

SPLITTINGS OF FINITELY GENERATED GROUPS OVER TWO-ENDED SUBGROUPS

BRIAN H. BOWDITCH

ABSTRACT. We describe a means of constructing splittings of a one-ended finitely generated group over two-ended subgroups, starting with a finite collection of codimension-one two-ended subgroups. In the case where all the two-ended subgroups have two-ended commensurators, we obtain an annulus theorem, and a form of the JSJ splitting of Rips and Sela. The construction uses ideas from the work of Dunwoody, Sageev and Swenson. We use a particular kind of order structure which combines cyclic orders and treelike structures. In the special case of hyperbolic groups, this provides a link between combinatorial constructions, and constructions arising from the topological structure of the boundary. In this context, we recover the annulus theorem of Scott and Swarup. We also show that a one-ended finitely generated group which contains an infinite-order element, and such that every infinite cyclic subgroup is (virtually) codimension-one is a virtual surface group.

0. INTRODUCTION

In this paper, we describe a means of constructing splittings of a finitely generated group starting with a finite collection of codimension-one two-ended subgroups. We obtain versions of the annulus theorem and JSJ splittings for such groups. As a corollary, we also obtain a characterisation of a virtual surface group as a one-ended finitely generated group which contains an infinite order element and such that every infinite order element is codimension-one.

In the case of (almost) finitely presented groups, stronger results have been obtained by different methods in [DSa] and [FP]. There is also a strong version of the annulus theorem for finitely generated groups in [DSw]. The methods in this paper differ in that they all apply directly to the finitely generated case, and that the splittings obtained are more canonical. In the case of hyperbolic groups, we recover many of the results of [Bo1], and we provide a link between the methods of that paper and those of [DSw] and [DSa].

The term “JSJ splitting” has been applied to a number of similar canonical splittings of a one-ended group over certain classes of subgroups. They are inspired by an analogous decomposition of irreducible 3-manifolds due to Waldhausen, Johannson, Jaco and Shalen. In the setting of group theory, JSJ splittings were introduced by Sela [Se]. Since then several accounts have been given in different contexts using a variety of different methods (see [RS], [Bo1], [DSa], [FP]). The “annulus theorems” are also modelled on a 3-manifold construction, originally due to Scott. The group theoretical analogues give criteria for recognising the existence of splittings

Received by the editors January 31, 2001 and, in revised form, July 01, 2001.
2000 *Mathematics Subject Classification*. Primary 20F65, 20E08.

over certain classes of subgroups. There are again a number of different approaches (see [ScS1], [Bo1], [DSw], [P]).

In this theory, virtual surface groups play a central role, and so therefore do geometric criteria for recognising them. A number of such criteria are described in [Bo4] and [P]. These rest ultimately on the characterisation of fuchsian groups as convergence groups acting on the circle [T1], [Ga], [CJ], which completed the proof of the Seifert conjecture [M]. As mentioned above, we give another criterion in this paper (Theorem 13.1). If one strengthens the finite generation hypothesis to almost finite presentability, then with bit of work, this result can deduced from the results and methods of [DSw] and [DSa], as observed by Graham Niblo. These methods are also relevant to [Bo4]. Many of the arguments of the present paper are inspired by [DSw] and [Sw].

We remark that a series of results relating splittings and end invariants have been obtained by Kropholler and Roller. See for example [K], [KR], and the end of Section 1 for some discussion of how this relates to the present paper.

Part of the material for this paper was prepared while visiting the Université de Toulouse III, and I am grateful to Michel Boileau and Gilbert Levitt for their invitation. Conversations with Martin Dunwoody, Ross Geoghegan and Graham Niblo have also been helpful.

1. TWO-ENDED SUBGROUPS OF ONE-ENDED GROUPS

Suppose that Γ is a group and that $G, H \leq \Gamma$. We say that G and H are *commensurate* if $G \cap H$ has finite index in both G and H . The *commensurator*, $\text{Comm}(G)$, of G in Γ is the set of elements, $g \in \Gamma$, such that gGg^{-1} is commensurate with G . Thus, $\text{Comm}(G)$ is a subgroup of Γ . We say that G is *self-commensurating* if $G = \text{Comm}(G)$. Note that if G has finite index in $\text{Comm}(G)$, then $\text{Comm}(G)$ is self-commensurating. (Note that, for two-ended subgroups, “self-commensurating” is the same as “almost malnormal” in the terminology of [ScS2].)

If $g, h \in \Gamma$, we say that g and h are commensurate if $\langle g \rangle$ and $\langle h \rangle$ are commensurate, in other words, if there exist $m, n \in \mathbf{Z} \setminus \{0\}$ such that $g^m = h^n$.

Suppose now that Γ acts properly discontinuously on a locally finite connected one-ended graph, X . We view this as a metric space with the combinatorial path metric. Suppose that $G \leq \Gamma$. Let $\mathcal{S}(G)$ be the set of G -invariant subgraphs of X with finite quotient under G . Thus, $\mathcal{S}(G)$ is a non-empty directed set under inclusion. Given $A \in \mathcal{S}(G)$, let $\mathcal{C}_\infty(A)$ be the set of complementary components of A that are not contained in any other element of $\mathcal{S}(G)$. If $A \subseteq B \in \mathcal{C}_\infty(A)$, then there is a natural surjective map from $\mathcal{C}_\infty(B)$ to $\mathcal{C}_\infty(A)$. We thus have an inverse limit system, $(\mathcal{C}_\infty(A))_{A \in \mathcal{S}(G)}$. We write $\mathcal{E}(G)$ for the inverse limit. In the case of interest to us here, G will be two-ended. In this case, the sets $\mathcal{C}_\infty(A)$ will all be finite (cf. the construction of axes in Section 9). Taking the inverse limit in the category of topological spaces, we therefore arrive at a compact totally disconnected space, $\mathcal{E}(G)$.

Definition. We refer to an element of $\mathcal{E}(G)$ as a *coend* of G (with respect to the action of G on X).

(The term “filtered end” has been used for an equivalent notion elsewhere.)

Note that we arrive at the same definition by using the (cofinal) subset, $\mathcal{S}_0(G) \subseteq \mathcal{S}(G)$ of non-empty connected subgraphs. We write $\bar{e}(\Gamma, G) = |\mathcal{E}(G)| \in \mathbf{N} \cup \{\infty\}$. We say that G is *codim-one* (“codimension 1”) if $\bar{e}(\Gamma, G) \geq 2$. By a *codim-one*

element of Γ , we mean an infinite order element, g , such that $\langle g \rangle$ is codim-one. In fact, we shall see that $\bar{e}(\Gamma, G)$ is the same as the invariant $\tilde{e}(\Gamma, G)$ defined in [KR] — see the discussion at the end of this section.

Suppose $G, H \in \Gamma$ are two-ended subgroups and that $A \in \mathcal{S}_0(G)$ and $B \in \mathcal{S}_0(H)$. If G and H are commensurate, then A and B are parallel (i.e. the Hausdorff distance between them is finite). Moreover, $\mathcal{S}_0(G)$ and $\mathcal{S}_0(H)$ are cofinal, and so there is a natural homeomorphism between $\mathcal{E}(G)$ and $\mathcal{E}(H)$. In particular, $\bar{e}(\Gamma, G) = \bar{e}(\Gamma, H)$.

If G, H are not commensurate, then A and B are divergent in the sense that the intersection of B with any element of $\mathcal{S}(G)$ is compact (and vice versa). In particular, each end of B determines a coend of G , which is easily seen to be independent of the choice of B . If these ends are different, we say that H crosses G . This clearly implies that $\bar{e}(\Gamma, G) \geq 2$, and it also implies that $\bar{e}(\Gamma, H) \leq 2$. If $\bar{e}(\Gamma, H) = 2$, then it follows that G crosses H , hence $\bar{e}(\Gamma, G) = 2$. In summary, for codim-one two-ended subgroups, the relation of crossing is symmetric, and implies that the subgroups concerned have precisely two coends. (For more details, see Section 10.)

Note that G acts by homeomorphism on $\mathcal{E}(G)$. It has a particular kind of dynamics in that if $x, y \in \mathcal{E}(G)$ are distinct, then there is a G -invariant clopen partition of $\mathcal{E}(G)$ which separates x and y (see Lemma 9.3).

If we restrict to actions of Γ which are cocompact, then the construction of $\mathcal{E}(G)$ is, up to natural homeomorphism, independent of the choice of X , as is the notion of crossing. This can be seen from the fact that the constructions are quasi-isometry invariant. We shall not give details of the proof here. It is essentially the same as the proof that shows that the number of ends of a finitely generated group is quasi-isometry invariant. This observation allows us to speak unambiguously of the space of coends of any two-ended subgroup of a finitely generated one-ended group.

In general, a two-ended group may have infinitely many coends. A trivial example is the \mathbf{Z} -factor in $\mathbf{Z} \times (\mathbf{Z} * \mathbf{Z})$, where the space of coends is a Cantor set. However, I don't know if this can happen for self-commensurating subgroups.

Question. Suppose that G is a self-commensurating two-ended subgroup of a one-ended finitely generated group, Γ . Is the number of coends of G necessarily finite?

The possibility of the existence of an infinite number of coends adds a few complications to our constructions. A partial answer to the above question has recently been given by Scott and Swarup [ScS2], as we discuss at the end of this section.

Codim-one subgroups arise out of group splittings — any subgroup over which a group splits non-trivially is codim-one. More to the point for us, we have the following observations. Suppose that Γ is finitely generated and one-ended and admits a representation a finite graph of groups, \mathcal{G} . In other words, Γ acts simplicially and without edge inversions on a “Bass-Serre” tree, Σ , with quotient \mathcal{G} . Let X be any locally finite connected graph on which Γ acts properly discontinuously and cocompactly, and let $f : X \rightarrow \Sigma$ be a Γ -equivariant map sending vertices to vertices. Suppose that v is a vertex of Σ , and that S is a branch of Σ based at v (i.e. the closure of a connected component of $\Sigma \setminus \{v\}$). If $H \leq \Gamma$ is a two-ended subgroup of Γ fixing v , then we say that a coend of H is *associated to* the branch, S , if we can find $A \in \mathcal{S}(H)$ so that the corresponding element of $\mathcal{C}_\infty(A)$ gets mapped into S under f . (This can be seen to be independent of the choice of space X .) We note for example that if the stabiliser, G , of v is two-ended, then there are only finitely many branches of Σ based at v . This gives rise to a G -invariant clopen partition of

$\mathcal{E}(G)$, where different elements of the partition are associated to different branches of Σ .

The splittings we construct in this paper will be bipartite. We can write the vertex set, $V(\mathcal{G}) = V_0(\mathcal{G}) \sqcup V_1(\mathcal{G})$. The vertex groups in $V_0(\mathcal{G})$ will be of one of two types. Either they will be (conjugate to) fuchsian groups, with the adjacent edge groups as maximal parabolic subgroups, or else they will be commensurators of codim-one elements of Γ .

We shall describe how to construct such splittings starting with a finite collection of codim-one subgroups, possibly with the additional information of a clopen partition of the space of coends of each such subgroup. Subgroups which cross each other will end up in the same fuchsian vertex group of our splitting. The splitting we obtain will be canonical based on the starting data. Detailed statements will be given in Sections 14 and 15.

By taking larger and larger (finite) collections of codim-one subgroups we obtain more and more refined splittings. If Γ is finitely presented, then standard accessibility results (cf. [BeF]) tell us that the process terminates, modulo the addition of “reducible” vertices. In the case where the commensurator of any (codim-one) infinite-order element is two-ended, then we recover the usual JSJ splitting of Γ over two-ended subgroups, as a maximal splitting. In particular, this applies to hyperbolic groups, and so provides a link between the constructions of [DSw] and those of [Bo1].

In general, the situation is more complicated. Consider the case of the \mathbf{Z} -factor of $\mathbf{Z} \times (\mathbf{Z} * \mathbf{Z})$. Here, Γ splits over this subgroup in many different ways. There is, however, no canonical way to represent Γ as a graph of groups with this group as a vertex group, so our methods will not detect such a splitting. For this one needs other, non-canonical constructions. Such constructions can be found in [DSw] or [FP] for example.

We finish this section with some comments on other results about ends of pairs of groups. Connections between various end invariants are discussed in the recent book by Geoghegan [Geo]. I am indebted to G. A. Swarup and Ross Geoghegan for their comments on what follows. Suppose that Γ is finitely generated, and that $G \leq \Gamma$ is a subgroup. Let X be any Cayley graph of Γ . The number of ends, $e(\Gamma, G)$ of the pair (Γ, G) , can be defined as the number of ends of the space X/G . This again is independent of the choice of X , and can be given various homological characterisations (see, for example [K]). In our set-up where G is two-ended, $e(\Gamma, G)$ can be defined as the maximal cardinality of \mathcal{P}/G , where \mathcal{P} ranges over G -invariant partitions of the space of coends, $\mathcal{E}(G)$ (or as the cardinality of the hausdorffification of $\mathcal{E}(G)/G$). In the case where G is two-ended, we see that $\bar{e}(\Gamma, G) = \max\{e(\Gamma, H)\}$, as H ranges over finite index subgroups of G .

Another end invariant, $\bar{e}(\Gamma, G)$, associated to a pair (Γ, G) is studied in [KR]. This notion makes sense for any pair of groups $G \leq \Gamma$. If Γ is finitely generated and one-ended and G is two-ended, this reduces to the invariant $\bar{e}(\Gamma, G)$ defined earlier. The quantity $\bar{e}(\Gamma, G)$ may be defined as follows. Let $\mathcal{P}(\Gamma)$ be the power set of Γ , thought of as a \mathbf{Z}_2 -vector space admitting a linear Γ -action by right multiplication. Let $\mathcal{F}_G(\Gamma)$ be the Γ -invariant subspace of $\mathcal{P}(\Gamma)$ consisting of those $Q \subseteq \Gamma$ for which there exists a finite subset, $F \subseteq \Gamma$, such that $A \subseteq GF$. We define $\bar{e}(\Gamma, G)$ to be the dimension of $\mathcal{W} = (\mathcal{P}(\Gamma)/\mathcal{F}_G(\Gamma))^\Gamma$, i.e. the set of Γ -fixed points of the quotient space $\mathcal{P}(\Gamma)/\mathcal{F}_G(\Gamma)$.

Now suppose that Γ is finitely generated, and let X be any Cayley graph of Γ . Thus, $V(X) = \Gamma$ and $g, h \in \Gamma$ are adjacent if $g^{-1}h$ is a generator. Given $Q \subseteq \Gamma$, let $\text{bd}(Q)$ be the set of elements of Q adjacent in X to an element of the complement $\Gamma \setminus Q$. Now it is easily seen that Q represents an element, say $[Q]$, of \mathcal{W} if and only if $\text{bd}(Q) \in \mathcal{F}_G(\Gamma)$.

We explain why $\tilde{e}(\Gamma, G) = \bar{e}(\Gamma, G)$ in the case where $\tilde{e}(\Gamma, G) < \infty$, and leave the infinite case to the reader. (This applies to any subgroup G of a one-ended finitely generated group, Γ .) If $A \in \mathcal{S}(G)$ and P is a connected component of $X \setminus A$, then $[V(P)] \in \mathcal{W}$. Moreover, $[V(P)]$ is non-zero if and only if $P \in \mathcal{C}_\infty(A)$. We choose $A \in \mathcal{S}(G)$ so that $\mathcal{C}_\infty(A)$ is identified with $\mathcal{E}(G)$. We thus get a map $f: \mathcal{E}(G) \rightarrow \mathcal{W}$ by sending P to $[V(P)]$. Now, if $P_1, \dots, P_n \in \mathcal{C}_\infty(A)$ are distinct, and $n > 0$, then $V(P_1) \cup \dots \cup V(P_n) \notin \mathcal{F}_G(\Gamma)$. This shows that the elements $f(P)$ for $P \in \mathcal{E}(G)$ are distinct and linearly independent in \mathcal{W} . To see that they span, suppose Q is a subset of Γ with $\text{bd}(Q) \in \mathcal{F}_G(\Gamma)$. Choose $B \in \mathcal{S}(G)$ so that $A \cup \text{bd}(Q)$ lies in the interior of B . We again get a natural identification of $\mathcal{C}_\infty(B)$ with $\mathcal{E}(G)$. Let $\mathcal{C}_Q \subseteq \mathcal{C}_\infty(B)$ consist of those elements, P , for which $V(P) \subseteq Q$. Now we see that the symmetric difference of Q and $\bigcup_{P \in \mathcal{C}_Q} V(P)$ lies in $\mathcal{F}_G(\Gamma)$, showing that $[Q] = \sum_{P \in \mathcal{C}_Q} f(P)$ in \mathcal{W} . Thus $\{f(P) \mid P \in \mathcal{E}(G)\}$ is a basis for \mathcal{W} as required.

In [K], Kropholler shows that if Γ and G are both one-ended, and G is malnormal in Γ , then Γ splits over G if and only if $e(\Gamma, G) \geq 2$. Also, Γ splits over a subgroup of G if and only if $\tilde{e}(\Gamma, G) \geq 2$.

In [KR], it was shown that if G is a finitely generated subgroup of a finitely generated group Γ , and G has infinite index in its commensurator, then $\tilde{e}(\Gamma, G) \in \{1, 2, \infty\}$. On the other hand, in [ScS2], Scott and Swarup show that if Γ is either finitely presented or torsion-free, and $G \leq \Gamma$ is one- or two-ended and almost malnormal, then $e(\Gamma, G)$ is finite. In particular, this applies to self-commensurating two-ended subgroups. They go on to show that if Γ is finitely presented, and $G \leq \Gamma$ is two-ended and malnormal, then $\bar{e}(\Gamma, G)$ is finite.

2. PRETREES

The structures we use in this paper combine treelike structures and cyclic orders. We describe these in turn in the next two sections, beginning with a brief review of the notion of a “pretree”.

A *pretree* consists of a set, T , together with a ternary relation on T , denoted xyz for $x, y, z \in T$, satisfying the following axioms for $x, y, z, w \in T$: (T0): $xyz \Rightarrow x \neq z$, (T1): $xyz \Rightarrow zyx$, (T2): $xyz \Rightarrow \text{not } xzy$, (T3): xyz and $w \neq y \Rightarrow xyw$ or wyz . The intuitive interpretation of the relation xyz is that y lies “strictly between” x and z . The axioms express the treelike nature of this relation. Indeed, if $F \subseteq T$ is any finite subset, then one can embed F in a finite simplicial tree, Σ , such that if $x, y, z \in F$, then xyz holds in T if and only if x and z lie in different components of $\Sigma \setminus \{y\}$. In fact, this property serves as an equivalent definition of a pretree. (Sets F of cardinality at most 4 suffice.)

The notion of a pretree (under different names) appears in [W] and [AN], and is developed further in [Bo3]. Only the following observations are directly relevant to this paper.

Suppose T is a pretree, and $a \in T$. The relation, \sim , defined on $T \setminus \{a\}$ by $x \sim y$ if not xya is an equivalence relation. We refer to the equivalence classes as *branches* of T based at a .

A *closed interval* in T is a set of the form $[x, y] = \{x, y\} \cup \{z \in T \mid xzy\}$. We say that T is *discrete* if every interval is finite. A discrete pretree can be canonically embedded as a subset of the vertex set of a simplicial tree, Σ , so that the pretree relation on T agrees with the betweenness relation in Σ . In fact, we can take Σ to be a bipartite tree with vertex set $V(\Sigma) = V_0(\Sigma) \sqcup V_1(\Sigma)$, where $V_0(\Sigma)$ can be identified with T , and $V_1(\Sigma)$ can be identified with the set of “stars” of T . A subset $A \subseteq T$ is a *star* if there are no relations of the form xzy with $x, y \in A$ and $z \in T$ and if, for each $w \in T \setminus A$ there exist $x, y \in T$ with xyw . A star is deemed to be adjacent in Σ to each element of T that it contains. Every terminal vertex of Σ lies in $V_0(\Sigma)$. If $A \in V_1(\Sigma)$, then the branches of T based at A are in bijective correspondence with the adjacent vertices of Σ , and hence with the elements of A .

3. CYCLIC ORDERS AND PATTERNS

By a *cyclic order* on a set, Θ , we mean a quaternary relation which intuitively tells us when a pair of points are “linked with” (or “separated by”) another pair of points. It can be defined by the property that if $F \subseteq \Theta$ is any finite subset, then F can be embedded in the circle, S^1 , such that the quaternary relation restricted from Θ to F agrees with the standard cyclic order on S^1 . (In fact, it suffices to consider subsets of S^1 of cardinality at most 5.) Note that our cyclic order is “unoriented” unless otherwise stated. An *oriented cyclic order* can be similarly defined as a ternary relation on Θ which tells us when three points are ordered positively around Θ . Any cyclic order admits two compatible orientations.

Given a cyclically ordered set, Θ , we write $\Phi(\Theta)$ for the set of subsets of cardinality 2. (Thus, $\Phi(S^1)$ is an open Möbius band.) If $\lambda = \{x, y\}$ and $\mu = \{z, w\}$ are disjoint, we say that λ and μ *cross* if λ separates z from w , and we write $\lambda \times \mu$. Otherwise, we write $\lambda : \mu$.

By a *pattern* in Θ , we mean a subset, Π , of pairwise disjoint elements of $\Phi(\Theta)$. It is *full* if $\bigcup \Pi = \Theta$. We define the *nerve*, $N(\Pi)$, of Θ , as the graph whose vertex set is Π and where $\lambda, \mu \in \Pi$ are adjacent in $N(\Pi)$ if $\lambda \times \mu$. Following [Sw], we say that $N(\Pi)$ is *cross-connected* if $N(\Pi)$ is connected. In general, not every finite connected graph arises in this way, nor does $N(\Pi)$ determine Π . (Though for $|\Pi| \leq 4$, there is a bijection between isomorphism classes of cross-connected patterns and connected graphs.) In general, to recover the cyclic order, we need additional structure which we describe later (Lemma 3.2). It is often convenient to think of a finite pattern, Π , as a finite set of geodesics in the hyperbolic plane, \mathbf{H}^2 , whose ideal endpoints in $S^1 \equiv \partial \mathbf{H}^2$ are disjoint.

Suppose Π is a pattern on Θ . Given $\lambda \in \Pi$ and $x, y \in \Theta \setminus \lambda$, we write $x - \lambda - y$ to mean that λ separates x from y . Note that this partitions $\Theta \setminus \lambda$ into two intervals, so that $x - \lambda - y$ holds if and only if x and y lie in different intervals. We refer to these intervals as “sides” of λ . If Π is full, then we refer to the above ternary relation as a “cactaceous structure” on Θ . In the cross-connected case, we shall see that this structure determines the cyclic order on Θ (Lemma 3.2). First, we note:

Lemma 3.1. *Suppose that Π is a full cross-connected pattern on a finite cyclically ordered set, Θ . Then there is some $\lambda \in \Pi$ such that $\Pi \setminus \{\lambda\}$ is cross-connected.*

Proof. Choose any $\lambda \in \Pi$ which is not a global cut point of the nerve $N(\Pi)$. \square

Lemma 3.2. *Suppose that Θ and Θ' are cyclically ordered sets, and that $\Pi \subseteq \Phi(\Theta)$ and $\Pi' \subseteq \Phi(\Theta')$ are full cross-connected patterns. Suppose that $f : \Theta \rightarrow \Theta'$ is a bijection which sends Π to Π' , and such that for all $\lambda \in \Theta$ and $x, y \in \Theta$, $x - \lambda - y$ hold if and only if $f(x) - f(\lambda) - f(y)$ holds. Then f is a cyclic order isomorphism from Θ to Θ' .*

Proof. It's enough to verify this for finite subsets of Π and Π' . By enlarging these subsets if necessary, we can assume them to be cross-connected. We may therefore assume that Π and Π' are finite. We can imagine them as embedded in the circle. By Lemma 3.1, there is some $\lambda = \{x, y\} \in \Pi$ such that $\Pi \setminus \{\lambda\}$ (and hence $\Pi' \setminus \{\lambda'\}$) is cross-connected. By induction on $|\Pi|$, we can assume that $f|_{\Theta \setminus \lambda}$ is an order isomorphism. Now, if $\mu \in \Pi \setminus \{\lambda\}$, then the relations of the form $x - \lambda - z$ for $z \in \Theta \setminus (\lambda \cup \mu)$ determine on which side of μ the point x must lie. As μ varies over $\Theta \setminus \{\lambda\}$, these sides determine a complementary interval of $\Theta \setminus \lambda$ in the circle. The corresponding relations in Θ' determine the corresponding complementary interval of $\Theta' \setminus f(\lambda)$. We see that $f|_{\Theta \setminus \{y\}}$ is an order isomorphism of $\Theta \setminus \{y\}$ onto $\Theta' \setminus \{f(y)\}$. The same goes for $f|_{\Theta \setminus \{x\}}$. But now, since Π is cross-connected, x and y lie in different complementary intervals of $\Pi \setminus \lambda$. Thus, the quaternary cyclic order relations involving both x and y are determined by those involving just one of x or y . It follows that f is an order isomorphism as required. \square

We note that Lemma 2.2 fails if Π is not cross-connected. Suppose, for example, that Π is finite and has two cross-connected components. If we imagine $\Theta = \bigcup \Pi$ as embedded in S^1 , then we can find $x, y \in S^1$ such that $\{x, y\}$ separates these components. Now, we can reverse the orientation of one of these components without changing the cactaceous structure. In fact, we can more naturally imagine Θ as embedded in the “figure-of-eight” obtained by identifying x and y . This is sufficient to recover the cyclic order on either component, but not on Θ as a whole. This construction gives us an example of a “cactus” as described in the next section.

We introduce the following discreteness condition on patterns, which will be elaborated on, in a more general context, in Section 6.

Definition. Suppose that Π is a pattern on a cyclic order Θ . We say that Π is *discrete* if, given any pairs $\lambda, \mu \in \Phi(\Theta)$, only finitely many elements of Π separate λ from μ .

If $\Theta = S^1$, then this is equivalent to asserting that Π is topologically discrete as a subset of the Möbius band, $\Phi(S^1)$. Note also that if $\bigcup \Pi$ is dense in Θ , then it is sufficient to apply the definition to pairs λ and μ in $\bigcup \Pi$ — in other words discreteness in this case is intrinsic to Π .

The notion of discreteness gives us the following criterion recognising convergence actions on the circle. It is a slight rephrasing of the criterion due to Swenson [Sw]. (See [GerM], [T2] or [Bo3] for a discussion of convergence groups.) Suppose that a group, Γ , acts by homeomorphism on the circle, S^1 . Given a Γ -invariant pattern, $\Pi \subseteq S^1$, and $\lambda \in \Pi$, we write $\Gamma(\lambda)$ for the setwise stabiliser of λ .

Theorem 3.3. *Suppose that Π is a discrete cross-connected Γ -invariant pattern with $\bigcup \Pi$ dense in S^1 and with Π/Γ finite. Suppose that if $\lambda, \mu \in S^1$ with $\lambda \times \mu$, then $\Gamma(\lambda) \cap \Gamma(\mu)$ is finite. Suppose also that, for all $\lambda \in \Pi$, there then is some infinite order element $g \in \Gamma(\lambda)$. Then, Γ acts as a convergence group in S^1 .*

Proof. If $\lambda \in \Pi$ then there is some $\mu \in \Pi$ with $\lambda \times \mu$. By discreteness, it is easy to see that $g^2\mu : \mu$ and that $g^{2n}\mu$ converges on one point of λ , while $g^{-2n}\mu$ converges on the other. Again, applying discreteness it is now easy to see that $\langle g \rangle$ has finite index in $\Gamma(\lambda)$ and that $\Gamma(\lambda)$ acts as a convergence group on S^1 . The main result of [Sw] now applies. \square

In fact, in the above result, we could replace the circle by any dense cyclic order, Θ , and with Γ acting by order isomorphism on Θ and satisfying the same hypotheses. This follows since Θ can be canonically completed to a circle. We shall apply the result, in this form, in Sections 7 and 13.

The interest in convergence actions on the circle follows from the following central result:

Theorem 3.4 ([T1], [Ga], [CJ]). *Suppose that Γ acts as a convergence group on S^1 . Then, Γ also admits a properly discontinuous action on the hyperbolic plane, \mathbf{H}^2 , so that the induced action on the boundary $\partial\mathbf{H}^2$, is topologically conjugate to the original.*

If Γ is finitely generated, then it is virtually the fundamental group of a surface whose peripheral subgroups are precisely the (conjugacy classes of) maximal parabolic subgroups of this action. We shall refer to such a group, together with its peripheral structure, as a “fuchsian group”. Note that if Γ is one-ended, then there are no peripheral subgroups, and Γ is a virtual (closed) surface group. We refer to such a group as a “planar group”.

4. PATTERNS IN GRAPHS

We can think of a finite pattern as defined in the last section as a set of pairs of vertices of a graph that happens to be homeomorphic to a circle, i.e. every vertex has degree 2. In this section, we generalise this to arbitrary connected graphs.

Let Υ be a finite connected graph, with vertex set $V = V(\Upsilon)$. (We do not allow Υ to have loops or multiple edges.) Let $\Phi(\Upsilon)$ be the set of unordered pairs of elements of V .

Definition. A subset, $\Pi \subseteq \Phi(\Upsilon)$ of pairwise disjoint pairs of vertices of Υ is a *pattern* if no element of $\bigcup \Pi$ is a cut vertex of Υ .

We say that Π is a *proper pattern* if, in addition, $\Upsilon \setminus \lambda$ is disconnected for all $\lambda \in \Pi$.

Note that this is automatically satisfied if Υ happens to be a circle, so that the patterns discussed in Section 3 are all examples of proper patterns.

We write $V_0 = \bigcup \Pi \subseteq V$ and $V_1 = V \setminus V_0$. Given $x \in V_0$, we write x^* for the vertex of V_0 such that $\{x, x^*\} \in \Pi$. We write $\hat{x} = \{x, x^*\}$.

Given $x, y \in V_0$, and $\lambda \in \Pi$, we write $x - \lambda - y$ to mean that x and y lie in different components of $\Upsilon \setminus \lambda$. If $\mu = \hat{x} \in \Pi$, we write $\mu \times \lambda$ to mean that $x - \lambda - x^*$, and we say that μ *crosses* λ .

Lemma 4.1. *Suppose that Π is a proper pattern. If $\lambda, \mu \in \Pi$ and $\mu \times \lambda$, then $\lambda \times \mu$.*

Proof. This is a simple exercise. It also follows from Proposition 4.4 below. \square

Thus, as in Section 2, we can speak of a proper pattern as being “cross-connected” and of “cross-connected components” of proper patterns.

We can think of the pattern Π together with the ternary relation $. - . - .$ defined above as defining a kind of structure on V_0 , which we shall refer to as a “cactaceous structure”. There are a number of simplifying operations we can perform on the graph, Υ , without changing this structure.

First note that we can insert a vertex of degree 2 between any two adjacent vertices of V_0 . Moreover, we can contract any component of $\Gamma \setminus V_0$ to a point (and collapse any set of multiple edges arising to a single edge). In this way, there is no loss of generality in supposing that Υ is bipartite, with the vertices partitioned as $V = V_0 \sqcup V_1$.

More generally, we can define a relation \approx on V_1 by writing $x \approx y$ if, for all $\lambda \in \Pi$, x and y lie in the same component of $\Upsilon \setminus \lambda$. Now, Π can be considered as a pattern on the graph Υ/\approx , and the induced cactaceous structure remains the same. Replacing Υ by Υ/\approx , we can thus assume that every pair of points of V_0 are separated by some element of Π . We say that a pattern on a bipartite graph with this property is *efficient*. Note that all the above operations preserve the propriety of a pattern.

Suppose that Π is a pattern on Υ . Given $\lambda \in \Pi$, we write $\mathcal{C}(\lambda)$ for the set of connected components of $\Upsilon \setminus \lambda$. If Υ is bipartite, $C \in \mathcal{C}(\lambda)$ and $x \in \lambda$, then all the vertices adjacent to x in C lie in V_1 , and they are all equivalent under the relation \approx . Thus, if Π is efficient, there is only one such vertex. We deduce:

Lemma 4.2. *Suppose Π is an efficient pattern on Υ and $\lambda = \{x, y\} \in \Pi$. If $C \in \mathcal{C}(\lambda)$, then C has precisely one edge incident on x and one edge incident on y .*

In this case, we write $n(\lambda) = |\mathcal{C}(\lambda)| = \deg x = \deg y$ (where \deg denotes “degree”). Thus, Π is proper if $n(\lambda) \geq 2$ for all $\lambda \in \Pi$.

Lemma 4.3. *Suppose Π is a proper efficient pattern. Suppose that $x, y \in V_0$ with $y \notin \hat{x}$. If $\deg x \geq 3$, then x and x^* lie in the same element of $\mathcal{C}(\hat{y})$.*

Proof. Since $|\mathcal{C}(\hat{x})| = \deg x \geq 3$, there is some $C \in \mathcal{C}(\hat{x})$ with $y, y^* \notin C$. Thus, C connects x to x^* in $\Upsilon \setminus \hat{y}$. □

Suppose that Υ is (any) connected graph. Recall that a *block* of Υ is a maximal 2-vertex-connected subgraph. Let $\mathcal{D} = \mathcal{D}(\Upsilon)$ be the set of blocks of Υ . Let $\Sigma(\Upsilon)$ be the bipartite graph with vertex set $V(\Sigma) = V(\Upsilon) \sqcup \mathcal{D}(\Upsilon)$ such that $x \in V(\Upsilon)$ is deemed to be adjacent to $D \in \mathcal{D}(\Upsilon)$ in Σ if $x \in D$ in Υ . A well-known fact of elementary graph theory tells us that Σ is a tree. Moreover, Υ is bipartite if and only if each of its blocks is bipartite.

Suppose now that Π is a pattern on the finite connected graph Υ . Since no element of V_0 is a global cut point, two blocks of Υ meet, if at all, in a single vertex of V_1 . Note that if $x, y \in V_0$ lie in different blocks, then $\Upsilon \setminus \{x, y\}$ is connected. If Π is proper, we deduce that if $\lambda \in \Pi$, $D \in \mathcal{D}(\Upsilon)$ and $\lambda \cap D \neq \emptyset$, then $\lambda \subseteq D$. We write Π_D for the set of such λ . Thus $\Pi = \bigsqcup_{D \in \mathcal{D}} \Pi_D$. It is easily checked that Π_D is a pattern on D . Moreover, Π is proper (in Υ) if and only if Π_D is proper in D for all $D \in \mathcal{D}$. Similarly Π is efficient if and only if each block is efficient.

We have the following structure theorem for efficient proper patterns:

Proposition 4.4. *Suppose that Π is an efficient proper pattern on a finite connected graph, Υ . Then, each block, D , of Υ is either a circle (with alternating vertices in V_0 and V_1) or else $V_0 \cap D$ consists of a single element of Π and D is the complete bipartite graph on $V_0 \cap D$ and $V_1 \cap D$.*

Proof. From the discussion above, we may as well assume that Υ is 2-vertex-connected. We want to show that Υ has one of the two types described.

Suppose first that there is some $x \in V_0$ with $\deg x \geq 3$. Thus $|\mathcal{C}(\hat{x})| \geq 3$. Given $A \in \mathcal{C}(\hat{x})$, let $a, b \in V_1(\Upsilon)$ be the vertices adjacent to x and x^* respectively. We claim that $a = b$. To see this, suppose that $\lambda \in \Pi \setminus \{\hat{x}\}$. By Lemma 4.3, $x, x^* \in B \in \mathcal{C}(\hat{y})$. It follows that $a, b \in B$. In other words, a and b are not separated by λ . Since $\lambda \in \Pi \setminus \{\hat{x}\}$ and since Π is efficient, it follows that $a = b$. We deduce that each element of V_1 is adjacent to both x and x^* . Thus, in this case, Υ is a complete bipartite graph on $V_0 = \hat{x}$ and V_1 .

We can now suppose that every vertex of V_0 has degree 2. Let Ψ be the graph (with vertex set V_1) obtained by deleting every vertex of V_0 and merging the two adjacent edges. (We allow the possibility of Ψ having multiple edges.) It is easily seen that Ψ satisfies the hypotheses of Lemma 4.5 below, and is hence a circle. It follows that Υ is a circle as required. \square

Lemma 4.5. *Suppose that Ψ is a finite 2-vertex-connected graph (possibly with multiple edges). If every pair of distinct vertices of Ψ are separated by some pair of edges of Ψ , then Ψ is a circle.*

Proof. If not, then we can find two vertices x and y which are connected by three arcs which intersect pairwise precisely in $\{x, y\}$. Clearly x and y are not separated by any pair of edges. \square

We can draw a number of immediate corollaries from Proposition 4.4. Firstly we deduce Lemma 4.1, since if $\lambda, \mu \in \Pi$ with $\mu \times \lambda$, then λ and μ must lie in the same circular block, and it follows that $\lambda \times \mu$. We can thus define the nerve, $N(\Pi)$, of Π exactly as for patterns on the circle, and we can speak of cross-connected components of Π . Clearly each cross-connected component must lie in a subpattern, Π_D , for some $D \in \mathcal{D}$. In fact:

Lemma 4.6. *Let Π be an efficient proper pattern on Υ and let \mathcal{D} be the set of blocks of Υ . Then, for all $D \in \mathcal{D}$, the subpattern Π_D is cross-connected.*

Proof. We need only check this for a circular block, D . It is easily seen that if Π_D is not cross-connected, then there must be two vertices of $V_1 \cap D$ not separated by any element of Π_D and hence by any element of Π , thereby contradicting efficiency. \square

We see that the subpatterns, Π_D , for $D \in \mathcal{D}$ are precisely the cross-connected components of Π . The set of such components embeds in the finite tree $\Sigma(\Upsilon)$, and so has a natural treelike structure.

The following observations are also simple consequences of the structure theorem, Proposition 4.4.

Suppose $\lambda \in \Pi$. The relation defined on $V_0 \setminus \lambda$ by $x \sim y$ if not $x - \lambda - y$ is an equivalence relation. If some element of Π crosses λ , then there are precisely two equivalence classes. Suppose that $x, y, z \in V_0$ and that $x - \hat{y} - z$ and $x - \hat{z} - y$ hold, then $\hat{y} \times \hat{z}$. If $x, y, z \in V_0$ and $\hat{x} : \hat{z}$, $\hat{x} : \hat{y}$ and $y - \hat{x} - z$ hold, then $\hat{y} : \hat{z}$ holds. (Recall that $\lambda : \mu$ means that λ and μ do not cross.)

It will sometimes be convenient to abbreviate $x - \hat{y} - z$ to $x - y - z$. In this way, we can view $.-.-.$ as a ternary relation on V_0 . We similarly write $x \times y$ for $\hat{x} \times \hat{y}$ and $x : y$ for $\hat{x} : \hat{y}$.

We can think of the graph described by Proposition 4.4 as a tree of circles, and hence the terminology “cactaceous structures” — it is meant to resemble a cactus

of the genus *Opuntia* for example. In the next section, we explain how this serves as a basis for a more general definition of cacti.

Much of the above discussion can be generalised to patterns which are not necessarily proper. In this case the crossing relation on pairs need not be symmetric, so we cannot speak of “cross-connected components” as such. If Π is an efficient pattern, then the set of terminal vertices of Υ is precisely the union of pairs, λ , for which $\Upsilon \setminus \lambda$ is connected. If we throw away all such vertices and their incident edges then we would be left with an efficient proper pattern. This tells us that Υ has the form described by Proposition 4.4 except that we need to allow for a number of terminal vertices of V_0 attached at certain vertices of V_1 . This time, there is no reason for a non-separating pair to lie in a single block of Υ .

Suppose that Π is any pattern in Υ and $\lambda \in \Pi$. If λ does not cross any other element of Π , then we can adjoin another vertex of V_1 which we deem to be adjacent to both elements of λ , without altering the ternary relation on V_0 . In particular we see that if Π has the property that, for each $\lambda \in \Pi$ which crosses some other element of Π , there exist $x, y \in \Pi$ with $x - \lambda - y$, then we can enlarge Υ so that Π becomes proper.

5. CACTI

In this section, we define the notion that a “cactus” generalises the cactaceous structures of the last section.

Let Ω be a set, and let Π be a partition of Ω into disjoint ordered pairs. We have a fixed point free involution, $[x \mapsto x^*]$, on Ω such that for any $x \in \Omega$, $\hat{x} = \{x, x^*\} \in \Pi$.

Suppose we have a ternary relation $.-.-.$ on Ω so that $y - x - z$ holds if and only if $y - x^* - z$ holds. We shall also write $x - \lambda - y$ where $\lambda = \hat{x}$.

Definition. A *cactus* is a set, Ω , with a partition Π and ternary relation, $.-.-.$, as above, with the property that for any finite subset $\Pi_0 \subseteq \Pi$, we can realise Π_0 as a pattern on a finite graph, Υ , such that the ternary relation induced on $V_0(\Upsilon) = \bigcup \Pi_0$ (as described in Section 4) agrees with the restriction of the ternary relation on Ω .

We say that Ω is a *proper cactus* if we can always take Π_0 to be a proper pattern on Υ .

In other words, any finite $*$ -invariant subset of Ω is cactaceous in the sense defined in the previous section. Given a $*$ -invariant subset $P \subseteq \Omega$, we write $\Pi_P = \{\lambda \in \Pi \mid \lambda \subseteq P\}$.

If $\lambda \in \Pi$, then it is easily seen that the relation \sim_λ defined on $\Omega \setminus \lambda$ by $x \sim_\lambda y$ if not $x - \lambda - y$ is an equivalence relation on $\Omega \setminus \lambda$. We refer to an equivalence class as a *side* of λ . More generally, if $P \subseteq \Omega$ is any $*$ -invariant subset, we define a relation \sim_P on $\Omega \setminus P$ by writing $x \sim_P y$ to mean that $x - z - y$ does not hold for any $z \in P$.

As for patterns, if $\lambda, \mu \in \Omega$, we write $\mu \times \lambda$ to mean that μ crosses λ , i.e. that $x - \lambda - x^*$, where $\mu = \hat{x}$. The following is an immediate consequence of the discussion at the end of Section 4:

Lemma 5.1. *Suppose $(\Omega, \Pi, .-.-.)$ is a cactus. Then the following are equivalent:*

- (1) Ω is proper.
- (2) For all $\lambda, \mu \in \Pi$ with $\lambda \times \mu$, we have $\mu \times \lambda$.
- (3) If $\lambda, \mu \in \Pi$ with $\lambda \times \mu$, then there exist $x, y \in \Omega$ with $x - \lambda - y$.

Let us suppose, for the moment, that Ω is proper. We can thus speak of “cross-connected” components of Ω (or of Π) as usual. Let $\mathcal{T} = \mathcal{T}(\Omega)$ be the set of cross-connected components. Given $P, Q, R \in \mathcal{T}$, we write QPR to mean that there exist $x \in P$, $y \in Q$ and $z \in R$ such that $y - x - z$. Note that if $w \in R$, then we must also have $y - x - w$. (It is sufficient to verify this in the case where $z \times w$. But if not $y - x - w$, then we must have $z - x - w$. Since $z : x$ and $w : x$, we would derive the contradiction that $z : w$.) We see that, in fact, $y - x - z$ holds for all $y \in Q$ and all $z \in R$.

It now follows easily that:

Lemma 5.2. *With the ternary relation defined above, \mathcal{T} is a pretree.*

We shall say that an element $\lambda \in \Pi$ is *isolated* if it is not crossed by any other element of Π ; in other words if $\lambda \in \mathcal{T}$. We write of \mathcal{T}_0 for the subset of \mathcal{T} consisting of such isolated pairs. We write $\mathcal{T}_\infty = \mathcal{T} \setminus \mathcal{T}_0$. The following result shows how we can put a natural cyclic order on each element of \mathcal{T}_∞ .

Lemma 5.3. *Suppose that $P \in \mathcal{T}_\infty$. Then there is a unique cyclic order on P such that the cactaceous structure induced on the pattern Π_P agrees with that coming from Ω .*

Proof. The corresponding statement for any finite cross-connected subpattern of Ω follows from Lemma 3.2 (uniqueness) and Proposition 4.4 (existence). But the set of finite cross-connected subpatterns is cofinal in the set of all finite subsets of P . We thus get a unique cyclic order on P as required. \square

Given $P \in \mathcal{T}$, let \sim_P be the equivalence relation on $\Omega \setminus P$ defined above. We write $\mathcal{B}(P) = (\Omega \setminus P) / \sim_P$ for the set of equivalence classes. We see that there is a natural bijection between $\mathcal{B}(P)$ and the set of branches of the pretree \mathcal{T} based at P . More precisely, if B is such a branch, then $\bigcup B \in \mathcal{B}(P)$. Let $S(P)$ be the abstract disjoint union $S(P) = P \sqcup \mathcal{B}(P)$. There is a natural projection $\phi : \Omega \rightarrow S(P)$ which is the identity on P and which sends each point of $\Omega \setminus P$ to its \sim_P -class.

Suppose now that $P \in \mathcal{T}_\infty$. We can extend the cyclic order on P to $S(P)$:

Lemma 5.4. *There is a unique cyclic order on $S(P)$ with the property that if $x \in P$ and $y, z \in \Omega \setminus \hat{x}$, then $y - x - z$ holds if and only if $\{\phi(y), \phi(z)\}$ crosses \hat{x} in this cyclic order.*

Proof. Let $P_0 \subseteq P$ and $\mathcal{B}_0(P) \subseteq \mathcal{B}(P)$ be finite subsets. We can find a finite cross-connected pattern $\Pi_0 \subseteq \Pi_P$ with the property that $\Pi_0 \subseteq P'_0$ where $P'_0 = \bigcup \Pi_0$, and such that two elements of $\mathcal{B}_0(P)$ are separated by some element of Π_0 . For each $B \in \mathcal{B}_0(P)$, choose any $\lambda_B \in \Pi$ with $\lambda_B \subseteq B$. Let $\Pi' = \Pi_0 \cup \{\lambda_B \mid B \in \mathcal{B}_0(P)\}$, and let $P' = \bigcup \Pi'$. Thus, $P' \supseteq P$ is finite. We now apply Proposition 4.4 to the finite pattern Π' on P' . The resulting graph, Υ , consists of a central circular block, C , with $V_0(C) \equiv P'_0$, and with a number of square blocks attached to certain vertices of $V_1(C)$. These square blocks each contain one pair, λ_B , and are hence in bijective correspondence with $\mathcal{B}_0(P)$. In this way, we put a cyclic order on $S_0 = P'_0 \cup \mathcal{B}_0(P)$, which has the required properties for this subset. We know (Lemma 2.2) that the order restricted to P'_0 is unique, and it is easily deduced that the order on S_0 is at least determined by the choice of λ_B .

Suppose we chose a different set of pairs, $\{\lambda'_B \mid B \in \mathcal{B}_0(P)\}$. Let $\Pi'' = \Pi_0 \cup \{\lambda'_B, \lambda''_B \mid B \in \mathcal{B}_0(P)\}$ and let $P'' = \bigcup \Pi''$. Applying Proposition 4.4 to the pattern Π'' , we get a graph, Υ' , consisting of a central circular block, C' , again

with $V_0(C') \equiv P'_0$, and with connected graphs attached at various points of $V_1(C')$. Each such graph (which is either an octagon or two squares joined at a vertex) contains both λ_B and λ'_B for some $B \in \mathcal{B}_0(P)$. In this way, we see that the order induced on S_0 is independent of the choice of the λ_B .

In summary, for each subset of the type $S_0 \subseteq S(P)$, we have a well-defined cyclic order on S_0 . Now, the set of such subsets is cofinal in the set of all finite subsets of $S(P)$, so we get a well-defined cyclic order on $S(P)$. \square

In fact, the conclusions of Lemmas 5.2 and 5.4 could be used to give an alternative formulation of the notion of a proper cactus. It is also possible to write down a complete set of axioms for a proper cactus in terms of the partition Π and the ternary relation $\cdot - \cdot - \cdot$, though we shall not describe this here. We note that any set with free involution, $*$, and a ternary relation is a cactus if and only if every finite $*$ -invariant subset is a cactus. (However, unlike the analogous cases of pretrees and cyclic orders, it is not sufficient to restrict to subsets of any given finite cardinality: it is possible to find, for arbitrarily large $n \in \mathbf{N}$, a structure on a set Ω of $2n$ points with the property that deleting any element of Π gives us a cactus, whereas Ω is not itself a cactus.)

EXAMPLES

We finish off this section by describing how cacti arise from continua. In Sections 9 and 10, we shall see further examples arising in group theory. In the case of hyperbolic groups, these constructions are closely related (cf. [Bo1]).

Recall that a *continuum*, M , is a compact connected hausdorff topological space. We say that two points, $x, y \in M$ are *separated* by a closed subset, $F \subseteq M$, if they lie in different quasicomponents of $M \setminus F$, i.e. we can write $M \setminus F$ as a disjoint union of two open subsets, $M \setminus F = O \sqcup U$, with $x \in O$ and $y \in U$. If $F = \lambda$ is a pair of points, we write $x - \lambda - y$. A point of M is a (*global*) *cut point* if it separates two points of M (i.e. if its complement is disconnected).

Proposition 5.5. *Suppose that M is a continuum and that $\Pi \subseteq \Phi(M)$ is a set of pairwise disjoint unordered pairs of M . Let $\Omega = \bigcup \Pi$. Suppose that no point of Ω is a global cut point. With the ternary relation described above, Ω is a cactus. Moreover, if $M \setminus \lambda$ is disconnected for all $\lambda \in \Pi$, then Ω is proper.*

Proof. It is enough to verify this in the case where Ω is finite. It is a consequence of the following observation (see [Bo2]). Let $F \subseteq M$ be any finite subset. Let Υ be the graph with vertex set F where $x, y \in F$ are deemed to be adjacent if they are not separated by $F \setminus \{x, y\}$. Then, if $x, y \in F$ and $C \subseteq F \setminus \{x, y\}$ is any subset, then x and y are separated by C in M if and only if they are separated by C in Υ . Taking $C = \emptyset$, we see that Υ is necessarily connected.

To verify that Ω is a cactus, we take $F = \Omega$, or indeed any finite subset of M containing Ω . Now, from the hypothesis on Ω , we see that no point of $F \equiv V_0(\Upsilon)$ is a cut vertex of Υ . Thus, Π is a pattern on Υ . Moreover, the ternary cactaceous relation on Υ agrees with the separation relation coming from M . This shows that Ω is a cactus.

For the second statement, we need to choose F so that each element of Π separates some pair of points of F in M . It follows that this is also true in Υ , and so Π is a proper pattern in Υ as required. \square

Note that we do not need to assume that M is locally connected for this to work. In the case where $M = S^1$, the hypotheses on Π are necessarily satisfied. We therefore recover the notion of a pattern described in Section 3. All such patterns are proper.

6. DISCRETENESS CONDITIONS

Suppose that $(\Omega, \Pi, \cdot - \cdot - \cdot)$ is a cactus. Given $\lambda \in \Pi$ and subsets $P, Q \in \Omega$, we write $P - \lambda - Q$ to mean that for all $x \in P$ and all $y \in Q$, we have $x - \lambda - y$. If $\mu, \nu \in \Pi$, then $\mu - \lambda - \nu$ implies that λ, μ and ν are pairwise non-crossing. Note also that if $\lambda \in \Pi, x, y \in \Omega$, then the tree conditions $x - \lambda - y$ and $\hat{x} : \lambda$ and $\hat{y} : \lambda$ together imply $\hat{x} - \lambda - \hat{y}$.

We can impose the following discreteness conditions on Ω .

(F0) If $x, y, z, w \in \Omega$ are distinct, then $\{\lambda \in \Pi \mid \{x, y\} - \lambda - \{z, w\}\}$ is finite.

(F1) If $\mu, \nu \in \Pi$ satisfy $\mu : \nu$, then $\{\lambda \in \Pi \mid \mu - \lambda - \nu\}$ is finite.

(F2) If $\mu, \nu \in \Pi$ are distinct, then $\{\lambda \in \Pi \mid \lambda \times \mu, \lambda \times \nu\}$ is finite.

Clearly (F0) implies both (F1) and (F2). Note that (F0) corresponds to the definition of “discreteness” for patterns on the circle given in Section 3.

Lemma 6.1. *If Ω is a proper cactus, then (F1) and (F2) implies (F0).*

Proof. Suppose that Ω is a proper cactus satisfying (F1) and (F2). Suppose that $x, y, z, w \in \Omega$ are distinct. First, we deal with the case where $z = x^*$, so that $\mu = \{x, z\} \in \Pi$.

Now if $\lambda \in \Pi$ satisfies $\{x, y\} - \lambda - \{z, w\}$, then at least one of the following possibilities hold:

- (a) $\lambda \times \hat{y}$,
- (b) $\lambda \times \hat{w}$, or
- (c) $\lambda : \hat{y}$ and $\lambda : \hat{w}$.

Now, since $\lambda \times \mu$, by (F2) only finitely many such λ can satisfy (a) or (b). Moreover, condition (c) implies that $\hat{y} - \lambda - \hat{w}$ holds. Thus, by (F1), this can again arise for only finitely many λ . In other words, we have shown that, $\{\lambda \in \Pi \mid \{x, y\} - \lambda - \{z, w\}\}$ is finite.

To deal with the general case (x, y, z, w arbitrary) we repeat the above argument, this time using the result already obtained, in place of (F2), to deal with cases (a) and (b). □

Definition. We say that a cactus is *discrete* if it satisfies (F0).

Note that this implies that the pretree, $\mathcal{T}(\Omega)$ as defined in Section 5 is also discrete.

EXAMPLES

Let M be a continuum. Recall that a point $x \in M$ is said to be a *weak cut point* if there exist $y, z \in M \setminus \{x\}$ such that every subcontinuum of M containing both y and z also contains x . Let $\Phi(M)$ be the space of unordered pairs of M with the obvious topology (i.e. as a quotient of $M \times M$ minus the diagonal).

Proposition 6.2. *Suppose that M is a continuum with no weak cut point. Suppose that $\Pi \subseteq \Phi(M)$ is a set of pairwise disjoint pairs of M . Let $\Omega = \bigcup \Pi$ (so that Ω is a cactus by Proposition 5.5). Suppose that Π is discrete as a subset of $\Phi(M)$. Then, Ω is discrete (as a cactus).*

Proof. Suppose, for contradiction, that there are distinct points $x, y, z, w \in \Omega$ such that $\Pi_0 = \{\lambda \in \Pi \mid \{x, y\} - \lambda - \{z, w\}\}$ is infinite. It follows that there is some point $a \in M$ such that every neighbourhood of a contains an element of Π_0 . Without loss of generality, $a \notin \{x, z\}$. Since a is not a weak cut point, there is a continuum, $K \subseteq M \setminus \{x\}$, with $x, z \in K$. But now we can find $\lambda \in \Pi_0$ with $\lambda \cap K = \emptyset$. Thus, x and z are not separated by λ , contrary to the assumption that $\{x, y\} - \lambda - \{z, w\}$. \square

Note that every (global) cut point of a continuum is a weak cut point, though the converse fails in general. However, these notions coincide for Peano continua (i.e. locally connected continua). Thus, Proposition 6.2 applies to Peano continua with no global cut points, and hence to the set-up described in [Bo1].

7. GROUP ACTIONS

In this section, we consider group actions on discrete proper cacti. We show, under the hypotheses (1)–(4) below, that such an action gives rise to a kind of JSJ splitting. In the case of cross-connected patterns this reduces to the result of [Tu], [Ga], [CJ], [Sw]. In the general case, we use this result to recognise the “hanging fuchsian” subgroups of the splitting.

Suppose that the group, Γ , acts on a cactus, $(\Omega, \Pi, . - . - .)$. Given $\lambda \in \Pi$, we write $\Gamma(\lambda)$ for the setwise stabiliser of λ , and $\Gamma_0(\lambda)$ for the pointwise stabiliser of λ . Clearly, $\Gamma_0(\lambda)$ has index at most 2 in $\Gamma(\lambda)$. We shall assume:

- (1) Π is countable, discrete and proper,
- (2) for all $\lambda \in \Pi$, $\Gamma(\lambda)$ contains an infinite-order element,
- (3) if $\lambda, \mu \in \Pi$, then $\Gamma(\lambda) \cap \Gamma(\mu)$ is finite, and
- (4) Π/Γ is finite.

As in Section 5, we shall say that an element of Π is “isolated” if it is not crossed by any other element of Π . We shall see (by Lemma 7.2) that if λ is not isolated, then $\Gamma(\lambda)$ is two-ended. We shall not, in general, assume this of isolated pairs.

Let \mathcal{T} be the set of cross-connected components of Ω . Let $\mathcal{T}_1 \subseteq \mathcal{T}$ be the set of isolated pairs, and let $\mathcal{T}_\infty = \mathcal{T} \setminus \mathcal{T}_1$. Clearly, each element of \mathcal{T}_∞ is infinite. We shall show:

Proposition 7.1. *Γ admits a (possibly trivial) canonical splitting as a finite bipartite graph of groups, \mathcal{G} , with vertex set $V(\mathcal{G}) = V_0(\mathcal{G}) \sqcup W(\mathcal{G})$. The vertex groups of $V_0(\mathcal{G})$ are in bijective correspondence with the conjugacy classes of (set-wise) cross-connected component stabilisers of Ω , where the corresponding groups are isomorphic. Those corresponding to infinite components (in \mathcal{T}_∞) are fuchsian relative to the incident edge groups.*

By the last statement, we mean that such a vertex group, G , admits a properly discontinuous action on the hyperbolic plane, whose limit set is the whole circle, and such that each incident edge group, H , is the stabiliser of an ideal point, and is either finite or parabolic. Since parabolic subgroups are two-ended, it follows that all incident edge groups are finitely generated. If Γ is finitely generated, it therefore follows that G is finitely generated. In this case, the action on the hyperbolic plane must have finite coarea. If we assume, in addition, that Γ is one-ended, none of edge groups can be finite, so each of the incident edge groups must be two-ended and parabolic. Moreover, every parabolic subgroup must arise in this way. (Otherwise, G would split over a finite group relative to the incident edge groups, thereby giving rise to a splitting of Γ over a finite group.) We conclude:

Proposition 7.2. *If, in addition to the properties (1)–(4), we assume:*

(5): Γ is finitely generated and one-ended,

then each of the fuchsian vertex groups (corresponding to elements of \mathcal{T}_∞) is finitely generated (hence finite coarea), and the incident edge groups are precisely the maximal parabolic subgroups.

In such a case, we refer to such a subgroup as a *hanging fuchsian subgroup* and an incident edge groups as a *peripheral subgroup*.

Note that the assumption that Γ is finitely generated implies automatically that Ω is countable.

If Ω is infinite and cross-connected, then the splitting is trivial, consisting of a single fuchsian group with no peripheral subgroups, i.e. Γ is a planar group.

We now go on to describe the construction of the above splitting.

We saw in Section 5, that the set, \mathcal{T} , of cross-connected components has the structure of a pretree. In this case, \mathcal{T} is discrete, and so embeds canonically in a bipartite simplicial tree, Σ , as the vertex set of one colour, $V_0(\Sigma) \subseteq V(\Sigma)$. The action of Γ on \mathcal{T} extends to a simplicial action on Σ . Since $V_0(\Sigma)/\Gamma$ is finite, and since no vertex of $V_1(\Sigma) = V(\Sigma) \setminus V_0(\Sigma)$ is terminal, it is easily seen that Σ/Γ is finite. We shall see that Σ/Γ is our required graph of groups.

Given $P \in \mathcal{T}$, we write $\Gamma(P)$ for the setwise stabiliser of P in Γ . The main content of Propositions 7.1 and 7.2 therefore amounts to the following:

Lemma 7.3. *If $P \in \mathcal{T}_\infty$, then $\Gamma(P)$ admits a properly discontinuous action on the hyperbolic plane such that each adjacent edge group fixes a point and is either finite or parabolic.*

Now, the edges of Σ adjacent to P are in bijective correspondence with the components of $\Sigma \setminus \{P\}$, and hence with the branches of the pretree, \mathcal{T} , based at P . Let $S(P) = P \sqcup \mathcal{B}(P)$. We saw in Section 5 that $S(P)$ admits a canonical cyclic order. This order is preserved by $\Gamma(P)$. Let $\Gamma_0(P)$ be the subgroup of $\Gamma(P)$ (of index at most 2) which preserves either of the orientations on $S(P)$.

Suppose $\lambda \in \Pi_P$ (i.e. $\lambda \in \Pi$ and $\lambda \subseteq P$). Now λ has precisely two sides (i.e. equivalence classes under the relation defined by $x \sim y$ if not $x - \lambda - y$). By hypothesis, there is some infinite order element $g \in \Gamma(\lambda)$. By replacing g by g^2 , if necessary, we can suppose that $g \in \Gamma_0(P) \cap \Gamma_0(\lambda)$. In particular, g preserves the two sides of λ in P .

Lemma 7.4. *$\langle g \rangle$ has finite index in $\Gamma(\lambda)$. Moreover, if $\mu \in \Pi_P$ is any pair crossing λ , then $g\mu : \mu$, and $g^n\mu$ and $g^{-n}\mu$ converge, respectively to the two points of λ in the cyclic order on $S(P)$.*

Writing $\lambda = \{x, y\}$, we interpret the last statement to mean that, for any $z, w \in S(P)$ on opposite sides of λ , then for all sufficiently large n , we have $g^n\mu$ and $g^{-n}\mu$ on the same side of $\{z, w\}$ as x and y respectively (possibly interchanging the roles of x and y , or equivalently, replacing g by g^{-1}). From the definition of the cyclic order on $S(P)$, it is sufficient to prove this for $z, w \in P$.

Proof. If $g\mu \times \mu$, then it follows inductively that $g^n\mu \times \mu$ for all $n \geq 0$. Since $g^n\mu \times \lambda$ for all n , we contradict discreteness. Thus, $g\mu : \mu$, and without loss of generality, we can suppose that $\mu - g\mu - y$. It follows that $g^m\mu - g^n\mu - y$ whenever $m < n$.

Now suppose $z, w \in \Omega \setminus \{y\}$ with $z - \lambda - w$. If $\{z, w\} - g^n\mu - y$ does not hold, then either $\mu - g^n\mu - \{z, y\}$ or $\mu - g^n\mu - \{w, y\}$ must hold. By discreteness, this can

happen for only finitely many n . Thus, we have $\{z, w\} - g^n\mu - y$ for all sufficiently large n . Similarly, if $z, w \neq x$, then $\{z, w\} - g^{-n}\mu - x$ for all sufficiently large n .

It remains to show that $\langle g \rangle$ has finite index in $\Gamma(\lambda)$. Since $\Gamma(\mu) \cap \Gamma(\lambda)$ is finite, it suffices to show that there are finitely many $\Gamma(\lambda)$ -orbits of μ modulo the action of $\langle g \rangle$. To this end, suppose that $\nu = \{a, b\}$ is a $\Gamma(\lambda)$ -image of μ . Up to the action of $\langle g \rangle$, we can suppose that a lies between μ and $g\mu$, i.e. $x - \mu - a$ and $a - g\mu - y$. Now, $\nu \times \lambda$ and one of $\nu \times \mu$, $\nu \times g\mu$ or $\mu - \nu - g\mu$ must hold. There are thus only finitely many possibilities for ν as required. \square

Since $P = \bigcup \Pi_P$, this shows, in particular, that no point of P can be adjacent to any element of $S(P)$ in the cyclic order. Moreover, from the construction of $S(P)$, any pair of distinct elements of $\mathcal{B}(P) \subseteq S(P)$ are separated by a pair of points of P (namely an element of Π_P). Thus, now two points of $\mathcal{B}(P)$ are adjacent. Since $S(P) = P \sqcup \mathcal{B}(P)$, it follows that the cyclic order on $S(P)$ is dense. Since it is countable, it can be canonically completed to a circle, S^1 .

Now, Γ acts by homeomorphism on S^1 , and $\Pi_P/\Gamma(P)$ is finite. We see that the hypotheses of Theorem 3.3 are satisfied. This action is thus homeomorphically conjugate to the action induced on $S^1 \equiv \partial\mathbf{H}^2$ by a properly discontinuous isometric action of $\Gamma(P)$ on the hyperbolic plane, \mathbf{H}^2 . Using Lemma 6.4, it is easy to see that if $\lambda \in \Pi_P$, then $\Gamma(\lambda)$ acts as a loxodromic group on S^1 with limit set $\lambda \subseteq S^1$.

Now if $H \leq \Pi(P)$ is an incident edge group to P , then H is the (setwise) stabiliser of a branch of Σ , and hence an element of $\mathcal{B}(P)$, which we can identify with a point $x \in S^1$. Now, $H = \Gamma(x)$ must be finite, parabolic or loxodromic. To complete the proof of Lemma 7.3, we need to rule out the last possibility. In other words, we want to show:

Lemma 7.5. *If $x \in \mathcal{B}(P) \subseteq S^1$ and $y \in S^1 \setminus \{x\}$, then $\Gamma(x) \cap \Gamma(y)$ is finite.*

Proof. Suppose, to the contrary, that $\Gamma(x) \cap \Gamma(y)$ is infinite. Then it must contain an infinite order element, say g . Since Π_P is cross-connected, there must be some $\lambda \in \Pi_P$ which separates x from y in S^1 . From Lemma 7.4, we see that $\Gamma(\lambda) \cap \Gamma(x)$ finite. Since $g \in \Gamma(x)$, it follows that there are infinitely many distinct $\langle g \rangle$ -images of λ . By discreteness, these cannot all cross each other. Thus, there is some $n \in \mathbf{Z}$ with $g^n\lambda : \lambda$. Since λ and $g^n\lambda$ both separate x from y , without loss of generality (replacing g by g^{-1} if necessary) we can assume that $g^n\lambda$ separates λ from x . It follows that $g^{mn}\lambda$ separates λ from x for all $m \geq 1$. Now, if μ is any element of Π lying in the branch of \mathcal{T} corresponding to $x \in S^1$, then it follows that $\lambda - g^{mn}\lambda - \mu$ for all $m \geq 1$, contrary to the discreteness of Π . \square

Propositions 7.1 and 7.2 now follow from the constructions described earlier.

EXAMPLE

Suppose that M is a metrisable continuum with no weak cut point. Suppose that Γ acts on M as a convergence group. Suppose that $G \leq \Gamma$ is a maximal loxodromic subgroup with limit sets ΛG . Thus, ΛG is a pair of points of M . The Γ -orbit of ΛG is necessarily discrete as a subset of the space of pairs of M . Thus, by Proposition 6.2, the Γ -orbit has the structure of a discrete cactus. If we suppose, in addition, that $M \setminus \Lambda G$ is disconnected, then this cactus is proper. We see that all the conditions (1)–(4) are satisfied, and so Proposition 7.1 applies to give us a splitting of Γ . Note that we could equally well start with any finite collection of maximal loxodromic subgroups whose limit sets individually separate M .

Suppose now that Γ is a one-ended hyperbolic group. Then Γ acts as a convergence group on $\partial\Gamma$. Moreover, Γ is a metrisable Peano continuum with no global cut point. Any two-ended subgroup of Γ acts as a loxodromic group on $\partial\Gamma$. Moreover its limit set disconnects $\partial\Gamma$ if and only if it is (virtually) codimension one. We therefore get the annulus theorem for hyperbolic groups (see [ScS1], [Bo1]), namely that the existence of a codimension-one two-ended subgroup implies the existence of a non-trivial splitting over a two-ended subgroup. Moreover, from the results of [BeF] there is a bound on the complexity of such splittings. From this is not hard to see that there is a canonical maximal such splitting so we effectively recover the JSJ splitting as described in [Bo1] (though one needs to work harder to see that every local cut point plays a role in this splitting).

8. STRIPS

In this section, we explain how cacti arise from collections of two-ended subsets of a one-ended metric space. We shall apply this to groups in subsequent sections.

Let (X, ρ) be a proper (complete locally compact) one-ended path-metric space. It will be technically convenient to assume that X is a locally finite graph, and that ρ is the combinatorial metric on X (which assigns every edge a unit length). We shall also tacitly assume that all subsets we work with are subgraphs or complements thereof. We shall speak loosely of the “ends” of a subset, $A \subseteq X$, by which we mean, more precisely, the ends of its closure, \bar{A} in X . Given $A \subseteq X$, we shall write $A^C = X \setminus A$. We shall write $\mathcal{C}(A)$ for the set of connected components of A^C . Given $r \geq 0$, we write $N(A, r) = \{y \in X \mid \rho(y, A) \leq r\}$.

Definition. A *strip* in X consists of a pair, (A, I) , of two-ended subsets $A \subseteq I \subseteq X$, with A lying in the interior of I and such that for all $P \in \mathcal{C}(A)$, P is one-ended and $I \cap P$ is two-ended.

We say that A is *proper* if $|\mathcal{C}(A)| \geq 2$.

In the cases of interest to us, $\mathcal{C}(A)$ will be finite, and it will be convenient to take this as an additional assumption in this section, though it is not strictly necessary.

We shall usually denote a strip simply as A , and assume that it has associated with it an appropriate subset $I = I_A$. We shall refer to an element of $\mathcal{C}(A)$ as a *side* of A . If P is a side of A , we shall refer to the set $P \cap I_A$ as the *edge* of P or of A . We shall refer to the ends of a strip as *half-strips* and to the ends of an edge as *half-edges*. We write $\Omega(A)$ for the (unordered) pair of half-strips associated to A .

Definition. Suppose A and B are strips. We say that A and B are *separated* if $I_A \cap I_B$ is compact.

Suppose that \mathcal{A} is a set of pairwise separated strips in X . Let $\Omega(\mathcal{A}) = \bigcup_{A \in \mathcal{A}} \Omega(A)$ be the set of half-strips. Thus, $\Pi = \Pi(\mathcal{A}) = \{\Omega(A) \mid A \in \mathcal{A}\}$ is a partition of \mathcal{A} into disjoint pairs. Suppose that $\lambda = \Omega(A) \in \Pi$ and $x, y \in \Omega(A)$. We write $x - \lambda - y$ to mean that the half-strips x and y lie in different sides of A .

Lemma 8.1. *With this structure, $\Omega(\mathcal{A})$ is a cactus.*

Proof. We can assume that \mathcal{A} is finite. Given any $A \in \mathcal{A}$, let K_A be a compact set that separates the ends of I_A . Let $K \subseteq X$ be a compact set which contains K_A for all $A \in \mathcal{A}$ and $I_A \cap I_B$ for all distinct pairs $A, B \in \mathcal{A}$. We construct a graph $\Upsilon = \Upsilon_K(\mathcal{A})$ as follows.

Let $V_0(\Upsilon) = \Omega(\mathcal{A})$ and $V_1(\Upsilon) = \mathcal{C}(K \cup \bigcup \mathcal{A})$ and let $E(\Upsilon)$ be the set of half-edges of elements of \mathcal{A} . We let Υ be the bipartite graph with vertex set $V(\Upsilon) = V_0(\Upsilon) \sqcup V_1(\Upsilon)$ and edge set $E(\Upsilon)$, where an element of $E(\Upsilon)$ is deemed to be incident on the corresponding half-strip of $V_0(\Upsilon)$ as well as to the unique element of $V_1(\Upsilon)$ in which it is contained.

We need to verify that Υ is connected, and that no element of $V_0(\Upsilon) = \bigcup \Pi(\mathcal{A})$ is a cut vertex.

Suppose that $\lambda = \Omega(A) \in \Pi(\mathcal{A})$ and that $P \in \mathcal{C}(A)$. Let $\Upsilon^P \subseteq \Omega$ be the full subgraph whose vertex set consists of λ together with all those half-strips of \mathcal{A} lying in P , and all elements of $\mathcal{C}(K \cup \bigcup \mathcal{A})$ lying in P (so that the edge set consists of all half-edges lying in P). Since P is by hypothesis one-ended, it follows easily that Υ^P is connected. Since $\lambda \subseteq \Upsilon^P$ for all $P \in \mathcal{C}(A)$ and $\Upsilon = \bigcup_{P \in \mathcal{C}(A)} \Upsilon^P$, it follows that Υ is connected, and that neither element of λ can be a global cut point, as required. \square

Note that, in the above argument, if $|\mathcal{C}(A)| \geq 2$, then $\Upsilon \setminus \lambda$ is disconnected. We therefore deduce:

Lemma 8.2. *If $|\mathcal{C}(A)| \geq 2$ for all $A \in \mathcal{A}$, then $\Omega(\mathcal{A})$ is a proper cactus.*

In such a case, we say that \mathcal{A} is *proper*. In all cases, we see that if $\lambda = \{x, y\} \in \Pi$, then $\deg x = \deg y$ and this is equal to the number of components of $\Upsilon \setminus \lambda$.

In Section 4, we defined an equivalence relation, \approx , on $V_1(\Upsilon_K(\mathcal{A}))$ by deeming two vertices to be equivalent if they are not separated by any element of Π . If $L \supseteq K$ is another compact set, then there is an obvious canonical map from $\Upsilon_K(\mathcal{A})$ to $\Upsilon_L(\mathcal{A})$ which is the identity on $V_0(\Upsilon_K) = V_0(\Upsilon_L) \equiv \Omega(\mathcal{A})$, and which may identify certain vertices of $V_1(\Upsilon_K)$ lying in the same \approx -class. If Υ_K is efficient, then this must be the identity. Note that $|V_0(\Upsilon_K)|$ is bounded by $2 \sum_{A \in \mathcal{A}} |\mathcal{C}(A)|$. Thus, by taking larger and larger compact sets K , we see that $\Upsilon_K(\mathcal{A})$ must stabilise on some graph which we denote by $\Upsilon(\mathcal{A})$. If $\Upsilon(\mathcal{A})$ is efficient, then $\Upsilon_K(\mathcal{A}) = \Upsilon(\mathcal{A})$ for all compact sets K .

If $A, B \in \mathcal{A}$, we say that B *crosses* A if the corresponding elements of $\Omega(\mathcal{A})$ cross. In other words, the ends of B lie on different sides of A .

Suppose that \mathcal{A} is proper. Then we can speak of “cross-connected” components of \mathcal{A} .

Lemma 8.3. *If \mathcal{A} is finite and cross-connected, then $\Upsilon(\mathcal{A})$ is efficient (hence a circle if \mathcal{A} has more than one element).*

Proof. Suppose that $|\mathcal{A}| \geq 2$. We show directly that $\Upsilon = \Upsilon(\mathcal{A})$ is a circle. If $\lambda = \{x, y\} \in \Pi$ then $\Upsilon \setminus \lambda$ must have precisely two components (otherwise no element of Π could separate the two points of λ). Thus $\deg x = \deg y = 2$. Since $V_0 = \bigcup \Pi$, we see that each element of V_0 has degree 2. Moreover, no element of V_1 can be a cut-vertex (otherwise it would cut Π into at least two cross-connected components). Let Φ be the graph obtained by deleting each vertex in V_0 and amalgamating the two incident edges. The result follows from the following characterisation of the circle. \square

Lemma 8.4. *Let Φ be a connected graph (possibly with multiple edges) with no cut-vertex. Suppose that for every pair of edges $e, f \in E(\Phi)$ there is a chain of pairs of distinct edges, $\{e_i, e'_i\}_{i=1}^n$ with $e_0 = e$ and $e_n = f$, such that $\Phi \setminus \{e_0, e'_0\}$ is*

disconnected and such that for all $i < n$, e_i and e'_i lie in different components of $\Phi \setminus \{e_{i+1}, e'_{i+1}\}$. Then, Φ is a circle.

Proof. If not, then we can find three distinct arcs α , β and γ in Φ which intersect pairwise precisely in their endpoints. Suppose that e_0, e'_0 are distinct edges of Φ with e_0 contained in α . Since $\Phi \setminus \{e_0, e'_0\}$ is disconnected, we see that e'_0 also lies in α . Moreover, if e_0 and e'_0 lie in different components of $\Phi \setminus \{e_1, e'_1\}$, then at least one of e_1 or e'_1 lies in α . It now follows inductively that, if the sequence $\{e_i, e'_i\}$ is as in the hypotheses, the e_i must lie in α for all i , giving a contradiction. \square

It now follows that if \mathcal{A} is finite proper and cross-connected, then $\Upsilon_K(\mathcal{A}) = \Upsilon(\mathcal{A})$ for all compact K for which Υ_K is defined. In particular, we can take $K = \bigcup\{(A \cap I_B \mid A, B \in \mathcal{A}, A \neq B)\}$, so that $K \cup \bigcup \mathcal{A} = \bigcup \mathcal{A}$. In particular, each element of $\mathcal{C}_\infty(\bigcup \mathcal{A})$ contains precisely two half-edges.

Given a pairwise separated collection of strips, \mathcal{A} , we shall use the notations $A \times B$, $A : B$, $A - B - C$ etc. to mean $\Omega(A) \times \Omega(B)$ etc. in the cactaceous structure on $\Omega(\mathcal{A})$.

There is a variation on the construction we have described in this section which we remark upon for future reference. Suppose that \mathcal{A} is a pairwise separated collection of strips. Suppose that for each $A \in \mathcal{A}$, we have a partition, $\mathcal{R}(A)$, of $\mathcal{C}(A)$. Let \mathcal{A} be the set of half-strips as before. This time, we define the relation $x - \lambda - y$ for $x, y \in \Omega(\mathcal{A})$ and $\lambda = \Omega(A) = \Pi(\mathcal{A})$ to mean that the half-strips x, y lie in different elements of the partition, $\mathcal{R}(A)$. We see, by essentially the same construction, that the resulting structure on $\Omega(\mathcal{A})$ is again a cactus. We shall denote this cactus by $\Omega_{\mathcal{R}}(\mathcal{A})$.

9. AXES

Suppose that G is a two-ended group acting properly discontinuously on X .

Definition. An *axis*, A , of G is a G -invariant strip.

We interpret this to mean that both A and I_A are G -invariant.

Lemma 9.1. *Any two-ended group, G , has an axis.*

Proof. We start with any connected G -invariant set, A_0 with A_0/G compact. Now $\mathcal{C}(A_0)/G$ is finite. Let A be the union of A_0 and all those elements of $\mathcal{C}(A_0)$ which lie in some uniform neighbourhood of A_0 . We see that A/G is compact, and $\mathcal{C}(A)$ is finite. We first claim that every element of $P \in \mathcal{C}(A)$ is preserved by some infinite order element.

To see this, let $g \in G$ be any infinite order element. Since X is one-ended, $\mathcal{B}(P) \cap A$ is unbounded. Thus, infinitely many $\langle g \rangle$ -translates of A must meet some given compact set. Since $\mathcal{C}(A)$ is locally finite, the result follows.

Since $\mathcal{C}(A)/\Gamma$ is finite, it follows that there is a fixed infinite order element, $h \in G$, such that $hP = P$ for all $P \in \mathcal{C}(A)$. Again, by local finiteness, it follows that $\mathcal{C}(A)$ is finite.

Now suppose that $P \in \mathcal{C}(A)$ and that $H \leq G$ is the setwise stabiliser of P (which has finite index in G). The next step is to construct a closed H -invariant subset J_P with $A \subseteq J_P \subseteq A \cup P$, such that J_P/H is compact, $J_P \setminus A$ is connected, and such that the closure of $P \setminus J_P$ in X does not meet A . If we drop the condition that $J_P \setminus A$ be connected, this is trivial. To ensure that it is connected, let $h_0 \in H$ be any infinite order element, and let h_1, \dots, h_p be a coset transversal of $\langle h \rangle$ in H .

Let x be any point of $J_P \setminus A \subseteq P$, and let α_i be a path in P connecting x to $g_i x$. We now add to J_P all H -images of all paths $\alpha_0, \dots, \alpha_p$.

We can clearly carry out this construction G -equivariantly for all $P \in \mathcal{C}(A)$. Let $I_A = \bigcup_{P \in \mathcal{C}(A)} J_P$. We claim that (A, I_A) is an axis of G . It remains to verify that if $P \in \mathcal{C}(A)$, then \bar{P} is one-ended.

We first claim that \bar{P} has at most two ends. For suppose $K \subseteq \bar{P}$ is compact and is such that $\bar{P} \cap K^C$ has at least three unbounded components. Now $J_P \cap \bar{P}$ is H -invariant with compact quotient. Thus, $J_P \cap \bar{P}$ is two-ended. Each end of $J_P \cap \bar{P}$ must be contained in an unbounded component of $\bar{P} \cap K^C$. There is therefore some unbounded component, say C , which meets $J_P \cap \bar{P}$ in a compact set, say L . Now the boundary of P in X lies in $A \subseteq J_P$ and so the boundary of C in X lies in $K \cup L$. Since $K \cup L$ is compact and C is unbounded, we contradict the fact that X is one-ended. This proves the claim.

Now, if \bar{P} were two-ended, it would follow that \bar{P}/H must be compact. Thus \bar{P} would lie in a bounded neighbourhood of A , contrary to our construction. We deduce that \bar{P} is one-ended as required. \square

Note that our argument also shows that any axis, A , of G satisfies $|\mathcal{C}(A)| < \infty$. Moreover, our construction gives the following:

Lemma 9.2. *If K is compact, and \mathcal{E} is a finite set of coends of H , then we can choose the axis A to contain K , and to separate any pair of coends of \mathcal{E} .*

Proof. Choose $A_0 \supseteq K$ to separate all pairs of coends of G . \square

By Lemma 9.2, we see that any codim-one two-ended group acting properly discontinuously on X has a proper axis, A (i.e. $|\mathcal{C}(A)| \geq 2$).

If A is any axis of G , and $P \in \mathcal{C}(A)$, then the set of coends of G lying in P is a clopen subset of $\mathcal{E}(G)$. Thus, A determines a G -invariant clopen partition of $\mathcal{E}(G)$, i.e. a partition of $\mathcal{E}(G)$ into clopen subsets. (Such a partition is necessarily finite.) As a result, we conclude:

Lemma 9.3. *If $x, y \in \mathcal{E}(G)$ are distinct, then there is a G -invariant clopen partition of $\mathcal{E}(G)$ with x and y contained in different elements of the partition.*

10. GROUP ACTIONS ON AXES

In this section, we prove a lemma concerning the geometry of axes of a group acting on X . First, we give some definitions:

Definition. We say that two strips, A and B , are *parallel* if the hausdorff distance between them is finite (i.e. each lies in a uniform neighbourhood of the other).

We say that A and B are *divergent* if for all $r \geq 0$, there is a compact subset $K \subseteq X$, such that $\rho(A \cap K^C, B \cap K^C) \geq r$.

We say that A and B are *strictly separated* if $I_A \cap I_B = \emptyset$.

Clearly, if A and B are either divergent or strictly separated, then they are separated ($I_A \cap I_B$ is compact).

Suppose that a group, Γ , acts properly discontinuously isometrically on X . Suppose that $G, H \leq \Gamma$ are two-ended subgroups, and that A and B are axes of G and H respectively. Then, either G and H are commensurate, and A and B are parallel, or else G and H are non-commensurate, and A and B are divergent. In the latter case, we say that H *crosses* G if we can choose A and B so that B crosses A . This

is clearly independent of the choice of B . Moreover, provided we choose A to be a proper axis, it is also independent of the choice of A (see below). This notion of crossing agrees with that defined in Section 1.

Suppose that $G, H \leq \Gamma$ are non-commensurate two-ended subgroups. Given any axes, A, B of G, H respectively, we can construct the cactus, $\Omega = \Omega(\{A, B\})$. If B crosses A , then A must be proper, and $|\mathcal{C}(A)| = 2$. Using Lemma 9.2, we see that G has precisely two coends.

Suppose that H is also codim-one. Let B be any proper axis of H . If B crosses A , then A crosses B (see Lemma 3.1), and so G crosses H . In other words, for codim-one two-ended subgroups, the notion of crossing is symmetric.

If A is an axis of G , we say that $g \in G$ is *side-preserving* if $gP = P$ for all $P \in \mathcal{C}(A)$. If G has precisely two coends, then this is independent of any choice of proper axis, and the side-preserving subgroup of G has index at most 2.

The following is a slight variation on Lemma 4.5 of [DSw].

Lemma 10.1. *Suppose that $g \in \Gamma$ has infinite order and that A is an axis of g . Suppose that g is side-preserving and that B is a strip which diverges from A . Then it cannot happen that gB is separated from B and crosses B .*

Proof. Suppose, to the contrary, that gB and B are separated and that gB crosses B . By Lemma 9.2, we can assume that all the bounded components of $I_B \setminus (B \cup gB)$ lie in the interior of A . Let A^-, A^+ and B^-, B^+ be the ends of A and B respectively. We can assume that g translates A in the direction of A^+ . The set of ends, $\{A^\pm, B^\pm, gB^\pm\}$ has the structure of a cross-connected cactus, $\Omega(\{A, B, gB\})$. They thus carry a cyclic order, which without loss of generality we can take to be $A^+, B^+, gB^+, A^-, B^-, gB^-, A^+$.

By Lemma 8.3, the graphs $\Upsilon(\{A, B\})$ and $\Upsilon(\{A, B, gB\})$ are both circles. Let $P \in \mathcal{C}(A \cup B)$ be the unbounded component of $(A \cup B)^C$ lying between A^- and B^+ . We claim that $gP \subseteq P$. To see this, let $Q \in \mathcal{C}(A \cup B \cup gB)$ be the unbounded component of $(A \cup B \cup gB)^C$ lying between A^- and gB^+ . Since $P \cap Q \neq \emptyset$, it is clear that $Q \subseteq P$. But now $Q \cap B = \emptyset$ (otherwise Q would meet an unbounded component of $I_B \setminus (B \cup gB)$, and hence contain a half-edge of B). It follows that $Q \in \mathcal{C}(A \cup gB)$. Thus, $Q = gP$, and so $gP \subseteq P$ as claimed.

Now let α be any path connecting a point $x \in P$ to a point $y \in A$. Note that $g^n x \in P$ for all n . Also, $g^n y$ tends to the end A^+ , which lies on the opposite side of B from the component P . It follows that $g^n \alpha \cap B \neq \emptyset$ for all sufficiently large n . But now this contradicts the assumption that B and A are divergent. \square

11. VERIFICATION OF DISCRETENESS CONDITIONS

Suppose that Γ acts properly discontinuously isometrically on X . Let Γ act properly discontinuously on X . Let \mathcal{A} be a pairwise divergent Γ -invariant set of axes in X with \mathcal{A}/Γ finite. In this section, we aim to show:

Proposition 11.1. *The cactus $\Omega(\mathcal{A})$ is discrete.*

We shall prove this by verifying property (F1) and then deducing (F2). We begin with some preliminary observations. Firstly note that \mathcal{A} is locally finite (as a collection of subsets of X). Moreover, modulo Γ , there are only finitely many pairs $A, B \in \mathcal{A}$ such that $I_A \cap I_B \neq \emptyset$. In particular, it follows that if A, B do not cross, then B can only penetrate a bounded distance on the “wrong side” of A . More precisely:

Lemma 11.2. *There is some constant $r \geq 0$ such that if $A, B \in \mathcal{A}$ do not cross and both ends of B lie in $P \in \mathcal{C}(A)$, then $B \subseteq P \cup N(A, r)$.*

We also note that if $A \in \mathcal{A}$, then, modulo its stabiliser, $\Gamma(A)$, only finitely many elements of \mathcal{A} can cross A .

Lemma 11.3. $\Omega(\mathcal{A})$ satisfies (F1).

Proof. Suppose $A, B \in \mathcal{A}$ with $A : B$. We want to verify that $\{C \in \mathcal{A} \mid A - C - B\}$ is finite.

To see this, let α be any path in X connecting some point $x \in A$ to some point $y \in B$. Suppose $C \in \mathcal{A}$ satisfies $A - C - B$. We claim that $\rho(\alpha, C) \leq r$. For suppose $\rho(\alpha, C) > r$. It follows in particular that $\alpha \cap C = \emptyset$ so there is some $P \in \mathcal{C}(C)$ with $\alpha \subseteq P$, and $\rho(x, C) > r$. It follows that $\rho(x, C \cup Q) > r$ for all $Q \in \mathcal{C}(C) \setminus \{P\}$. By Lemma 11.2 it therefore follows that both ends of B lie in P . Similarly both ends of A lie in P . But this contradicts the assumption that $A - C - B$. We conclude that $\rho(\alpha, C) \leq r$. By local finiteness of \mathcal{A} , there are only finitely many such C . \square

Lemma 11.4. $\Omega(\mathcal{A})$ satisfies (F2).

Proof. Suppose that $A, B \in \mathcal{A}$ are distinct, hence divergent. Suppose, for contradiction that infinitely many elements of \mathcal{A} cross both A and B . In particular, $|\mathcal{C}(A)| = |\mathcal{C}(B)| = 2$. Let g, h be infinite-order side-preserving elements of $\Gamma(A)$ and $\Gamma(B)$ respectively. Modulo $\Gamma(A)$, and hence modulo $\langle g \rangle$, only finitely many elements of \mathcal{A} can cross A . Thus, there is some $C \in \mathcal{A}$ and an infinite sequence, n_i , of natural numbers such that $g^{n_i}C$ crosses both A and B for all i . We can assume that $n_0 = 0$.

Now, by Lemma 10.1, $g^m C : g^n C$ for all $m, n \in \mathbf{N}$. In particular, it follows that if $n_i \leq n \leq n_{i+1}$, then $g^{n_i}C - g^n C - g^{n_{i+1}}C$, and so $g^n C$ crosses both A and B . (This observation is easily made using the fact that $\{A, B\} \cup \{g^n C \mid n \in \mathbf{N}\}$ is cross-connected, so that its set of ends admits a cyclic order.) Thus, by interpolation, we see, in fact, that $g^n C \times A$ and $g^n C \times B$ for all $n \in \mathbf{N}$.

Now, modulo the action of $\langle h \rangle$, only finitely many elements of \mathcal{A} can cross B . In particular, for each $p \in \mathbf{N}$, there is some $q \in \mathbf{N}$ such that $h^{-q}g^p C$ lies in some fixed finite subset of \mathcal{A} . Thus there is some fixed $D \in \mathcal{A}$ and sequences p_i, q_i with $g^{p_i}C = h^{q_i}D$ for all i . We can assume that $p_i \rightarrow \infty$ but that $p_{i+1} - p_i \not\rightarrow \infty$. Thus, $p_{i+1} - p_i$ equals some fixed constant, say m , for infinitely many i . For any such i , we have $h^{q_i}D \times A$ and $h^{q_{i+1}}D \times A$, and so using Lemma 10.1 again, we see that $h^{q_i}D - h^{q_i}D - h^{q_{i+1}}D$ for all q lying strictly between q_i and q_{i+1} . Thus, $C - g^{-p_i}h^q D - g^m C$ (since $g^m C = g^{-p_i}h^{q_{i+1}}D$). But now, since C and $g^m C$ diverge, and A and B diverge, the intersections of $h^{q_i}D = g^{p_i}C$ and $h^{q_{i+1}}D = g^{p_i+m}C$ with B are arbitrarily far apart as $i \rightarrow \infty$. Thus $|q_{i+1} - q_i| \rightarrow \infty$. (This is the only geometric input into the argument.) Thus, by choosing i sufficiently large, we can find arbitrarily many axes of the form $g^{-p_i}h^q D$ between C and $g^m C$. Thus, $\{E \in \mathcal{A} \mid C : E : g^m C\}$ is infinite, contrary to Lemma 11.3. \square

By Lemma 6.1, this proves Proposition 11.1.

12. VIRTUALLY ABELIAN GROUPS

In this section, we give a characterisation of groups which are virtually $\mathbf{Z} \oplus \mathbf{Z}$ (Proposition 12.2). We first need to define some general notions.

Suppose that Γ acts properly discontinuously on a locally finite one-ended graph, X . Given $g \in \Gamma$, let $D(g) = \min\{\rho(x, gx) \mid x \in X\}$. Clearly $D : \Gamma \rightarrow [0, \infty)$ is conjugacy invariant, and $D(g^n) \leq nD(g)$ for $n \geq 0$. We say that g is *quasigeodesic* if $D(g^n)$ is bounded below by an increasing linear function of n . (If X/Γ is compact, this is the same as saying that $\langle g \rangle$ is quasigeodesically embedded in Γ .)

Suppose that g is quasigeodesic and $h \in \text{Comm}(g)$. Then $hg^mh^{-1} = g^n$ for some $m, n \in \mathbf{Z} \setminus \{0\}$. Thus, for all $p \in \mathbf{N}$, g^{m^p} is conjugate to g^{n^p} , and so $D(g^{m^p}) = D(g^{n^p})$. It follows that $|m|^p$ is bounded above and below by linear functions of $|n|^p$, and so $m = \pm n$. We see that h^2 commutes with g^n .

Given a two-sided strip, A , we write $\mathcal{C}(A) = \{P_A, Q_A\}$. (There is no reason to prefer P_A or Q_A .) If A, B are two-sided and strictly separated, then they are “nested” in the sense that precisely one of the statements $I_A \cup P_A \subseteq P_B$, $I_A \cup Q_A \subseteq P_B$, $I_A \cup P_A \subseteq Q_A$, $I_A \cup Q_A \subseteq Q_B$ must hold. Moreover, $I_A \cup P_A \subseteq P_B$ is equivalent to $I_B \cup Q_B \subseteq Q_A$ etc. If we have three parallel strictly separated two-sided strips, A, B, C , then they must also be “nested”, i.e. up to permuting A, B, C and choosing P_A, P_B, P_C appropriately we can assume that $I_A \cup P_A \subseteq P_B$ and $I_B \cup P_B \subseteq P_C$.

Suppose that A is a two-sided strip and $h \in \Gamma$, with hA strictly separated from A . We can assume that either $P_A \cap hP_A = \emptyset$ or that $h(I_A \cup P_A) \subseteq P_A$. In the latter case, h has infinite order and any axis of h must cross A (and hA). Moreover, we have $\bigcap_{n=0}^{\infty} h^n P_A = \emptyset$ (since the sets $h^n P_A$ for $n \in \mathbf{N}$ are nested and their boundaries form a locally finite family). Similarly $\bigcap_{n=0}^{\infty} h^{-n} Q_A = \emptyset$. It follows easily that $\bigcup_{n=0}^{\infty} Q_A = \bigcup_{n=0}^{\infty} P_A = X$. We also see that h is quasigeodesic (since any path connecting x to $h^n x$ must cross n copies of I_A).

Conversely, if A is two-ended and h is an infinite-order element whose axis crosses A , then the first case (namely $P_A \cap hP_A = \emptyset$) cannot arise (since one end of any axis of h must lie in P_A).

Lemma 12.1. *Suppose $g, h \in \Gamma$ are codim-one elements which cross. Then g and h are quasigeodesic.*

Proof. Replacing h by h^2 if necessary, we can assume that h is side-preserving. Let A be any proper axis of g . We claim that there is some $n > 0$ such A and $h^n A$ are strictly separated. (For if not, it would follow by local finiteness that the collection $\{h^n A \mid n \geq 0\}$ lies in finitely many $\langle g \rangle$ -orbits of A . Since $A/\langle g \rangle$ is compact and the action of Γ is properly discontinuous, we see that for each n , there is some m such that $g^m h^n$ lies in some fixed finite subset of Γ . It in turn follows that g and h are commensurable. This contradicts the assumption that they cross, and proves the claim.) From the earlier discussion we conclude that h is quasigeodesic. The same argument applies to g . \square

Proposition 12.2. *Suppose that Γ acts properly discontinuously on X . Suppose that there exist codim-one elements $g, h \in \Gamma$ such that h crosses g , and such that $[\text{Comm}(g) : \langle g \rangle] = \infty$. Then Γ is virtually $\mathbf{Z} \oplus \mathbf{Z}$, and X/Γ is compact.*

Proof. Let A be a proper axis of g (so that A is two-sided). Since $[\text{Comm}(g) : \langle g \rangle]$ is infinite, we can find $k_1, k_2 \in \text{Comm}(g)$ so that $A, k_1 A$ and $k_2 A$ are pairwise strictly separated. Since they are parallel they are nested (in the sense described before Lemma 12.1). It therefore follows that there is some $k \in \text{Comm}(g)$ such that $k(I_A \cup P_A) \subseteq P_A$ where $\mathcal{C}(A) = \{P_A, Q_A\}$ (here, $k = k_1, k_2$ or $k_1^{-1} k_2$). From the earlier discussion, it follows that k has infinite order and crosses g (though we don't a priori know that k is codim-one). Given the existence of h , it follows by Lemma

12.1 that g is quasigeodesic. Thus (after replacing g and k by suitable powers) we can suppose that g and k commute. Since k crosses g , it is not commensurable with g and so $\langle g, k \rangle \cong \mathbf{Z} \oplus \mathbf{Z}$. It therefore remains to show that $X/\langle g, k \rangle$ is compact.

By Lemma 8.2, there is a proper axis, B , of g such that $I_A \cup kI_A \subseteq B$. We claim that $X = \bigcup \langle k \rangle B$ from which the result follows.

Suppose, to the contrary, that C is a nonempty component of $X \setminus \bigcup \langle k \rangle B$. Let $\mathcal{C}(B) = \{P_B, Q_B\}$, with $P_B \subseteq P_A$ and $Q_B \subseteq Q_A$. Since $kA \subseteq B$, and since $Q_A \cap kP_A = \emptyset$, we also have $P_B \subseteq kP_A$ and $Q_B \subseteq kQ_A$. Thus, $X = B \cup Q_B \cup P_B = B \cup Q_A \cup kP_A$. For each n , $C \cap k^n A = \emptyset$ and so C lies in either $k^n P_A$ or $k^n Q_A$. Now, $k^n P_A$ and $k^{-n} Q_A$ are both increasing sequences whose union is the whole of X . It follows that there is some n with $C \subseteq k^n P_A \cup k^{n+1} Q_A$, so that $C \cap k^{n+1} P_A = C \cap k^n Q_A = \emptyset$. But $X = k^n (B \cup Q_A \cup kP_A) = k^n B \cup k^n Q_A \cup k^{n+1} P_A$ and $C \cap k^n B = \emptyset$, giving the contradiction that $C = \emptyset$. \square

13. PLANAR GROUPS

In this section, we give the following characterisation of a planar group (i.e. a group which contains a finite index subgroup which is the fundamental group of a closed surface other than the 2-sphere or projective plane).

Theorem 13.1. *Let Γ be a finitely generated one-ended group which contains an infinite order element and such that every infinite order element is codim-one. Then, Γ is planar.*

Proof. We let X be any Cayley graph of Γ . Suppose first that there is some infinite order $g \in \Gamma$ with $[\text{Comm}(g) : \langle g \rangle] = \infty$. Now the argument of Proposition 12.2 shows that there is an infinite order element $h \in \text{Comm}(g)$ which crosses g . By hypothesis, g and h are both codim-one, and so Proposition 12.2 applies showing that Γ is virtually $\mathbf{Z} \oplus \mathbf{Z}$.

We can therefore suppose that for any infinite order $g \in \Gamma$, $\text{Comm}(g)$ is two-ended and self-commensurating. We can thus find an axis, A , of $\text{Comm}(g)$. Since g is assumed to be codim-one, we can also take A to be proper. Note that any Γ -image of A is either equal to A or divergent from A .

Now, suppose \mathcal{A} is any Γ -invariant collection of pairwise divergent axes, with \mathcal{A}/Γ finite. (Such a collection exists by the discussion of the previous paragraph.) By Lemma 5.2 and Proposition 11.1, the set of cross-connected components of \mathcal{A} has the structure of a discrete pretree, and so can be canonically identified as a subset $V_0(\Sigma)$, of the vertex set, $V(\Sigma)$, of a simplicial tree, $\Sigma = \Sigma(\mathcal{A})$. Now, Γ acts simplicially without edge inversions and with finite quotient graph (see Section 7).

If Γ consists of a single vertex, then \mathcal{A} is cross-connected, and Γ is planar (see Proposition 7.2 and the subsequent discussion). We therefore assume that Σ is non-trivial.

We claim that Σ has no terminal vertices. To see this, note that by the construction of Σ , any terminal vertex would have to correspond to a cross-connected component, \mathcal{B} , of \mathcal{A} . There are two cases to consider. Firstly, suppose that $\mathcal{B} = \{A\}$ for some $A \in \mathcal{A}$. Now, all the other elements of \mathcal{A} lie on one side of A . By Lemma 11.2, there is a bound on how far any element of \mathcal{A} can penetrate into any other side of A . But this contradicts the fact that every point of X lies a bounded distance from some element of \mathcal{A} . The second case is where \mathcal{B} is infinite. We can similarly arrive at a contradiction. For example, we know that the setwise stabiliser of \mathcal{B} is conjugate to a fuchsian group with each incident edge stabiliser corresponding to

a maximal parabolic subgroup. There are clearly infinitely many such subgroups. This proves the claim.

It now follows that the action of Γ on Σ is minimal. Since the quotient is finite, a simple argument from Bass-Serre theory shows that any two vertices of Σ lie in the axis (i.e. minimal invariant subtree) of some hyperbolic element of Γ (i.e. an infinite order element of Γ that does not fix any vertex of Σ).

Now let \mathcal{B} be any cross-connected component of \mathcal{A} , and let g_1, \dots, g_n be a finite set of generators of Γ . For each $i = 1, \dots, n$, choose any hyperbolic element, h_i , such that the vertices corresponding to \mathcal{B} and $g_i\mathcal{B}$ are both contained in the axis of h_i in Σ . (If $g_i\mathcal{B} = \mathcal{B}$, then we can ignore this.) It is easily seen that h_i must cross some element of \mathcal{B} and some element of $g_i\mathcal{B}$. (Note that there is a branch of Σ based at \mathcal{B} containing all forward iterates $h_i^p(\mathcal{B})$, and another distinct branch containing all backward iterates $h_i^{-p}\mathcal{B}$. From the construction of Σ , there is some $B \in \mathcal{B}$ that separates $h_i^p\mathcal{B}$ from $h_i^{-p}\mathcal{B}$ in the cactaceous structure on the set of ends of \mathcal{A} . Thus, any axis of h must meet two distinct sides of B in non-compact sets, i.e. h crosses B . The same argument applies to $g_i\mathcal{B}$.) Now, let H_1, \dots, H_m be a collection of self-commensurating pairwise non-conjugate two-ended subgroups of Γ , such that each h_i is conjugate into some H_j . Let A_i be an axis of H_i , and let $\mathcal{A}' = \bigcup \Gamma\mathcal{B} \cup \Gamma A_1 \cup \dots \cup \Gamma A_m$.

Thus, \mathcal{A}' is a Γ -invariant set of pairwise divergent axes with \mathcal{A}'/Γ finite. This time, we have arranged that \mathcal{A}' is cross-connected and so Γ is planar as required. \square

In the almost finitely presented case (i.e. FP_2 over \mathbf{Z}_2), the above result can be obtained using the results and methods of [DSa] and [DSw]. Graham Niblo has also suggested ways in which these methods might be adapted to deal with the finitely generated case.

14. CANONICAL SPLITTINGS

Let Γ be a finitely generated one-ended group. In this section, we show how a finite collection of self-commensurating codim-one two-ended subgroups give rise to a splitting of Γ in a canonical way. If we drop the condition on self-commensurability, the situation becomes more complicated, as we explain in the next section.

Firstly, we show:

Proposition 14.1. *Suppose that Γ is a finitely generated one-ended group. Suppose that \mathcal{H} is a collection of self-commensurating codim-one two-ended subgroups of Γ , consisting of a finite union of Γ -conjugacy classes. Then, there is a representation of Γ as a finite bipartite graph of groups, \mathcal{G} , with vertex set $V(\mathcal{G}) = V_0(\mathcal{G}) \sqcup V_1(\mathcal{G})$. Every edge group is two-ended, and every element of \mathcal{H} is conjugate to a vertex group in $V_0(\mathcal{G})$. Each vertex group in $V_0(\mathcal{G})$ is either (the conjugacy class of) an element of \mathcal{H} , or else (conjugate to) a finitely generated fuchsian group, where the incident edge groups are precisely the maximal parabolic subgroups.*

Proof. The proof consists of piecing together our earlier constructions. First, choose a proper axis for each element of \mathcal{H} . We do this equivariantly, so as to give us a Γ -invariant collection, \mathcal{A} , of pairwise divergent axes. By Proposition 11.1, this gives us a discrete cactus $\Omega(\mathcal{A})$. Now, Γ acts on this cactus, and the properties (1)–(5) laid out in Section 7 are satisfied. Proposition 7.2 now gives us the required splitting of Γ . In this splitting, the subset $V_0(\mathcal{G})$ corresponds to the (Γ -orbits of) cross-connected components of \mathcal{A} . If $H \in \mathcal{H}$, then H preserves an element of \mathcal{A} ,

and hence a cross-connected component of \mathcal{A} . In other words, H is conjugate into a vertex stabiliser in $V_0(\mathcal{G})$ as required. \square

We have not explained to what extent this construction is canonical. A priori, it might depend on the choice of axes, \mathcal{A} . If each $H \in \mathcal{H}$ has only finitely many coends, then we can choose the corresponding axes to separate any pair of such coends. Given this, the cactus, $\Omega(\mathcal{A})$, is determined independently of the choice of axes. In fact, it's not too hard to see that it also does not depend on the choice of the space X , since the relations that determine the structure of the cactus are clearly quasi-isometry invariant.

As mentioned in the introduction, I don't know in general if it is possible for a self-commensurating two-ended subgroup to have infinitely many coends. If we want to allow for this possibility, the statement becomes a little more complicated:

Proposition 14.2. *Suppose that, in the set-up described in Proposition 14.1, we have associated to each $H \in \mathcal{H}$ a non-trivial H -invariant clopen partition, $\mathcal{P}(H)$, of the space of coends of H , and that the collection of these partitions is Γ -invariant. Then, the splitting described by Proposition 14.1 can be assumed to be canonical. Moreover, suppose that $H \in \mathcal{H}$ is conjugate into the vertex group, G , of the splitting. Then distinct elements of the partition $\mathcal{P}(H)$ are associated to distinct branches of the Bass-Serre tree based at the vertex which stabilises G .*

(The meaning of the last statement is elaborated on in Section 1. Of course, if H has only finitely many coends, we could take each element of $\mathcal{P}(H)$ to be a singleton.)

Proof. Using Lemma 9.2, for each $H \in \mathcal{H}$, we choose the corresponding axis, A , to separate different elements of $\mathcal{P}(H)$. We define a partition, $\mathcal{R}(A)$, of $\mathcal{C}(A)$ by deeming two sides of A to be equivalent if they correspond to the same element of $\mathcal{P}(H)$. Performing this construction Γ -equivariantly for all $H \in \mathcal{H}$, we get a Γ -cactus, $\Omega_{\mathcal{R}}(\mathcal{A})$, as described at the end of Section 8. Since $\Omega_{\mathcal{R}}(\mathcal{A})$ has fewer relations than $\Omega(\mathcal{A})$, it is discrete by Proposition 11.1. The splitting described by Proposition 14.1 now has the required properties. \square

The vertex groups of the second type are called “hanging fuchsian” groups. In general, we can define a *hanging fuchsian* group as one that features as a vertex group of this type in some splitting of Γ .

As a corollary of Proposition 14.1, we get the following (weak) version of the Annulus Theorem.

Proposition 14.3. *Suppose that Γ is a one-ended finitely generated group which contains a two-ended codim-one self-commensurating subgroup. Then, either Γ is planar, or else Γ splits non-trivially over a two-ended subgroup.*

Proof. Apply Proposition 14.1 to the set of Γ -conjugates of the subgroup, H . Either the corresponding graph of groups, \mathcal{G} , is trivial and Γ is planar, or else Γ splits over a two-ended subgroup. \square

Note that Proposition 14.3 applies, in particular, to hyperbolic groups, giving another proof of the result of [ScS1] and [Bo1].

15. MORE GENERAL SPLITTINGS

If we allow for two-ended subgroups with large commensurators (i.e. not two-ended), then the situation becomes more complicated. However, we do get some kind of canonical splitting as we now describe:

Proposition 15.1. *Suppose that Γ is a finitely generated one-ended group. Suppose that \mathcal{H} is a collection of codim-one two-ended subgroups of Γ , consisting of a finite union of Γ -conjugacy classes. Suppose that we have associated to each $H \in \mathcal{H}$ a non-trivial H -invariant clopen partition, $\mathcal{P}(H)$, of the space of coends of H , and that the collection of such partitions is Γ -invariant. Then, associated to \mathcal{H} , we have a canonical representation of Γ as a finite bipartite graph of groups, \mathcal{G} , with vertex set $V(\mathcal{G}) = V_0(\mathcal{G}) \sqcup V_1(\mathcal{G})$. Every element of \mathcal{H} is conjugate into a vertex group in $V_0(\mathcal{G})$. A vertex group in $V_0(\mathcal{G})$ is of one of two kinds. Either it is (conjugate to) a finitely generated fuchsian group where the incident edge groups are precisely the maximal parabolic subgroups, or else it is the commensurator of some element of \mathcal{H} . Moreover, suppose that $H \in \mathcal{H}$ is conjugate into the vertex group, G , of the splitting. Then, distinct branches of the Bass-Serre tree based at G correspond to distinct elements of the partition, $\mathcal{P}(H)$.*

(There is no reason to suppose that a vertex group of the second kind must be finitely generated.)

The splitting may be constructed as follows. For each $H \in \mathcal{H}$, we choose an axis which separates the elements of the partition, $\mathcal{P}(H)$. This gives an H -invariant partition, $\mathcal{R}(A)$, of the set $\mathcal{C}(A)$. Performing this construction Γ -equivariantly, we get a Γ -invariant collection, \mathcal{A} , of axes (as in the proof of Proposition 14.2). However, this time, they need not be pairwise divergent.

Consider the set of half-strips of elements of \mathcal{A} . We deem two half-strips to be equivalent if they are parallel (i.e. at finite Hausdorff distance). Let Ω be the set of such parallel classes. The equivalence relation respects the involution which swaps the ends of any given strip. Therefore we get a partition, Π , of Ω into pairs. Given $\lambda \in \Pi$, we write $\mathcal{A}[\lambda]$ for the set of axes in \mathcal{A} in the parallel class corresponding to λ .

Given $x, y \in \Omega$ and $\lambda \in \Pi$, we write $x - \lambda - y$ to mean that there is some $A \in \mathcal{A}[\lambda]$ such that the half-strips x and y lie in different elements of the partition, $\mathcal{R}(A)$, of $\mathcal{C}(A)$. (This is clearly independent of the choice of representative half-strips for x and y .)

Lemma 15.2. *With this structure, Ω is a discrete cactus.*

Proof. We first verify that Ω is a cactus.

Suppose Π_0 is a finite subset of Π . Let $\Omega_0 = \bigcup \Pi_0 \subseteq \Omega$. Suppose $\lambda \in \Pi_0$. For each pair of elements, $x, y \in \Omega_0$ such that the relation $x - \lambda - y$ holds, we choose some axis, $A \in \mathcal{A}[\lambda]$ such that x and y lie in different elements of the partition, $\mathcal{R}(A)$, of $\mathcal{C}(A)$ (as in the definition of the ternary relation). Let $\mathcal{A}_0[\lambda] \subseteq \mathcal{A}[\lambda]$ be the finite set of axes arising this way. Now, all such axes are parallel, so the corresponding elements of \mathcal{H} are commensurate. It follows that there is a two-ended subgroup, $H(\lambda)$, of Γ , which preserves setwise each element of $\mathcal{A}_0[\lambda]$. By Lemma 9.2, there is an axis, $B(\lambda)$, of $H(\lambda)$ that contains I_A for all $A \in \mathcal{A}_0[\lambda]$. Note that each element of $\mathcal{C}(B)$ lies in a unique element of $\mathcal{C}(A)$ for each $A \in \mathcal{A}_0[\lambda]$. We define an equivalence relation on $\mathcal{C}(B)$, by deeming two elements, $P, Q \in \mathcal{C}(B)$, to be equivalent if, for all $A \in \mathcal{A}_0[\lambda]$, the sides of A containing P and Q lie in

the same element of $\mathcal{R}(A)$. This equivalence relation defines a partition, $\mathcal{R}(B)$, of $\mathcal{C}(B)$. From the construction, it follows that if $x, y \in \Omega_0$, then $x - \lambda - y$ holds if and only if x and y lie in different elements of $\mathcal{R}(B)$.

We perform this construction for each $\lambda \in \Pi_0$. Let $\mathcal{B} = \{B(\lambda) \mid \lambda \in \Pi_0\}$. Thus, \mathcal{B} is a finite pairwise divergent set of axes, and each element of \mathcal{B} comes equipped with a partition of its sides. It follows from the discussion at the end of Section 8 that Ω_0 is a cactus. We deduce that Ω is a cactus as claimed.

To prove that Ω is discrete, we will need the following observation. Suppose that $\lambda, \mu \in \Pi$ with $\mu \times \lambda$. In other words, there is some $A \in \mathcal{A}$ such that (any representative of) μ crosses A . It follows that each element of $\mathcal{A}[\lambda]$ has precisely two sides. It in turn follows that, if $x, y \in \Omega$ with $x - \lambda - y$, then x and y lie on different sides of B for all $B \in \mathcal{A}[\lambda]$.

The proof of discreteness will now follow from the arguments of Section 11. However, since the collection \mathcal{A} need not be pairwise divergent, we cannot apply Proposition 11.1 directly.

First, we verify property (F1). Suppose that $\lambda, \mu, \nu \in \Pi$ satisfy $\lambda - \nu - \mu$. Let $\lambda = \{x, y\}$ and $\mu = \{z, w\}$. Since the relation $x - \nu - z$ holds, there is some $A \in \mathcal{A}[\nu]$ such that x and z lie in different elements of $\mathcal{R}(A)$. We claim that A that y and z must also lie in different elements of $\mathcal{R}(A)$. For if not, then λ must cross ν . Now, since $y - \nu - z$ holds, there is some $B \in \mathcal{A}[\nu]$ with y and z in different elements of $\mathcal{R}(B)$. But, from the above discussion, we see that they must also lie in different elements of $\mathcal{R}(A)$, thereby giving us a contradiction, thus proving the claim. Continuing with the same argument, we show that the same applies to the pair x and w and to the pair y and w . We write $\lambda - A - \mu$ (meaning that there do not exist an element of λ and an element of μ lying in the same element of $\mathcal{R}(A)$).

Now, suppose that $\lambda, \mu \in \Pi$ and let $\Pi' = \{\mu \in \Pi \mid \lambda - \nu - \mu\}$. For each $\nu \in \Pi'$, we choose $A(\nu) \in \mathcal{A}[\nu]$ with $\lambda - A(\nu) - \mu$ as in the previous paragraph. The proof of Proposition 11.3 now shows that $\{A(\nu) \mid \nu \in \Pi'\}$, and hence Π' itself, are finite. (Note that Lemma 11.2 did not require that \mathcal{A} be pairwise divergent.) Thus, Ω satisfies (F1).

Finally, suppose $\lambda, \mu \in \Pi$ are distinct. Let $\Pi'' = \{\nu \in \Pi \mid \nu \times \lambda, \nu \times \mu\}$. If $\nu \in \Pi''$, then we know that λ and μ cross A for all $A \in \mathcal{A}[\nu]$. Choosing such an element for all $\nu \in \Pi''$, we obtain a pairwise divergent set of axes. By the argument of Lemma 11.4, this set, and hence Π' , must be finite.

This shows that Ω satisfies (F2), and is hence discrete by Lemma 6.1. \square

Suppose that $\lambda \in \Pi$, and that $A \in \mathcal{A}[\lambda]$ corresponds to $H \in \mathcal{H}$. Then, the setwise stabiliser, $\Gamma(\lambda)$, of λ , is precisely the commensurator, $\text{Comm}(H)$, of H in Γ . Suppose H has infinite index in $\Gamma(\lambda)$, and that λ is not isolated (i.e. is crossed by some other element of Π). Then, by Proposition 12.1, Γ is virtually $\mathbf{Z} \oplus \mathbf{Z}$, and we set \mathcal{G} to be a single point. If we rule out this case, then we see that each cross-connected component of Ω is of one of two types. Either it consists of a single pair, $\lambda \in \Pi$, whose stabiliser is equal to $\text{Comm}(H)$ for some $H \in \mathcal{H}$, or else it is infinite, with each pair stabiliser two-ended. In the latter case, the setwise stabiliser of the component is a fuchsian group.

The remainder of Proposition 15.1 now follows as with Proposition 14.1.

REFERENCES

- [AN] S.A.Adeleke, P.M.Neumann, *Relations related to betweenness: their structure and automorphisms* *Memoirs Amer. Math. Soc.* Volume 131, no. 623 (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
- [Bo1] B.H.Bowditch, *Cut points and canonical splitting of hyperbolic groups*, Acta Math. **180** (1998) 145–186. MR **99g**:20069
- [Bo2] B.H.Bowditch, *Convergence groups and configuration spaces*, in “Group Theory Down Under”, (ed. J.Cossey, C.F.Miller III, W.D.Neumann, M.Shapiro), de Gruyter (1999) 23–54. MR **2001d**:20035
- [Bo3] B.H.Bowditch, *Treelike structures arising from continua and convergence groups*, Memoirs Amer. Math. Soc. Volume 139, no. 662 (1999). MR **2000c**:20061
- [Bo4] B.H.Bowditch, *Planar groups and the Seifert conjecture*, preprint, Southampton (1999).
- [CJ] A.Casson, D.Jungreis, *Convergence groups and Seifert fibered 3-manifolds*, Invent. Math. **118** (1994) 441–456. MR **96f**:57011
- [DSa] M.J.Dunwoody, M.E.Sageev, *JSJ-splittings for finitely presented groups over slender subgroups*, Invent. Math. **135** (1999) 25–44. MR **2000b**:20050
- [DSw] M.J.Dunwoody, E.L.Swenson, *The algebraic torus theorem*, Invent. Math. **140** (2000) 605–637. MR **2001d**:20039
- [FP] K.Fujiwara, P.Papasoglu, *JSJ decompositions of finitely presented groups and complexes of groups*, preprint (1997).
- [Ga] D.Gabai, *Convergence groups are fuchsian groups*, Ann. of Math. **136** (1992) 447–510. MR **93m**:20065
- [Geo] R.Geoghegan, *Topological methods in group theory*, manuscript, Binghamton (2000).
- [GerM] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I*, Proc. London Math. Soc. **55** (1987) 331–358. MR **88m**:30057
- [Gr] M.Gromov, *Hyperbolic groups*, in “Essays in Group Theory” (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263. MR **89e**:20070
- [K] P.H.Kropholler, *A group theoretic proof of the torus theorem*, in “Geometric Group Theory, Volume 1”, London Math. Soc. Lecture Note Series No. 181, (ed. G.A.Niblo, M.A.Roller), Cambridge University Press (1993) 138–158. MR **94i**:57029
- [KR] P.H.Kropholler, M.A.Roller, *Relative ends and duality groups*, J. Pure Appl. Algebra **61** (1989) 197–210. MR **91b**:20069
- [M] G.Mess, *The Seifert conjecture and groups that are coarse quasiisometric to planes*, preprint, UCLA (1988).
- [P] P.Papasoglu, *Quasi-isometry invariance of group splittings*, preprint, Orsay (2000).
- [RS] E.Rips, Z.Sela, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. **146** (1997) 53–109. MR **98m**:20044
- [ScS1] P.Scott, G.A.Swarup, *An algebraic annulus theorem*, Pacific J. Math. **196** (2000) 461–506. CMP 2001:05
- [ScS2] P.Scott, G.A.Swarup, *The number of ends of a pair of groups*, preprint, Ann Arbor/Melbourne (2000).
- [Se] Z.Sela, *Structure and rigidity in (Gromov) hyperbolic groups and discrete, groups in rank 1 Lie groups II* Geom. Funct. Anal. **7** (1997) 561–593. MR **98j**:20044
- [Sw] E.L.Swenson, *Axial pairs and convergence groups on S^1* , Topology **39** (2000) 229–237. MR **2001c**:20097
- [T1] P.Tukia, *Homeomorphic conjugates of fuchsian groups*, J. Reine Angew. Math. **391** (1988) 1–54. MR **89m**:30047
- [T2] P.Tukia, *Convergence groups and Gromov’s metric hyperbolic spaces*, New Zealand J. Math. **23** (1994) 157–187. MR **96c**:30042
- [W] L.E.Ward, *Axioms for cutpoints*, in “General topology and modern analysis”, Proceedings, University of California, Riverside (ed. L.F.McAuley, M.M.Rao), Academic Press (1980) 327–336. MR **82g**:54053

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

E-mail address: B.H.Bowditch@maths.soton.ac.uk