to be reflection in the origin if g swaps x and y.

We now perform this construction for one jump in each Γ-orbit. To extend over a given Γ-orbit of jumps, we conjugate by appropriate elements of Γ.

This extends the action over S^1. It's now a simple exercise to verify that this is a convergence action. $\diamondsuit$

(In fact, the construction of the extension to S^1 probably isn't really necessary — one should be able to adapt the arguments of [T1] directly to this case though I've not worked through the details.)

Note that we have not assumed that Γ is finitely generated. In this special case, we obtain:

Proposition 4.9 : *With the hypotheses of Theorem 4.8, if, in addition, Γ is finitely generated without parabolics, then Γ is conjugate to a bounded fuchsian group.* $\diamondsuit$

Here "bounded" means convex cocompact but not cocompact, as described in the introduction. Note that the peripheral subgroups are precisely the stabilisers of jumps. In particular, each such stabiliser is two-ended, and there are finitely many conjugacy classes of such.

5. Convergence actions on Peano continua.

In this section we continue in a similar vein to Section 3, but now introducing convergence actions into the picture.

Let M be a metrisable Peano continuum, and let $\Phi = \Phi(M)$ be the space of distinct triples in M. Let Γ be a group acting as a uniform convergence group on M, i.e. acting properly discontinuously and cocompactly on Φ. (One can probably avoid the metrisability assumption, but this would raise technical complications we don't really want to be bothered with here.) We shall shortly see (Proposition 5.4) that M has no global cut point, so from that point on, we can bring the results of Section 3 into play.

By Lemma 4.4, we know that Γ is finitely generated and one-ended. We shall not use the finite generation result explicitly here, only the fact that Γ does not split over any finite subgroup. The fact that Γ is countable is used for Lemma 5.31.

Lemma 5.1 : *There are finite collections, $(U_i)_{i=1}^p$, $(V_i)_{i=1}^p$ and $(W_i)_{i=1}^p$ of open connected sets $U_i, V_i, W_i \subseteq M$ such that $\bar{U}_i \cap \bar{V}_i = \bar{V}_i \cap \bar{W}_i = \bar{W}_i \cap \bar{U}_i = \emptyset$ for all $i \in \{1, \ldots, p\}$, and such that if $x, y, z \in M$ are all distinct, then is some $\gamma \in \Gamma$ and $i \in \{1, \ldots, p\}$ such that $\gamma x \in U_i$, $\gamma y \in V_i$ and $\gamma z \in W_i$.*

Proof : Let $\Phi_0 \subseteq \Phi$ be compact with $\Phi = \bigcup \Gamma \Phi_0$. Given any $(x, y, z) \in \Phi_0$, we can certainly find open connected sets $U \ni x$, $V \ni y$ and $W \ni z$ whose closures are all disjoint. We now cover Φ_0 with finitely many sets of the form $U \times V \times W$. $\diamondsuit$

Lemma 5.2 : *There are finite collections $(U_i)_{i=1}^p$ and $(O_i)_{i=1}^p$ of open connected subsets of M such that $\bar{U}_i \cap \bar{O}_i = \emptyset$ for all i, and such that if $K \subseteq M$ is closed and $x \in M \setminus K$, then there is some $\gamma \in \Gamma$ and $i \in \{1, \ldots, p\}$ such that $\gamma x \in U_i$ and $\gamma K \subseteq O_i$.*

Proof : Let U_i, V_i, W_i be as described by Lemma 5.1. Let O_i be an open connected set such that $\bar{V}_i \subseteq O_i$ and $\bar{O}_i \cap \bar{U}_i = \emptyset$.

Choose any $y \in M \setminus \{x\}$ and choose a sequence of points $z_n \in M \setminus \{x, y\}$ with $z_n \to x$. Now, applying Lemma 5.1, we can find a sequence $(\gamma_n)_{n \in \mathbf{N}}$ of elements of Γ such that $\gamma_n x \in U_{i(n)}$, $\gamma_n y \in V_{i(n)}$ and $\gamma_n z_n \in W_{i(n)}$, for some $i(n) \in \{1, \ldots, p\}$. Passing to a subsequence, we can suppose that $i(n)$ is constant, equal to k say, and that the γ_n are all distinct. Passing to a further subsequence, we can suppose that $\gamma_n x \to a \in \bar{U}$, $\gamma_n y \to b \in \bar{V}$ and $\gamma_n z_n \to c \in \bar{W}$. Note that a, b, c are all distinct. Suppose that for all $n \in \mathbf{N}$, $\gamma_n K \not\subseteq O_k$. Then we can find $w_n \in K$ such $\gamma_n w_n \in M \setminus O_k$.

We now apply the convergence property to $(\gamma_n)_{n \in \mathbf{N}}$. Passing to a subsequence, we can find points $\lambda, \mu \in M$ such that $\gamma_n | (M \setminus \{\lambda\})$ converges uniformly to μ. Now since $\gamma_n x \to a$ and $\gamma_n y \to b$, we see that μ must equal a or b.

Suppose that $\mu = a$. Since $\gamma_n y \not\to a$, we must have $\lambda = y$. Since $z_n \to x$, the points z_n remain in a compact subset of $M \setminus \{y\}$. Thus, $\gamma_n z_n \to a$ contradicting the fact that $\gamma_n z_n \to c \neq a$.

Suppose that $\mu = b$. Since $\gamma_n x \not\to b$, we must have $\lambda = x$. Now, $w_n \in K \subseteq M \setminus \{x\}$, and K is compact. Thus, $\gamma_n w_n \to b$. Now $b \in \bar{U}_k$ and $\gamma_n w_n \in M \setminus O_k$. But $\bar{U}_k$ and $M \setminus O_k$ are disjoint closed sets, so we again get a contradiction.

In conclusion, we deduce that there must be some $n \in \mathbf{N}$ such that $\gamma_n K \subseteq O_k$. Also $\gamma_n x \in U_k$ as required. ◊

In the last section, we described a process associating an annulus $(B^-, B^+)'$ to any pair of disjoint closed subsets, B^- and B^+, of M. Given $i \in \{1, \ldots, p\}$, set $A_i = (\bar{O}_i, \bar{U}_i)'$. Thus, $O_i \subseteq A_i^-$ and $U_i \subseteq A_i^+$. This gives a finite collection of annuli, $(A_i)_{i=1}^p$. The basic property of this collection may be expressed as follows:

Lemma 5.3 : *Suppose $K \subseteq M$ is closed, and $x \in M \setminus K$. Then, there is some $\gamma \subset \Gamma$ and $i \in \{1, \ldots, p\}$ such that $K < \gamma A_i < x$.*

Proof : This follows immediately from Lemma 5.2 (with γ replacing γ^{-1}). ◊

We now get one of the principal results:

Proposition 5.4 : *M has no global cut point.*

Proof : Suppose $x \in M$ is a global cut point. Choose any $y, z \in M$ separated by x. Applying Lemma 5.3, with $K = \{y, z\}$, we get an annulus $A = \gamma A_i$, with $K < A < x$. But now, $y, z \in A^-$ which is connected, and $x \notin A^-$, contradicting the fact that x separates y from z. ◊

We are now in the situation described in Section 3. We thus have the definition of valency of a point, the sets $M(n)$ and $M(n+)$, and the relations, $\sim$ and $\approx$ etc. We immediately have the following result.

Let $N = \max\{N(A_i) \mid 1 \leq i \leq p\}$.

Proposition 5.5 : *If $x \in M$, then $\mathrm{val}(x) \leq N$.*

Proof : Suppose $\mathrm{val}(x) > N$. By Lemma 3.11 and Lemma 3.12, there is an annulus (K, x) with $N(K, x) > N$. But now, Lemma 5.3 gives us some $i \in \{1, \ldots, p\}$ and some $\gamma \in \Gamma$ such that $\gamma A_i \ll (K, x)$. By Corollary 3.10, we get $N(A_i) = N(\gamma A_i) \geq N(K, x) > N$, contradicting the definition of N. $\diamondsuit$

In particular, we see that $M(\infty) = \emptyset$.

Now, exactly as in Section 1, if $G \leq \Gamma$ is a two-ended subgroup, we define $e(G)$ to be the number of components of $M \setminus \Lambda G$. Similarly, if $\gamma \in \Gamma$ is loxodromic, set $e(\gamma) = e(\langle \gamma \rangle)$. Thus, $e(\gamma) = N(x, y)$, where $\{x, y\} = \mathrm{fix}(\gamma) = \Lambda(G)$.

Lemma 5.6 : *If $\gamma \in \Gamma$ is loxodromic with $\mathrm{fix}(\Lambda) = \{x, y\}$, then $\mathrm{val}(x) = \mathrm{val}(y) = N(x, y) = e(\gamma)$.*

Proof : As in Section 1, after raising γ to some power, if necessary, we can suppose that γ fixes each component $U \in \mathcal{U}(x, y)$. Now, $U/\langle \gamma \rangle$ is compact. (It is a connected component of $(M \setminus \mathrm{fix}(\gamma))/\langle \gamma \rangle$). Since $\langle \gamma \rangle$ is two-ended, it follows that U has precisely two ends. Since $\bar{U} = U \cup \{x, y\}$ is compact, we see that these ends are compactified by the points x and y. Thus, $M \setminus \{x\} = \bigcup \{U \cup \{y\} \mid U \in \mathcal{U}(x, y)\}$ has precisely $N(x, y) = \mathrm{card}\, \mathcal{U}(x, y)$ ends. Thus $\mathrm{val}(x) = N(x, y) = e(\gamma)$. Similarly for y. $\diamondsuit$

Corollary 5.7 : *If $e(\gamma) \geq 3$ and $\mathrm{fix}(\gamma) = \{x, y\}$, then $x \approx y$.* $\diamondsuit$

We shall see (Proposition 5.13), that there is a converse to Corollary 5.7.

We are assuming that Γ acts cocompactly on the space of distinct triples, so it certainly acts cocompactly (though not properly discontinuously) on the space of distinct unordered pairs, Π. (In other words, there is a compact set $\Pi_0 \subseteq \Pi$ such that $\Pi = \bigcup \Gamma \Pi_0$.)

Recall that $\Pi(3+) \subseteq \Pi$ is the set of $\approx$-pairs. Clearly this is Γ-invariant. Lemma 3.15 tells us that this is discrete, and so:

Lemma 5.8 : *$\Pi(3+)/\Gamma$ is finite.* $\diamondsuit$

Putting this together with Corollary 5.7, we obtain:

Proposition 5.9 : *There are finitely many conjugacy classes of maximal two-ended subgroups $G \leq \Gamma$ for which $e(G) \geq 3$.* $\diamondsuit$

The same argument, using Lemma 3.16 in place of Lemma 3.15, will enable us to establish that there are only finitely many Γ-orbits of $\sim$-classes in $M(2)$. For this to work, however, we will first need to establish that there are no singleton $\sim$-classes. This will form part of a more substantial analysis later on. We first need to make a few more general observations.

¿From Lemma 4.1, we note:

Lemma 5.10 : *Suppose that A is an annulus, and $\gamma \in \Gamma$ with $A < \gamma A$. Then, γ is loxodromic, and $\mathrm{fix}^-(\gamma) < A$ and $A < \mathrm{fix}^+(\gamma)$.* ♢

Extending Lemma 5.3, we get the following nesting property of annuli:

Lemma 5.11 : *Suppose $K \subseteq M$ is closed, and $x \in M \setminus K$. Then, there is some $i \in \{1, \ldots, p\}$, and a sequence of elements $(\gamma_n)_{n \in \mathbf{N}}$ such that*

$$K < \gamma_0 A_i < \gamma_1 A_i < \gamma_2 A_i < \cdots < x.$$

Proof : Lemma 5.3 gives us $\gamma_0 \in \Gamma$ and $i(0) \in \{1, \ldots, p\}$ such that $K < \gamma_0 A_{i(0)} < x$. We now apply Lemma 5.3 again, with $M \setminus \gamma_0 \,\mathrm{int}\, A^+_{i(0)}$ replacing K. This gives γ_1 and $i(1)$ such that $M \setminus \gamma_0 \,\mathrm{int}\, A^+_{i(0)} < \gamma_1 A_{i(1)} < x$. But the first relation is equivalent to $\gamma_0 A_{i(0)} < \gamma_1 A_{i(1)}$. We now continue inductively. On passing to a subsequence, we can suppose that $i(n)$ is constant. ♢

Lemma 5.12 : *Every point of $M(3+)$ is the fixed point of some loxodromic in Γ.*

Proof : Suppose $x \in M(3+)$, so $3 \le \mathrm{val}(x) < \infty$. Lemma 3.11 gives an annulus (K, x) with $N(K, x) \ge 3$. Lemma 5.11 now gives us some $i \in \{1, \ldots, p\}$ and a sequence, $(\gamma_n)_{n \in \mathbf{N}}$ of elements of Γ with $K < \gamma_0 A_i < \gamma_1 A_i < \cdots < x$. For each n, we have $\gamma_n A_i \ll (K, x)$, and so $A_i \ll (\gamma_n^{-1} K, \gamma_n^{-1} x)$. Applying Lemma 3.13, we can find $m < n \in \mathbf{N}$ with $\gamma_n^{-1} x = \gamma_m^{-1} x$. Let $\gamma = \gamma_n \gamma_m^{-1}$ so that $\gamma x = x$. Now $\gamma_m A_i < \gamma_n A_i = \gamma(\gamma_m A_i)$, so by Lemma 5.10, γ is loxodromic. ♢

Proposition 5.13 : *Suppose $x, y \in M$. Then $x \approx y$ if and only if there is a loxodromic $\gamma \in \Gamma$ with $e(\gamma) \ge 3$ and $\mathrm{fix}(\gamma) = \{x, y\}$. Moreover $M(3+)$ is a disjoint union of such $\approx$-pairs.*

Proof : By Corollary 5.7, if $\gamma \in \Gamma$ is loxodromic and $e(\gamma) \ge 3$, then $\mathrm{fix}(\gamma)$ is a $\approx$-pair. Suppose $x \in M(3+)$, then Lemma 5.12 gives us such a γ with $x \in \mathrm{fix}(\gamma)$. By Lemma 3.8, if $x \approx y$, it follows that $\mathrm{fix}(\gamma) = \{x, y\}$. ♢

Lemma 5.14 : *Every isolated point of a $\sim$-class in $M(2)$ is the fixed point of a loxodromic in Γ.*

Proof : Suppose $\sigma \subseteq M(2)$ is a $\sim$-class. To say that $x \in \sigma$ is isolated means that there is an open set $U \subseteq M$ such that $\sigma \cap U = \{x\}$. Let $F = M \setminus U$. By Lemma 3.11, there is an annulus (K, x) with $F < (K, x)$ and $N(K, x) = 2$. We now argue as in the proof of Lemma 5.12, using Lemma 3.14 in place of Lemma 3.13, to obtain $\gamma_m, \gamma_n \in \Gamma$ with

$F < \gamma_m A_i < \gamma_n A_i < x$ and $\gamma_m^{-1} x \sim \gamma_n^{-1} x$. Let $\gamma = \gamma_n \gamma_m^{-1}$, so that $x \sim \gamma x$, and γ is loxodromic. Now, $\gamma_m A_i < x$, so $\gamma_n A_i = \gamma(\gamma_m A_i) < \gamma x$. Thus $K < \gamma_n A_i < \gamma x$, and so $\gamma x \in M \setminus K = U$. Since $x \sim \gamma x$ and $\sigma \cap U = \{x\}$, we deduce that $\gamma x = x$. ♢

Corollary 5.15 : *If σ is a $\sim$-class in $M(2)$ which contains an isolated point, then there is a loxodromic $\gamma \in \Gamma$ such that $\mathrm{fix}(\gamma) = \sigma$.*

Proof : Let $x \in \sigma$ be an isolated point of σ. By Lemma 5.14, x is fixed by some loxodromic $\gamma \in \Gamma$. We can suppose that $x = \mathrm{fix}^+(\gamma)$. Let $y = \mathrm{fix}^-(\gamma)$. By Lemma 5.6, $N(x, y) = 2$ and $y \in M(2)$, and so $x \sim y$. In other words, $\mathrm{fix}(\gamma) \subseteq \sigma$. Suppose there is some $z \in \sigma \setminus \mathrm{fix}(\gamma)$. Then $\gamma^n z \in \sigma$ for all n. But $\gamma^n z \to x$, contradicting the assumption that x is isolated in σ. We deduce that $\sigma = \mathrm{fix}(\gamma)$. ♢

In particular, we see immediately that there are no singleton $\sim$-classes. As a consequence, we deduce:

Lemma 5.16 : *There are finitely many Γ-orbits of $\sim$-classes in $M(2)$.*

Proof : As in the proof of Lemma 5.8, using Lemma 3.16 in place of Lemma 3.15. ♢

Definition : We shall say that a subset $\Delta \subseteq M$ is a *necklace* if it is the closure of a $\sim$-class, σ, in $M(2)$, and $\mathrm{card}(\Delta) \geq 3$.

Note that by Corollary 3.6, Δ determines σ uniquely. We write $\sigma = \sigma(\Delta)$. By Lemma 2.2, Δ is a cyclically separating set, and by Corollary 5.15, Δ is perfect.

We shall later describe necklaces as the limit sets of the "MHF" subgroups of Γ. The first objective will be to show that the setwise stabiliser of Δ in Γ is quasiconvex.

Suppose $\Phi_0 \subseteq \Phi$ is compact. Lemma 3.16 and the subsequent remarks tell us that the set of necklaces, Δ, such that $\Phi_0 \cap \Phi_M(\Delta) \neq \emptyset$ is finite.

Proposition 5.17 : *Suppose Δ is a necklace, and $Q \leq \Gamma$ is the setwise stabiliser of Δ. Then, Q is quasiconvex, and Δ is the limit set of Q.*

Proof : Since Δ is perfect, it's enough to show that $\Phi_M(\Delta)/Q$ is compact.

Let $\Phi_0 \subseteq \Phi$ be a compact set such that $\Phi = \bigcup \Gamma \Phi_0$. Let $\{\gamma_1 \Delta, \ldots, \gamma_m \Delta\}$ where $\gamma_i \in \Gamma$ be the set of Γ-images of Δ such that $\Phi_0 \cap \Phi_M(\gamma_i \Delta) \neq \emptyset$. This is a finite set, as noted above. Let $\Psi = \bigcup_{i=1}^m \gamma_i^{-1} \Phi_0$. Thus, $\Psi \subseteq \Phi$ is compact.

Suppose $\rho \in \Phi_M(\Delta)$. Now, there is some $\gamma \in \Gamma$ such that $\gamma\rho \in \Phi_0$. Now $\gamma\rho \in \Phi_M(\gamma\Delta)$ and so $\Phi_0 \cap \Phi_M(\gamma\Delta) \neq \emptyset$. Thus, $\gamma\Delta = \gamma_i \Delta$ for some $i \in \{1, \ldots, m\}$, and so $\gamma^{-1}\gamma_i \in Q$. Now, $\rho \in \gamma^{-1}\Phi_0 = (\gamma^{-1}\gamma_i)\gamma_i^{-1}\Phi_0 \subseteq \bigcup Q\Psi$. This shows that $\Phi_M(\Delta) \subseteq \bigcup Q\Psi$, and so $\Phi_M(\Delta)/Q \subseteq (\bigcup Q\Psi)/Q$. Since $\Phi_M(\Delta)$ is a closed subset of Φ, it follows that $\Phi_M(\Delta)/Q$ is compact, as claimed. ♢

Note that the cyclic order on a necklace, Δ, is defined purely in terms of the topology on M, and is thus Q-invariant. In particular, the set of jumps, $J(\Delta)$, is Q-invariant.

Proposition 5.18 : *If M is not homeomorphic to a circle, then every necklace in M is a cantor set.*

Proof : Let $\Delta \subseteq M$ be a necklace. Since M is metrisable and hence separable, then so is Δ.

If $J(\Delta) = \emptyset$, then Δ is a compact separable cyclically ordered set with no jumps. Thus Δ is homeomorphic to a circle. By Lemma 2.3, we see that $\mathcal{U}(\Delta) = \emptyset$, and so $M \setminus \Delta = \bigcup \mathcal{U}(\Delta) = \emptyset$. Thus, $M = \Delta$ is itself a circle.

Suppose that $J(\Delta) \neq \emptyset$. Let Δ' be the closure of $\bigcup J(\Delta)$. Thus, Δ' is a non-empty closed subset of Δ. Since it is canonically defined, it must be invariant under the setwise stabiliser of Δ. Applying Proposition 5.17 and Lemma 4.6, we see that $\Delta' = \Delta$. In other words, $\bigcup J(\Delta)$ is dense in Δ. Since Δ is separable and perfect, it follows that it must be a cantor set. $\diamondsuit$

Note that by Theorem 4.7, we see that if M is homeomorphic to a circle, then Γ is a cocompact fuchsian group. Moreover, the action of Γ on M is topologically conjugate to the action of the fuchsian group in $\partial \mathbf{H}^2$.

¿From now on, we shall assume that M is not homeomorphic to a circle.

Definition : A subgroup, $Q \leq \Gamma$ is an *MHF subgroup* if it is the setwise stabiliser of a necklace $\Delta \subseteq M$. A *peripheral subgroup* of Q is the stabiliser, in Q, of a jump of Δ.

Thus, an MHF subgroup, Q, is quasiconvex, and its limit set $\Delta = \Lambda Q$ is a cantor set. Peripheral subgroups are either finite or two-ended. In fact, by Theorem 4.8, we see that Q is conjugate to a fuchsian group, and the peripheral subgroups have their usual meaning. We make no explicit use of this fact in this section. We do not yet know that MHF subgroups are finitely generated.

The term "MHF" is meant to be an abbreviation of "maximal hanging fuchsian". This is essentially what is called a "maximal quadratically hanging" subgroup in Sela's terminology, as discussed in the introduction. In Section 6, we show, under the assumption that M is the boundary of a hyperbolic group, that MHF subgroups are finitely generated, and hence bounded fuchsian groups. In this case, the terminology agrees with the notion of "maximal hanging fuchsian" as described in the introduction.

Let's make a few more observations about MHF subgroups in general.

Lemma 5.19 : *Suppose Q is an MHF subgroup, and $\Delta = \Lambda(Q)$. Then $J(\Delta)/Q$ is finite.*

Proof : Recall that Γ acts cocompactly on the space, Π, of unordered pairs in M, i.e. there is a compact set $\Pi_0 \subseteq \Pi$ such that $\Pi = \bigcup \Gamma \Pi_0$. Using Lemma 3.16, we see that only finitely many Γ-images, $\gamma_1 \Delta, \ldots, \gamma_m \Delta$, of Δ meet Π_0. By Lemma 2.4, for each i, the set $J_i = J(\Delta) \cap \gamma_i^{-1} \Pi_0$ is finite. Now, we see easily, that the set $\bigcup_{i=1}^m J_i$ gives a finite transversal to the Q-action on $J(\Delta)$. $\diamondsuit$

We see immediately that:

Corollary 5.20 : *An MHF subgroup, Q, has finitely many Q-conjugacy classes of peripheral subgroups.* $\diamondsuit$

Recall that a necklace, Δ, it the closure of a unique $\sim$-class, denoted $\sigma(\Delta)$.

Lemma 5.21 : *Suppose that $Q \subseteq \Gamma$ is an MHF subgroup. Then $\Delta \setminus \sigma(\Delta) = \bigcup J_0(\Delta)$, where $J_0(\Delta)$ is some Q-invariant subset of $J(\Delta)$. Moreover, if G is the peripheral subgroup corresponding to some jump in $J_0(\Delta)$, then G is two-ended, and $e(G) \geq 3$.*

Proof : Suppose $x \in \Delta \setminus \sigma(\Delta)$. By Lemma 3.7, $\mathrm{val}(x) \geq 3$. Thus, by Lemma 5.12, there is a loxodromic $\gamma \in \Gamma$, with $x \in \mathrm{fix}(\gamma)$. Let y be the other fixed point of γ. Thus, by Lemma 5.6, $N(x,y) = \mathrm{val}(y) = e(\gamma) = \mathrm{val}(x) \geq 3$. Since $\mathrm{val}(y) > 2$, we have that $y \notin \sigma(\Delta)$.

Now, we can assume that $y = \mathrm{fix}^+(\gamma)$. Now, if the necklaces $\gamma^n \Delta$ for $n \geq \mathbf{N}$ were all distinct, they would accumulate at both x and y violating Lemma 3.16. We must therefore have $\gamma^m \Delta = \gamma^n \Delta$ for some $m \neq n$. Thus, replacing γ by γ^{n-m} we can suppose that $\gamma\Delta = \Delta$, and so $\gamma \in Q$. We conclude that $y \in \Delta$. Now since $N(x,y) > 2$, we see that no pair of points of M, and so in particular of $\sigma(\Delta)$, can separate x from y. Thus, $\{x,y\}$ is a jump of Δ. Note that $\{x,y\} \subseteq \Delta \setminus \sigma(\Delta)$.

We now let $J_0(\Delta)$ to be the set of jumps arising in this way. $\diamondsuit$

Corollary 5.22 : *Suppose that Q and Q' are distinct MHF subgroups of Γ. Let $\Delta = \Lambda(Q)$ and $\Delta' = \Lambda(Q')$. Then, either $\Delta \cap \Delta' = \emptyset$, or $\Delta \cap \Delta' \in J_0(\Delta) \cap J_0(\Delta')$. In the former case, $Q \cap Q'$ is finite. In the latter case, $Q \cap Q'$ is two-ended, and of finite index in peripheral subgroups of both Q and Q'. Moreover, $e(Q \cap Q') \geq 3$.*

Proof : Note that by Lemma 3.5, $\mathrm{card}(\Delta \cap \Delta') \leq 2$. If $\Delta \cap \Delta' = \emptyset$, then $Q \cap Q'$ is finite. (Otherwise, we would easily arrive at a contradiction, on applying the convergence hypothesis to a sequence of elements in $Q \cap Q'$.)

Now, $\sigma(\Delta) \cap \sigma(\Delta') = \emptyset$. Thus, if $x \in \Delta \cap \Delta'$, then without loss of generality $x \in \Delta \setminus \sigma(\Delta)$. Thus, by Lemma 5.21, we see that $\mathrm{val}(x) > 3$. It follows that x also lies in $\Delta' \setminus \sigma(\Delta')$. Again, by Lemma 5.21, we see that $x \in \mathrm{fix}(\gamma) \cap \mathrm{fix}(\gamma')$, where $\gamma \in Q$ and $\gamma' \in Q'$ are loxodromics. Since they share a common fixed point, γ and γ' lie in a common two-ended subgroup of Γ. It follows easily that they have a common power, $g \in Q \cap Q'$. Now, $\mathrm{fix}(g) \in J_0(\Delta) \cap J_0(\Delta')$, so $\mathrm{fix}(g) \subseteq \Delta \cap \Delta'$. Since $\mathrm{card}(\Delta \cap \Delta') \leq 2$, we see that $\mathrm{fix}(g) = \Delta \cap \Delta'$. $\diamondsuit$

At the end of Section 3, we described how to put a pretree structure on the set, T, consisting of all $\approx$-pairs and $\sim$-classes. We saw (Proposition 3.20) that T is in fact a discrete pretree, and so embeds naturally in a discrete median pretree, Θ, using the construction outlined in Section 2. We can identify Θ as the vertex set of a simplicial tree, Σ. We shall write $E = E(\Sigma)$ for the set of edges of Σ.

It is natural to partition T into two subsets as follows. We write $\Theta_2 \subseteq T$ for the set consisting of all infinite $\sim$-classes, and write $\Theta_1 = T \setminus \Theta_2$. Thus, Θ_1 consists of all $\approx$-pairs and $\sim$-pairs. (By the latter we mean a $\sim$-class consisting of two points.) We write

$\Theta_3 = \Theta \setminus T$. Thus $\Theta = \Theta_1 \sqcup \Theta_2 \sqcup \Theta_3$. We shall sometimes speak of this as defining a "3-colouring" of the vertex set.

¿From the construction of Θ from T, it's not hard to see that every vertex of Θ_3 has degree at least 3 in Σ, and that no two vertices of Θ_3 are adjacent. (In fact, we shall see later that no two adjacent vertices of Θ have the same colour, so that this is indeed a 3-colouring in the usual sense).

Since the construction of Σ is canonical, we see that Γ acts simplicially on Σ, preserving the 3-colouring. Given $\theta \in \Theta$ and $e \in E$, we write $\Gamma(\theta)$ and $\Gamma(e)$ for the vertex and edge stabilisers respectively. Thus, if $\theta \in \Theta_1$, then $\Gamma(\theta)$ is a two-ended subgroup, and if $\theta \in \Theta_2$, then $\Gamma(\theta)$ is an MHF subgroup.

Reinterpreting Lemma 5.8 and 5.16, we get:

Lemma 5.23 : $(\Theta_1 \cup \Theta_2)/\Gamma$ *is finite.* ◊

We shall see in fact that Σ has finite quotient under Γ. To be able to speak of the quotient graph we need to know that there are no edge inversions, which will follow when we have established the claim that adjacent vertices have different colours. We shall also see that all edge stabilisers are elementary, and thus two-ended given that Γ does not split over a finite subgroup.

Given $\theta \in \Theta$, let $E_\theta \subseteq E$ be the set of edges incident on θ. We can define an equivalence relation, $\simeq$, on $\Theta \setminus \{\theta\}$ by writing $\zeta \simeq \eta$ if and only if not $\zeta\theta\eta$. Thus, there is a natural bijection between E_θ and $(\Theta \setminus \{\theta\})/\simeq$. In fact, since no vertex of Θ_3 is terminal, we see that $T = \Theta_1 \cup \Theta_2$ must intersect every $\simeq$-class in $\Theta \setminus \{\theta\}$. We can thus restrict the relation $\simeq$ to $T \setminus \{\theta\}$. In this way we get a natural bijection between E_θ and $(T \setminus \{\theta\})/\simeq$.

Given a pair of distinct points, $\theta = \{x, y\} \subseteq M$, we write $\mathcal{U}(\theta) = \mathcal{U}(x, y)$.

Lemma 5.24 : *If $\theta \in \Theta_1$, then there is natural bijection between E_θ and $\mathcal{U}(\theta)$.*

Proof : By Lemma 3.18, if $\zeta \in T \setminus \{\theta\}$, then $\zeta \subseteq U$ for some $U \in \mathcal{U}(\theta)$. Since the action of Γ on M is minimal, we see that each such U must contain some such ζ. By the definition of the pretree structure on T, we see that $\zeta \simeq \eta$ if and only if they lie in the same element U. ◊

Lemma 5.25 : *Suppose that $\sigma \in \Theta_2$. Let Δ be the necklace $\bar{\sigma}$. Then, there is a natural bijection between E_σ and $J(\Delta)$.*

Proof : Suppose $\zeta \in T \setminus \{\sigma\}$. Using Lemma 5.21, we see that either $\zeta \cap \Delta = \emptyset$ or $\zeta \in J_0(\Delta)$. If $\zeta \cap \Delta = \emptyset$, then there is some $U \in \mathcal{U}(\Delta)$ with $\zeta \in U$. (For if ζ intersected two components of $M \setminus \Delta$, then it would be separated by some pair of points of σ, contradicting Lemma 3.18.) Now by Lemma 2.3, there is a unique $\theta \in J(\Delta)$ such that $U \in \mathcal{U}_\Delta(\theta)$. We can thus define a map $h : T \setminus \{\sigma\} \longrightarrow J(\Delta)$ so that either $h(\zeta) = \zeta \in J_0(\Delta)$, or $h(\zeta) \subseteq U \in \mathcal{U}_\Delta(h(\zeta))$. It follows easily from the minimality of the action of Γ on M and Lemma 3.18 that h is surjective. We claim that $h(\zeta) = h(\eta)$ if and only if $\zeta \simeq \eta$.

Firstly, suppose $h(\zeta) \neq h(\eta)$. If $\zeta, \eta \notin J_0(\Delta)$, then they lie in different components of $M \setminus \Delta$ and are thus separated by some pair of points of σ. Thus $\zeta \not\simeq \eta$. Similarly, if

$\zeta \in J_0(\Delta)$, so $h(\zeta) = \zeta$, we can suppose that $\eta \notin J_0(\Delta)$. Thus, $\zeta \subseteq U$ where $U \in \mathcal{U}_\Delta(\theta)$ and $\theta \in J(\Delta)$ and $\theta \neq \zeta$. Again, we see that $\zeta \not\simeq \eta$.

Conversely, suppose that $h(\zeta) = h(\eta) = \theta$, say. We distinguish two cases.

If $\theta \in J(\Delta) \setminus J_0(\Delta)$, then $\theta \subseteq \sigma$, and so $\operatorname{card} \mathcal{U}(\theta) = 2$. Thus, $\mathcal{U}_\Delta(\theta)$ consists of a single element, U. Since $\zeta, \eta \neq \theta$, we have $\zeta, \eta \subseteq U$, and so they are not separated by any pair of points of σ. Thus $\zeta \simeq \eta$.

Finally, suppose that $\theta \in J_0(\Delta)$. Thus $\theta \subseteq \Delta \setminus \sigma$. Let $K = \theta \cup \bigcup \mathcal{U}_\Delta(\theta)$, so that K is closed and connected. Also $M \setminus K$ is the component of $M \setminus \theta$ which contains σ. In particular, $K \cap \sigma = \emptyset$. By the definition of h, we have $\zeta, \eta \subseteq K$, and so again we see that $\zeta \simeq \eta$. ♢

We immediately get:

Lemma 5.26 : *If $\theta \in \Theta_1$, then E_θ is finite. If $\theta \in \Theta_2$, then $E_\theta / \Gamma(\theta)$ is finite.*

Proof : If $\theta \in \Theta_1$, then $\mathcal{U}(\theta)$ is finite, so we apply Lemma 5.24. If $\theta \in \Theta_2$, then θ is an MHF subgroup. In this case, Lemma 5.19 tells us that $J(\Delta)/\Gamma(\theta)$ is finite. ♢

Since we don't yet know that there are no edge inversions, we define $\Gamma_0(e)$ to be the "directed edge stabiliser" of an edge e. In other words, $\Gamma_0(e) = \Gamma(\zeta) \cap \Gamma(\eta)$, where $\zeta, \eta \in \Theta$ are the endpoints of e. Thus, $\Gamma_0(e)$ has index at most 2 in $\Gamma(e)$.

Suppose $\theta \in \Theta_1$. Now $\Gamma(\theta)$ is two-ended, and $e(\Gamma(\theta)) = \operatorname{card} \mathcal{U}(\theta) = \operatorname{card} E_\theta$. In other words, $e(\Gamma(\theta))$ equals the degree of θ in Σ. If $e \in E_\theta$, then $\Gamma_0(e)$ has finite index in $\Gamma(\theta)$ and is thus a two-ended group.

Suppose $\sigma \in \Theta_2$, so that $\Gamma(\sigma)$ is an MHF group. If $e \in E_\sigma$, then $\Gamma_0(e)$ is the stabiliser of a jump of $\bar{\sigma} = \Lambda(\Gamma_0(e))$. In other words, $\Gamma_0(e)$ is a peripheral subgroup of $\Gamma(\sigma)$, and is thus elementary. Since Γ doesn't split over any finite subgroup, it follows that $\Gamma(e)$ is two-ended. (Note that, from the construction of Σ, every edge, e, lies in an arc connecting two points of $\Theta_1 \cup \Theta_2$. We know corresponding vertex groups to be infinite, so if $\Gamma(e)$ were finite, we would obtain a non-trivial splitting.)

Now every edge $e \in E$ is incident on some vertex of $\Theta_1 \cup \Theta_2$. It follows that $\Gamma_0(e)$ and hence $\Gamma(e)$ is elementary.

We next want to show that no two adjacent vertices have the same colour. We have already noted that this is true of Θ_3.

Suppose that $\zeta, \eta \in \Theta_1$ are the endpoints of some edge $e \in E$. Now, $\Gamma_0(e) = \Gamma(\zeta) \cap \Gamma(\eta)$ is of finite index in both ζ and η. But these groups are both two-ended with limit sets ζ and η respectively. We get the contradiction that $\zeta = \eta$, showing that this situation is not possible.

Suppose now that $\sigma, \sigma' \subseteq \Theta_2$ are the endpoints of the edge $e \in E$. Let $\Delta = \Lambda(\Gamma(\sigma))$ and $\Delta' = \Lambda(\Gamma(\sigma'))$. Since $\Gamma_0(e) = \Gamma(\sigma) \cap \Gamma(\sigma')$ is two-ended, we see, by Corollary 5.22, that $\theta = \Delta \cap \Delta' \in J(\Delta) \cap J(\Delta')$. Now, $e(\Gamma_0(e)) \geq 3$, and so $\theta \in \Theta_1$. But now, we have $\sigma\theta\sigma'$ contradicting the assumption that σ and σ' are adjacent.

We have shown that no two points of Θ_1 and no two points of Θ_2 are adjacent. Thus, the partition, $\Theta = \Theta_1 \sqcup \Theta_2 \sqcup \Theta_3$ is indeed a 3-colouring in the usual sense. It follows that there are no edge inversions, i.e. $\Gamma_0(e) = \Gamma(e)$ for all $e \in E$. We can thus construct the quotient graph Σ/Γ. Clearly Σ/Γ is connected.

Now we can identify the vertex set, $V(\Sigma/\Gamma)$, of Σ/Γ with Θ/Γ. Thus $V(\Sigma/\Gamma) = W_1 \sqcup W_2 \sqcup W_3$ gives a 3-colouring of Σ/Γ, where $W_i = \Theta_i/\Gamma$. By Lemma 5.23, $W_1 \cup W_2$ is finite. By Lemma 5.26, each vertex of $W_1 \cup W_2$ has finite degree in Σ/Γ. Also every vertex of W_3 is adjacent of a vertex of $W_1 \cap W_2$. We conclude that Σ/Γ is finite.

Note that no vertex of Σ is terminal (has degree 1). Since Σ/Γ is finite, we see easily that the action of Γ on Σ is minimal (i.e. there is no proper invariant subtree).

There a couple more observations we can make concerning the structure of the vertex groups.

Lemma 5.27 : *If $\theta \in \Theta_3$, then $\Gamma(\theta)$ is non-elementary.*

Proof : Suppose, for contradiction, that $\Gamma(\theta)$ is elementary. We know that θ has degree at least 3 in Σ. Since $\Gamma(\theta)$ contains each incident edge group, we see that it contains a loxodromic element, and is therefore two-ended. If $\zeta \in \Theta_1$ is adjacent to θ, then, since the edge group, $\Gamma(\zeta) \cap \Gamma(\theta)$ is two ended, we see that $\Gamma(\theta)$ and $\Gamma(\zeta)$ are commensurable, and so we see that $\zeta = \Lambda(\Gamma(\zeta)) = \Lambda(\Gamma(\theta))$. It follows that at most one vertex of Θ_1 can be adjacent to θ. We can thus find two distinct vertices, $\sigma, \tau \in \Theta_2$ adjacent to θ. But now, $\Gamma(\theta)$ is commensurable with peripheral subgroups of $\Gamma(\sigma)$ and $\Gamma(\tau)$. Applying Corollary 5.22, we see that $\eta = \Lambda(\Gamma(\sigma)) \cap \Lambda(\Gamma(\tau)) \in \Theta_1$, and the the betweenness relation $\sigma\eta\tau$ holds. We thus get the contradiction that $\theta = \eta \in \Theta_1$. We conclude that $\Gamma(\theta)$ is non-elementary as claimed. $\diamondsuit$

It's also clear that $\Gamma(\theta)$ cannot be an MHF subgroup. (If it were, it would be equal to $\Gamma(\theta')$ for some $\theta' \in \Theta_2$. If e is any edge in the arc connecting θ to θ', then we get the contradiction that $\Gamma(e) = \Gamma(\theta)$ is not two-ended.)

Note that it follows from Lemma 5.27 that Θ_1 is precisely the set of vertices of Σ of finite degree.

Also, if $e, e' \in E$ and $\Gamma(e)$ and $\Gamma(e')$ are commensurable, then $\Lambda\Gamma(e) = \Lambda\Gamma(e')$ is a $\sim$-pair or a $\approx$-pair, and so $\Gamma(e), \Gamma(e') \subseteq \Gamma(\theta)$ for some $\theta \in \Theta_1$ incident on e and e'.

Finally, note that if $\theta \in \Theta_1$ is adjacent to some vertex, $\sigma \in \Theta_2$, then it follows from Lemma 5.21 that $e(\Gamma(e)) \geq 3$, where e is the connecting edge. Since $\Gamma(\theta)$ is commensurable with $\Gamma(e)$, we see that $e(\Gamma(\theta)) = e(\Gamma(e)) \geq 3$. On the other hand, if $e \in E(\Sigma)$ connects a vertex in Θ_2 to a vertex in Θ_3, then $\Gamma(e)$ is maximal two-ended and $e(\Gamma(e)) = 2$.

We want to summarise these properties in group theoretical terms, with a view to describing the uniqueness of the splitting in Section 6. We will get a slightly cleaner statment if we, somewhat artificially, insert a vertex of degree 2 at the midpoint of each edge of Σ which has one endpoint in Θ_2 and the other in Θ_3. With these extra vertices, we can partition $V(\Sigma)$ as $V_1(\Sigma) \sqcup V_2(\Sigma) \sqcup V_3(\Sigma)$, where $V_2(\Sigma) = \Theta_2$, $V_3(\Sigma) = \Theta_3$, and $V_1(\Sigma)$ consists of the set Θ_1 together with the set of new degree-2 vertices we have just introduced.

Summarising the above discussion, we have:

Theorem 5.28 : *Γ acts minimally and simplicially on Σ without edge-inversions and with finite quotient Σ/Γ. The edge stabilisers are all two-ended subgroups. The action of*

Γ preserves the 3-colouring, $(V_1(\Sigma), V_2(\Sigma), V_3(\Sigma))$, of the vertex set of Σ. No two vertices of $V_1(\Sigma)$ are adjacent, and no two vertices of $V_2(\Sigma) \cup V_3(\Sigma)$ are adjacent.

If $v \in V_1(\Sigma)$, then the vertex stabiliser, $\Gamma(v)$, is a maximal two-ended subgroup of Γ, and $e(\Gamma(v))$ equals the degree of v in Σ. This degree is finite and at least 2. If the degree equals 2, then at least one of the incident vertices lies in $V_3(\Sigma)$.

If $v \in V_2(\Sigma)$, then $\Gamma(v)$ is an MHF subgroup of Γ. If $e \in E(\Sigma)$ is incident on v, then $\Gamma(e) \leq \Gamma(v)$ is a peripheral subgroup of $\Gamma(v)$.

If $v \in V_3(\Sigma)$, then $\Gamma(v)$ is non-elementary (and not a torsion group) and is not an MHF subgroup. ♢

We also note that any pair of commensurable edge stabilisers lie inside a common incident vertex stabiliser, $\Gamma(v)$ for $v \in V_1(\Sigma)$.

The term "MHF subgroup", as we have defined it, makes reference to our particular construction. We shall give an algebraic reinterpretation of this term in Section 6.

Clearly this construction has been canonical. It is describes a "maximal" splitting over two-ended subgroups. There are various ways to describe this maximality. Below we give a topological and a group theoretical formulation.

Proposition 5.29 : *Every local cut point of M lies in the limit set of some vertex stabiliser, $\Gamma(v)$, where $v \in V_1(\Sigma) \cup V_2(\Sigma)$.*

Proof : This follows since $M(2)$ is a union of $\sim$-classes and $M(3+)$ is a union of $\approx$-pairs. The set of local cut points is, by definition, $M(2) \cup M(3+)$. ♢

Proposition 5.30 : *Suppose $G \leq \Gamma$ is a two-ended group with $e(G) \geq 2$. Then $G \leq \Gamma(v)$ for some $v \in V_1(\Sigma) \cup V_2(\Sigma)$.*

Proof : If $e(G) \geq 3$, then $\Lambda(G)$ is a $\approx$-pair $\theta \in \Theta_1$, and so $G \leq \Gamma(\theta)$. If $e(G) = 2$, then either $\Lambda(G)$ is a $\sim$-pair $\theta \in \Theta_1$ so again $G \leq \Gamma(\theta)$, or else $\Lambda(G)$ is a subset of some infinite $\sim$-class $\sigma \in \Theta_2$. In the latter case, σ is G-invariant, and so $G \leq \Gamma(\sigma)$. ♢

Now, since Γ is countable and $\Pi(3+)/\Gamma$ is finite (Lemma 5.8), it follows that $M(3+)$ is countable. Using this, we deduce:

Proposition 5.31 : *Suppose $\Delta \subseteq M$ is a closed perfect cyclically separating set. Then, Δ lies in the limit set of $\Gamma(v)$ for some $v \in V_2(\Sigma)$.*

Proof : In other words, we claim that Δ lies in some necklace. Now, $\Delta \subseteq M(2) \cup M(3+)$, and since Δ is perfect, it is locally uncountable. Thus, $\Delta \cap M(2)$ is dense in Δ. But if $x, y \in \Delta \cap M(2)$, then $x \sim y$. Thus, $\Delta \cap M(2) \subseteq \sigma$ for some $\sim$-class σ. Thus $\Delta \subseteq \bar{\sigma}$, which, by definition, is a necklace. ♢

Note that another way of expressing the maximality of the splitting is to say that no vertex group of the form $\Gamma(v)$ for $v \in V_3(\Sigma)$ splits over a two-ended subgroup relative the incident edge groups (i.e. in such a way that the incident edge groups are conjugate into one of the vertex groups in the supposed splitting of $\Gamma(v)$).

6. Conclusion.

In this section, we reintroduce some geometry, and give a summary of our results in the context in which we are really interested, namely when Γ is a hyperbolic group in the sense of Gromov [Gr]. For an exposition of such groups, see, for example, [GhH].

Dunwooody's accessibility theorem for finitely presented groups [Du] tells us that any hyperbolic group can be split as a finite graph of groups where the edge groups are all finite, and where all the vertex groups are all finite or one-ended. Note that the vertex groups are all quasiconvex (see Proposition 2.1) and hence themselves hyperbolic. This can be thought of as analogous to the splitting of a 3-manifold into irreducible components along 2-spheres. Here we are concerned with the second stage of splitting, which is over two-ended subgroups. This is analogous to the characteristic submanifold construction for irreducible 3-manifolds described in [JaS] and [Jo] following ideas of Waldhausen.

Suppose, then, that Γ is a one-ended hyperbolic group, so that $\partial\Gamma$ is a continuum. In this case, Bestvina and Mess [BeM] conjectured that $\partial\Gamma$ has no global cut point and showed that this implies that $\partial\Gamma$ is locally connected. The converse is given by Proposition 5.4. As discussed in the introduction, this is now known for all one-ended hyperbolic groups, though we shall not explicitly use this fact here.

If $\partial\Gamma$ does contain a global cut point, then the set of all global cut points has the structure of a pretree (Section 2). ¿From this information, one can construct an equivariant quotient of $\partial\Gamma$ which is a non-trivial dendrite [Bo1]. (The equivalence relation on $\partial\Gamma$ can be defined by deeming two points to be not equivalent if there exists a set of cut points points individually separating them which is order isomorphic to the rational numbers in the natural linear order.) The group Γ acts as a convergence group on this dendrite, and one can adapt the Rips machinery for $\mathbf{R}$-trees to this context to show that Γ splits over a two-ended subgroup [Bo2]. (One can give an alternative argument using Levitt's generalisation of [Bo2] in [L].) One can continue along these lines to rule out this possibility altogether [Sw] (see also [Bo5]).

We shall now assume that there is no global cut point, so that $\partial\Gamma$ is a metrisable Peano continuum. Now, Γ acts as a uniform convergence group on $\partial\Gamma$, so we can quote the results of the last section. In particular, assuming that Γ is not a cocompact fuchsian group, we get a splitting of Γ as a finite graph of groups, Σ/Γ, as described by Theorem 5.28.

Now, all the edge groups in this splitting are two-ended and hence quasiconvex. It follows, by Proposition 1.2, that all the vertex groups are also quasiconvex, in the usual geometric sense as described in Section 1.

In fact, it's not hard to see that the notion of quasiconvexity defined in Section 4 agrees with the usual geometric notion in the case of a hyperbolic group acting on its boundary. We thus get two reasons why MHF subgroups are quasiconvex — either using Proposition 1.2 as above, of directly from Proposition 5.17. In particular, MHF subgroups are finitely generated in this case. By Proposition 4.9, we see that an MHF group is conjugate to a bounded fuchsian group by a conjugacy that sends incident edge groups to peripheral subgroups. Recall that a "bounded fuchsian group", as described in the introduction, is a non-elementary finitely generated group that acts properly discontinuously without parabolics on $\mathbf{H}^2$, and such that the quotient, $\mathbf{H}^2/\Gamma$ is not compact. We do not assume that the action is effective — only that the kernel is finite. Note that such an action is necessarily "convex cocompact", i.e. geometrically finite without parabolics. Thus, formally, a bounded fuchsian group consists of a virtually free group with a preferred collection of peripheral subgroups.

We next define the notion of a "hanging fuchsian" subgroup. In the torsion free case, this coincides with what Sela calls a "quadratically hanging" subgroup.

Definition : A subgroup, $Q \leq \Gamma$, is a *hanging fuchsian* subgroup if it occurs as the vertex group of a finite splitting of Γ, in such a way that Q admits an isomorphism with a bounded fuchsian group so that the incident edge groups in the splitting are precisely the peripheral subgroups of Q.

Without loss of generality, we can assume that every other vertex group of the splitting is adjacent to Q. In this case every edge group in the splitting is two-ended. Thus, by Proposition 1.2, every vertex group is quasiconvex. In particular, this shows that every hanging fuchsian subgroup is quasiconvex.

¿From the discussion of the boundary given in Section 1, it's not hard to see that the limit set, $\Lambda(Q)$, of Q is a perfect cyclically separating set. Thus, by Proposition 5.31, we see that $\Lambda(Q) \subseteq \Lambda(\Gamma(v))$ for some $v \in V_2(\Sigma)$. Since $\Gamma(v)$ is the setwise stabiliser of $\Gamma(v)$, we see that $Q \leq \Gamma(v)$. We conclude:

Proposition 6.1 : *Every hanging fuchsian subgroup of* Γ *is quasiconvex, and lies inside one of the MHF subgroups of the JSJ splitting.* $\diamondsuit$

We have already noted that an MHF subgroup is a bounded fuchsian group, and hence itself hanging fuchsian. This justifies the terminology — the MHF subgroups of Γ are precisely the maximal hanging fuchsian subgroups.

We have observed that our splitting is canonical in that it arises explicitly from the action of Γ on $\partial\Gamma$. In fact, the uniqueness of the splitting can be characterised in purely group theoretical terms. To describe exactly how this works, we digress for a moment to consider precisely what we mean by a splitting of an arbitrary group,

Usually a group splitting is described in terms of a presentation of a group as the fundamental group of a graph of groups. However, to describe uniqueness in these terms is a little clumsy. Formally it is more convenient to view a splitting as an additional structure associated to the group satisfying certain axioms. These can be laid out explicitly as follows. (Our formulation rules out the possibility of an edge group being equal to the two incident vertex groups, but this situation never arises with JSJ splittings.)

Suppose Γ is any group. We view a splitting of Γ as consisting of a collection, $\mathcal{V}$, of subgroups of Γ, together with a symmetric binary relation on $\mathcal{V}$ satisfying the following conditions. Firstly, $\mathcal{V}$ is closed under conjugacy, and each element of $\mathcal{V}$ is equal to its normaliser in Γ. This defines an action of Γ on $\mathcal{V}$, such that the stabiliser of any $G \in \mathcal{V}$ is equal to G itself. We suppose that the binary relation on $\mathcal{V}$ is Γ-invariant. Moreover, if we view $\mathcal{V}$ as the vertex set of a graph, Σ, with this adjacency relation, then we assume that Σ is a simplicial tree. Thus, Γ acts simplicially on Γ. We go on to assume that there are no edge inversions and that Σ/Γ is finite, giving us a graph of groups in the usual sense.

In the case where Γ is a one-ended hyperbolic group (with locally connected boundary), we have constructed such a splitting as described in Section 5. We can extract from our the construction those algebraic features which determine its uniqueness. (Of course, our construction has been entirely canonical throughout, so this is of interest primarily in relating our construction to others that have appeared elsewhere.) Firstly, Σ admits vertex 3-colouring $(V_1(\Sigma), V_2(\Sigma), V_3(\Sigma))$, satisfying the conclusion of Theorem 5.28. In this context, we should interpret the term "MHF subgroup" as a maximal hanging fuchsian subgroup as formulated in this section. Moreover, this splitting is maximal in the sense described by Proposition 5.30. Recall, in particular, from the conclusion of Theorem 5.28, that every edge group $\Gamma(e)$ is a finite index subgroup of an incident vertex group of the form $\Gamma(v)$ for $v \in V_1(\Sigma)$. Moreover $e(\Gamma(e)) = e(\Gamma(v))$ equals the degree of v in Σ. Other properties of the splitting can be deduced from this assumption; for example, the fact that any pair of commensurable edge groups are subgroups of a common incident vertex group of this type. Bringing the conclusion of Proposition 5.30 into play, one can go on to derive the fact that no vertex group of $V_3(\Sigma)$ splits over a two-ended subgroup relative to its incident edge groups. Moreover, every hanging fuchsian group lies inside a vertex group of $V_2(\Sigma)$. Thus, every maximal hanging fuchsian group occurs as such a vertex group. Finally we note that the adjacency relation in the tree, Σ, is determined by the observation that v is adjacent to w if and only if $\Gamma(v) \cap \Gamma(w)$ is infinite and $\{v, w\} \cap V_1(\Sigma)$ has exactly one element. (We shall not give proofs of these observations — we already knew all these things about our particular splitting, so, if we wish, we could simply add them to our list of requirements.) We now have enough information to see easily that our splitting is unique. We leave the details of these assertions as an exercise.

We should note that our formulation of the JSJ splitting differs slightly from that in [Se]. Apart from questions of torsion, the main difference is due to the fact that some of our vertex groups of type (1) are omitted from Sela's splitting. This has two consequences we remark upon. Firstly, in Sela's account, edge groups incident on MHF groups are only assumed to be of finite index in peripheral subgroups (and not equal to peripheral subgroups as here). Secondly, Sela's splitting is only canonical up to certain "sliding operations" whereas ours is unique.

Note that since it is canonical, any outer automorphism must respect this splitting. Thus, there is a finite index subgroup of $\mathrm{Out}(\Gamma)$ which fixes each vertex in the graph of groups. One can modify the arguments of [P] to show that each vertex group, G, of type (3) is rigid relative to the incident edge groups (i.e. the subgroup $\mathrm{Out}(G)$ preserving the conjugacy classes of these edge groups is finite). Thus, with a bit more work, one arrives at Sela's result (given in the torsion-free case) that $\mathrm{Out}(\Gamma)$ is virtually a direct product of infinite cyclic groups and orbifold mapping class groups. With these techniques at our disposal, it should not be hard, for example, to give a precise description of when $\mathrm{Out}(\Gamma)$ is infinite, though we shall not pursue this question here. (See [MNS] for further discussion.)

There are also a few topological consequences to our construction. ¿From Proposition 5.29, every local cut point plays a role in the splitting. In particular, the existence of a local cut point implies that the splitting is non-trivial. We conclude:

Theorem 6.2 : *Suppose Γ is a one-ended hyperbolic group which is not a cocompact fuchsian group. Then Γ splits over a two-ended subgroup if and only if $\partial\Gamma$ has a local cut point.*

In particular, we see that, modulo fuchsian groups, this property is quasiisometry invariant. (In fact, to see this, we only need that the existence of a global cut point would give rise to a splitting [Bo1,Bo2].)

For completeness, we note that a cocompact fuchsian group splits over a two-ended subgroup if and only if it is not a virtual semitriangle group, as discussed in the introduction. (A "semitriangle group" has a presentation of the form $\langle a, b \mid a^p = b^q = (ab)^r = 1 \rangle$, where $p, q, r \in \mathbf{N}$ satisfy $p^{-1} + q^{-1} + r^{-1} < 1$.)

References.

[AN] S.A.Adeleke, P.M.Neumann, *Relations related to betweenness: their structure and automorphisms* : to appear in Memoirs Amer. Math. Soc.

[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, *Group actions on trees and dendrons* : to appear in Topology.

[Bo3] B.H.Bowditch, *Boundaries of strongly accessible hyperbolic groups* : preprint, Melbourne (1996).

[Bo4] B.H.Bowditch, *Convergence groups and configuration spaces* : to appear in "Group Theory Down Under", (ed. J.Cossey, C.F.Miller, W.D.Neumann, M.Shapiro), de Gruyter.

[Bo5] B.H.Bowditch, *Connectedness properties of limit sets* : preprint, Melbourne (1996).

[Bo6] B.H.Bowditch, *A topological characterisation of hyperbolic groups* : preprint, Southampton (1996).

[CJ] A.Casson, D.Jungreis, *Convergence groups and Seifert fibered 3-manifolds* : Invent. Math. **118** (1994) 441–456.

[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.

[DuSa] M.J.Dunwoody, M.E.Sageev, *JSJ-splittings for finitely presented groups over slender subgroups* : preprint, Southampton (1996).

[DuSw] M.J.Dunwoody, E.L.Swenson, *The algebraic annulus theorem* : preprint, Southampton (1996).

[Fr] E.M.Freden, *Negatively curved groups have the convergenve property* : Ann. Acad. Sci. Fenn. Ser. A Math. **20** (1995) 333-348.

[FuP] K.Fujiwara, P.Papasoglu, *JSJ decompositions of finitely presented groups and complexes of groups* : preprint (1997).

[Ga] D.Gabai, *Convergence groups are fuchsian groups* : Ann. of Math. **136** (1992) 447–510.

[GeM1] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I* : Proc. London Math. Soc. **55** (1987) 331–358.

[GeM2] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups II* : handwritten notes.

[GhH] E.Ghys, P.de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Maths. No. 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.

[HY] J.G.Hocking, G.S.Young, *Topology* : Addison-Wesley (1961).

[JaS] W.H.Jaco, P.B.Shalen, *Seifert fibered spaces in 3-manifolds* : Amer. Math. Soc. Mem. **220** (1979).

[Jo] K.Johannson, *Homotopy equivalences of 3-manifolds with boundary* : Springer Lecture Notes in Mathematics, No. 761, Springer Verlag, Berlin (1979)

[K] P.H.Kropholler, *A group theoretic proof of the torus theorem* : in "Geometric Group Theory, Volume 1", (ed. G.A.Niblo, M.A.Roller), London Math. Soc. Lecture Notes Series No. 181, Cambridge University Press (1993) 138–158.

[L] G.Levitt, *Non-nesting actions on real trees* : to appear in Bull. London Math. Soc.

[MNS] C.F.Miller III, W.D.Neumann, G.A.Swarup, *Some examples of hyperbolic groups* : preprint, Melbourne (1996).

[P] F.Paulin, *Outer automorphisms of hyperbolic groups and small actions on* $\mathbf{R}$*-trees* : in "Arborial group theory" (ed. R.C.Alperin) M.S.R.I. Publications No. 19, Spinger-Verlag (1991) 331–343.

[RS] E.Rips, Z.Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition* : Ann. of Math. **146** (1997) 53–109.

[ScS] P.Scott, G.A.Swarup, *An algebraic annulus theorem* : preprint.

[Se] Z.Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II* : to appear in GAFA.

[Sha] P.B.Shalen, *Dendrology and its applications* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 543–616.

[Sho] 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).

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

[T1] P.Tukia, *Homeomorphic conjugates of fuchsian groups* : J. reine angew. Math. **391** (1988) 1–54.

[T2] P.Tukia, *Convergence groups and Gromov's metric hyperbolic spaces* : New Zealand J. Math. **23** (1994) 157–187.

[W] L.E.Ward, *Axioms for cutpoints* : in "General topology and modern analysis", Proceedings, University of California, Riverside (ed. L.F.McAuley, M.M.Rao), Academic Press (1980) 327–336.