

PERIPHERAL SPLITTINGS OF GROUPS

B. H. BOWDITCH

ABSTRACT. We define the notion of a “peripheral splitting” of a group. This is essentially a representation of the group as the fundamental group of a bipartite graph of groups, where all the vertex groups of one colour are held fixed—the “peripheral subgroups”. We develop the theory of such splittings and prove an accessibility result. The theory mainly applies to relatively hyperbolic groups with connected boundary, where the peripheral subgroups are precisely the maximal parabolic subgroups. We show that if such a group admits a non-trivial peripheral splitting, then its boundary has a global cut point. Moreover, the non-peripheral vertex groups of such a splitting are themselves relatively hyperbolic. These results, together with results from elsewhere, show that under modest constraints on the peripheral subgroups, the boundary of a relatively hyperbolic group is locally connected if it is connected. In retrospect, one further deduces that the set of global cut points in such a boundary has a simplicial treelike structure.

0. INTRODUCTION

In this paper, we study particular kinds of splittings of groups, which we shall term *peripheral splittings*. The theory can be developed in a general context, and fits into a broad scheme of studying various properties of groups (in this case splittings) relative to preferred classes of subgroups. However, the main applications we have in mind here are to relatively hyperbolic groups, where the preferred subgroups are the maximal parabolic subgroups. A peripheral splitting of such a group is reflected in the topology of its boundary; in particular, the existence of global cut points. It is well established that global cut points play an important role in the connectedness properties of boundaries (cf. [BeM]). In particular, the results we give here are central to the main result of [Bo6], namely that the boundary of a relatively hyperbolic group is locally connected if it is connected—given some mild constraints on the class of parabolic subgroups (see Theorem 1.5).

Suppose that Γ is a group with a preferred set, \mathcal{G} , of subgroups, which we call *peripheral subgroups*. Formally, a peripheral splitting of Γ can be defined as a presentation of Γ as a finite bipartite graph of groups, where the vertex groups of one colour (i.e. in one of the sets in the partition of vertices) are precisely the peripheral subgroups. They can also be thought of in terms of splittings of the group over subgroups of the peripheral subgroups, relative to peripheral subgroups. However, we shall see that this latter formulation is less natural for our purposes. One of the main results of this paper will be an accessibility result for such splittings; see Theorem 6.1.

Received by the editors November 19, 1999 and, in revised form, January 31, 2001.
2000 *Mathematics Subject Classification*. Primary 20F67, 20E08.

In order to tie these results in naturally with the theory of relatively hyperbolic groups, we shall take a similar approach to that described in [Bo5]. In that paper it was suggested that a natural way to view properties of a group, Γ , “relative to” a preferred class of subgroups, \mathcal{G} , is to consider actions of Γ on a set, V , for which the point stabilisers are precisely the elements of \mathcal{G} . The geometry of such actions can be analysed by extending the action to a connected graph with vertex set V and with finite quotient, thereby generalising the notion of a Cayley graph in the non-relative case. These ideas are outlined in Section 4. The notion of a relatively hyperbolic group has a clean formulation in these terms, and we describe here how to view relative splittings of Γ in similar fashion.

I am indebted to the referee for many helpful comments, in particular, for a simplification of the proof of Theorem 6.1, and some corrections to Section 7.

1. SUMMARY OF RESULTS AND APPLICATIONS

We summarise the main results of this paper as they apply to relatively hyperbolic groups, and some of their consequences concerning connectedness of boundaries.

Let Γ be a group. By a *splitting* of Γ , over a given class of subgroups, we mean a presentation of Γ as a finite graph of groups, where each edge group belongs to this class. Such a splitting is said to be *relative to* another class, \mathcal{G} , of subgroups if each element of \mathcal{G} is conjugate into one of the vertex groups. We usually take \mathcal{G} to be invariant under conjugacy, and refer to elements of \mathcal{G} as *peripheral* subgroups. A splitting is said to be *trivial* if it is relative to $\{\Gamma\}$, i.e. at least one of the vertex groups is equal to Γ . A *peripheral splitting* is a representation of Γ as a finite bipartite graph of groups, where \mathcal{G} consists precisely of the (conjugacy classes of) vertex groups of one colour. (It will be convenient to rule out the possibility of a non-peripheral vertex group of degree 1 being contained in the adjacent peripheral group—although this restriction need not concern us for the moment.) We say that one peripheral splitting is a *refinement* of another if there is a colour-preserving folding of the first splitting onto the second. (For a more precise definition, see Section 2.) Obviously, any peripheral splitting is relative to \mathcal{G} and over subgroups of elements of \mathcal{G} . There is also a somewhat weaker converse to this statement; see Section 5.

Suppose now that Γ is hyperbolic relative to \mathcal{G} . We write $\partial\Gamma$ for its boundary. The peripheral subgroups in this case are precisely the maximal parabolic groups. We begin by recalling the following result from [Bo5]:

Proposition 1.1. *The boundary, $\partial\Gamma$, is disconnected if and only if Γ splits non-trivially over a finite group relative to \mathcal{G} . Moreover, each vertex group in any such splitting is hyperbolic relative to the elements of \mathcal{G} that it contains.*

(Note that the kind of splitting referred to in Proposition 1.1 will not, in general, be a peripheral splitting of Γ .)

There is also an accessibility result for splittings of this type [Bo5], which leads us naturally to restrict to the case of connected boundaries. Peripheral splittings represent the next simplest means of cutting up a relatively hyperbolic group. In particular, we shall show:

Theorem 1.2. *Suppose that $\partial\Gamma$ is connected. If Γ admits a non-trivial peripheral splitting, then $\partial\Gamma$ contains a global cut point.*

We note that Γ admits a non-trivial peripheral splitting if and only if Γ splits non-trivially relative to \mathcal{G} over a parabolic subgroup (i.e. an infinite subgroup of an element of \mathcal{G}). The fixed point of this parabolic subgroup will give us the global cut point.

There is a kind of converse to Theorem 1.2. Under certain constraints on the peripheral subgroups (as in Theorem 1.5), one can show that every global cut point is a parabolic fixed point, though this result is much deeper (see [Bo4]). This is another essential ingredient in the proof of local connectedness (see Theorem 1.5).

A *component* of a peripheral splitting is a non-peripheral vertex group. We show:

Theorem 1.3. *Suppose Γ is a relatively hyperbolic group (not necessarily with connected boundary). Suppose that H is a component of some peripheral splitting of Γ . Then, H is hyperbolic relative to the set of infinite groups of the form $H \cap G$, where G ranges over the peripheral subgroups of Γ . Moreover, if $\partial\Gamma$ is connected, then the boundary of any such component, H , is connected. In this case, if G is a peripheral subgroup, then $G \cap H$ is infinite if and only if G is adjacent to H in the Bass-Serre tree corresponding to the splitting.*

In fact, we shall see in Sections 7 and 8 how the boundary of $\partial\Gamma$ can be recovered by piecing together the boundaries of components in a treelike manner. As a consequence, we deduce that if the boundary of each component is (locally) connected, then the boundary of Γ is (connected).

The main specific result of this paper is the following accessibility result:

Theorem 1.4. *Suppose that Γ is relatively hyperbolic with connected boundary. Then Γ admits a (possibly trivial) peripheral splitting which is maximal in the sense that it is not a refinement of any other peripheral splitting.*

As a consequence, the boundaries of the components of such a maximal splitting do not contain any global cut point.

Now, in [Bo6], it is shown that if the boundary of a relatively hyperbolic group is connected and contains no global cut point, then it is locally connected (generalising the argument given in [BeM]). Gathering all these facts together as outlined in Section 9, we deduce [Bo6]:

Theorem 1.5. *Suppose that Γ is relatively hyperbolic, and that each peripheral subgroup is finitely presented, one- or two-ended, and contains no infinite torsion subgroup. If $\partial\Gamma$ is connected, then it is locally connected.*

The local connectedness has many potential consequences. To begin with, one can go on to recover the maximal peripheral splitting (as in Theorem 1.4) from the topology of $\partial\Gamma$ (Theorem 9.2). Beyond this, one might hope to obtain the JSJ splitting from studying local cut points as in the non-relative case [Bo2]. In [Gu], Guralnik has already begun an analysis of the local cut point structure under the assumption of local connectedness.

Theorems 1.2, 1.3 and 1.4, will all be proven in this paper, although we shall formulate them in terms of actions on sets, as mentioned in the introduction. This will be readily translated back into the more familiar language of group splittings used above (see Section 4 and [Bo5]).

2. EXAMPLES

In this section, we describe a few examples peripheral splittings to illustrate the results of this paper. Typical examples of relatively hyperbolic groups are geometrically finite groups acting on pinched Hadamard manifolds (see [Bo1]). In this case, the boundary can be naturally identified with the limit set [Bo5]. Moreover the peripheral subgroups are all finitely generated virtually nilpotent, and hence, in particular, either one-ended or two-ended. If the manifold is a hyperbolic n -space, we refer to such groups as “kleinian groups”. In this case, the peripheral subgroups are virtually abelian.

Suppose S is a closed orientable surface of genus at least 2. Consider a fuchsian representation of $\pi_1(S)$ into the isometry group of hyperbolic 3-space. The limit set is a round circle, and the quotient of each complementary disc gives us a Riemann surface structure on S . There is a well developed deformation theory for 3-dimensional kleinian groups which tells us, in particular, that we can deform this group by varying these two complex structures. Suppose that C is an essential simple closed curve on S . We hold one of the Riemann surface structures on S fixed, and in the other, we shrink C to a point. In the limit, we arrive at a geometrically finite representation of $\pi_1(S)$, where C corresponds to the conjugacy class of an (infinite cyclic) peripheral subgroup. The limit set can be described topologically as the quotient of the original circular limit set after identifying the pair of endpoints in every lift of C . It thus consists of a union of circles joined together in a treelike fashion, and compactified by adjoining the set of ideal points of this tree (cf. Section 7). Each parabolic point is a global cut point which connects two such circles. The stabiliser of each circle is a representation of a punctured surface group (namely the fundamental group of a component of $S \setminus C$), which one can arrange to be fuchsian by choosing a suitable deformation. The tree described above is the Bass-Serre tree of a peripheral splitting of $\pi_1(S)$, and the punctured surface subgroups are the components of the splitting.

One can arrive at the same example by starting with fuchsian representations of punctured surfaces, and gluing them together along parabolic subgroups. Applying a combination theorem, we can ensure that the result will be discrete and geometrically finite. Indeed, one can do the same thing one dimension higher, starting with finite covolume 3-dimensional kleinian groups, in place of fuchsian groups. In this way we construct a 4-dimensional geometrically finite kleinian group which admits a splitting over rank-2 abelian peripheral subgroups. The limit set will be a compactified tree of 2-spheres. In fact, one can perform this construction by taking an amalgamated free product of an arbitrary finite number, n , of copies of a finite covolume kleinian group amalgamated over a common maximal parabolic subgroup. In this case, there will be n -spheres meeting at each of the corresponding parabolic points.

For another example, start with a fuchsian representation, G , of a punctured surface group in hyperbolic 3-space. Let g be the generator of a parabolic subgroup of G , and let h be another parabolic isometry with the same fixed point. We choose h so that the parabolic group $\langle g, h \rangle$ corresponds to a rectangular torus. By taking this torus sufficiently elongated in the h -direction, we can arrange that the group $\langle G, h \rangle$ will be discrete and geometrically finite. We again have a peripheral subgroup which is rank-2 free abelian, although this time Γ splits over an infinite cyclic subgroup of this group. The components of the peripheral splitting will be

the conjugates of G . The limit set is a compactified tree of circles, with infinitely many circles meeting at each parabolic point.

Returning to the first example, suppose we choose another essential curve, C' , on S . By collapsing C and C' in the two different Riemann surfaces, we again get a geometrically finite representation of $\pi_1(S)$. If C and C' are disjoint, then this example is similar to the first. However, if C and C' intersect each other non-trivially, the resulting group will have as a limit set a kind of circle packing, where the parabolic points will be local but not global cut points. Abstractly, the group splits over each (conjugacy class of) peripheral subgroup, but not both simultaneously. Moreover, such splittings will not be relative to the class of peripheral subgroups. This group does not admit any non-trivial peripheral splitting.

Examples of the type described suggest various conjectures one might make regarding the role of local and global cut points in splittings of relatively hyperbolic groups. These might be amenable to analysis once one assumes local connectedness (cf. [Gu]). However, since one of our main objectives is to prove local connectedness, we cannot take that as an assumption in this paper.

3. PRETREES

In this section, we shall consider discrete pretree structures on a set V , and we give an alternative formulation of this in terms of “arboreal structures”. Such a structure will serve as a formal definition of a “peripheral splitting” when we introduce group actions.

Let V be a set. Suppose $\mathcal{T} \subseteq V \times V \times V$. We view \mathcal{T} as a ternary relation on V and write xyz to mean that $(x, y, z) \in \mathcal{T}$. Given $x, y \in V$, we write $(x, y) = \{z \in V \mid xzy\}$ and $[x, y] = (x, y) \cup \{x, y\}$. We refer to $[x, y]$ as a *closed interval*. The following notion was defined in [Bo3]:

Definition. (V, \mathcal{T}) is a *pretree* if it satisfies the following axioms for all $x, y, z \in T$:

- (T0): $[x, x] = \{x\}$,
- (T1): $[x, y] = [y, x]$,
- (T2): If $y \in (x, z)$, then $z \notin (x, y)$,
- (T3): $[x, y] \subseteq [x, z] \cup [z, y]$.

The intuitive interpretation of xyz is that y “lies strictly between” x and z . The axioms express the idea that this betweenness relation has a treelike structure. The axioms of a pretree have been explored previously. They can be found (under different names) in [W] and [AN].

We note that any interval $[x, y]$ is totally ordered by the relation \leq defined by $a \leq b$ if $a \in [x, b]$. A *median* of $x, y, z \in V$ is a point of $[x, y] \cap [y, z] \cap [z, x]$. If it exists, a median is unique.

Definition. A *discrete* pretree is one for which $[x, y]$ is finite for all $x, y \in V$.

A *median* pretree is one for which a median exists for any three points of V .

Median pretrees, under a variety of names, have been studied for some time (see for example [Sh]). Discrete median pretrees are even more familiar, as they are precisely the pretrees that arise as vertex sets of simplicial trees. In fact, if V is a discrete median pretree, then V has the structure of a \mathbf{Z} -tree, where the \mathbf{Z} -distance between x and y is defined as $\text{card}([x, y]) - 1$. Thus, \mathbf{Z} -trees, discrete median pretrees, and (vertex sets of) simplicial trees, are all reformulations of the same underlying structure.

Suppose V is a pretree. A subset, $F \subseteq V$ is *null and full* if there are no relations of the form xyz for any $x, y \in F$ and $z \in V$. It is a *star* if it is a maximal subset with this property. In other words, $F \subseteq V$ is a star if and only if it has the property that, if $z \in V$, then $z \notin F$ if and only if there exist $x, y \in F$ with xyz . Given a pretree, \mathcal{T} , we write $\mathcal{R}(\mathcal{T})$ for the set of stars in V .

Now, suppose \mathcal{R} is any set of subsets of a set V . We construct a bipartite graph, $T(\mathcal{R})$, with vertex set $V \sqcup \mathcal{R}$, by deeming $x \in V$ to be adjacent to $R \in \mathcal{R}$ if $x \in R$. Clearly, the degree of R in $T(\mathcal{R})$ is equal to $\text{card}(R)$.

Definition. We say that \mathcal{R} is an *arboreal structure* on V if every element of \mathcal{R} has cardinality at least 2, and $T(\mathcal{R})$ is a simplicial tree.

Note that it follows that $\bigcup \mathcal{R} = V$, and that the intersection of two distinct elements of \mathcal{R} is either empty or a singleton. Also, the clause that every element of \mathcal{R} must contain at least two points ensures that every terminal vertex of $T(\mathcal{R})$ lies in V .

We also see that $V \cup \mathcal{R}$ is a discrete median pretree, and so this induces a discrete pretree structure on V , which we denote by $\mathcal{T}(\mathcal{R})$.

Conversely, if \mathcal{T} is a discrete pretree structure on V , then it's not hard to verify that $\mathcal{R}(\mathcal{T})$ is an arboreal set. (This is a special case of the completion process of general pretrees discussed in [Bo3]. A variation of this can also be found in [AN].) Indeed, we see that $\mathcal{T}(\mathcal{R}(\mathcal{T})) = \mathcal{T}$. Moreover, if \mathcal{R} is any arboreal structure, then $\mathcal{R}(\mathcal{T}(\mathcal{R})) = \mathcal{R}$ (using the fact that all terminal elements of $T(\mathcal{R})$ lie in V). We conclude:

Lemma 3.1. *Given a set V , there is a natural bijective correspondence between discrete pretree structures on V and arboreal structures on V .*

From now on, we will pass freely between these two concepts.

Given pretree structures \mathcal{T} and \mathcal{T}' on V , we say that \mathcal{T}' is a *refinement* of \mathcal{T} , or that \mathcal{T} is *subordinate* to \mathcal{T}' if $\mathcal{T} \subseteq \mathcal{T}'$. We note:

Lemma 3.2. *If $\mathcal{T}, \mathcal{T}'$ are pretree structures on V which admit a common refinement, then $\mathcal{T} \cup \mathcal{T}'$ is a pretree structure. Moreover, if \mathcal{T} and \mathcal{T}' are discrete, then so is $\mathcal{T} \cup \mathcal{T}'$.*

Proof. Directly from the axioms. The existence of a common refinement is needed for axiom (T2). \square

This can be reinterpreted in terms of arboreal structures. Suppose \mathcal{R} and \mathcal{R}' are arboreal structures on V . We write $\mathcal{R} \leq \mathcal{R}'$ to mean that for all $R \in \mathcal{R}'$, there is some $S \in \mathcal{R}$ such that $R \subseteq S$. This is equivalent to the refinement relation on pretrees, i.e. $\mathcal{R} \leq \mathcal{R}'$ if and only if $\mathcal{T}(\mathcal{R}) \subseteq \mathcal{T}(\mathcal{R}')$. The “only if” statement is immediate. The “if” statement follows from the descriptions of \mathcal{R} and \mathcal{R}' as the maximal full and null subsets with respect to the pretree relations, $\mathcal{T}(\mathcal{R})$ and $\mathcal{T}(\mathcal{R}')$ respectively. Note that if \mathcal{R}' is a strict refinement of \mathcal{R} , then we can find some $R \in \mathcal{R}'$ which contains at least two elements of \mathcal{R} . Thus there exist $x, y, z \in R$ so that the relation yxz holds in $\mathcal{T}(\mathcal{R}')$. In fact, we can make a stronger statement:

Lemma 3.3. *Suppose that \mathcal{R} is an arboreal structure, and that \mathcal{T}' is a pretree structure on V (not necessarily discrete) such that $\mathcal{T}(\mathcal{R}) \subseteq \mathcal{T}'$. Suppose that $x, y, z \in V$ and that the relation yxz holds in \mathcal{T}' but not in $\mathcal{T}(\mathcal{R})$. Then there is some $R \in \mathcal{R}$ containing x , and $b, c \in R$ such that the relation bxc holds in \mathcal{T}' .*

Proof. Let $b \in V$ be the point adjacent to x in the interval $[x, y]$ in $\mathcal{T}(\mathcal{R})$. Thus b also lies in the interval $[x, y]$ in \mathcal{T}' . Since yxz holds in \mathcal{T}' , we deduce that bxz holds in \mathcal{T}' . However, bxz cannot hold in $\mathcal{T}(\mathcal{R})$, otherwise we could deduce the relation yxz .

Now let c be the point adjacent to x in the interval $[x, z]$ in $\mathcal{T}(\mathcal{R})$. By the same argument (with b replacing z and c replacing b), we deduce that bxz holds in \mathcal{T}' but not in $\mathcal{T}(\mathcal{R})$. Now, $\{b, x, c\}$ is null and full in $\mathcal{T}(\mathcal{R})$ and hence lies in some $R \in \mathcal{R}$. □

Given two arboreal structures, \mathcal{R} and \mathcal{R}' , we write $\mathcal{R} \vee \mathcal{R}' = \{R \cap R' \mid R \in \mathcal{R}, R' \in \mathcal{R}'\} \setminus (\{\{x\} \mid x \in V\} \cup \{\emptyset\})$. If \mathcal{R} and \mathcal{R}' admit a common refinement, then $\mathcal{R} \vee \mathcal{R}'$ is arboreal, and $\mathcal{T}(\mathcal{R} \vee \mathcal{R}') = \mathcal{T}(\mathcal{R}) \cup \mathcal{T}(\mathcal{R}')$.

Familiar examples of arboreal structures arise in elementary graph theory. Suppose that K is a connected graph with vertex set V . A *block* is a maximal 2-vertex-connected subgraph of K (where we view a single edge as 2-vertex-connected). By a *block set*, we mean the set of vertices in some block. Let $\mathcal{B}(K)$ be the set of block sets of K . An elementary result of graph theory tells us that $\mathcal{B}(K)$ is an arboreal structure on V . The corresponding discrete pretree relation on V is the obvious one— xyz holds if y separates x from z in K .

More generally, suppose K is a connected graph. Given a subset $A \subseteq V$, we write $K(A)$ for the full subgraph of K on vertex set A . Suppose \mathcal{R} is an arboreal structure on V . We say that K *respects* \mathcal{R} if K is a union of the subgraphs $K(R)$ as R ranges over \mathcal{R} . This is equivalent to saying that $\mathcal{B}(K)$ is a refinement of \mathcal{R} .

Lemma 3.4. *Suppose that K is a graph with vertex set V , which respects an arboreal structure, \mathcal{R} . Then K is connected if and only if $K(R)$ is connected for all $R \in \mathcal{R}$.*

Proof. Suppose $K(R)$ is connected for each R . Given $x, y \in V$, let $x = x_0, x_1, \dots, x_n = y$ be the pretree interval $[x, y]$ in the natural order, in the corresponding pretree structure. Now, for each i , there is some $R \in \mathcal{R}$ such that $x_i, x_{i+1} \in R$. We connect x_i to x_{i+1} by a path in $K(R)$. Concatenating these paths gives us a path from x to y in K .

Conversely, suppose K is connected. Suppose $x, y \in R \in \mathcal{R}$. Let $x = x_0, x_1, \dots, x_n = y$ be a path connecting x to y in K . Let $x = y_0, y_1, \dots, y_m = y$ be the sequence obtained by deleting those x_i which do not lie in R . Now, if y_i and y_{i+1} are not consecutive in the sequence $(x_i)_i$, then it is easily verified that $y_i = y_{i+1}$. Thus, for all i , we see that y_i and y_{i+1} are either equal or adjacent in $K(R)$. We have thus connected x to y by a path in $K(R)$. (An alternative argument can be given along the lines of Lemma 5.3.) □

We end this section with an observation about constructing simplicial trees needed in Section 5.

Suppose T is a simplicial tree with vertex set $V(T)$. Let \mathcal{A} be a set of subtrees of T with the property that each edge of T lies in a unique element of \mathcal{A} . We construct a bipartite graph, $S(\mathcal{A})$, with vertex set $V(T) \sqcup \mathcal{A}$, where we deem $x \in V$ to be adjacent to $\tau \in \mathcal{A}$ if $x \in \tau$. We note:

Lemma 3.5. *$S(\mathcal{A})$ is simplicial tree.*

Proof. It is clear that $S(\mathcal{A})$ is connected. Suppose that $x_1\tau_1x_2\tau_2\dots x_n\tau_n$ is a circuit in $S(\mathcal{A})$, where $x_i \in V$ and $\tau_i \in \mathcal{A}$. Let α_i be the arc in τ_i connecting x_i to x_{i+1} (taking subscripts mod n). We see that $\alpha_1 \cup \alpha_2 \cup \dots \cup \alpha_n$ is a circuit in T . \square

4. PERIPHERAL SPLITTINGS

We now introduce group actions into the picture. We use the terminology of group actions on sets, as described in more detail in [Bo5].

Let Γ be a group. A Γ -set is a set, V , together with a Γ -action on V . We refer to the points of V as *vertices*. We say that V is *cofinite* if V/Γ is finite. If $x \in V$, we write $\Gamma(x)$ for its stabiliser. If $x, y \in V$ with $x \neq y$, we refer to the group $\Gamma(x) \cap \Gamma(y)$ as a *pair stabiliser*. By a (Γ, V) -graph, K , we mean a connected Γ -invariant graph with vertex set V , and with finitely many Γ -orbits of edges. We do not allow loops or multiple edges (unlike a graph of groups). We write $E(K)$ for the edge set of K . We say that V is *connected* (or *0-connected*) if it admits a (Γ, V) -graph. Clearly, a connected Γ -set is cofinite. We say that a Γ -set is *doubly connected* if it admits a 2-vertex-connected (Γ, V) -graph.

The following is easily verified [Bo5].

Lemma 4.1. *Suppose that V is a cofinite Γ -set, and that $W \subseteq V$ is a Γ -invariant subset. If W is connected and non-empty, then V is connected. Conversely, if V is connected, and W contains every point of V with an infinite stabiliser, then W is connected.*

We remark that Γ may itself be viewed as a Γ -set under left multiplication. In this case, Γ is 0-connected if and only if it is finitely generated. (In Section 4, we define a notion of 1-connectedness which corresponds to finite presentability.) In this set-up, a (Γ, V) -graph plays the role of a Cayley graph. Another example of a Γ -set, more directly relevant to the ideas of this paper, is obtained by considering the action of Γ by conjugation on some preferred set of subgroups of Γ . (Such subgroups are termed “peripheral” in [Bo5], in reference to their appearance as maximal parabolic subgroups of relatively hyperbolic groups. This accounts for some of the terminology of this section.)

We return to general setting of a Γ -set V .

Definition. A *peripheral splitting* of V is a Γ -invariant arboreal structure on V .

As discussed in Section 1, this is essentially the same as a Γ -invariant discrete pretree structure on V . We can adopt the terminology of “refinements” in reference to peripheral splittings.

Suppose that \mathcal{R} is a peripheral splitting of V . We can view \mathcal{R} , itself, as a Γ -set. We say that \mathcal{R} is *cofinite* if \mathcal{R}/Γ is finite. If $R \in \mathcal{R}$ we write $\Gamma(R)$ for the (setwise) stabiliser of R . Thus, R is a $\Gamma(R)$ -set. As such, we refer to it as a *component* of the peripheral splitting.

Lemma 4.2. *Suppose that \mathcal{R} is a cofinite peripheral splitting of a Γ -set, V . If each component of \mathcal{R} is connected, then V is connected.*

Proof. Given $R \in \mathcal{R}$, we can find a $(\Gamma(R), R)$ -graph, $K(R)$. We choose these graphs equivariantly with respect to the action of Γ . Let K be the graph with vertex set V and edge set $E(K) = \bigcup_{R \in \mathcal{R}} E(K(R))$. We see easily, using Lemma 3.4, that K is a (Γ, V) -graph. \square

In fact, the construction of Lemma 4.2 gives us a (Γ, V) -graph which respects the splitting (in the sense defined in Section 3).

We have the following converse of Lemma 4.2:

Lemma 4.3. *Suppose that V is a connected Γ -set, and \mathcal{R} is a peripheral splitting of V . Then \mathcal{R} is cofinite, and each component of \mathcal{R} is connected.*

Proof. Let K be a (Γ, V) -graph. Suppose $e \in E(K)$. Let $x, y \in V$ be the endpoints of e , and let $[x, y]$ be the closed pretree interval in the discrete pretree structure on V associated with \mathcal{R} . Now, $[x, y]$ is finite, and has a natural linear order $x = x_0, x_1, \dots, x_n = y$ (so that $x_i x_j x_k$ holds whenever $i < j < k$). Note that for each i , x_i and x_{i+1} lie in some component of \mathcal{R} . Let $E(e)$ be the set of edges $\{x_0 x_1, x_1 x_2, \dots, x_{n-1} x_n\}$. We now let L be the graph with vertex set V and edge set $E(L) = \bigcup_{e \in E(K)} E(e)$. Clearly, L is a (Γ, V) -graph which respects the peripheral splitting, \mathcal{R} . By Lemma 3.4, $L(R)$ is connected for all $R \in \mathcal{R}$. It is easily checked that $E(L(R))/\Gamma(R)$ is finite. Thus, $L(R)$ is a $(\Gamma(R), R)$ -graph, so R is connected as claimed.

Now, no two distinct subgraphs $L(R)$ can share an edge. Since $E(L)/\Gamma$ is finite, we deduce that \mathcal{R}/Γ is finite. □

Recall that $T(\mathcal{R})$ is a bipartite simplicial tree, with vertex set $V \sqcup \mathcal{R}$. We may form the quotient graph $T(\mathcal{R})/\Gamma$. We note:

Lemma 4.4. *If \mathcal{R} is a peripheral splitting of a connected Γ -set, V , then $T(\mathcal{R})/\Gamma$ is finite.*

Proof. We just need to show that $E(T(\mathcal{R}))/\Gamma$ is finite. Let K be a (Γ, V) -graph respecting the splitting. Now, if $x \in R \in \mathcal{R}$, we can find an edge, $e \in K(R)$, with endpoint at x (since $K(R)$ is connected). If we view it as a directed edge, then it determines the pair (x, R) uniquely. Since there are only finitely many directed edges of K up to the action of Γ , there are only finitely many such pairs, and hence finitely many edges of $T(\mathcal{R})$ up to the action of Γ . □

Recall that a Γ -set, V , is “doubly connected” if it admits a 2-vertex-connected (Γ, V) -graph.

Definition. A peripheral splitting of a connected Γ -set is *full* if every component is doubly connected.

Thus, \mathcal{R} is a full peripheral splitting if and only if V admits a (Γ, V) -graph, K , such that $\mathcal{R} = \mathcal{B}(K)$, where $\mathcal{B}(K)$ is the set of block sets of K . We note:

Lemma 4.5. *Any peripheral splitting of a connected Γ -set admits a full refinement.*

Proof. Let \mathcal{R} be a peripheral splitting of V . By the construction of Lemma 4.2, there is a (Γ, V) -graph, K , which respects \mathcal{R} . Now, $\mathcal{B}(K)$ is a refinement of \mathcal{R} . □

We say that a peripheral splitting is *maximal* if it admits no strict refinement (as a peripheral splitting). An immediate consequence of Lemma 4.5 is that any maximal peripheral splitting is full.

By the *minimal* splitting of a Γ -set, V , we mean the splitting $\{V\}$. (In other words, it is minimal with respect to refinement, and has minimal complexity. The term “trivial” has traditionally had a broader meaning in the terminology of group splittings, as we discuss shortly.) We note:

Lemma 4.6. *Each component of a maximal peripheral splitting admits no non-trivial peripheral splitting.*

Proof. Suppose \mathcal{R} is a peripheral splitting of a Γ -set, V . Suppose $R \in \mathcal{R}$, and that \mathcal{S} is a non-trivial peripheral splitting of R as a $\Gamma(R)$ -set. If $\gamma \in \Gamma$, we write $\gamma\mathcal{S} = \{\gamma S \mid S \in \mathcal{S}\}$. Thus, $\gamma\mathcal{S}$ is a peripheral splitting of the $\Gamma(\gamma R)$ -set γR . Also, if $\gamma \in \Gamma(R)$, then $\gamma\mathcal{S} = \mathcal{S}$. We now construct a peripheral splitting, \mathcal{R}' , of V , by replacing each $\{\gamma R\} \subseteq \mathcal{R}$ by $\gamma\mathcal{S}$. In other words, $\mathcal{R}' = \mathcal{R} \cup \{\gamma S \mid \gamma \in \Gamma, S \in \mathcal{S}\} \setminus \{\gamma R \mid \gamma \in \Gamma\}$. It is easily verified that $T(\mathcal{R}')$ is a tree, so that \mathcal{R}' is indeed a peripheral splitting of V . It is clearly a strict refinement of \mathcal{R} , so \mathcal{R} cannot be maximal. \square

In more familiar terms, a peripheral splitting of V can be viewed as a presentation of Γ as the fundamental group of a finite bipartite graph of groups, where all the vertex groups of one colour are determined—as the conjugacy representatives of the vertex stabilisers of V . More precisely, the graph arises as the quotient $T(\mathcal{R})/\Gamma$, as described by Lemma 4.4. This has vertex set $(V/\Gamma) \sqcup (\mathcal{R}/\Gamma)$. The vertex groups are thus (the conjugacy classes of) groups of the form $\Gamma(x)$ and $\Gamma(R)$ for $x \in V$ and $R \in \mathcal{R}$. The vertex groups corresponding to the elements of V are called *peripheral*. Note that we must rule out the possibility of a non-peripheral vertex group of degree 1 being equal to the adjacent edge group. Such a group would correspond to an element of \mathcal{R} which is terminal in the tree, $T(\mathcal{R})$, and hence a singleton—a situation we have explicitly ruled out.

Suppose that \mathcal{R}' is a refinement of \mathcal{R} . We get a natural map, f , from \mathcal{R}' to \mathcal{R} , so that $R \subseteq f(R')$ for all $R' \in \mathcal{R}'$. Defining f to be the identity on V , we get a Γ -equivariant map from $T(\mathcal{R}')$ to $T(\mathcal{R})$ which sends edges to edges. Thus, f is a “folding” of the tree $T(\mathcal{R}')$ (see, for example, [Du4]). This descends to a folding on the level of graphs of groups. In other words, we see that passing to a refinement can be interpreted as an unfolding. Of course, these are a particular kind of unfoldings, in that, one set of vertex groups are held fixed. We shall consider sequences of such unfoldings in Section 6.

The minimal peripheral splitting $\{V\}$ is obviously the unique splitting which does not refine any other, hence the terminology. If V is a singleton, then the corresponding graph of groups is also a singleton with no edges. Otherwise, the corresponding graph of groups consists of one “central” vertex group equal to Γ which is connected to a finite number of peripheral groups, each by a single edge group equal to the peripheral group. Following the (somewhat unfortunate) terminology of group splittings, we shall say that a peripheral splitting, \mathcal{R} , is *trivial* if some element of V or of \mathcal{R} is preserved by Γ . This is more general than minimal. However, if a trivial splitting fails to be minimal, then some point stabiliser of V must be contained in another point stabiliser—a situation one can rule out in many cases.

5. RELATIVE SPLITTINGS

In this section, we describe how peripheral splittings are related to other kinds of relative splittings. A general reference on group splittings is [DiD]. The terminology we use here is described in Section 1.

Suppose Γ is a group, and \mathcal{G} a conjugacy invariant collection of subgroups of Γ , consisting of finitely many conjugacy classes. We shall suppose that each element

of \mathcal{G} is equal to its normaliser in Γ , and that no element of \mathcal{G} is a subgroup of a different element of \mathcal{G} . Now Γ acts on \mathcal{G} by conjugacy. Thought of as a Γ -set in this way, we shall denote \mathcal{G} by V . In other words, the map $[x \mapsto \Gamma(x)]$ gives us the identification of V with \mathcal{G} . Note that under the conditions we have imposed, the notions of minimal and trivial splittings coincide.

Proposition 5.1. *The Γ -set, V , admits a non-trivial peripheral splitting if and only if Γ splits non-trivially relative to \mathcal{G} over a subgroup of an element of \mathcal{G} .*

Proof. A peripheral splitting of V corresponds to a presentation of Γ as a bipartite graph of groups in the manner described in Section 3. Such a splitting clearly has the required properties.

For the converse, we can suppose that the graph of groups has just one edge, i.e. is either an amalgamated free product or an HNN extension. The two cases are similar, so for simplicity of exposition, we will assume the former. Thus, Γ can be expressed as an amalgamated free product of subgroups K and K' over H . By assumption, H is a subgroup of some $G \in \mathcal{G}$ which we can, in turn, take to be a subgroup of K . We now represent K (trivially) as a graph of groups, consisting of a central vertex with group K , and additional vertices corresponding to those (conjugacy classes of) peripheral subgroups which lie in K . Each of these additional vertices is joined to the central vertex by an edge whose edge group equals the peripheral subgroup. (If it happens that $K = G$, then we just leave it alone.) We perform the same construction for K' . We now construct a bipartite graph of groups by connecting the vertex corresponding to G in the first graph (representing K) to the central vertex of the second graph (representing K') by an edge with group H . This gives rise to a proper peripheral splitting of V . \square

We note that the construction of the first paragraph is canonical, whereas, that of the second paragraph is somewhat artificial; for example, the edge group H might be a subgroup of more than one peripheral subgroup. This is one reason for preferring the former formulation.

In the case of relatively hyperbolic groups, for example, one has further constraints on the set, \mathcal{G} . We shall say that \mathcal{G} is a *peripheral structure* on Γ if (as before) it consists of finitely many conjugacy classes and each element of \mathcal{G} is equal to its normaliser, and if, in addition, every element of \mathcal{G} is infinite, and each pair of distinct elements of \mathcal{G} intersect in a finite group.

We remark that if \mathcal{G} is a peripheral structure, and if \mathcal{G} does not split over any finite subgroup relative to \mathcal{G} , then the correspondence between peripheral splittings and of relative splitting of Γ over subgroups of peripheral subgroups can be made more canonical. Note that any edge group in a splitting of the latter type must be infinite, and hence contained in a unique peripheral subgroup. Now the construction of the second paragraph of Proposition 5.1 can be applied to any graph of groups of this type, and in this case does not involve any arbitrary choices.

We shall say that a group is *one-ended* if it does not split non-trivially over any finite subgroup. In view of Stallings's Theorem, this coincides with the usual notion for finitely generated groups.

Proposition 5.2. *Suppose that Γ is a group, and \mathcal{G} is a peripheral structure with every peripheral subgroup one-ended. If Γ splits over a subgroup of a peripheral subgroup, then it splits relative to \mathcal{G} over a subgroup of a peripheral subgroup.*

Proof. Let T be a simplicial tree with a minimal action of Γ , with no edge inversions and with one orbit of edges, such that some (hence every) edge stabiliser of T is a subgroup of some element of \mathcal{G} . Let $e \in E(T)$, and let $H = \Gamma(e)$ be the edge stabiliser. Thus, $H \leq G$ for some $G \in \mathcal{G}$.

Suppose $G' \in \mathcal{G}$ is not a conjugate of G . We claim that G' fixes a vertex of T . To see this, note that every edge stabiliser of T is a conjugate of H and hence contained in a conjugate of G . It must therefore meet G' in a finite group. Since G' , by hypothesis, does not split over a finite group, its action on T must be “trivial” in the sense that it must fix a vertex.

Now, if G fixes a vertex of T , we are done. If not, let τ be the (unique) minimal G -invariant subtree of T . Suppose that $e \in E(g\tau)$ for some $g \in \Gamma$. Now $\Gamma(g^{-1}e) = g^{-1}Hg$. Since $g^{-1}e \in E(\tau)$ and G is one-ended, we see that $G \cap g^{-1}Hg$ is infinite. Since $H \leq G$, we get that $G \cap g^{-1}Gg$ is infinite. Thus, $G = g^{-1}Gg$, so $g \in G$ (since G is assumed to be equal to its normaliser). In particular, we see that $g\tau = \tau$. Now, since τ is non-trivial, we can certainly find some $g \in \Gamma$ so that $g^{-1}e \in E(\tau)$, so we conclude that $e \in E(g\tau)$. We also note that if $h \in \Gamma$ with $h\tau = \tau$, then $he \in E(\tau)$, and so $h \in G$. This shows that G is precisely the setwise stabiliser of τ .

Now, let \mathcal{A} be the set of Γ -images of τ . The argument of the last paragraph shows that each edge of T is contained in a unique element of \mathcal{A} . Thus, \mathcal{A} satisfies the hypotheses of Lemma 3.5, so we get a minimal action of Γ on the bipartite simplicial tree $S(\mathcal{A})$. Now, each element of \mathcal{G} fixes a vertex of $S(\mathcal{A})$. In fact, since G is precisely a vertex stabiliser, we see that Γ splits over a subgroup of G . Also, since τ is non-trivial, the corresponding vertex in $S(\mathcal{A})$ is incident to infinitely many edges, so the splitting is non-trivial. \square

The constraint on peripheral subgroups in Proposition 5.2 is fairly natural in the context of geometrically finite groups acting on pinched Hadamard manifolds, where all the peripheral subgroups are finitely generated and one- or two-ended. With a bit more work, one can adapt the argument of Proposition 5.2 to deal with the case where all the peripheral subgroups are one-ended, except for at most one conjugacy class of two-ended subgroups. However, in general, the result fails if we allow more than one conjugacy class of two-ended peripheral subgroups, as the final example in Section 2 shows. Indeed the existence of examples of this type causes some amount of complication in the subject (cf. [Bo4]).

The relative splittings we have described in this section have been of a particular type. However, it is possible to describe all relative splittings in this language. Suppose Γ is a group and V is a cofinite Γ -set. Suppose that Γ splits relative to $\{\Gamma(x) \mid x \in V\}$. This gives rise to an action of Γ on a simplicial tree, T , such that, for each $x \in V$, $\Gamma(x)$ fixes a vertex of T . We therefore get a Γ -equivariant map from V to $V(T)$ (though this map need not be canonical without further hypotheses). Let \mathcal{S} be the set of non-empty preimages of vertices of T under this map. Thus, \mathcal{S} is a Γ -invariant partition of V . Moreover, it carries the structure of a discrete pretree, arising from its embedding in T . In other words, viewed as a Γ -set, \mathcal{S} admits a peripheral splitting. One can easily invert the above process (though without further hypotheses, one need not in general recover the original splitting). We therefore see that one can view relative splittings of Γ in terms of peripheral splittings of invariant partitions of V . Thus, the machinery developed here can be applied more generally. As an example, we note that it is possible to reprove Proposition 1.1 by adapting the arguments of Theorems 1.2 and 1.3, though we shall not describe this here.

It is worth remarking that a splitting of this type can often be recognised geometrically. We say that a connected graph, K , is *one-ended* if, when we remove any finite set of edges, precisely one of the complementary components is unbounded. (Here we can equivalently interpret “unbounded” to mean with respect to either the intrinsic path-metric, or the metric on K .)

Definition. We say that a connected Γ -set with finite pair stabilisers is *one-ended* if every (Γ, V) -graph is one-ended.

It is not hard to see that it is sufficient that one (Γ, V) -graph be one-ended. In fact, the following is a consequence of the main result of [Du1]:

Theorem 5.3. *Suppose V is a connected Γ -set with finite pair stabilisers. Then, Γ is one-ended if and only if Γ does not split non-trivially over any finite subgroup relative to $\{\Gamma(x) \mid x \in V\}$.*

In the case of relatively hyperbolic groups, one can give a simpler proof using the arguments of [Du2]; see [Bo5].

6. ACCESSIBILITY

In this section, we prove an accessibility result for peripheral splittings of groups. We give a condition which ensures that a sequence of increasingly refined peripheral splittings must stabilise. This can also be interpreted in terms of unfolding bipartite graphs of groups. It fits into a general class of questions which ask: when do chains of unfoldings have upper bounds. There are a number of results of this type in the literature, although the version we need here does not seem to follow directly from these. It does, however, yield to methods of the type introduced in [Du2].

Let V be a connected Γ -set. In this section, it will be convenient to think of a peripheral splitting of V as a discrete Γ -invariant pretree structure, $\mathcal{T} \subseteq V \times V \times V$ on V . Thus \mathcal{T}' is a “refinement” of \mathcal{T} if $\mathcal{T} \subseteq \mathcal{T}'$.

Suppose we have an infinite sequence of increasingly refined peripheral splittings of V , say, $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \mathcal{T}_3 \subseteq \dots$. Their union, $\mathcal{T}_\infty = \bigcup_{n=1}^\infty \mathcal{T}_n$, is a Γ -invariant pretree structure on V , although a priori there is no reason to suppose it will be discrete. We say that $(\mathcal{T}_n)_{n \in \mathbb{N}}$ *stabilises* if $\mathcal{T}_n = \mathcal{T}_\infty$ for all sufficiently large n .

It seems to be natural to ask when such a sequence must stabilise, or when the limit \mathcal{T}_∞ must be discrete, or indeed whether every refinement of a particular splitting must be discrete. A large class of examples of finitely generated groups acting on non-discrete protrees and hence also pretrees are described in [Du3]. Here we shall confine ourselves to proving a positive result, namely Theorem 6.1.

As mentioned earlier, we can interpret refinement in terms of unfolding. As discussed in Section 3, if $n \geq m$, then there is a Γ -invariant folding map, $f : T_n \rightarrow T_m$, where T_i is the bipartite simplicial tree associated to \mathcal{T}_i . This descends to a folding of graphs of groups, $f : T_n/\Gamma \rightarrow T_m/\Gamma$. In this folding, the groups associated to V/Γ are held fixed. We remark that one can give a simple argument, via Grushko’s Theorem, to bound the complexity of the graphs T_n/Γ , under the assumption that Γ is finitely generated. In this case, it follows that the combinatorial types of the graphs T_n/Γ must stabilise. However, there are folding moves which do not alter the combinatorial type of the underlying graph (see [Du2], [BeF]), so this is not sufficient for our purposes.

To state the main result, we need another definition. Recall that the notion of “connectedness” (or “0-connectedness”) is analogous to finite generation for groups.




FIGURE 1.

We want another notion which is analogous to finite presentability. By a *2-complex* we mean a 2-dimensional simplicial complex. It is *1-connected* if it is connected and simply connected.

Definition. We say that a Γ -set, V , is *1-connected* if it can be represented as the vertex set of a 1-connected 2-complex, Σ , such that Σ/Γ is finite.

By saying that Σ/Γ is finite, we mean that there are finitely many orbits of edges and 2-simplices. To get a well-defined quotient it may be necessary to pass to the first barycentric subdivision, Σ' . We can think of Σ'/Γ as an orbihedron with fundamental group Γ , as discussed in [H].

Note that the 1-skeleton, K , of Σ is a (Γ, V) -graph (so that 1-connected implies 0-connected).

We shall only be concerned here with 1-connected Γ -sets which have finite pair stabilisers. (In the case of relatively hyperbolic groups, Σ will be locally finite away from V , though we need not assume that here.)

Suppose that V is connected with finite pair stabilisers. Recall that V is one-ended if and only if Γ does not split non-trivially over any finite group relative to $\{\Gamma(x) \mid x \in V\}$ (see Theorem 5.3).

We aim to prove:

Theorem 6.1. *Suppose V is a 1-connected Γ -set with finite pair stabilisers, and that V is one-ended. Then any sequence of increasingly refined peripheral splittings of V must stabilise.*

The argument will be in terms of tracks on our 2-complex, Σ . A “track” is a connected component of a pattern. A “pattern” is a closed subset of a 2-complex which meets every 1-simplex in a finite set of points, and meets every 2-simplex in one of a number of specific combinatorial possibilities; see Figure 1.

This is slightly more general than a “pattern” as defined in [Du2]. Specifically, we are allowing for a track to pass through vertices of Σ and for it to branch into the interior of a 2-simplex. Thus, a component of an intersection of a pattern with a 2-simplex might be a vertex of the simplex, an interval connecting interior points of two distinct edges of the simplex, an interval connecting a vertex of the simplex to an interior point on the opposite edge, or a tripod which meets each edge in a single interior point. Note that a track is an embedded graph. We can give a more formal presentation of this in terms of pretrees.

Suppose V satisfies the hypotheses of Theorem 6.1. We fix a 1-connected 2-complex, as given by the definition. Let K be the 1-skeleton of Σ . We write $C(\Sigma)$

for the set of 2-simplices of Σ . Thus, $E(\Sigma)/\Gamma$ and $C(\Sigma)/\Gamma$ are both finite. Given $\sigma \in C(\Sigma)$, we write $E(\sigma) \subseteq E(\Sigma)$ for the set of edges of σ .

Suppose that $\mathcal{T} \subseteq V \times V \times V$ is a Γ -invariant pretree relation on V . Suppose that the action of Γ on \mathcal{T} is *non-nesting*, i.e., no element of Γ sends any closed interval of \mathcal{T} into a proper subset of itself. (This is automatically true for any Γ -invariant discrete pretree or any increasing union of such.)

Given any edge, $e \in E(\Sigma)$ with endpoints $x, y \in V$, we choose an order-preserving embedding of the pretree interval $[x, y]$ in e , sending x to x and y to y . We write $\Pi(e) \subseteq e$ for the image of this embedding, and $\pi_e : \Pi(e) \rightarrow [x, y]$ for its identification with $[x, y]$. Using the non-nesting assumption, we can perform this construction Γ -equivariantly for all $e \in E(\Sigma)$. Let $\Pi = \bigcup_{e \in E(\Sigma)} \Pi(e)$, and let $\pi : \Pi \rightarrow V$ be the union of all the maps π_e . Note that $V \subseteq \Pi$, and that π restricts to the identity on V .

Now suppose that $\sigma \in C(\Sigma)$. We define a relation, \sim_σ , on Π as follows. Given $x, y \in \Pi$, we write $x \sim_\sigma y$ to mean that either $x = y$ or that there exist $e, e' \in E(\sigma)$ such that $x \in \Pi(e)$, $y \in \Pi(e')$ and $\pi(x) = \pi(y)$. Let \sim_σ be the equivalence relation generated by the union of the relations \sim_σ as σ ranges over $C(\Sigma)$. We write $\Phi = \Pi/\sim$. Note that π induces a natural map from Φ to V , which we also denote by π . The inclusion of V in Π , induces an inclusion, $j : V \hookrightarrow \Phi$, such that $\pi \circ j$ is the identity.

In the case where \mathcal{T} is discrete, we can interpret all this in terms of patterns on Σ . Suppose $\sigma \in C(\Sigma)$. Let $E(\sigma) = \{e_1, e_2, e_3\}$ and $n_i = \text{card}(\Pi(e_i))$. Now the numbers n_1, n_2 and n_3 are all at least 2, and satisfy strict triangle inequalities ($n_i < n_j + n_k$ for i, j, k distinct). They thus determine precisely one of the diagrams described by Figure 1, i.e., a 1-dimensional subset of σ determined up to isotopy relative to Π . The union of Π , together with these subsets, as σ ranges over $C(\Sigma)$, gives us a *pattern* on Σ . A *track* is a connected component of this pattern. We see that if $x, y \in \Pi$, then $x \sim y$ if and only if they lie in the same track. We can thus identify $\Phi = \Pi/\sim$ with the set of tracks on Σ .

We note that the pattern descends to a pattern on the orbihedron Σ'/Γ , where Σ' is the first barycentric subdivision of Σ .

We also remark that, since Σ is simply connected, the set of tracks on Σ separate Σ in a treelike fashion. In particular, they induce a discrete pretree relation on Φ . The pretree relation thus induced on V under the embedding $j : V \hookrightarrow \Phi$ is precisely \mathcal{T} .

We now return to our objective of proving Theorem 6.1.

Suppose $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \mathcal{T}_3 \subseteq \dots$ is an increasing sequence of peripheral splittings of V . Let $\mathcal{T}_\infty = \bigcup \mathcal{T}_n$. Suppose $\sigma \in C(\Sigma)$ has vertices $x, y, z \in V$. Now the points x, y, z may or may not have a median in the pretree structure \mathcal{T}_∞ . Let $C_0(\Sigma) \subseteq C(\Sigma)$ be the set of 2-simplices σ so that such a median exists. Now if $\sigma \in C_0(\Sigma)$, this median exists in the structure \mathcal{T}_n for all sufficiently large n . Thus (since $C(\Sigma)/\Gamma$ is finite) we can suppose that it exists for all n .

Now, given $n \in \mathbf{N} \cup \{\infty\}$ and $\sigma \in C(\Sigma)$, we write $\sim_{\sigma, n}$ for the relation \sim_σ defined in terms of the pretree structure \mathcal{T}_n . We write \sim_n for the equivalence relation generated by the $\sim_{\sigma, n}$. We shall denote the pretree interval between a and b in \mathcal{T}_n by $[a, b]_n$. We write $\Phi_n = \Pi_n/\sim_n$ for the spaces constructed earlier.

Lemma 6.2. *Suppose $x, y \in \Pi_\infty$ with $x \sim_\infty y$. If $x \in \Pi_n$ for $n \in \mathbf{N}$, then $y \in \Pi_n$ and $x \sim_n y$.*

Proof. It is sufficient to verify that if $\sigma \in C(\Sigma)$ and $x, y \in \Pi_\infty$ with $x \sim_{\sigma, \infty} y$, then $y \in \Pi_n$ and $x \sim_{\sigma, n} y$. Let the vertices of σ be a, b, c , so that $E(\sigma) = \{ab, bc, ca\}$. Without loss of generality, we can suppose that $x \in ab$ and $y \in ac$. Thus, $\pi(x) \in [a, b]_n \subseteq [a, c]_n \cup [b, c]_n$. It follows that there is some $z \in \Pi_n(ac) \cup \Pi_n(bc)$ such that $x \sim_{\sigma, n} z$. If x is the median of $\{a, b, c\}$ in \mathcal{T}_∞ , then by hypothesis, it is the median $\{a, b, c\}$ in \mathcal{T}_n . Thus, $y \in \Pi_n$ and $x \sim_{\sigma, n} y$. Otherwise, we must have $y = z$ and again the result follows. \square

We can thus identify each pattern, Φ_n , as a subset of Φ_∞ . In other words, we have an increasing sequence, $\Phi_1 \subseteq \Phi_2 \subseteq \Phi_3 \subseteq \dots$ of patterns on Σ . There is a natural Γ -equivariant map, $\pi : \Phi \rightarrow V$. If $t \in \Phi_\infty$, we write $\Gamma(t)$ for the stabiliser of t (in other words, the setwise stabiliser of t , thought of as a track on Σ). Clearly $\Gamma(t) \subseteq \Gamma(\pi t)$.

Let Σ' be the first barycentric subdivision of Σ . We can view Σ'/Γ as a finite orbihedron. If $t \in \Phi_\infty$, then t projects to a 1-complex, $s(t)$, embedded in Σ'/Γ . We can assume that $s(t)$ meets each simplex of Σ' in one of the combinatorial pictures described earlier, except that we need to allow for the possibility of edges of Σ' to be included in $s(t)$. However, this possibility can only arise for finitely many t .

Suppose that the sequence $(\Phi_n)_n$ does not stabilise. We can then find a sequence, $(t_i)_{i \in \mathbb{N}}$ of distinct tracks in Φ_∞ , which project to disjoint 1-complexes, $s_i = s(t_i)$, in Σ'/Γ . Since Σ'/Γ is a finite complex, we can assume that no s_i contains an edge of Σ'/Γ , so that each s_i is a track in the sense described earlier (Figure 1). In fact, we can eliminate pictures (2) and (3) (where s_i passes through a vertex, or branches into the interior of a 2-simplex). We are thus left with picture (1), so that s_i is a track in the traditional sense [Du2].

We now apply the standard Kneser-Dunwoody argument. Recall that a *band* in a simplicial 2-complex is a “thickened-up track”, or more precisely, a compact connected subset disjoint from the vertex set which meets each 1-simplex in a disjoint union of intervals, and each 2-simplex in a disjoint union of rectangles, each such rectangle having two opposite sides in different faces of the 2-simplex. A band is *twisted* if its boundary is connected. Two disjoint tracks are *parallel* if their union is the boundary of an (untwisted) band. If we have a finite number, n , of disjoint tracks embedded in a finite 2-complex, then all but a bounded number of complementary components must be bands. Moreover, for homological reasons, there is also a bound on the number of disjoint twisted bands we can embed. Thus, if n is sufficiently large, the set of tracks must contain a parallel pair.

In our set-up, it follows that we can find $i \neq j$, with $s_i \sqcup s_j$ the boundary of a band $A \subseteq \Sigma'/\Gamma$. We now lift A to a connected subset $B \subseteq K' \cong K$. We can suppose that $\partial B = t_i \sqcup t_j$. Now, $\Gamma(t_i) = \Gamma(t_j)$ is the setwise stabiliser of B , which we denote by $\Gamma(B)$. Thus, $\Gamma(B) \subseteq \Gamma(\pi t_i) \cap \Gamma(\pi t_j)$. But now, t_i and t_j both meet some edge of K . From the construction of π , it follows that $\pi t_i \neq \pi t_j$. Since we are assuming that V has finite pair stabilisers, it follows that $\Gamma(B)$ is finite, and that B is compact. Now the set of edges of K meeting B is finite and separates Σ into two unbounded components, contradicting one-endedness.

This proves Theorem 6.1.

In fact, we can see immediately that Γ splits relative to $\{\Gamma(x) \mid x \in V\}$ without using Theorem 5.3. If we collapse each Γ -image of B to an interval, and each connected component of $\Sigma \setminus (\bigcup \Gamma B)$ to a point, then we arrive at a non-trivial Γ -action on a simplicial tree with finite edge stabilisers. Each element $x \in V$ lies

inside some complementary component, and hence fixes the corresponding vertex of the tree.

As a corollary to Theorem 6.1, we note:

Proposition 6.3. *Suppose that V is a non-multiended 1-connected Γ -set with finite pair stabilisers. Suppose that \mathcal{T} is a Γ -equivariant pretree structure on V . Then, the set of peripheral splittings of V subordinate to \mathcal{T} is either empty or has a unique maximal element.*

Proof. This follows from Theorem 6.1, on noting that if \mathcal{T}_1 and \mathcal{T}_2 are peripheral splittings subordinate to \mathcal{T} , then so is $\mathcal{T}_1 \cup \mathcal{T}_2$. \square

This raises a number of questions. For example, is \mathcal{T} necessarily discrete, and hence equal to the maximal peripheral splitting? Also, under the hypotheses of Theorem 6.1, does there necessarily exist a pretree, \mathcal{T} , which is a refinement of every peripheral splitting (so that there exists a unique maximal peripheral splitting)? We shall see that this is the case if we add the hypothesis that V is a hyperbolic Γ -set. However, the argument in this case has some topological input, and so cannot immediately be reinterpreted in combinatorial terms.

7. TREES OF METRIC SPACES

In this section, we consider decompositions of continua into subcontinua glued together in treelike fashion. Such a situation may arise as the boundary of a Γ -set admitting a peripheral splitting. One of the main objectives (Proposition 7.4) will be to show that if each of the subcontinua (corresponding to the components of the peripheral splitting) is locally connected, then the whole space is locally connected. Since much of the proof is routine continuum topology, we shall motivate the argument by explaining how we expect the boundary to look with reference to particular examples, and we shall leave the details of the general argument to the reader.

Although it is not an essential part of the logic of the argument, a useful idea to keep in mind is the following decomposition of continua analogous to the block decomposition of graphs. It is essentially the simplicial case of the construction of Swenson [Swe].

Let M be a compact metric space. We can give M the structure of a pretree by writing yxz if y and z lie in different quasicomponents of $M \setminus \{x\}$. (In other words, if we can write $M \setminus \{x\} = O \sqcup U$, where O and U are open subsets of M containing y and z respectively.) We shall refer to the closure in M of a quasicomponent of $M \setminus \{x\}$ as a *branch* of M *rooted at* x . More generally, a closed (connected) subset of M is *branchlike* if its boundary is a singleton. Any intersection of branchlike sets is connected. Indeed such an intersection meets any subcontinuum of M in a connected set [Bo3].

Now, suppose $V \subseteq M$ is (for the moment) any subset. Given $x, y \in M$, write $x \sim y$ to mean that there does not exist $z \in V$ with xzy . We write $[x] = \{z \in M \mid z \sim x\}$. Thus $[x]$ is a subcontinuum of M (being an intersection of branches). If $x \notin V$, we refer to $[x]$ as a *block*. Two distinct blocks intersect, if at all, in a single point of V . Note that if $x, y \in V$ with $x \sim y$, then $[x] \cap [y]$ is a subcontinuum of M containing both x and y . Thus, if V is countable (or if every point of V is a global cut point of M [Swe]), then $[x] \cap [y]$ must contain an element $z \notin V$. Thus, x and y lie in a common block, namely $[z]$.

Now, V inherits a pretree structure from M . Let us suppose that this pretree structure is discrete. We construct a bipartite graph with vertex set $V \sqcup W$, where W is the set of blocks of T , by deeming $x \in V$ to be adjacent to $P \in W$ if $x \in P$.

From the discussion of the previous paragraph, it is easy to see that T is a tree. Moreover, the pretree structure induced on V from T agrees with the structure induced from M .

We now add the assumption that for any block, P , and for any $\epsilon > 0$, there are only finitely many branches rooted in P (i.e. rooted at a point of P but not containing P) with diameter greater than ϵ . (This is necessarily the case for locally connected continua, for example, and in the constructions that follow later.) In this situation, any union of branches rooted at the same point is branchlike.

Associated to T is its *ideal boundary*, ∂T (i.e. the set of cofinality classes of rays in T). An element $x \in \partial T$ determines a subcontinuum of M (not meeting V) as an intersection of branches in the obvious way. At this point, we need to add another assumption (again valid in the cases of interest), namely that each subset of this type is a singleton. In this way, we get an identification of ∂T with a subset of M . We refer to such points as *ideal points* of M (relative to V).

Suppose S is a subtree of T , so that $\partial S \subseteq \partial T$. We can associate to S a subset $\Phi(S) \subseteq M$, namely the union of all blocks and ideal points associated to S . Thus, $\Phi(S)$ is a subcontinuum of M (being an intersection of branchlike sets). Note that if S is a branch of T rooted at $x \in V$, then $\Phi(S)$ is a branch of M rooted at x , and conversely.

As an example, consider an amalgamated free product, Γ , of 3 copies of a finite covolume kleinian group, G , amalgamated over a common maximal parabolic subgroup, H (as discussed in Section 2). Suppose we take a maximal peripheral splitting of Γ , where the graph of groups is a tripod with central vertex group H and terminal vertex groups G . In this case, the boundaries of the components are precisely the blocks of $\partial\Gamma$. However, we shall need to consider more general situations. For example, we can also write Γ as an amalgamated free product of two relatively hyperbolic groups, namely $\Gamma \cong G *_H (G *_H G)$. In this case, the boundary of the component $G *_H G$ is the closure of a union of blocks as defined above. In what follows, this is an example of a “generalised block”.

To derive the main results of this section, we need to apply the above construction in reverse. Suppose we are given a compact metrisable space, M , and a bipartite simplicial tree, T , with vertex set $V(T) \sqcup W(T)$. To each element $x \in V(T) \cup \partial T$ we associate a point $p(x) \in M$, and to each $x \in W(T)$, we associate a closed subset $P(x) \subseteq M$. We refer to a subset of the form $P(x)$ as a *generalised block*, and to elements of $p(\partial T)$ as *generalised ideal points*. If $S \subseteq T$ is a subtree, we write $\Phi(S) = p(V(S) \cup \partial S) \cup \bigcup_{x \in W(S)} P(x)$. If S is a branch of T rooted at $x \in V$, we refer to $\Phi(S)$ as a *generalised branch* of M rooted at $p(x)$.

We also need the following hypotheses. First, we assume that the map $p : V(T) \cup \partial T \rightarrow M$ is injective, and that no generalised block contains any generalised ideal points. We assume that two distinct generalised blocks meet, if at all, in an element of $p(V(T))$, and that $x \in V(T)$ is adjacent to $y \in W(T)$ if and only if $p(x) \in P(y)$. We assume that $M = \Phi(T)$. Finally, we assume that if P is any generalised block, then the diameters of the generalised branches emerging from P (i.e. rooted in P but not containing P) must tend to 0 (i.e. for any $\epsilon > 0$, only finitely many have diameter greater than ϵ).

We begin with a few preliminary observations. First, note that if $F \subseteq P$ is any closed subset, then the union of F and the union of any subset of generalised branches emerging from F must be closed. We also note any generalised ideal point has a neighbourhood base consisting of generalised branches of M . In particular, the complement of the set of ideal points is dense in M . Finally, note that if $S \subseteq T$ is any subtree, then $\Phi(S)$ is closed.

An easy consequence of these observations is:

Proposition 7.1. *M is connected if and only if every generalised block of M is connected.*

We assume henceforth that M is connected.

The same argument applied to any subtree, S , of T , shows that $\Phi(S)$ is connected. In particular, any generalised branch of M is connected.

Now $V(T)$ inherits a pretree structure, \mathcal{T}_M , from its embedding $p : V(T) \hookrightarrow M$. It also inherits a pretree structure, \mathcal{T}_T , from its embedding in T .

From the earlier observations it is easily verified that:

Proposition 7.2. *\mathcal{T}_M is a refinement of \mathcal{T}_T .*

Suppose that P is a generalised block of M . Now P has a pretree structure as a subpretree of M , as well as an intrinsic pretree structure as a continuum. It is easily checked that these structures agree:

Proposition 7.3. *If P is a generalised block of M , then P with its intrinsic pretree structure is a subpretree of M .*

In particular, we see that any global cut point of P is also a global cut point of M .

Finally, we want to consider local connectedness. This is usually defined by demanding that every point has a base of open connected neighbourhoods. In fact, it is sufficient that every point has a base of connected neighbourhoods (see [K]).

Proposition 7.4. *M is locally connected if and only if every generalised block of M is locally connected.*

Proof. First suppose that M is locally connected. Let P be a generalised block of M . If F is any subcontinuum of M , the $F \cap P$ is connected. (This can be seen directly, or using the fact that P is an intersection of branchlike sets.) Thus, if $x \in P$ and F is a connected neighbourhood of x in M , then $F \cap P$ is a connected neighbourhood of x in P .

Conversely, suppose that every block of M is locally connected. We shall explicitly construct a base of connected neighbourhoods for each point $a \in M$. We distinguish three cases, namely $a \in p(\partial T)$, $a \in M \setminus p(V(T) \cup \partial T)$ and $a \in p(V(T))$.

If $a \in p(\partial T)$, then we have already observed that a has a neighbourhood base consisting of generalised branches of M . Moreover, such branches are necessarily connected.

We can thus assume that $a \notin p(\partial T)$, so that a lies in some generalised block, P , of M . Suppose first that $a \notin p(V(T))$. Given $\epsilon > 0$, we want to find a connected neighbourhood of a contained in $N(a, \epsilon)$.

Let \mathcal{B} be the set of generalised branches of M emerging from P (so that $M = P \cup \bigcup \mathcal{B}$). Let $\mathcal{B}_0 \subseteq \mathcal{B}$ be the (finite) subset of such branches of diameter at least $\epsilon/2$. Let $\delta = \frac{1}{2} \min(\{d(a, B) \mid B \in \mathcal{B}_0\} \cup \{\epsilon\}) > 0$. Now, let $F \subseteq P$ be a closed

connected neighbourhood of a in P contained in $N(a, \delta)$. If $B \in \mathcal{B}$ is a generalised branch rooted in F , then $d(a, B) \leq \epsilon/2$ and $\text{diam}(B) \leq \epsilon/2$. Thus, $B \subseteq N(a, \epsilon)$. Let F' be the union of F and the union of all such branches. It follows that F' is connected and contained in $N(a, \epsilon)$. Now let G be the closure of $P \setminus F$, and let G' be the union of G and the union of all elements of \mathcal{B} rooted in G . We see that G' is closed, that $a \notin G'$, and that $M = F' \cup G'$. It follows that F' is a neighbourhood of a in M .

Finally, we consider the case where $a \in p(V(T))$. Suppose $\epsilon > 0$. Let \mathcal{C} be the set of generalised branches of M rooted at a , and let \mathcal{C}_0 be the (finite) set of such branches of diameter at least ϵ . For each $C \in \mathcal{C}_0$, we use the same argument as in the previous paragraph to construct a neighbourhood, $A(C)$, of a in C contained in $N(a, \epsilon)$. We see that $\bigcup_{C \in \mathcal{C}_0} A(C) \cup \bigcup (C \setminus \mathcal{C}_0)$ is a connected neighbourhood of a in M contained in $N(a, \epsilon)$. \square

8. HYPERBOLICITY

In this section, we consider peripheral splittings of relatively hyperbolic groups, or more precisely, hyperbolic Γ -sets. We begin by recalling some definitions from [Bo5].

Let K be a connected graph. We can view K as a path-metric space by assigning each edge a length 1. We say that K is *hyperbolic* if it is (Gromov) hyperbolic with this metric [Gr]. We say that K is *fine* if, given any $n \in \mathbf{N}$ and any edge, e , of K , there are finitely many circuits of length n containing the edge e . It is a simple exercise to verify that a connected graph is fine if and only if each of its blocks is fine. Likewise, a connected graph is hyperbolic if and only if its blocks are uniformly hyperbolic.

Suppose V is a connected Γ -set with finite pair stabilisers. One can show that if one (Γ, V) -graph is fine, then they all are. Since all (Γ, V) -graphs are quasi-isometric, the same goes for hyperbolicity. This leads naturally to:

Definition. A Γ -set, V , is hyperbolic if it has finite pair stabilisers and admits a 2-vertex-connected (Γ, V) -graph which is fine and hyperbolic.

The constraint that the (Γ, V) -graph be 2-vertex-connected is essentially equivalent to demanding that the vertex stabilisers are all finitely generated. Also, given the finiteness condition, having finite pair stabilisers is equivalent to having finite edge stabilisers.

It was shown in [Bo5] that if V is a cofinite Γ -set, and $W \subseteq V$ is a non-empty Γ -invariant subset such that each point of $V \setminus W$ has finite stabiliser, then V is hyperbolic as a Γ -set if and only if W is. This fact will be used in the proof of Proposition 8.7. Thus, for most purposes one can reduce either to the case where all vertex stabilisers are finite (the standard case of word hyperbolic groups) or to the case where all vertex stabilisers are infinite. The last hypothesis leads to simplifications in certain cases, although we shall not explicitly assume it for the moment.

Recall that a peripheral splitting, \mathcal{R} , of V is “full” if and only if V admits a (Γ, V) -graph, K , with the property that \mathcal{R} is precisely the set of block sets of K .

An immediate consequence of the preceding observations is:

Proposition 8.1. *Suppose V is a cofinite Γ -set with finitely generated vertex stabilisers. Suppose \mathcal{R} is a full peripheral splitting of V . Then, V is a hyperbolic Γ -set if and only if R is a hyperbolic $\Gamma(R)$ -set for all $R \in \mathcal{R}$.*

To relate this to boundaries of Γ -sets, we need the following equivalent, more geometric, formulation of hyperbolicity [Bo5]:

Proposition 8.2. *A Γ -set, V , is hyperbolic if and only if we can represent V as the vertex set, $V = V(\Sigma)$, of a Γ -invariant simplicial 2-complex, such that Γ acts on Σ with finite quotient and finite edge-stabilisers, and such that Σ is simplicially hyperbolic and has no global cut vertices.*

Here “simplicially hyperbolic” means that every cycle, β , in the 1-skeleton of Σ , bounds a simplicial disc where the number of 2-simplices is bounded by a linear function of the length of β . In particular, Σ is simply connected, so V is 1-connected as defined in Section 6. As observed in [Bo5], Σ can be assumed to contain any given Γ -invariant 2-complex with finite quotient.

We may construct, from Σ , a geometric 2-complex, $X(\Sigma)$, by giving each 2-simplex the structure of an ideal triangle in such a way that the union of two adjacent triangles is isometric to an ideal hyperbolic square. We thus obtain a complete locally compact Γ -invariant path-metric on $X(\Sigma) \cong X \setminus V$. In [Bo5], it was shown that $X(\Sigma)$ is hyperbolic in the usual Gromov sense. We write $\partial V = \partial X(\Sigma)$ for the ideal boundary of $X(\Sigma)$. This is well-defined up to Γ -equivariant homeomorphism, independently of Σ . Thus, ∂V is compact metrisable, with V naturally embedded in ∂V . The isolated points of ∂V are precisely the points of V which have finite stabiliser. The group Γ acts on ∂V as a geometrically finite convergence group (see [T]). (We note that this construction makes sense if Σ has a global cut vertex, except that in this case we end up with a disjoint union of hyperbolic spaces.)

More generally, suppose that (X, ρ) is a proper (i.e. complete locally compact) hyperbolic space. Suppose that Γ acts isometrically on X , and that there is a Γ -invariant embedding of V in ∂X . An *invariant system of horoballs*, $(B(x))_{x \in V}$, consists of a choice of closed horoballs, $B(x)$, about $x \in V \subseteq \partial X$ for each $x \in V$ such that $B(\gamma x) = \gamma B(x)$ for all $\gamma \in \Gamma$. It is *r -separated* if $\rho(B(x), B(y)) \geq r$ whenever $x \neq y$. If it is r -separated for some $r > 0$, then the action of Γ is geometrically finite if and only if $(X \setminus \bigcup_{x \in V} \text{int } B(x))/\Gamma$ is compact. Moreover, ∂X is Γ -equivariantly homeomorphic to ∂V . For further elaboration, see [Bo5], [T]. Note that V (thought of as a subset of ∂X) is precisely the set of parabolic points for the action of Γ on ∂X .

We now want to interpret these constructions in the case where the hyperbolic Γ -set, V , admits a full peripheral splitting, \mathcal{R} . For each $R \in \mathcal{R}$, let $\Sigma(R)$ be a 2-complex as given by Proposition 8.2, using the fact that R is a hyperbolic $\Gamma(R)$ -set (Proposition 8.1). We choose these complexes Γ -equivariantly to give us a Γ -equivariant complex Σ_0 , with vertex set V , and with finite quotient. We can now embed Σ_0 as a subcomplex of another 2-complex, Σ , with the properties given by Proposition 8.2. Let $X(\Sigma)$ be the hyperbolic complex constructed from Σ , and let $X(\Sigma_0)$ be the subcomplex corresponding to Σ_0 . Thus, $X(\Sigma_0)$ is a disjoint union of subcomplexes $X(\Sigma(R))$ as R varies over \mathcal{R} .

Given $t \geq 0$, we shall construct a Γ -invariant subset, $Y(t)$, of $X(\Sigma)$, with $X(\Sigma_0) \subseteq Y(t) \subseteq X(\Sigma)$ as follows.

Suppose $x \in V$ and σ is a 2-simplex of Σ , with x as a vertex. Let $s(\sigma, x, t) \subseteq \sigma$ be the “spike” in σ consisting of the intersection of σ (thought of as an ideal triangle in the hyperbolic plane) with a horoball (in the classical hyperbolic sense) about x which is at a distance $t + 1$ from the centre of the triangle. Let $B(x, t)$ be the union of all the spikes $s(\sigma, x, t)$ as σ ranges over the set of 2-simplices of Σ with vertex x . Now let $Y(t) = X(\Sigma_0) \cup \bigcup_{x \in V} B(x, t)$. Thus $Y(t)$ is closed in $X(\Sigma)$, and complete and locally compact in the induced path-metric.

Now, a simple variation on the argument of [Bo5] (used in proving that $X(\Sigma)$ is hyperbolic) shows that the space $Y(t)$ is a proper hyperbolic space, with V naturally embedded as a dense subset of $\partial Y(t)$. Moreover, the collection $(B(x, t))_{x \in V}$ forms a $(2t)$ -separated invariant system of horoballs for this action. The quotient $(Y(t) \setminus \bigcup_{x \in V} \text{int } B(x, t))/\Gamma$ is compact, so the action of Γ on V is geometrically finite. Also, the subcomplexes $X(\Sigma(R))$ are uniformly quasiconvex in $Y(t)$. The constants involved can all be fixed independently of t . Setting $X = Y(t)$, $X(R) = X(\Sigma(R))$ and $B(x) = B(x, t)$ we have shown:

Proposition 8.3. *Suppose V is a hyperbolic Γ -set, and that \mathcal{R} is a full peripheral splitting of V . Suppose $r > 0$. Then, we can find a proper hyperbolic space, X , with a properly discontinuous isometric action of Γ and a Γ -equivariant embedding of V in ∂X , such that the action is geometrically finite. Moreover, for each $R \in \mathcal{R}$, we can find a closed $\Gamma(R)$ -invariant subset, $X(R)$, of X , which is quasiconvex and intrinsically hyperbolic. (In fact, we can assume that the induced metric on $X(R)$ is already a path metric.) The collection $\{X(R) \mid R \in \mathcal{R}\}$ is locally finite in X . Moreover, $X(\gamma R) = \gamma X(R)$ for all $\gamma \in \Gamma$, and $X(R) \cap X(R') = \emptyset$ if $R \neq R'$. We can also find an r -separated invariant system of horoballs, $(B(x))_{x \in V}$, for X , such that $X = \bigcup_{R \in \mathcal{R}} X(R) \cup \bigcup_{x \in V} B(x)$. Moreover, if $x \in V$ and $R \in \mathcal{R}$, then $X(R) \cap B(x) \neq \emptyset$ if and only if $x \in R$. If $x \in R \in \mathcal{R}$, then $X(R) \cap B(x)$ is a horoball about x in $X(R)$. (In fact, we can also assume that the metric on each $B(x)$ induced from X is already a path metric.) Finally, all the constants (of hyperbolicity and quasiconvexity) involved can be fixed independently of r .*

We note that if $x, y \in V$, then any path in X connecting $B(x)$ to $B(y)$ must meet $B(z)$ for all z in the pretree interval $[x, y] \subseteq V$.

Suppose that $R \in \mathcal{R}$. We have $R \subseteq V \subseteq \partial V$. We write $P(R)$ for the closure of R in ∂V . Thus, $P(R)$ may be identified with $\partial X(R)$. We know that R is a hyperbolic $\Gamma(R)$ -set by Proposition 8.1. We claim that we can identify ∂R with $P(R)$. In the explicit construction of X as a 2-complex we gave earlier, this is clear from the definitions. It also follows directly in the general set-up described by Proposition 8.3. For this, it's sufficient to see that the action of $\Gamma(R)$ on $X(R)$ is geometrically finite. To this end, we note that $(X \setminus \bigcup_{x \in V} \text{int } B(x))/\Gamma$ is compact. Moreover, it can be expressed as a finite disjoint union of sets of the form $(X(R) \setminus \bigcup_{R \in \mathcal{R}} \text{int } B(x))/\Gamma(R)$ as R ranges over a Γ -transversal of \mathcal{R} . In particular, we see that each of these subsets is compact. Since $X(R) \cap B(x)$ is a horoball in $X(R)$ for all $x \in R$, it follows that the action of $\Gamma(R)$ on $X(R)$ is geometrically finite. By the results of [Bo5], it now follows that $\partial X(R)$ is $\Gamma(R)$ -equivariantly homeomorphic to ∂R .

Now let $T = T(\mathcal{R})$ be the bipartite tree with vertex set $V \sqcup \mathcal{R}$ associated to \mathcal{R} , as defined in Section 1. We define a map $p : V \sqcup \partial T \rightarrow \partial V$ as follows. If $x \in V$, set $p(x) = x$. If $x \in \partial T$, let β be a ray tending to x , and let $(x_n)_n$ be the sequence of vertices of V in β . We claim that $(x_n)_n$ is convergent in ∂V . To see this, fix

any basepoint, $a \in B(x_0)$, and let α_n be a shortest path from a to $B(x_n)$. Now the geodesics α_n must subconverge on some geodesic ray, α . We see that α must intersect $B(x_n)$ for all n . It follows easily that x_n must converge to the ideal point of α , proving the claim. Now this limit is clearly independent of the choice of α , and we denote it by $p(x)$.

We need to verify that the maps p and P defined above satisfy the hypotheses of a generalised block decomposition given in Section 7. Here, the generalised blocks are sets of the form $P(R) \equiv \partial X(R)$ for $R \in \mathcal{R}$. The facts that p is injective and that no block meets $p(\partial T)$ are easy consequences of the definitions. Also, note that any geodesic ray that does not remain within a bounded distance of any $X(R)$ must pass through an infinite sequence of horoballs corresponding to a ray in T . From this it follows that $\partial V = \Phi(T)$. Now, if $R, R' \in \mathcal{R}$ are disjoint, then $X(R)$ and $X(R')$ are a distance at least r apart. Since r can be chosen large in relation to the constants of hyperbolicity and quasiconvexity, it follows that $P(R) \cap P(R') = \emptyset$. On the other hand, if $R \cap R' = \{x\}$, then if $a \in X(R)$ and $b \in X(R')$, then the distance from a or b to $B(x)$ is less than the distance between them. Again choosing r appropriately, we see that $P(R) \cap P(R') = \{x\}$.

Suppose that $S \subseteq T$ is a subtree. Let $X(S) = \bigcup_{x \in V \cap S} B(x) \cup \bigcup_{R \in \mathcal{R}(S)} X(R)$. From the manner in which the sets $X(R)$ and $B(x)$ are connected, it is easily verified that $X(S)$ is a uniformly quasiconvex subset of X , and that $\partial(X(S)) \cong \Phi(S)$. It remains to verify that given $R \in \mathcal{R}$, there are only finitely many generalised branches rooted in R with diameter greater than any positive constant. Suppose, to the contrary, that S_n is a sequence of distinct branches rooted at $x_n \in R$ with $\text{diam}(\Phi(S_n))$ bounded below. We can assume that the points x_n are either all distinct or else constant. In the former case, we see that the sets $X(S_n)$ are disjoint, locally finite and uniformly quasiconvex, giving the contradiction that $\text{diam}(\partial X(S_n)) \rightarrow 0$. We can thus assume that $x_n = x$ is constant. In this case, let $R_n \in \mathcal{R}$ be the vertex of S_n adjacent to x . The sets R_n are thus distinct. On passing to a subsequence, we can find $y_n \in \Phi(S_n)$ with $y_n \rightarrow y \neq x$. Let α_n and α be geodesics in X connecting x to y_n and y respectively. We can assume that α_n converges to α . Let z_n (respectively z) be the point where α_n (respectively α) leaves the horoball $B(x)$. From the combinatorial structure of X , we see that $z_n \in X(R_n)$. But $z_n \rightarrow z$, contradicting the local finiteness of the collection of spaces $X(R_n)$.

Proposition 8.4. *Suppose V is a hyperbolic Γ -set and that \mathcal{R} is a peripheral splitting of V . Then ∂V is connected if and only if ∂R is connected for all $R \in \mathcal{R}$. Moreover, in this case, the pretree structure on V induced by the continuum ∂V is a refinement of the pretree structure on V associated to \mathcal{R} .*

Proof. The first statement follows from Proposition 7.1, and the second from Proposition 7.2 on noting that the pretree structure associated to \mathcal{R} agrees with that induced from the tree $T(\mathcal{R})$. □

The next proposition follows from Proposition 7.4 in the same way Proposition 8.4 followed from Proposition 7.1:

Proposition 8.5. *Suppose that V is a hyperbolic Γ -set and that \mathcal{R} is a peripheral splitting of V . Then ∂V is a Peano continuum, if and only if ∂R is a Peano continuum for each $R \in \mathcal{R}$.*

Now Theorems 1.2 and 1.3 are essentially reformulations of Propositions 8.1 and 8.4 in the language of relatively hyperbolic groups. Suppose that Γ is a group and

that \mathcal{G} is a peripheral structure on Γ , as defined in Section 4. We identify \mathcal{G} with a Γ -set, V , obtained by taking the action of Γ on \mathcal{G} by conjugation. Thus, \mathcal{G} is precisely the set of point stabilisers of V . We say that Γ is *hyperbolic relative to \mathcal{G}* if V is hyperbolic as a Γ -set. We can assume that $\mathcal{G} \neq \emptyset$, otherwise the subject is vacuous. Since every element of \mathcal{G} is, by assumption, infinite, we can identify ∂V with the boundary, $\partial\Gamma$, of Γ . More generally, if V is any hyperbolic Γ -set, then Γ will be hyperbolic relative to the infinite point stabilisers of V . In this case, $\partial\Gamma$ is equal to ∂V minus the set of isolated points.

Suppose that \mathcal{R} is a peripheral splitting of V , and that $R \in \mathcal{R}$. Then $H = \Gamma(R)$ is a component of Γ , as we defined in Section 1. Suppose the stabiliser, $H(x) = H \cap \Gamma(x)$, of x in H is infinite. If $g \in \Gamma$ normalises $H(x)$, then $H(gx) = H(x)$. In particular, $\Gamma(gx) \cap \Gamma(x) \supseteq H(x)$ is infinite, and so $gx = x$. Thus, g normalises $\Gamma(x)$, so $g \in \Gamma(x)$, and $g \in H(x) = H \cap \Gamma(x)$. Thus, $H(x)$ is equal to its normaliser. We see that the set of infinite $H(x)$ as x ranges over R is a peripheral structure on H . Now, by Lemma 8.1, R is a hyperbolic H -set. Thus H is hyperbolic relative to this peripheral structure.

Suppose now that $\partial V = \partial\Gamma$ is connected. It follows by Proposition 8.4, that ∂R is connected. In particular, ∂R has no isolated points. Thus every point stabiliser, $H(x)$ for $x \in R$ is infinite. Thus, by definition, ∂H can be identified with ∂R .

Finally, note that if \mathcal{R} is non-trivial, then the pretree structure induced on $\partial\Gamma$ arising from its topology is also non-trivial. In other words, $\partial\Gamma$ has a global cut point.

This proves Theorems 1.2 and 1.3.

9. CONCLUSION

In this section, we put together the results of previous sections to give us an accessibility result for hyperbolic Γ -sets with connected boundaries (Theorem 9.1). Reinterpreted in terms of a relatively hyperbolic group, this proves Theorem 1.4. We explain how this relates the programme of proving local connectedness of such boundaries (Theorem 1.5). In the non-relative case, local connectedness follows from previous work (see [BeM], [Bo3], [Swa], [Bo4]). In the case of geometrically finitely kleinian groups, the argument presented below can be simplified somewhat (see [BoS]).

Suppose then, that Γ is a group, and that V is a hyperbolic Γ -set. Let's suppose that ∂V is connected, so that, by Proposition 1.1, V is one-ended. Since V is hyperbolic, it is certainly 1-connected, and has finite pair stabilisers. Let $\mathcal{T}_{\partial V}$ be the pretree structure on V induced from the topology of ∂V (see Section 7). Suppose that \mathcal{R} is a peripheral splitting, and that $\mathcal{T}(\mathcal{R})$ is the corresponding pretree structure on V . By Proposition 8.4, $\mathcal{T}(\mathcal{R})$ is subordinate to $\mathcal{T}_{\partial V}$. It thus follows from Proposition 6.3 that V admits a unique maximal peripheral splitting. We have shown:

Theorem 9.1. *Suppose that V is a hyperbolic Γ -set with connected boundary. Then V admits a unique maximal peripheral splitting.*

Let \mathcal{R}_0 be the maximal peripheral splitting of V . (Note that this splitting might be trivial.) By Lemma 4.6, no component of \mathcal{R}_0 admits any non-trivial peripheral splitting. We also note that by Proposition 8.4, every component of \mathcal{R}_0 has connected boundary.

Note that the pretree structure $\mathcal{T}(\mathcal{R}_0)$ is subordinate to $\mathcal{T}_{\partial V}$. At the moment, there is no reason to suppose that these are equal, or even that $\mathcal{T}_{\partial V}$ is discrete. Things become clearer once we know that ∂V is locally connected, as we shall explain later. First, we explain the connection with the proof of local connectedness.

Suppose then that ∂V is hyperbolic with ∂V connected. The proof given in [Bo6] uses the notion of a “separating horoball”. An example of a separating horoball would be a horoball of the type $B(x)$ in Proposition 8.3, where $x \in V$ is a non-terminal point of the peripheral splitting. We shall not give a formal definition here since only the logic of the argument is relevant. It is sufficient to note that the “centre” of a separating horoball is always a global cut point of ∂V .

We recall two facts from [Bo6]. First, we note that if V has no separating horoball, then ∂V is locally connected (cf. [BeM]). Second, if ∂V has a separating horoball centred on a point of V , then V admits a non-trivial peripheral splitting. We also need a third fact, proven in [Bo4], which tells us that, under the same hypotheses as Theorem 1.5, every global cut point of ∂V is centred on a point of V (i.e. a parabolic point). This last fact is by far the most difficult part of the argument, and calls for considerable input from elsewhere.

Suppose, for contradiction, that V satisfies the hypotheses of Theorem 1.5, but that ∂V is not locally connected. Let \mathcal{R}_0 be the maximal peripheral splitting given by Theorem 9.1. By Propositions 8.4 and 8.5, some component, $R \in \mathcal{R}_0$, is such that ∂R is connected but not locally connected, hence R admits a separating horoball. The centre, x , of this horoball is a global cut point of ∂R , and hence, by Proposition 7.3, also a global cut point of ∂V . Thus, $x \in V$. Since $R = V \cap \partial R$, we see that $x \in R$. In other words, R has a separating horoball centred on a point of R , and hence admits a non-trivial peripheral splitting. By Lemma 4.6, this contradicts the maximality of \mathcal{R}_0 , and hence proves Theorem 1.5. For more details, see [Bo6].

From hereon, life becomes much easier. The way is open to a more detailed analysis of boundaries. To begin with, it was shown in [Bo6] that if the boundary of a relatively hyperbolic group is connected and locally connected, then every parabolic global cut point is the centre of a separating horoball, and hence gives rise to a peripheral splitting. From this, we can deduce:

Theorem 9.2. *Suppose V is a hyperbolic Γ -set such that ∂V is connected and locally connected. Let $\mathcal{T}_{\partial V}$ be the pretree structure on V induced from the topology of ∂V . Then, $\mathcal{T}_{\partial V}$ is precisely the maximal peripheral splitting of V .*

Proof. Let $\mathcal{T}(\mathcal{R}_0)$ be the maximal peripheral splitting of V . By Proposition 8.4, $\mathcal{T}_{\partial V}$ is a refinement of $\mathcal{T}(\mathcal{R}_0)$. Suppose, for contradiction, that it is a strict refinement. By Lemma 3.3, there is some $R \in \mathcal{R}_0$ and points $x, y, z \in R$ such that yxz holds in $\mathcal{T}_{\partial V}$.

Now by Propositions 8.4 and 8.5, ∂R is connected and locally connected. By Proposition 7.3, the relation yxz holds in the pretree structure associated to ∂R . Thus, x is a global cut point of ∂R . But now, by the observation above, R admits a non-trivial peripheral splitting, contradicting Lemma 4.6. \square

In particular, it follows that $\mathcal{T}_{\partial V}$ is discrete.

If we assume that V satisfies the hypotheses of Theorem 1.5, then we know (by the result of [Bo4]), that every global cut point of ∂V lies in V . Thus, the pretree relation on the whole of ∂V is discrete. In other words, any two cut points of ∂V can be separated by only finitely many other cut points.

REFERENCES

- [AN] S.A. Adeleke, P.M. Neumann, *Relations related to betweenness: their structure and automorphisms*, Memoirs Amer. Math. Soc. No. 623, American Mathematical Society, Providence, Rhode Island (1998). MR **98h**:20008
- [BeF] M. Bestvina, M. Feighn, *Bounding the complexity of simplicial actions on trees*, Invent. Math. **103** (1991) 449–469. MR **92c**:20044
- [BeM] M. Bestvina, G. Mess, *The boundary of negatively curved groups*, J. Amer. Math. Soc. **4** (1991) 469–481. MR **93j**:20076
- [Bo1] B.H. Bowditch, *Geometrical finiteness with variable negative curvature*, Duke Math. J. **77** (1995) 229–274. MR **96b**:53056
- [Bo2] B.H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta. Math. **180** (1998) 145–186. MR **99b**:20069
- [Bo3] B.H. Bowditch, *Treelike structures arising from continua and convergence groups*, Memoirs Amer. Math. Soc. No. 662, American Mathematical Society, Providence, Rhode Island (1999). MR **2000c**:20061
- [Bo4] B.H. Bowditch, *Connectedness properties of limit sets*, Trans. Amer. Math. Soc. **351** (1999) 3673–3686. MR **2000d**:20056
- [Bo5] B.H. Bowditch, *Relatively hyperbolic groups*, preprint, Southampton (1997).
- [Bo6] B.H. Bowditch, *Boundaries of geometrically finite groups*, Math. Z. **230** (1999) 509–527. MR **2000b**:20049
- [BoS] B.H. Bowditch, G.A. Swarup, *Cut points in the boundaries of hyperbolic groups*, in preparation.
- [DiD] W. Dicks, M.J. Dunwoody, *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics No. 17, Cambridge University Press (1989). MR **91b**:20001
- [Du1] M.J. Dunwoody, *Cutting up graphs*, Combinatorica **2** (1982) 15–23. MR **84k**:05050
- [Du2] M.J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985) 449–457. MR **87d**:20037
- [Du3] M.J. Dunwoody, *Groups acting on protrees*, J. London Math. Soc. **56** (1997) 125–136. MR **98f**:20008
- [Du4] M.J. Dunwoody, *Folding sequences*, in “The Epstein Birthday Schrift”, Geometry and Topology Monographs Volume 1, (ed. I. Rivin, C. Rourke, C. Series), International Press, 139–158. MR **2000f**:20037
- [Gr] M. Gromov, *Hyperbolic groups*, in “Essays in Group Theory” (S.M. Gersten, ed.) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263. MR **89e**:20070
- [Gu] D.P. Guralnik, *Constructing a splitting-tree for a cusp-finite group acting on a Peano continuum* (Hebrew) : M.Sc. Dissertation, Technion, Haifa (1998).
- [H] A. Haefliger, *Complexes of groups and orbihedra*, in “Group theory from a geometrical viewpoint” (E. Ghys, A. Haefliger, A. Verjovsky, eds.), World Scientific (1991) 504–540. MR **93m**:20048
- [K] J.L. Kelley, *General topology*, Graduate Texts in Mathematics, 27 Springer-Verlag, 1975 (reprint of Van Nostrand edition 1955). MR **51**:6681
- [Sh] M. Sholander, *Trees, lattices, order and betweenness*, Proc. Amer. Math. Soc. **3** (1952) 369–381. MR **14**:9b
- [Swa] G.A. Swarup, *On the cut point conjecture*, Electron. Res. Announc. Amer. Math. Soc. **2** (1996) 98–100 (Electronic). MR **97f**:20048
- [Swe] E.L. Swenson, *A cutpoint tree for a continuum*, in “Computational and Geometric Aspects of Modern Algebra”, London Math. Soc. Lecture Note Series, No. 275, (M. Atkinson, N. Gilbert, J. Howie, S. Linton, E. Robertson, eds.), Cambridge University Press (2000), 254–265. CMP 2001:01
- [T] P. Tukia, *Conical limit points and uniform convergence groups*, J. Reine. Angew. Math. **501** (1998) 71–98. MR **2000b**:30067
- [W] L.E. Ward, *Axioms for cutpoints*, in “General Topology and Modern Analysis”, Proceedings, University of California, Riverside (L. F. McAuley, M.M. Rao, eds.) Academic Press (1980) 327–336. MR **82g**:54053

FACULTY OF MATHEMATICAL STUDIES, UNIVERSITY OF SOUTHAMPTON, HIGHFIELD, SOUTHAMPTON SO17 1BJ, GREAT BRITAIN

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