

Groups acting on Cantor sets and the end structure of graphs

Brian H. Bowditch

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

B.H.Bowditch@maths.soton.ac.uk

[First draft: February 2000. Revised: January 2001]

0. Introduction.

In this paper, we describe some of the interconnections between the end structure of graphs, groups acting on protrees, and convergence actions on Cantor sets. Our work ties in with recent work of Gerasimov and Dunwoody and earlier work of Bergman.

It was shown by Hopf in the 1930s that the space of ends of (any Cayley graph of) a finitely generated group, Γ , either consists of 0, 1 or 2 points, or else is a Cantor set. In the late 1960s, Stallings used topological methods to show that in the last case Γ splits non-trivially over a finite subgroup. Shortly afterwards, Bergman [Ber] gave a different proof of the same result. One can interpret his construction as giving us a Γ -invariant nested subset of the boolean algebra of clopen sets of the space of ends, and hence a splitting of the group via Bass-Serre theory. In fact, one can think of this algebra combinatorially in terms of cuts of the Cayley graph, i.e. finite sets of edges which separate the graph into more than one infinite subset (cf. [Du1]). Bergman's proof uses a certain "norm" defined on the set of such cuts. Recently Dunwoody and Swenson showed how one can use Bergman's norm to construct nested generating sets of arbitrary invariant subalgebras of the algebra of cuts of a locally finite graph (see [DuS] and Section 7). In this paper we generalise Bergman's norm, and some its consequences, to certain non-locally finite graphs (see Section 10). One of the main applications we describe here will be to convergence group actions on Cantor sets. Similar results have already been obtained by Gerasimov [Ger], by other methods. One can also use these ideas to give a simplified proof of the Almost Stability Theorem of [DiD] (see Section 15).

The notion of a convergence group was defined by Gehring and Martin [GehM]. It extracts the essential dynamical features of a kleinian group acting on the boundary of (classical) hyperbolic space. This generalises to boundaries of proper hyperbolic spaces in the sense of Gromov [Gr], see for example, [T2,F,Bo3]. The celebrated result of [T1,Ga,CJ] tells us that every convergence action on the circle arises from a fuchsian group. (In particular, the work of Tukia [T1] deals with the case of cyclically ordered Cantor sets.) There has also been a study of convergence actions on more general continua and their applications to hyperbolic and relatively hyperbolic groups (see for example [Bo2] and the references therein). Here we turn our attention to the opposite extreme, namely convergence actions on Cantor sets. An extensive study of such actions in relation to median algebras can be found in [Ger]. An example of such an action is that of a finitely generated group acting on its space of ends, which is a Cantor set if the group is not virtually cyclic and splits non-trivially over a finite subgroup. We shall see that, under some finiteness assumptions on the group, this type of example is typical.

Via Stone duality, a convergence action of a group, Γ , on a Cantor set is equivalent to a certain kind of action on a Boolean algebra. To obtain a splitting of Γ from this action, there are two issues to be addressed. The first is to find a Γ -invariant nested generating set for the boolean algebra, and the second is to arrange that this set is cofinite (in other words to show that the algebra is finitely generated as a Γ -boolean algebra). To get a handle on these issues, we assume that Γ is (relatively) finitely generated, and work with the algebra of cuts of a connected Γ -graph given by this hypothesis. The nested generating set is obtained using the Bergman norm (see Sections 7 and 10). In [Ger], a similar result is obtained via the theory of median algebras. For the second part, we need to assume that Γ is (relatively) almost finitely presented. We use a result of [DiD] (or [BesF]) to obtain a cofinite generating set (Section 12). This is really an issue of accessibility. (We remark that it is shown in [DiD] that a finitely generated group, Γ , is accessible if and only if the boolean algebra of almost invariant subsets is finitely generated as a Γ -boolean algebra.)

I am indebted to Victor Gerasimov for explaining to me some of his work in this area. I have benefited greatly from discussions with Martin Dunwoody, in particular for introducing me to the work of Bergman, as well as explaining to me how accessibility results could be use to obtain finitely generated algebras (see Sections 11 and 12). Some of the work for the present article was carried out while visiting the Centre de Recerca Matemàtica in Barcelona. I am grateful for the support and hospitality of this institution, and for helpful discussions with Warren Dicks while there.

1. Summary of results about convergence groups.

We shall introduce some of the results of this paper by giving a series of results, which are, in some sense, increasingly general, but which require greater elaboration in their formulation. They will be proven in Section 13. As mentioned in the introduction, most of these results have been obtained, in some form, by Gerasimov using other methods [Ger]. The results will be made more precise in later sections. We shall assume throughout that the convergence actions are minimal, i.e. that there is no discontinuity domain. For definitions regarding convergence actions, see Section 8.

The simplest example of a convergence action on a Cantor set is that of a (virtually) free group acting on its space of ends (which is the same as its boundary as a hyperbolic group). Such an action has no parabolic points. One result says that this is typical:

Theorem 1.1 : *If Γ is an almost finitely presented group acting as a minimal convergence group on a Cantor set, M , without parabolic points, then Γ is virtually free, and there is a Γ -equivariant homeomorphism from M to the space of ends of Γ .*

By *almost finitely presented* we mean that Γ admits a cocompact properly discontinuous action on a connected 2-complex, Σ , with $H_1(\Sigma; \mathbf{Z}_2) = 0$. (Without loss of generality, one can assume that the action on Σ is also free.) Clearly, finite presentability in the usual sense implies almost finite presentability.

More generally, it is well-known that any finitely generated group acts as a convergence

group on its space of ends (see Section 9). If Γ is finitely presented, then Dunwoody’s accessibility theorem [Du2] tells us that we can represent Γ as a finite graph of groups with finite edge groups and finite or one-ended vertex groups. The one-ended vertex groups are uniquely determined as the maximal one-ended subgroups of Γ , and are precisely the maximal parabolic subgroups with respect to this convergence action. We have a converse:

Theorem 1.2 : *Suppose that Γ is an almost finitely presented group and acts as a minimal convergence group on a Cantor set, M , with all maximal parabolic subgroups finitely generated and one-ended. Then, M is equivariantly homeomorphic to the space of ends of Γ .*

The hypothesis on parabolic subgroups in the above result is somewhat unnatural. Note that the splitting given by Dunwoody’s accessibility result gives us, via Bass-Serre theory, an action on a simplicial tree with finite edge stabilisers and finite quotient. To such a tree, we can associate a “boundary” as described in Section 5. The group acts as a convergence group on this boundary, with the infinite vertex stabilisers as the maximal parabolic subgroups. If these happen to be one-ended, then we recover the space of ends Γ . We also have a converse:

Theorem 1.3 : *Suppose that Γ is an almost finitely presented group and acts as a minimal convergence group on a Cantor set, M . Then, Γ has a representation as a finite graph of groups with finite edge groups such that M is equivariantly homeomorphic to the “boundary” of the associated Bass-Serre tree.*

In fact, we only really require that Γ be finitely presented relative to a set of parabolic subgroups. More precisely,

Theorem 1.4 : *Suppose that a group, Γ , acts as a minimal convergence group on a Cantor set, M . Suppose that \mathcal{G} is a finite collection of parabolic subgroups of Γ with respect to this action. Suppose that Γ is almost finitely presented relative to \mathcal{G} . Then, the conclusion of Theorem 1.3 holds.*

The definition of Γ being “almost finitely presented relative to \mathcal{G} ” will be elaborated on in Section 3.

Note that, from the conclusion of Theorem 1.3, one may deduce that there are only finitely many conjugacy classes of maximal parabolic subgroups. In fact, it follows that the action of Γ on M is geometrically finite, and that Γ is hyperbolic relative to the collection of maximal parabolic subgroups (see Section 8 for definitions). Moreover, the boundary of Γ as a relatively hyperbolic group may be identified with the boundary of the Bass-Serre tree.

We can further weaken the hypotheses, and assume only that Γ is finitely generated (relative to a class of parabolic subgroups). However, in this case we can only be assured of an action of Γ on a protree, as opposed to a simplicial tree. From this, one can deduce that the original action on M is an inverse limit of geometrically finite actions of the above

type. This is elaborated on in Section 14.

2. Results concerning boolean algebras

In Section 10, we give a generalisation of the result of Bergman concerning the boolean algebra of slices of a connected graph. By the argument of [DuS], this allows us to construct nested systems of generators for such algebras. We describe below how a special case of this relates to the results outlined in Section 1.

Suppose that K is a connected graph. By a (*directed*) K -*slice*, we mean a subset, A , of the vertex set, $V(K)$, of K , such that only finitely many edges of K have precisely one endpoint in A . (It will sometimes be more convenient to work with *undirected* K -slices, i.e. unordered pairs $\{A, A^*\}$, where A is a K -slice.) The set of K -slices is a boolean subalgebra of the power set of $V(K)$. We say that two K -slices are *nested* if one is contained in the other, or in the complement of the other.

Suppose that Γ acts on K with finite edge stabilisers and such that there are finitely many orbits of edges, and finitely many orbits of circuits of any given length. If K (or equivalently Γ) is countable, we shall show that any boolean subalgebra, \mathcal{A} , of the algebra of K -slices of K has a Γ -invariant (pairwise) nested generating set.

In fact, we can weaken the above hypotheses. What we really require is that the graph, K , has finite quotient, and that for any $A \in \mathcal{A}$, and any distinct $x, y \in V(K)$, only finitely many distinct Γ -images of A contain x but not y . This is automatically satisfied in the set-up of the previous paragraph.

In the situation described in Section 1, the finite generation hypotheses on the group, Γ , give us our graph, K , and the boolean algebra of clopen sets of M gives us the boolean subalgebra of K -slices. The nested system of generators supplied by the above result gives us a treelike structure, or more precisely a protree in the sense of Dunwoody [Du4], on which the group acts. If we assume in, an addition, that Γ is (relatively) almost finitely presented, then an accessibility result (see Section 12) tells us that our boolean subalgebra must be finitely generated, and we end up with a simplicial tree. This gives us our required splitting of Γ via Bass-Serre theory. Stone duality tells us that the “boundary” of this tree is equivariantly homeomorphic to the original space, M .

Some other applications of the constructions of Section 10 are outlined in Section 15.

3. Finiteness conditions.

In this section, we elaborate on the finiteness conditions featured in Section 1. These are most conveniently expressed in terms of group actions on sets (cf. [Bo4]).

Suppose Γ is a group. A Γ -*set*, V , is a set on which the group Γ acts. Given $x \in V$, we write $\Gamma(x) = \{g \in \Gamma \mid gx = x\}$. We write $V_\infty = \{x \in V \mid |\Gamma(x)| = \infty\}$. We can interpret a property of V as a property of the group, Γ “relative to” the set of infinite point stabilisers, $\{\Gamma(x) \mid x \in V_\infty\}$. We shall also speak of “ Γ -graphs”, “ Γ -trees” and “ Γ -boolean algebras” etc. for such objects admitting Γ -actions.

We shall say that the Γ -set, V , is *cofinite* if $|V/\Gamma|$ is finite. A *pair stabiliser* is a subgroup of the form $\Gamma(x) \cap \Gamma(y)$ for $x \neq y$. We shall frequently assume that V has finite pair stabilisers. We shall say that V is “0-connected” if it can be identified with the vertex set, $V(K)$, of a connected cofinite Γ -graph, K (or equivalently, as a Γ -invariant subset of $V(K)$ containing $V_\infty(K)$). We say the V is *1-connected* if it is the vertex set of a cofinite simply connected CW-complex (or equivalently, a Γ -invariant subset containing $V_\infty(\Sigma)$ of the vertex set of a cofinite simplicial complex, Σ .) Note that 1-connectedness is equivalent to the assertion that for some (or equivalently any) cofinite connected Γ -graph with vertex set V , there is some n such that $\Omega_n(K)$ is simply connected. Here $\Omega_n(K)$ is the 2-complex obtained by attaching a 2-cell along every circuit of length n in K . More generally, we say that V is *\mathbf{Z}_2 -homologically 1-connected* if $\Omega_n(K)$ is \mathbf{Z}_2 -acyclic, i.e. $H_1(\Omega_n(K); \mathbf{Z}_2) = 0$.

We say that a graph, K , is *fine* if, for each n , there are only finitely many circuits of length n containing any given edge. Clearly, this implies that the complex $\Omega_n(K)$ described above is locally finite away from $V = V(K)$ (and can thus be subdivided to give a simplicial complex that is locally finite away from V). We say that a Γ -set is *fine* if some (hence any) cofinite connected Γ -graph, K , with vertex set V , is fine. Here, the fineness of K is equivalent to saying that there are finitely many circuits of any given length modulo Γ .

Note that if V is fine and (\mathbf{Z}_2 -homologically) 1-connected we can embed V equivariantly in the vertex set, $V(\Sigma)$, of a cocompact simply connected (\mathbf{Z}_2 -acyclic) 2-dimensional simplicial complex, Σ , which is locally finite away from $V(\Sigma)$, and such that the stabiliser of each element of $V(\Sigma) \setminus V$ is finite. Indeed, we can easily arrange that each such stabiliser is trivial.

Suppose that \mathcal{G} is a non-empty collection of self-normalising subgroups of Γ , which is a finite union of conjugacy classes, and such that the intersection of any two distinct elements of \mathcal{G} is finite. We may view \mathcal{G} as a Γ -set with Γ acting by conjugation. We say that Γ is *finitely generated relative to \mathcal{G}* if \mathcal{G} is 0-connected. We say that Γ is *finitely presented relative to \mathcal{G}* if \mathcal{G} is 1-connected. In the case of interest to us, \mathcal{G} will always be fine. (Arguably it might be more natural to include the hypothesis of fineness in the definition of relative finite presentability, though this will not matter to us here.)

Suppose Γ acts as a convergence group on a Cantor set, and Π is a cofinite Γ -invariant subset of parabolic points. The corresponding collection, \mathcal{G} , of maximal parabolic subgroups of Γ is isomorphic to Π as a Γ -set. Moreover, we shall see (Lemma 8.2) that Π (or equivalently \mathcal{G}) is necessarily fine.

4. Boolean algebras.

For the next few sections, we forget about our group, Γ , and review a few of the connections between boolean rings (or algebras), topological spaces and treelike structures of various sorts.

Let \mathcal{B} be a boolean ring, i.e. a commutative ring with a one, 1, satisfying $x^2 = x$ for all $x \in \mathcal{B}$. We write $x^* = 1 + x$, $x \wedge y = xy$ and $x \vee y = x + y + xy$. Thus, $(\mathcal{B}, \wedge, \vee, *)$ is a boolean algebra. We write $x \leq y$ to mean that $xy = x$. Thus \leq is a partial order on \mathcal{B} ,

and $[x \mapsto x^*]$ is an order reversing involution.

Given any set, X , its power set, $\mathcal{P}(X)$ is a boolean algebra with $P^* = X \setminus P$, $P \wedge Q = P \cap Q$ and $P \vee Q = P \cup Q$, for $P, Q \in \mathcal{P}(X)$. If \mathcal{C} is a subalgebra of X and $Y \subseteq X$, we write $Y \wedge \mathcal{C}$ for the boolean subalgebra, $\{P \cap Y \mid P \in \mathcal{C}\}$, of $\mathcal{P}(Y)$. The map $[P \mapsto P \cap Y] : \mathcal{C} \rightarrow Y \wedge \mathcal{C}$ is an epimorphism of boolean algebras.

Suppose M is a compact totally disconnected topological space. The set, $\mathcal{B}(M)$ of clopen subsets of M is a boolean algebra.

The Stone duality theorem [St] tells us that every boolean algebra arises in this way (see for example [Si]). Suppose that \mathcal{B} is a boolean algebra. We associate to \mathcal{B} a compact totally disconnected space $\Xi = \Xi(\mathcal{B})$, called the *Stone dual*, such that $\mathcal{B}(\Xi) \cong \mathcal{B}$. This can be described in a number of equivalent ways. For example we can define Ξ as the set of boolean ring homomorphisms from \mathcal{B} to \mathbf{Z}_2 . This is a closed subset of the Tychonoff cube $\mathbf{Z}_2^{\mathcal{B}}$, and we topologise Ξ accordingly.

Alternatively, we define Ξ as the maximal ideal spectrum of the ring \mathcal{B} with the Zariski topology (or the prime ideal spectrum if we remove the zero ideal). Note that the complement of a maximal ideal in \mathcal{B} is an ultrafilter. We therefore get the same thing by taking the set of ultrafilters on \mathcal{B} . From this point of view, we can define a basis for the closed sets, by taking a typical basis element to be the set of all ultrafilters that contain a given element of \mathcal{B} .

If M is compact and totally disconnected, then $\Xi(\mathcal{B}(M)) = M$. The isomorphism from M to $\Xi(\mathcal{B}(M))$ is given by $[x \mapsto \mathcal{O}(x)]$, where $\mathcal{O}(x)$ is the ultrafilter $\{P \in \mathcal{B}(M) \mid x \in P\}$ (or, equivalently, the set of homomorphisms from $\mathcal{B}(M)$ to \mathbf{Z}_2 which send each set containing x to 1).

Note that if $f : \mathcal{B} \rightarrow \mathcal{B}'$ is a homomorphism, we get a continuous dual map, $f_* : \Xi(\mathcal{B}') \rightarrow \Xi(\mathcal{B})$. If f is surjective, then f_* is injective, so we can identify $\Xi(\mathcal{B}')$ as a closed subset of $\Xi(\mathcal{B})$.

We say that two non-zero elements $x, y \in \mathcal{B}$ are *nested* if xy is equal to 0, x, y or $1 + x + y$. This is equivalent to saying that one of xy, xy^*, x^*y or x^*y^* equals 0. We say that a subset, $\mathcal{E} \subseteq \mathcal{B}$ is *nested* if $0, 1 \notin \mathcal{E}$ and every pair of elements of \mathcal{E} are nested.

Suppose that \mathcal{B} is a subalgebra of $\mathcal{P}(X)$ for some set X . If $x \in X$, then $\mathcal{O}(x) = \{P \in \mathcal{B} \mid x \in P\}$ determines a point of $\Xi(\mathcal{B})$, so we get natural map from X to $\Xi(\mathcal{B})$.

5. Simplicial trees.

Let T be a simplicial tree with vertex set $V = V(T)$ and edge set $E(T)$. We write $\vec{E}(T)$ for the directed edge set. Given $\vec{e} \in \vec{E}(T)$, we write e for the underlying undirected edge, and $-\vec{e}$ for the same edge pointing in the opposite direction. We write $T(\vec{e})$ for the component of $T \setminus e$ containing the head of \vec{e} (where we think of \vec{e} as directed from its tail to its head), and write $V(\vec{e}) = T(\vec{e}) \cap V(T)$. If $\vec{e}, \vec{f} \in \vec{E}(T)$, we write $\vec{e} < \vec{f}$ to mean that \vec{f} points towards \vec{e} and \vec{e} points away from \vec{f} . (In some papers the opposite convention is used.) Note that this is equivalent to saying that $V(\vec{e})$ is strictly contained in $V(\vec{f})$. Clearly \leq is a partial order on $\vec{E}(T)$ with order reversing involution $[\vec{e} \mapsto -\vec{e}]$. Moreover, if $\vec{e}, \vec{f} \in \vec{E}(T)$, then precisely one of the statements $\vec{e} < \vec{f}$, $-\vec{e} < \vec{f}$, $\vec{e} < -\vec{f}$, $-\vec{e} < -\vec{f}$,

$\vec{e} = \vec{f}$ or $\vec{e} = -\vec{f}$ holds.

A subset, $F \subseteq \vec{E}(T)$ is a *transversal* if, for all $\vec{e} \in \vec{E}(T)$, precisely one of \vec{e} or $-\vec{e}$ lies in F . A transversal, F , is a *flow* on T if no two elements of F point away from each other (i.e. there do not exist $\vec{e}, \vec{f} \in F$ with $\vec{e} \leq -\vec{f}$). A flow must be of one of two types. Either there is some (unique) $v \in V(T)$ such that each element of F points towards v , or else there is some infinite ray, $\alpha \subseteq T$ such that all edges of $F \cap E(\alpha)$ point away from its basepoint, and all other elements of F point towards α . We can alternatively think of a flow of the second kind as a cofinality class of rays in T , where two rays are *cofinal* if they intersect in a ray. We can identify the set of such flows (as a set) with the boundary, ∂T , of T thought of as a hyperbolic space in the sense of Gromov [Gr]. Here, however, we shall want to think of the boundary of T as the set of all flows with an appropriate topology which we now go on to describe.

Given $W \subseteq V(T)$, we shall write $I(W) \subseteq E(T)$ for the set of edges with precisely one endpoint in T . Clearly, $I(W^*) = I(W)$. Let $\mathcal{B} = \mathcal{B}(T)$ be the boolean algebra consisting of those $W \in \mathcal{P}(V)$ for which $I(W)$ is finite. Let $\mathcal{E} = \mathcal{E}(T) = \{W \in \mathcal{B} \mid |I(W)| = 1\} = \{V(\vec{e}) \mid \vec{e} \in \vec{E}(T)\}$. The map $[\vec{e} \mapsto V(\vec{e})]$ therefore gives an identification of $\vec{E}(T)$ with \mathcal{E} . Note that the relation, \leq , defined on $\vec{E}(T)$ above agrees with the partial order defined in the boolean algebra \mathcal{B} . Moreover, the involution $[\vec{e} \mapsto -\vec{e}]$ corresponds to the involution $[W \mapsto W^*]$ on \mathcal{B} . We see, in fact, that \mathcal{E} is a nested set of generators for \mathcal{B} .

Let $\Xi(T) = \Xi(\mathcal{B}(T))$ be the Stone dual of $\mathcal{B}(T)$. If we think of an element of $\Xi(T)$ as an ultrafilter on $\mathcal{B}(T)$, then the elements of \mathcal{E} lying in that ultrafilter define a flow on T . In other words, we may identify $\Xi(T)$ with the set of flows on T , and hence with $V(T) \sqcup \partial T$.

Now, T is fine (as defined in Section 3) and hyperbolic (in the sense of Gromov). It thus has associated with it a compact space ΔT , as in [Bo4]. As a set, ΔT may be defined as $V(T) \sqcup \partial T$. Note that any two elements, $x, y \in \Delta T$ can be “connected” by a unique arc, $[x, y]$, which may be compact, a ray, or a bi-infinite geodesic depending on whether or not $x, y \in V(T)$. Given $x \in \Delta T$ and $I \subseteq E(T)$, let $B(x, I) = \{y \in \Delta T \mid I \cap E([x, y]) = \emptyset\}$. We define a topology on ΔT by taking a neighbourhood base of $x \in \Delta(T)$ to be the collection $\{B(x, I)\}_I$ as I runs over all finite subsets $E(T)$. In this topology, ΔT is compact and hausdorff. This construction can be used to define the “boundary” of a relatively hyperbolic group, as in [Bo4]. In our particular case, the topology can be equivalently defined as follows. First note that if $S \subseteq T$ is a subtree, then ΔS is a subset of ΔT . If $x \in V(T)$, we take a neighbourhood base of x to consist of sets of the form ΔS , where S runs over subtrees of T which consist of all but finitely many branches of T based at x . If $x \in \partial T$, then a neighbourhood base of x is given by $\{\Delta T(\vec{e}) \mid \vec{e} \in F\}$, where F is the flow on T corresponding to x . It now follows easily that the topologies on $\Xi(T)$ and ΔT agree. In other words, we can identify $\Xi(T)$ with the ideal boundary, ΔT , of T as defined in [Bo4].

We note that $V(T)$ is dense in ΔT , since if $x \in \partial T$ then the sequence of vertices in any ray representing x converges to x in the topology on ΔT . Note that it follows that the boolean algebra of sets of the form $A \cap V(T)$, as A runs over clopen subsets of ΔT , is isomorphic to the original boolean algebra $\mathcal{B}(T)$.

Note that each finite-degree vertex of T is isolated in ΔT . We shall define the *ideal boundary* $\Delta_0 T$, of the tree, T , as the space ΔT minus the vertices of T of finite degree.

6. Protrees and nested sets.

We begin by recalling the definition of a “protree” due to Dunwoody, see for example [Du4].

A *protree* is a set Θ , with an involution $[x \mapsto x^*]$ and a partial order, \leq , with the property that for any $x, y \in \Theta$, precisely one of the six relations $x < y$, $x^* < y$, $x < y^*$, $x^* < y^*$, $x = y$, $x = y^*$ holds. We refer to a $*$ -invariant subset of Θ as a *subprotree*. This clearly has itself the structure of a protree.

An example of a protree is the directed edge set of any simplicial tree, as described in the last section. In fact, any finite protree can be realised as the directed edge set of a finite tree (a property which could serve as an equivalent definition). More generally any “discrete” protree can be realised as the directed edge set of a simplicial tree. A protree is said to be *discrete* if, for all $x, y \in \Theta$, the set $\{z \in \Theta \mid x \leq z \leq y\}$ is finite. We can define transversals and flows on a protree, exactly as for simplicial trees. (However, we cannot in general classify flows in the same way.)

Suppose Θ is a protree. Let $\mathcal{F} = \mathcal{F}(\Theta)$ be the boolean ring with generating set Θ , and with relations $x + x^* = 1$ for all $x \in \Theta$ and $xy = 0$ for all $x, y \in \Theta$ satisfying $x < y^*$. Note that it follows that if $x, y \in \Theta$ are distinct, then xy must be equal to 0, x , y or $1 + x + y$. In particular, it follows that any element of \mathcal{F} can be written as a sum of finitely many elements of Θ . Note that if $x \in \Theta$, then we can define an epimorphism, $\theta : \mathcal{F} \rightarrow \mathbf{Z}_2$ by $\theta(x) = 1$ and $\theta(y) = 0$ for all $y \in \Theta \setminus \{x\}$. It therefore follows that $x \neq 0$ for all $x \in \mathcal{F}$.

Suppose that $F \subseteq \Theta$ is any transversal. We may identify F as a subset of $\mathcal{F}(\Theta)$, and as such, it generates $\mathcal{F}(\Theta)$ as a ring (indeed as an additive group). If $x, y \in F$, then precisely one of the relations $xy = 0$, $x^*y = 0$, $xy^* = 0$, $x^*y^* = 0$ holds. Thus, Θ is a nested set of generators for $\mathcal{F}(\Theta)$. Any element of $\mathcal{F}(\Theta)$ can be written in a *standard* form $\epsilon + \sum_{i=1}^n x_i$ where $\epsilon \in \{0, 1\}$ and x_1, \dots, x_n are distinct elements of F .

If Θ is a discrete protree, and T is the corresponding simplicial tree, then there is an epimorphism $\phi : \mathcal{F}(\Theta) \rightarrow \mathcal{B}(T)$ defined by $\phi(x) = V(\vec{e})$, where $\vec{e} \in \vec{E}(T)$ is the directed edge corresponding to $x \in \Theta$. Let $F \subseteq \Theta$ be any transversal. If $x_1, \dots, x_n \in F$ are distinct, then the corresponding elements of $\mathcal{B}(T)$ are distinct. If $n \neq 0$, it follows easily that their symmetric difference can be neither \emptyset nor $V(T)$. In other words, we see that any standard form of any element in the kernel of ϕ must be trivial. Hence, the kernel is trivial, so ϕ is, in fact, an isomorphism.

We now return to the case of a general protree, Θ , with transversal F . Suppose that $\Phi \subseteq \Theta$ is any subprotree. Then, there is a natural epimorphism, $\theta : \mathcal{F}(\Theta) \rightarrow \mathcal{F}(\Phi)$ defined by $\theta(x) = x$ if $x \in \Phi$, and $\theta(x) = 0$ if $x \in \Theta \setminus \Phi$. In particular, suppose x_1, \dots, x_n are distinct. We get an epimorphism from $\mathcal{F}(\Theta)$ to $\mathcal{F}(\Phi)$, where $\Phi = \bigcup_{i=1}^n \{x_i, x_i^*\}$. Now, $\mathcal{F}(\Phi)$ is isomorphic to the boolean algebra on a finite tree, as above, and so it follows that $x_1 + \dots + x_n \notin \{0, 1\}$. We see that the standard form of an element of Θ (with respect to a given transversal, F) is unique. We therefore have an explicit description of the ring $\mathcal{F}(\Theta)$.

If $\Phi \subseteq \Theta$ is again any subprotree, we also get a homomorphism from $\mathcal{F}(\Phi)$ to $\mathcal{F}(\Theta)$ which extends the inclusion of Φ in Θ . From the above description, it is clear that this is injective. We therefore get a surjective map from $\Xi(\Theta)$ to $\Xi(\Phi)$. If Θ is an increasing

union of subprotrees, $(\Theta_n)_{n \in \mathbf{N}}$, then it is easily seen that $\Xi(\Theta)$ is an inverse limit of the system $(\Xi(\Theta_n))_n$ of topological spaces.

Let $\Xi(\Theta) = \Xi(\mathcal{F}(\Theta))$ be the Stone dual. If we think of an element of $\Xi(\Theta)$ as an ultrafilter on $\mathcal{F}(\Theta)$, then its intersection with Θ is a flow on Θ . We may therefore identify $\Xi(\Theta)$ with the set of flows on Θ . For non-discrete protrees, however, we do not get a clear distinction between vertices and boundary points, as in the simplicial case.

Suppose that \mathcal{B} is any boolean algebra with a nested set of generators, $\mathcal{E} \subseteq \mathcal{B}$. Now, \mathcal{E} has the structure of a protree, with involution and partial order induced from \mathcal{B} . We can therefore construct the boolean algebra $\mathcal{F}(\mathcal{E})$ as above, and identify \mathcal{B} as a quotient of $\mathcal{F}(\mathcal{E})$. Note that $\Xi(\mathcal{B})$ can thus be identified as a closed subset of the space $\Xi(\mathcal{E})$.

In fact, we can say more than this. We can formally identify \mathcal{E} as a subset of \mathcal{F} . When composed with the canonical epimorphism from \mathcal{F} to \mathcal{B} , this gives the inclusion of \mathcal{E} into \mathcal{B} . Let \mathcal{I} be the kernel of the canonical epimorphism from \mathcal{F} to \mathcal{B} . Note that if $x, y \in \mathcal{E}$, then $x, y, x + y \notin \mathcal{I}$ (since, by the definition of a nested set, they correspond to distinct non-zero elements of \mathcal{B}). Now a combinatorial argument shows that \mathcal{I} is generated by elements of the form $1 + \sum_{i=1}^n x_i$, where $x_1, \dots, x_n \in \mathcal{E}$ have the property that $x_i x_j = 0$ for $i \neq j$, and if $y \in \mathcal{E}$ then for some $i \in \{1, \dots, n\}$, we have $y x_i \in \{x_i, 1 + x_i + y\}$. In other words, we can think of the elements x_1, \dots, x_n as a set of edges whose tails all meet at a ‘‘vertex’’ of degree n of the protree \mathcal{E} . (This statement is precise if the protree \mathcal{E} happens to be discrete, and hence the edge set of a simplicial tree.) Suppose that $a \in \Xi(\mathcal{E}) \setminus \Xi(\mathcal{B})$. Now a corresponds to a flow, F , on \mathcal{E} . This cannot be identically zero on \mathcal{I} and so must be non-zero on some generator of \mathcal{I} of the above form. From this, it is easy to see that F converges to some vertex of finite degree. But such a point is easily seen to be isolated in $\Xi(\mathcal{E})$. We have shown that every point of $\Xi(\mathcal{E}) \setminus \Xi(\mathcal{B})$ is isolated.

Finally, suppose that \mathcal{B} is a subalgebra of $\mathcal{P}(V)$ for some set, V . There is a natural map from V to $\Xi(\mathcal{B})$ as defined at the end of Section 4. We therefore get a map from V to $\Xi(\mathcal{E})$. The image of a point $x \in V$ in $\Xi(\mathcal{E})$ is defined by the flow $\{A \in \mathcal{E} \mid x \in A\}$ on \mathcal{E} .

7. Construction of nested generating sets.

It was shown in [DuS] how the Bergman norm can be used to construct invariant nested subsets of a boolean algebra. Dunwoody has observed how an elaboration of this argument in fact gives us nested generating sets. The central idea may be formulated in a general fashion as follows.

Let \mathcal{B} be a boolean ring. We say that two elements, $x, y \in \mathcal{B}$ are *disjoint* if $xy = 0$.

Suppose that S is an ordered abelian group (or cancellative semigroup). Suppose that to each disjoint pair, $x, y \in \mathcal{B}$, we have associated an element $\sigma(x, y) \in S$. We suppose that $\sigma(x, y) = \sigma(y, x) \geq 0$, and that $\sigma(x, y) > 0$ if $x, y \neq 0$. Moreover, if $x, y, z \in \mathcal{B}$ are pairwise disjoint, then $\sigma(x, y + z) = \sigma(x, y) + \sigma(x, z)$. Given any $x, y \in \mathcal{B}$, we write $\mu(x) = \mu(x, x^*)$ and $\mu(x|y) = \sigma(xy, x^*y)$. Clearly $\mu(x^*) = \mu(x)$ and $\mu(x^*|y) = \mu(x|y)$.

Suppose now that $x, y \in \mathcal{B}$ are non-nested, i.e. that xy, x^*y, xy^*, x^*y^* are all non-zero.

If $\mu(y|x) \leq \mu(x|y^*)$, then

$$\begin{aligned}
\mu(xy) &= \sigma(xy, (xy)^*) \\
&= \sigma(xy, x^* + xy^*) \\
&= \sigma(xy, x^*) + \sigma(xy, xy^*) \\
&\leq \sigma(xy, x^*) + \sigma(xy^*, x^*y^*) \\
&< \sigma(xy, x^*) + \sigma(xy^*, x^*y^*) + \sigma(xy^*, x^*y) \\
&= \sigma(xy, x^*) + \sigma(xy^*, x^*y^* + x^*y) \\
&= \sigma(xy, x^*) + \sigma(xy^*, x^*) \\
&= \sigma(x, x^*) \\
&= \mu(x).
\end{aligned}$$

Similarly, if $\mu(y|x) \leq \mu(x|y)$, then $\mu(xy^*) < \mu(x)$.

Now, if we allow ourselves to permute x, y, x^*, y^* , then we can always arrange that $\mu(y|x)$ is minimal among $\{\mu(x|y), \mu(y|x), \mu(x|y^*), \mu(y|x^*)\}$, so that $\max\{\mu(xy), \mu(xy^*)\} < \mu(x)$.

Lemma 7.1 : *If $\mu(\mathcal{B} \setminus \{0, 1\})$ is well-ordered (as a subset of S), then \mathcal{B} has a nested set of generators.*

Proof : Let $\mathcal{E} \subseteq \mathcal{B} \setminus \{0, 1\}$ be the set of $x \in \mathcal{B}$ such that x does not lie in the ring generated by $\{z \in \mathcal{B} \setminus \{0, 1\} \mid \mu(z) < \mu(x)\}$. Clearly, \mathcal{E} generates \mathcal{B} . Moreover, \mathcal{E} is nested. For if $x, y \in \mathcal{E}$ were not nested, then, without loss of generality, $\max\{\mu(xy), \mu(xy^*)\} < \mu(x)$. But $x = xy + xy^*$, giving a contradiction. \diamond

8. Convergence groups.

The notion of a convergence group was defined in [GehM]. For further discussion, see [T2,Bo3,T3].

Suppose that M is a compact metrisable space and that Γ is a group acting by homeomorphism on M . We say that this is a *convergence action* (or that Γ is a *convergence group*) if the induced action on the space of distinct triples of M (i.e. $M \times M \times M$ minus the large diagonal) is properly discontinuous. This is equivalent to the statement that if $(g_n)_{n \in \mathbf{N}}$ is any infinite sequence of distinct elements of Γ , then there are points, $a, b \in M$, and a subsequence $(g_{n_i})_i$ such that $g_{n_i}|_{M \setminus \{a\}}$ converges locally uniformly to b . We refer to the latter statement as the ‘‘convergence property’’ of Γ .

A subgroup, G , of Γ is *parabolic* if it is infinite, and fixes a unique point. Such a fixed point, x , is called a *parabolic point*, and its stabiliser, $\Gamma(x)$, is a maximal parabolic subgroup of Γ . We say that x is a *bounded* parabolic point if $(M \setminus \{x\})/\Gamma(x)$ is compact. (We allow for the possibility of a parabolic group being an infinite torsion group.)

A point $x \in M$ is a *conical limit point* if there is a sequence of elements $g_n \in \Gamma$, and distinct points, $a, b \in M$ such that $g_n x \rightarrow a$ and $g_n y \rightarrow b$ for all $y \in M \setminus \{x\}$. It is shown

in [T3] that a conical limit point cannot be a parabolic point. We say that the action of Γ on M is *geometrically finite* if every point of M is either a conical limit point or a bounded parabolic point. Such actions have been studied by Tukia [T3].

By the *space of distinct pairs* of M , we mean $M \times M$ minus the diagonal.

Lemma 8.1 : *Suppose that Γ acts on M as a convergence group, and that $x, y \in M$ are distinct and not conical limit points. Then, $\Gamma(x) \cap \Gamma(y)$ is finite, and the Γ -orbit of (x, y) is a discrete subset of the space of distinct pairs.*

Proof : Suppose, for contradiction, that $g_n \in \Gamma$ is a sequence of distinct elements of Γ with $g_n x \rightarrow a$ and $g_n y \rightarrow b$ with $a, b \in M$ distinct. The convergence property tells us that after passing to a subsequence, either $g_n|_{M \setminus \{x\}}$ converges (locally uniformly) to b , or else $g_n|_{M \setminus \{y\}}$ converges (locally uniformly) to a . Thus, either x or y is a conical limit point. \diamond

Recall the terms “connected” and “fine” as defined in Section 3.

Lemma 8.2 : *Suppose that Γ acts on M as a convergence group, and that $\Pi \subseteq M$ is a Γ -invariant subset. Suppose that no point of Π is a conical limit point. If Π is connected (as a Γ -set) then it is fine.*

Proof : We show that if K is any cofinite Γ -graph with vertex set $V(K) = \Pi$, then modulo Γ , there are only finitely many circuits of length n for any given n . Since Π is connected, we can take K to be connected, and we see that K is fine. Together with the first part of Lemma 8.1, this implies that Π is fine as claimed.

Suppose, for contradiction, that $(\beta^k)_{k \in \mathbf{N}}$ is an infinite sequence of circuits of length n in K , each lying in a different Γ -orbit. We write $\beta^k = x_1^k \dots x_n^k$, taking subscripts mod n . Passing to a subsequence, we can suppose that for all i each of the edges $\{x_i^k x_{i+1}^k\}_k$ lie in the same Γ -orbit. Thus, modulo Γ , we can suppose that $x_0^k = x_0$ and $x_1^k = x_1$ are independent of k . Now by Lemma 8.1, the set of pairs $\{(x_1, x_2^k) \mid k \in \mathbf{N}\}$ is discrete in the space of distinct pairs. Thus, again after passing to a subsequence, we can suppose that either $x_2^k = x_2$ is constant, or that $x_2^k \rightarrow x_1$. In the latter case, we can suppose that the x_2^k are all distinct, so since the set of pairs $(x_2^k, x_3^k)_k$ is discrete, we must also have $x_3^k \rightarrow x_1$. It follows inductively that $x_i^k \rightarrow x_1$ for all i , giving the contradiction that $x_0 = x_n^k \rightarrow x_1$. We can thus assume that $x_2^k = x_2$ is constant. But now, the same argument tells us that x_3^k is constant, so by induction, x_i^k is constant for all i . We derive the contradiction that β^k is constant. \diamond

Note that, by the result of Tukia [T3], Lemma 8.2 applies to a set of parabolic points.

A standard example of a convergence group is the action induced on the boundary, ∂X , by any properly discontinuous action of a group, Γ , on a proper (complete locally compact) hyperbolic space, X . If the action on ∂X is geometrically finite, we say that Γ is *hyperbolic relative* to the set, \mathcal{G} , of maximal parabolic subgroups of Γ . (In [Bo4], we impose the additional requirement that each element of \mathcal{G} be finitely generated, but this need not concern us here.) It is necessarily the case that \mathcal{G} is cofinite and connected (and hence fine) as a Γ -set. Relatively hyperbolic groups were introduced by Gromov [Gr]. It

turns out that they can be characterised dynamically as geometrically finite convergence groups [Y].

We note that, in the case where M is totally disconnected, we can express the convergence property in terms of the boolean algebra, $\mathcal{B} = \mathcal{B}(M)$ of clopen subsets of M . By a *ternary partition* of \mathcal{B} , we mean a triple of pairwise disjoint non-zero elements, $A, B, C \in \mathcal{B}$ such that $A + B + C = 1$. (In other words, $M = A \sqcup B \sqcup C$.) Note that $A \times B \times C$ is a compact subset of the space of distinct triples of M .

In general, if Γ acts by isomorphism on a boolean algebra, \mathcal{B} , then we say that the action is a *convergence action* if, for any two ternary partitions, (a_1, a_2, a_3) and (b_1, b_2, b_3) of \mathcal{B} , then $\{g \in \Gamma \mid b_1 \wedge ga_1 \neq 0, b_2 \wedge ga_2 \neq 0, b_3 \wedge ga_3 \neq 0\}$ is finite. To see that this agrees with the notion already defined for M , note that compact subset of the set of distinct triples of M can be finitely covered by sets of the form $A \times B \times C$, where A, B, C is a ternary partition of M .

We finally note that if M is an inverse limit of compact spaces, $(M_n)_{n \in \mathbf{N}}$, each admitting a Γ -action that commutes with the inverse limit system, then the action of Γ on M is a convergence action if and only if the action on each M_n is a convergence action.

9. Ends of graphs.

Let K be a connected graph, with vertex set $V = V(K)$, and edge set $E(K)$. Let $V_0 = V_0(K)$ be the set of vertices of finite degree. As in the case of simplicial trees, if $W \subseteq V(K)$ we write $I(W)$ for the set of edges with precisely one endpoint in W . We write $\mathcal{B} = \mathcal{B}(K)$ for the boolean algebra of K -slices (i.e. those $W \subseteq V(K)$ for which $I(W)$ is finite). We write $\Xi(K) = \Xi(\mathcal{B})$ for the Stone dual of \mathcal{B} .

There is a natural map, $\xi : V \rightarrow \Xi(K)$, defined by sending $x \in V$ to the ultrafilter of elements of \mathcal{B} containing x . Note that $\xi|_{V_0}$ is injective, and every point of $\xi(V_0)$ is isolated in $\Xi(K)$. We write $\Xi_0(K) = \Xi(K) \setminus \xi(V_0)$. This is a closed subset of $\Xi(K)$. Note that if K is locally finite, then $\Xi_0(K)$ is the space of ends of K . If T is a simplicial tree, then we can identify $\Xi_0(T)$ with the ideal boundary, $\Delta_0 T$, of T , as defined in Section 5.

Another way to define $\Xi_0(K)$ is as follows. Let \mathcal{I} be the ideal of \mathcal{B} consisting of finite subsets of V , and let $\mathcal{C}(K)$ be the quotient \mathcal{B}/\mathcal{I} . There is an inclusion of $\Xi(\mathcal{C})$ in $\Xi(\mathcal{B})$ whose image is precisely $\Xi_0(K)$.

We shall say that K is *one-ended* if $\Xi_0(K)$ consists of a single point, i.e. no finite set of edges separates K into two or more infinite subgraphs.

Lemma 9.1 : *If a group Γ acts on a connected graph, K , with finite edge stabilisers, then the induced action of Γ on $\Xi(K)$ is a convergence action.*

Proof : Note that if $F \subseteq E(K)$ is any finite set of edges, then $\{g \in \Gamma \mid F \cap gF \neq \emptyset\}$ is finite.

Suppose that (A_1, A_2, A_3) and (B_1, B_2, B_3) are ternary partitions of \mathcal{B} . Let $I = I(A_1) \cup I(A_2) \cup I(A_3)$ and $J = I(B_1) \cup I(B_2) \cup I(B_3)$. Let X and Y be connected finite subgraphs of K containing I and J respectively.

Suppose that $E(X) \cap E(Y) = \emptyset$. Now, $I \cap E(Y) = \emptyset$ and so all the vertices of Y must lie in the same element of $\{A_1, A_2, A_3\}$, say $V(Y) \subseteq A_i$. Similarly, $V(X) \subseteq B_j$ for some j . We claim that $V(K) \subseteq A_i \cup B_j$. For suppose $x \in V(K) \setminus (A_i \cup B_j)$. Let α be a shortest path connecting x to $A_i \cup B_j$. Let y, z be, respectively, the last and last but one vertices of α . Without loss of generality, $y \in A_i$. Now $z \notin A_i$, and so the edge connecting y and z must lie in $I \subseteq E(X)$. It follows that $z \in V(X) \subseteq B_j$, giving a contradiction, and hence proving the claim. Note that if $k \neq i, j$, then $A_k \cap A_i = B_k \cap B_j = \emptyset$, so that $A_k \cap B_k = \emptyset$.

Now, if (A_1, A_2, A_3) and (B_1, B_2, B_3) are ternary partitions of \mathcal{B} , and $g \in \Gamma$ with $B_i \cap gA_i \neq \emptyset$ for each $i = 1, 2, 3$, then, by the previous paragraph, we see that $E(Y) \cap gE(X) \neq \emptyset$. But the set of such g is finite. This shows that Γ is a convergence group as claimed. \diamond

In particular, we deduce the well-known fact that any finitely generated group acts a convergence group on its space of ends. (Take K to be any Cayley graph of Γ .)

Suppose that $f : K \rightarrow L$ is a contraction of a connected graph, L (i.e. such that the preimage of each edge of L is an edge of K , and the preimage of every vertex of L is a connected subgraph of K). There is a natural inclusion of $\mathcal{B}(L)$ into $\mathcal{B}(K)$. If the preimage of every finite-degree vertex of L is finite, then this descends to an injection from $\mathcal{C}(L)$ to $\mathcal{C}(K)$. If, in addition, the preimage of every infinite degree vertex of L is one-ended, then this is an isomorphism, so that $\Xi_0(K)$ and $\Xi_0(L)$ are canonically homeomorphic.

Suppose that Γ acts on a simplicial tree, T , with finite edge stabilisers, then the induced action on $\Xi(T) \equiv \Delta T$ is a convergence action. In fact, if the action of Γ on T is cofinite, then the action on ΔT is geometrically finite, with the infinite vertex groups precisely the maximal parabolic subgroups. Thus, Γ is hyperbolic relative to the infinite vertex groups, and its boundary can be identified with $\Xi_0(T) \equiv \Delta_0 T$ (see [Bo4]). Moreover, if Γ is finitely generated, and each infinite vertex group is finitely generated and one-ended, then from the discussion of the previous paragraph, we see that $\Delta_0 T$ is canonically homeomorphic to the space of ends of Γ .

10. A variation on the Bergman norm.

In this section, we give a generalisation of Bergman's result [Ber].

Let V be a set, and let $\mathcal{P}(V)$ be its power set. In this section, it will be more convenient to work with the set, $\mathcal{R} = \mathcal{R}(V)$ of binary partitions of V , i.e. pairs $\{A, B\} \subseteq \mathcal{P}(V)$ such that $V = A \sqcup B$. We say that $\{A, B\}$ is *non-trivial* if $A, B \neq \emptyset$. Note that \mathcal{R} has the structure of an abelian group, with addition defined by $\{A, A^*\} + \{B, B^*\} = \{(A \cap B) \cup (A^* \cap B^*), (A \cap B^*) \cup (A^* \cap B)\}$. The same structure can be obtained by quotienting the additive group of the boolean ring $\mathcal{P}(V)$ by the subgroup $\{0, 1\}$.

Let Ψ be the set of unordered pairs of V . If $\pi \in \Psi$ and $P = \{A, A^*\}$, we say that π *crosses* P if $\pi \cap A \neq \emptyset$ and $\pi \cap A^* \neq \emptyset$. Let $\Psi(P)$ be the set of $\pi \in \Psi$ such that π crosses P .

Suppose that $(\Psi_n)_{n \in \mathbf{N}}$ are subsets of Ψ with $\Psi = \bigcup_{n \in \mathbf{N}} \Psi_n$. Given $P \in \mathcal{R}$, we write $\Psi_n(P) = \Psi_n \cap \Psi(P)$.

Definition : We say that a partition, $P \in \mathcal{R}$, is an (*undirected*) *slice* (with respect to $(\Psi_n)_n$) if $\Psi_n(P)$ is finite for all n .

Let $\mathbf{N}^{\mathbf{N}}$ be the set of infinite sequences of natural numbers. This has the structure of an ordered abelian group with the lexicographic order. We define a map, $\mu : \mathcal{S} \rightarrow \mathbf{N}^{\mathbf{N}}$ by setting $\mu(P) = (\mu_n(P))_n$, where $\mu_n(P) = |\Psi_n(P)|$.

Now suppose that Γ acts on V preserving each Ψ_n , with V/Γ and Ψ_n/Γ finite for all n . Thus, Ψ_n determines a Γ -graph, K_n , with vertex set V and edge set Ψ_n .

Recall that V is *connected* if it is the vertex set of some connected Γ -graph with finite quotient (i.e. some finite union of the K_n is connected). We show:

Theorem 10.1 : *Let V be a connected cofinite Γ -set. Then $\mu(\mathcal{S}) \subseteq \mathbf{N}^{\mathbf{N}}$ is well-ordered, and the map μ is finite-to-one modulo the action of Γ .*

In the above, we could for example take $(\Psi_n)_n$ to be an enumeration of the Γ -orbits of Ψ , if V is countable. However, it will be more convenient for the proof to assume that sets Ψ_n are increasing, i.e. $\Psi_m \subseteq \Psi_n$ whenever $m \leq n$. There is no loss of generality in doing this, for if we set $\Psi'_n = \bigcup_{m \leq n} \Psi_m$, then if $(P^\alpha)_{\alpha \in \mathbf{N}}$ is an infinite sequence of slices which is non-increasing with respect to the order determined by $(\Psi_n)_n$, then some subsequence will be non-increasing with respect to the order determined by $(\Psi'_n)_n$. We may as well also assume that K_0 , and hence every K_n , is connected.

It is easily verified that if $n \in \mathbf{N}$, then $\mu_n(P + Q) \leq \mu_n(P) + \mu_n(Q)$, with equality precisely if no element of Ψ_n crosses both P and Q .

We now set about the proof of Theorem 10.1.

For the moment, we forget about the group, Γ . Let K be a connected graph, with vertex set $V(K) = V$ and edge set $E(K)$. A finite subset, $I \subseteq E(K)$ is *separating* if $K \setminus I$ is disconnected.

Definition : A *cut* is a finite non-empty subset, $I \subseteq E(K)$, such that every circuit (or equivalently, cycle) in K contains an even number of edges of I .

Thus, every cut is separating. A cut is *minimal* if it contains no proper subcut. We similarly define minimality for separating sets. It is easily verified that if $I \subseteq E(K)$ is finite, then the following three conditions are equivalent: I is a minimal cut, I is a minimal separating set, or $K \setminus I$ has precisely two connected components.

Note that each cut $I \subseteq K$ determines a non-trivial partition, $P(K, I) = \{A, A^*\}$, where each path from A to A^* contains an odd number of edges of I . Conversely, each non-trivial partition $P = \{A, A^*\}$ determines the subset, $I = I(K, P)$, of edges which cross from A to A^* . If I is finite, then it is a cut, and we say that P is a *K-slice*. There is thus a natural bijection between cuts and *K*-slices.

Given two cuts I and J , we write $I + J$ for their symmetric difference. This is also a cut, and the operation agrees with that already defined on the set of partitions. If $J \subseteq I$ is a subcut, then $I \setminus J = I + J$ is also a subcut, and $I = J + (I \setminus J)$. We shall measure the “size” of a cut by its cardinality.

We note:

Lemma 10.2 : *Suppose $e \in E(K)$ and $n \in \mathbf{N}$. There are finitely many minimal cuts of size n containing e .*

Proof : Suppose, for contradiction, that the set, \mathcal{I} , of such minimal cuts is infinite. Choose $I \subseteq E(K)$ maximal such that I is contained in infinitely many elements of \mathcal{I} . Let $\mathcal{J} = \{J \in \mathcal{I} \mid I \subseteq J\}$. Now, I cannot separate K (otherwise I would be a minimal cut, and $\mathcal{J} = \{I\}$). Let γ be a path in $K \setminus I$ connecting the endpoints of e . Now each element of \mathcal{J} contains some edge of γ , and so some infinite subset of elements of \mathcal{J} all contain the same edge, say f , of γ . But now $I \cup \{f\}$ is contained in infinitely many elements of \mathcal{J} and hence of \mathcal{I} , contradicting the maximality of I . \diamond

Definition : We say that a cut, I , is *blocklike* if every pair of elements of I lie in a minimal subset of I .

Clearly, in Lemma 10.2, one can replace “minimal cut” by “blocklike cut”. Every cut can be uniquely decomposed into maximal blocklike cuts. One way to see this is as follows.

Suppose Υ is a finite connected graph. A *block* of Υ is a maximal 2-vertex-connected subgraph. Two blocks intersect, if at all, in a single vertex. If $e, f \in E(\Upsilon)$ are distinct, then e, f lie in the same block if and only if they lie in some circuit in Υ , and if and only if they lie in some minimal separating set. Note that Υ is bipartite (i.e. $E(\Upsilon)$ is a cut) if and only if each of its blocks is bipartite.

Suppose that K is any connected graph, and $I \subseteq E(K)$ is a cut. Let $\Upsilon = \Upsilon(K, I)$ be the graph obtained by collapsing each connected component of $K \setminus I$ to a point. Thus, Υ is a finite connected bipartite graph. There is a canonical surjective map $\phi : K \rightarrow \Upsilon$. The preimage of every connected subgraph is connected.

Now it is easily checked that the preimage of every cut in Υ is a subcut of I . Moreover, every subcut, J , arises in this way: $J = \phi^{-1}(\phi J)$. We also see easily that J is minimal if and only if ϕJ is minimal (since $K \setminus J$ has the same number of components as $\phi(K \setminus J) = K \setminus \phi J$). Thus it follows that J is blocklike if and only if ϕJ is blocklike. We can thus decompose I canonically by taking the preimages of (the edge sets of) blocks of Υ .

If P is a K -slice, then we shall write $\Upsilon(K, P) = \Upsilon(K, I(K, P))$. The canonical decomposition of $I(K, P)$ gives us a canonical decomposition of P (depending on K).

Now suppose that K' is another graph on the same vertex set, V , with $K \subseteq K'$. Let $\phi : K \rightarrow \Upsilon = \Upsilon(K, P)$ and $\phi' : K' \rightarrow \Upsilon' = \Upsilon(K', P)$ be collapsing maps described above. We can obtain Υ' from Υ by identifying certain vertices of the same colour and/or adding edges between vertices of different colours. There is thus a natural map, $\psi : \Upsilon \rightarrow \Upsilon'$, such that $\psi \circ \phi = \phi' \circ \iota$, where ι is the inclusion of K in K' . If we measure the complexity, $c(\Upsilon)$, of a finite graph, Υ , by the number of edges in the complementary graph, i.e. $c(\Upsilon) = \frac{1}{2}|V(\Upsilon)|(|V(\Upsilon)| - 1) - E(\Upsilon)$, then we see that the map ψ cannot increase complexity. We have $c(\Upsilon') = c(\Upsilon)$ if and only if ψ is an isomorphism. In this case, the canonical decomposition of P with respect to K is identical to its canonical decomposition with respect to K' . Note also that in general, $|V(\Upsilon')| \leq |V(\Upsilon)| \leq |I(K, P)| + 1$.

Suppose now that V is a set with $(\Psi_n)_n$ an increasing collection of sets of pairs of V as described above. We thus get an increasing collection of graphs, $(K_n)_n$. We suppose that K_0 (and hence each of these graphs) is connected. If P is any slice, then the sequence of complexities, $c(\Upsilon(K_n, P))$, is non-increasing in n . Moreover, the cut $I(K_n, P)$ will be minimal for all sufficiently large n (depending on P). In this case, $\Upsilon(K_n, P)$ consists of a single edge (so its complexity is 0).

Suppose that $\mathcal{T} \subseteq \mathcal{S}$ is an infinite set of slices with $\mu_0(P) = |I(K_0, P)|$ bounded, by ν , say for $P \in \mathcal{T}$. Thus, for each $P \in \mathcal{T}$ and $n \in \mathbf{N}$, we have $|V(\Upsilon(K_n, P))| \leq |V(\Upsilon(K_0, P))| \leq \mu_0(P) + 1 \leq \nu + 1$, so there are only finitely many possibilities for the graph $\Upsilon(K_n, P)$. Given a graph Υ , and $n \in \mathbf{N}$, let $\mathcal{T}(n, \Upsilon) = \{P \in \mathcal{T} \mid \Upsilon(K_n, P) = \Upsilon\}$. For any fixed n , this partitions \mathcal{T} into finitely many subsets. We can therefore choose $\Upsilon = \Upsilon_0$ minimal complexity, $c(\Upsilon_0)$, with the property that for some $n = n_0$, say, the set $\mathcal{T}_0 = \mathcal{T}(n_0, \Upsilon_0)$ is infinite. For any $n \geq n_0$, $\Upsilon(K_n, P) = \Upsilon_0$ for all but finitely many $P \in \mathcal{T}_0$. (Since $c(\Upsilon(K_n, P)) \leq c(\Upsilon(K_{n_0}, P)) = c(\Upsilon_0)$, and if we had strict inequality for infinitely many P , then we would contradict the minimality of $c(\Upsilon_0)$.)

We now introduce the action of Γ . Note that by Lemma 10.2, K_n has only finitely many blocklike cuts of size s modulo Γ , for any $n, s \in \mathbf{N}$.

Suppose that $(P^\alpha)_{\alpha \in \mathbf{N}}$ is a sequence of cuts all lying in distinct Γ -orbits. Suppose that $\mu(P^\alpha)$ is non-increasing. We want to derive a contradiction.

First note that applying the construction of the first paragraph, with $\mathcal{T} = \{P^\alpha \mid \alpha \in \mathbf{N}\}$, we can assume that, after passing to a subsequence, $(P^\alpha)_\alpha$ has the following property. There exist $n_0 \in \mathbf{N}$ and a finite bipartite graph, Υ_0 , such that if $n_0 \leq n \leq \alpha$, then $\Upsilon(K_n, P^\alpha) = \Upsilon_0$. For notational convenience (replacing Ψ_0 by Ψ_{n_0}), we can assume that $n_0 = 0$.

Now for each α , we decompose the cut $I(K_0, P^\alpha)$ into its maximal blocklike subcuts, $I(K_0, P^\alpha) = I_1^\alpha + \cdots + I_p^\alpha$. Note that p is the number of blocks of Υ_0 , and therefore constant. After passing to a subsequence, we can assume that each cut, I_i^α , is the Γ -image of some fixed cut J_i , independent of α . Let $P_i^\alpha = P(K_0, I_i^\alpha)$ and $Q_i = P(K_0, J_i)$. Thus, $P^\alpha = P_1^\alpha + \cdots + P_p^\alpha$. Note that since the P^α all lie in different Γ -orbits, we have $p \geq 2$.

Now we know that, for each α , $I(K_n, P^\alpha)$ is a minimal cut for all sufficiently large n . Thus, there is some n (depending on α) such that the cuts $\{I(K_n, P_i^\alpha)\}_{i=1}^p$ are not disjoint. Let $n(\alpha)$ be the smallest such n , and let $m = \min\{n(\alpha) \mid \alpha \in \mathbf{N}\}$. (Thus $m > 0$.) If $n < m$, then for all $\beta \in \mathbf{N}$, the cuts $\{I(K_n, P_i^\alpha)\}_{i=1}^p$ are disjoint. Thus, for $n < m$, we have $\mu_n(P^\beta) = \mu_n(P_1^\beta) + \cdots + \mu_n(P_p^\beta) = \mu_n(Q_1) + \cdots + \mu_n(Q_p)$, which is independent of β .

Now for some α , the cuts $\{I(K_m, P_i^\alpha)\}_{i=1}^p$ are no longer disjoint, so this time we get strict inequality: $\mu_m(P^\alpha) < \mu_m(Q_1) + \cdots + \mu_m(Q_p)$. However, for $\beta \geq m$, we have $\Upsilon(K_m, P^\beta) = \Upsilon_0 = \Upsilon(K_0, P^\beta)$, and the natural map from $\Upsilon(K_0, P^\beta)$ to $\Upsilon(K_m, P^\beta)$ is an isomorphism. Thus, the canonical decompositions of P^β with respect to K_0 and K_m are identical. In particular, the cuts $\{I(K_m, P_i^\beta)\}_{i=1}^p$ are disjoint, and again, we have equality: $\mu_m(P^\beta) = \mu_m(Q_1) + \cdots + \mu_m(Q_p)$. Thus $\mu_m(P^\beta) > \mu_m(P^\alpha)$. Taking $\beta > \alpha$, we derive a contradiction to the assumption that μ is non-increasing.

This proves Theorem 10.1.

We finish this section with one corollary of this result.

Suppose that V is a Γ -set. By a (*directed*) *slice* of V , we mean a subset, $A \subseteq V$ with the property that only finitely many Γ -images of any pair in V can meet both A and A^* . Note that the set of slices forms a boolean subalgebra of $\mathcal{P}(V)$. By an *undirected slice*, we mean a pair, $\{A, A^*\}$, where A is a slice.

Suppose that V is countable. Let $(\Psi_n)_n$ be an enumeration of the Γ -orbits of the set of pairs, Ψ . The notion of a slice as defined above thus agrees with that defined earlier. Let \mathcal{B} be the boolean algebra of slices of V . Suppose $A, B \in \mathcal{B}$ are disjoint. Let $\sigma_n(A, B)$ be the number of pairs in Ψ_n with one element in A and one element in B . Let $\sigma(A, B) = (\sigma_n(A, B))_n \in \mathbf{N}^{\mathbf{N}}$. Thus, $\mu(A) = \sigma(A, A^*)$. We are thus in the situation described in Section 7. Applying Lemma 7.1, we deduce:

Corollary 10.3 : *Suppose Γ is a countable group, and V is a connected cofinite Γ -set. Let \mathcal{B} be the boolean algebra of slices of V (as defined above). Then, any Γ -invariant subalgebra of \mathcal{B} has a Γ -invariant nested set of generators.*

Proof : The construction of Section 7 was canonical, and hence Γ -invariant. ◇

11. A finiteness result for boolean algebras related to simplicial trees.

In this section, using a result of [DiD], we give a proof of the following:

Lemma 11.1 : *Suppose that Γ is a countable group, and that T is a cofinite simplicial Γ -tree with finite edge stabilisers. Suppose that \mathcal{A} is any Γ -invariant boolean subalgebra of the boolean algebra, $\mathcal{B}(T)$. Then \mathcal{A} is finitely generated as a Γ -boolean algebra.*

Here, $\mathcal{B}(T)$ is the boolean algebra of T -slices, as defined in Section 5. To say that \mathcal{A} is finitely generated as a Γ -boolean algebra means that it has a generating set which is a finite union of Γ -orbits.

Now, $V = V(T)$ has finite pair stabilisers, and so Corollary 10.3 gives us a Γ -invariant nested set, \mathcal{E} , of generators of \mathcal{A} . If we know already that \mathcal{A} is finitely generated, then we can assume that \mathcal{E} is cofinite. It follows that if $e \in E(T)$, the set $\{A \in \mathcal{E} \mid e \in I(A)\}$ is finite. We see that if $x, y \in V$, then $\{A \in \mathcal{E} \mid x \in A, y \notin A\}$ is finite. In particular, if $A, B \in \mathcal{E}$, then $\{C \in \mathcal{E} \mid A \subseteq C \subseteq B\}$ is finite. In other words, \mathcal{E} satisfies the finite interval condition. We can thus identify \mathcal{E} with the directed edge set, $\vec{E}(S)$, of a simplicial tree, S . Note that Γ acts on S with finite edge stabilisers.

If $x \in V(T)$, then $\{A \in \mathcal{E} \mid x \in A\}$ is a flow on \mathcal{E} , and hence determines a flow on S . Moreover, from the observation of the previous paragraph, there can be no infinite decreasing sequence in the flow (i.e. any strictly decreasing sequence, $A_1 \supset A_2 \supset A_3 \supset \dots$ with $A_i \in \mathcal{E}$ must terminate). The flow thus arises from a unique vertex of S . This therefore defines a Γ -equivariant map $\phi : V(T) \rightarrow V(S)$.

Suppose $y \in V(S)$. Let $\mathcal{E}(y) \subseteq \mathcal{E}$ be the set of elements of \mathcal{E} which correspond to directed edges with tail at y . Let $\mathcal{E}^*(y) = \{A^* \mid A \in \mathcal{E}(y)\}$. If $x \in \bigcap \mathcal{E}^*(y) \subseteq V(T)$, then

each edge corresponding to an element of $\mathcal{E}^*(y)$ must point towards $\phi(x) \in V(S)$. It follows that $\phi(x) = y$. From this observation, we see that if $y \notin \phi(V(T))$, then $\bigcap \mathcal{E}^*(y) = \emptyset$, so that $\bigcup \mathcal{E}(y) = V(S)$. If it also happens that $\mathcal{E}(y)$ is finite, then $\bigcup \mathcal{E}(y)$ is the sum (i.e. symmetric difference) of the elements of $\mathcal{E}(y)$. We deduce:

Lemma 11.2 : *If $y \in V(S) \setminus \phi(V(T))$ has finite degree in S , then $\sum_{i=1}^n A_i = 1$ in $\mathcal{B}(T)$, where $A_1 \dots A_n$ are the elements of \mathcal{E} which correspond to those edges with tails at y . \diamond*

To deduce Lemma 11.1, one can use an accessibility result of Dicks and Dunwoody. Following [DiD], we say that an edge of a Γ -tree is *compressible* if its endpoints lie in distinct Γ -orbits and if its stabiliser is equal to an incident vertex stabiliser. (Such an edge can be “compressed” in the corresponding graph of groups to give a smaller graph.) A Γ -tree is *incompressible* if it has no compressible edges.

The following is shown in [DiD] (III 7.5 page 92):

Proposition 11.3 : *Let Γ be a group. Suppose that S, T are cofinite simplicial Γ -trees with finite edge stabilisers, and that S is incompressible. If there is a Γ -equivariant map from $V(T)$ to $V(S)$, then $|V(S)/\Gamma| \leq |V(T)/\Gamma| + |E(T)/\Gamma|$. \diamond*

(In [DiD] is assumed also that T is incompressible. However, it is clear that one can always collapse T to give another tree, T' , with this property, and with $V(T')$ isomorphic as a Γ -set to a Γ -invariant subset of $V(T)$. This process can only decrease $|V(T)/\Gamma|$ and $|E(T)/\Gamma|$.)

Alternatively, we can use (the argument of) the “elliptic” case of the accessibility result of Bestvina and Feighn [BesF]. This gives a slightly different result:

Proposition 11.4 : *Suppose that Γ is a group and that S, T are cofinite simplicial Γ -trees without edge inversions. Suppose that S is incompressible, and that every edge stabiliser of S fixes a vertex of T . If there is a Γ -equivariant map from $V(T)$ to $V(S)$, then $|V(S)/\Gamma| \leq \max\{1, 5|E(T)/\Gamma|\}$.*

In particular, this applies to the case of finite edge-stabilisers. We shall sketch a proof below, which is condensed out of the relevant part of [BesF]. Our direct use of Grushko’s Theorem bypasses the use of folding sequences. We begin with some preliminary remarks.

Let t be a graph of groups, and let $\Gamma = \pi_1(t)$ be its fundamental group. Thus Γ acts on the corresponding Bass-Serre tree, T , with quotient the underlying graph $|t| = T/\Gamma$. Given $v \in V(t)$ or $e \in E(t)$, we write $\Gamma(v) = \Gamma_t(v)$ or $\Gamma(e) = \Gamma_t(e)$ for the corresponding vertex or edge groups. A subgroup of Γ is *elliptic* if it is conjugate into a vertex group. Let t_0 be the graph of groups with the same underlying graph, and with all vertex and edge groups trivial. Thus, $\pi_1(t_0) = \pi_1(|t|)$ is free of rank $\beta(t) = |E(t)| - |V(t)| + 1$. Moreover, there is a natural epimorphism from Γ to $\pi_1(t_0)$ whose kernel contains $\langle\langle \bigcup_{v \in V(t)} \Gamma(v) \rangle\rangle$, where $\langle\langle \cdot \rangle\rangle$ denotes normal closure. In particular, if Γ is the normal closure of some vertex group, $\Gamma(v)$, then $|t|$ is a tree. We see easily that if every other vertex group is Γ -conjugate into $\Gamma(v)$, then $\Gamma = \Gamma(v)$. Indeed, if every vertex group is conjugate into a subgroup, $G \leq \Gamma(v)$, then $\Gamma = G = \Gamma(v)$.

Proof of Proposition 11.4 : After subdividing the edges of T , we obtain a Γ -tree, T' , with $V(T) \subseteq V(T')$, and an equivariant morphism $\phi : T' \rightarrow S$ (which sends each edge of T' to an edge or vertex of S). This descends to a graph-of-groups morphism $\phi : t' \rightarrow s$ (where $|t'| = T'/\Gamma$ and $|s| = S/\Gamma$) which induces the identity map on $\Gamma = \pi_1(t') = \pi_1(s)$. We make a series of observations.

Claim 1 : If $v \in V(s) \setminus \phi(V(t))$ and $G \subseteq \Gamma(v)$ is T -elliptic, then G is $\Gamma(v)$ -conjugate into an incident edge group. This is easily seen by considering the arc connecting a lift of v to $V(S)$ (fixed by $\Gamma(v)$) to the ϕ -image of a fixed point of G in $V(T)$.

In fact, the same argument shows that if $v, w \in V(s) \setminus \phi(V(t))$ are the endpoints of an edge $e \in E(s)$, then any T -elliptic subgroup, G , of $\langle \Gamma(v) \cup \Gamma(w) \rangle$ is conjugate into an incident edge group adjacent to (but different from) e . In particular, from the T -ellipticity hypothesis of the proposition, this applies to $G = \Gamma(e) = \Gamma(v) \cap \Gamma(w)$.

We say that a vertex, $v \in V(s)$ is *dead* if $v \notin \phi(V(t))$ and if $\Gamma(v)$ is the normal closure of the incident edge groups. Otherwise it is *live*. We thus decompose $V(s) = V_D(s) \sqcup V_L(s)$ into dead and live vertices.

Claim 2 : $|V_L(s) \setminus \phi(V(t))| \leq \beta(t) - \beta(s)$. To see this, let \bar{s} be the graph of groups with underlying graph $|s|$ obtained by collapsing each edge group in $E(s)$ and each vertex group in $\phi(V(t))$ to the trivial group, and by collapsing each remaining vertex group, $\Gamma_{\bar{s}}(v)$, to the quotient of $\Gamma_s(v)$ by the normal closure of its incident edge groups. Thus if $v \in V(\bar{s})$, then $\Gamma_{\bar{s}}(v)$ is non-trivial if and only if $v \in V_L(s) \setminus \phi(V(t))$. It follows that $\pi_1(\bar{s})$ has at least $|V_L(s) \setminus \phi(V(t))| + \beta(s)$ non-trivial free factors. Now there is a natural epimorphism from $\pi_1(t_0)$ to $\pi_1(\bar{s})$. The former group is free of rank $\beta(t)$, and so Claim 2 follows by Grushko's Theorem.

Clearly, $|\phi(V(t))| \leq |V(t)|$ and so $|V_L(s)| \leq |E(t)| - \beta(s) + 1$.

Claim 3 : Suppose $v \in V_D(s)$ and G is an incident edge group. Suppose that every other incident edge group is $\Gamma(v)$ -conjugate into G . Then $\Gamma(v) = G$. To see this, consider the action of $\Gamma(v)$ on a minimal $\Gamma(v)$ -invariant subtree of T . Let r be the corresponding graph of groups. Now by the T -ellipticity hypothesis, G is conjugate into a vertex group, $\Gamma_r(w)$, where $w \in V(r)$. Moreover, if H is any other vertex group of r , then again H is T -elliptic, and hence, by Claim 1, is $\Gamma(v)$ -conjugate into an incident edge group to v in s . Thus H is $\Gamma(v)$ -conjugate into G . But now $\Gamma(v) = \langle\langle G \rangle\rangle = \langle\langle \Gamma(w) \rangle\rangle$, and so it follows by the discussion before the proof that $\Gamma(v) = G$. The claim follows.

As an immediate corollary, we see (by the incompressibility of s) that any such vertex must have degree at least 3 in s . In particular, all terminal vertices of s are live.

Claim 4 : We cannot have two adjacent dead vertices of degree 2 in s . For suppose to the contrary that $e \in E(s)$ has endpoints $v, w \in V_D(s)$, both of degree 2. From the remark after Claim 1, we see that, without loss of generality, $\Gamma(e)$ is conjugate into $\Gamma(f)$, where $f \in E(s)$ is the other edge incident on v . But now, from Claim 3, we derive the contradiction that v has degree at least 3.

We now have enough information to bound $|V(s)|$ in terms of the complexity of $|t|$. First note that since every terminal vertex of s is live (Claim 3), the number of such vertices is bounded by Claim 2. Moreover (Claim 2), we have $\beta(s) \leq \beta(t)$. This places a bound on the number of vertices of s of degree at least 3. Finally, Claim 4 together with the bound on live vertices places a bound on the number of vertices of degree 2. More careful

bookkeeping shows that if $|s|$ is not a point, then $|V(s)| \leq 5|E(t)| - |V(t)|$. Proposition 11.4 now follows. \diamond

Now since, $\beta(s) \leq \beta(t)$, we also get a bound on $|E(s)| = |E(S)/\Gamma|$. In fact, to obtain such a bound, we can weaken the hypotheses slightly:

Corollary 11.5 : *Let Γ be a group and that S, T be cofinite Γ -trees with finite edge stabilisers. Suppose that $\phi : V(T) \rightarrow V(S)$ is a Γ -equivariant map, and that each compressible edge of S is incident on some element of $\phi(V(T))$. Then there is a bound on $|E(S)/\Gamma|$ in terms of $|E(T)/\Gamma|$.*

Proof : To see this, note that we can obtain a Γ -tree, S' , by collapsing down a union of trees, F , in S/Γ , consisting of a union of compressible edges. Moreover, we can assume that if $x \in \phi(V(T))$, then the image of x in S/Γ is incident to at most one edge of F that is terminal S/Γ . (Since collapsing such an edge will be sufficient to render all the other edges incident on x incompressible.) Now, Proposition 11.3 or 11.4 gives a bound on the complexity of S'/Γ . This, in turn, gives a bound on the number of edges of F , and hence a bound on the complexity of S/Γ as claimed. \diamond

We can now set about the proof of Lemma 11.1. Suppose that Γ , T and \mathcal{A} are as in the hypotheses. By Corollary 10.3, there is a Γ -invariant nested set, \mathcal{E} , of generators of \mathcal{A} . Now, if \mathcal{A} is not finitely generated as a Γ -boolean algebra, then we can find an infinite sequence, $(A_n)_{n \in \mathbf{N}}$, of elements of \mathcal{E} such that A_n does not lie in the Γ -boolean algebra generated by $\{A_i \mid i < n\}$. Let $\Gamma(A_n)$ be the stabiliser of A_n . After passing to a subsequence, we can assume that $|\Gamma(A_n)|$ is non-decreasing in n . Let \mathcal{E}_n be the union of the Γ -orbits of $\{A_i, A_i^* \mid i \leq n\}$. Given $A \in \mathcal{E}_n$, we write $m(A) = i$ to mean that A or A^* lies in the Γ -orbit of A_i .

Now fix some n . As discussed after the statement of Lemma 11.1, we can identify \mathcal{E}_n with the directed edge set of a simplicial tree, S_n , and there is an equivariant map, $\phi : V(T) \rightarrow V(S_n)$. Note that $|E(S_n)/\Gamma| = n$. We claim that S_n satisfies the weakened hypotheses of Corollary 11.5. In fact, we show that any vertex whose stabiliser fixes an incident edge must lie in $\phi(V(T))$.

Suppose, to the contrary, that $y \in V(S_n) \setminus \phi(V(T))$ is incident on an edge $\vec{e} \in \vec{E}(S_n)$ with tail at y and with $\Gamma(e) = \Gamma(y)$. Note that $\Gamma(y)$ is finite, so that y has finite degree. Let $A \in \mathcal{E}_n$ be the element corresponding to \vec{e} , so that $\Gamma(A) = \Gamma(y)$. Let $\mathcal{E}_n(y)$ be the set of elements of \mathcal{E}_n which correspond to edges with tails at y . We can assume that \vec{e} is chosen so that $m(A)$ is maximal among those elements of $\mathcal{E}_n(y)$ with stabilisers equal to $\Gamma(y)$. Write $\mathcal{E}_n(y) = \{A, B_1, \dots, B_k\}$. By Lemma 11.2, we have that $A = 1 + \sum_{j=1}^k B_j$ in the boolean algebra $\mathcal{B}(T)$. In particular, A lies in the boolean algebra generated by $\{B_1, \dots, B_k\}$.

Now, for each j , $\Gamma(B_j) \leq \Gamma(A)$. Either $\Gamma(B_j) = \Gamma(A)$ so that, by the maximality of $m(A)$, we have $m(B_j) < m(A)$, or else $|\Gamma(B_j)| < |\Gamma(A)|$ so that, by the construction of the sequence $(A_i)_i$, we again have $m(B_j) < m(A)$. We therefore deduce that A lies in the Γ -boolean algebra generated by $\{A_i \mid i < m(A)\}$, contrary to the construction of $(A_i)_i$. This proves the claim.

Now applying Corollary 11.5, we get a bound on the complexity $n = |E(S_n)/\Gamma|$. But we could have chosen n arbitrarily large, thereby giving a contradiction.

This proves Lemma 11.1.

12. An application to 1-connected Γ -sets.

In this section, we apply Lemma 11.1 to show:

Lemma 12.1 : *Suppose Σ is a countable \mathbf{Z}_2 -acyclic simplicial 2-complex which is locally finite away from the vertex set $V(\Sigma)$. Suppose a group Γ acts on Σ with finite quotient and such that $\Gamma(x) \cap \Gamma(y)$ is finite for all distinct $x, y \in \Gamma$. Let $\mathcal{S}(V(\Sigma))$ be the boolean algebra of slices of $V(\Sigma)$. Then, any Γ -invariant subalgebra, \mathcal{A} , of $\mathcal{S}(V)$ is finitely generated as a Γ -boolean algebra.*

In particular, this applies to any countable fine \mathbf{Z}_2 -homologically 1-connected Γ -set, V , with finite pair stabilisers. As described in Section 3, such a set can be equivariantly embedded in a simplicial complex, Σ , which is locally finite away from V . Of course, V might be a proper subset of $V(\Sigma)$. In this case, we apply the result as stated to the subalgebra $\mathcal{A}' = \{W \subseteq V(\Sigma) \mid W \cap V \in \mathcal{A}\}$ of $\mathcal{S}(V(\Sigma))$. We deduce that \mathcal{A}' is finitely generated. Since $\mathcal{A} \subseteq \mathcal{S}(V)$ is a quotient of \mathcal{A}' , we deduce that \mathcal{A} is finitely generated.

We can reduce Lemma 12.1 to Lemma 11.1 using the machinery of patterns and tracks as in [Du2]. (The overall strategy for the proof is thus similar to that of the accessibility result of [BesF].) Let $\Sigma, \Gamma, V, \mathcal{A}$ be as in the hypotheses, and let K be the 1-skeleton of Σ . Recall that a *pattern*, t , on Σ is a compact subset of $\Sigma \setminus V(\Sigma)$ which meets each 1-simplex either in the empty set or a single point, and which meets each 2-simplex, σ , either in the empty set or in a single arc connecting two distinct faces of σ . It *represents* a subset, $A \subseteq V(\Sigma)$ if it meets precisely those edges of Σ which connect A to A^* . Every Σ -slice is represented by a pattern. A *track* is a connected pattern. Two disjoint tracks, s, t , are *parallel* if there is a closed subset of $\Sigma \setminus V(\Sigma)$ homeomorphic to $s \times [0, 1] \cong t \times [0, 1]$ with boundary $s \sqcup t$. If \mathcal{T} is a Γ -equivariant set of disjoint pairwise non-parallel tracks, then there is a bound on $|\mathcal{T}/\Gamma|$, (see [Du2]).

By Corollary 10.3 there is a Γ -invariant nested set of generators, \mathcal{E} , for \mathcal{A} . By a standard construction (cf. [Du2]), we can find a set of patterns $(t(A))_{A \in \mathcal{E}}$ such that $t(A)$ represents A , $t(A) = t(A^*)$, $t(A) \cap t(B) = \emptyset$ if $B \neq A, A^*$ and $t(gA) = gt(A)$ for all $g \in \Gamma$. By the observation of the previous paragraph, we can find a cofinite Γ -invariant set, \mathcal{T} , of tracks such that if $A \in \mathcal{E}$, then each connected component of the pattern $t(A)$ is parallel to some element of \mathcal{T} . Now \mathcal{T} determines a simplicial tree, T , whose edges are in bijective correspondence with \mathcal{T} , and whose vertices are in bijective correspondence with the connected components of $\Sigma \setminus \bigcup \mathcal{T}$. (It is here that we use the fact that Σ is \mathbf{Z}_2 -acyclic, so that every track separates Σ .) There is a canonical map $\phi : V(\Sigma) \rightarrow V(T)$, where two vertices of Σ get mapped to the same vertex of T if and only if they are not separated by any element of \mathcal{A} . Note that Γ acts with finite edge stabilisers on T .

Suppose that $A \in \mathcal{E}$. Let $I_A \subseteq E(T)$ be the set of edges of T that correspond to

the connected components of $t(A)$. Now I_A in turn determines an element, $B(A) \in \mathcal{B}(T)$, with the property that $I(B(A)) = I_A$ and such that $\phi(A) \subseteq B(A)$. It follows that $A = \phi^{-1}(B(A))$. Let $\mathcal{F} = \{B(A) \mid A \in \mathcal{E}\}$. Thus, \mathcal{F} is a nested subset of $\mathcal{B}(T)$. By Lemma 11.1, some cofinite Γ -invariant subset of \mathcal{F} is sufficient to generate the boolean algebra generated by \mathcal{F} . The corresponding elements of \mathcal{E} now generate \mathcal{A} . (This follows because $A = \phi^{-1}(B(A))$ for all $A \in \mathcal{E}$, so that any relation between the elements of \mathcal{F} also holds between the corresponding elements of \mathcal{E} .) We see that \mathcal{A} has a cofinite generating set, as required. This proves Lemma 12.1.

13. Convergence actions on Cantor sets in the finitely presented case.

In this section, we shall give proofs of the main results stated in Section 1.

Suppose the group Γ acts as a convergence group on the Cantor set, M . We write $\Pi \subseteq M$ for the set of parabolic points, and $\mathcal{B}(M)$ for the boolean algebra of clopen sets of M . We suppose that there is a cofinite Γ -invariant collection, \mathcal{G} , of parabolic subgroups of Γ such that Γ is almost finitely presented relative to \mathcal{G} . By Lemma 8.2, \mathcal{G} is fine as a Γ -set. Thus, Γ acts on a \mathbf{Z}_2 -acyclic simplicial 2-complex, Σ , with finite quotient and finite edge stabilisers and such that some Γ -invariant subset, $V \subseteq V(\Sigma)$ is isomorphic to \mathcal{G} as a Γ -set. Moreover, we can assume that Γ acts freely on $V(\Sigma) \setminus V$. We distinguish two cases.

Firstly, suppose that $\Pi \neq \emptyset$. We can find a Γ -equivariant map, $\phi : V(\Sigma) \rightarrow \Pi \subseteq M$. Now, if $A \in \mathcal{B}$ and $x, y \in M$, then applying Lemma 8.1, we see that $\{g \in \Gamma \mid gx \in A, gy \notin A\}$ is finite. It follows that $\phi^{-1}A$ is a slice of $V(\Sigma)$. Let $\mathcal{A} = \{\phi^{-1}A \mid A \in \mathcal{B}(M)\}$. Since $\phi^{-1}(V(\Sigma))$ is dense in M , we see that the map $[A \mapsto \phi^{-1}A]$ from $\mathcal{B}(M)$ to \mathcal{A} is an isomorphism of Γ -boolean algebras. It follows that $\Xi(\mathcal{A})$ is Γ -equivariantly homeomorphic to $\Xi(\mathcal{B}(M)) \cong M$.

Now, by Corollary 10.3 and Lemma 12.1, \mathcal{A} has a cofinite Γ -invariant nested set of generators, \mathcal{E} . Moreover, as in Section 12, we see that \mathcal{E} is discrete, and can thus be identified as the directed edge set of a cofinite simplicial tree, T , with finite edge stabilisers. We can identify $\Xi(\mathcal{A})$ as a closed subset of $\Xi(T)$. Moreover, there is a canonical map from $V(T)$ to $\Xi(T)$, and each point of $\Xi(T) \setminus \Xi(\mathcal{A})$ is the image of a vertex of finite degree under this map. Each such point is isolated. Since $\Xi(\mathcal{A}) \cong M$ is a Cantor set, we see in fact that $\Xi(T) \setminus \Xi(\mathcal{A})$ consists precisely of these vertices. Thus, we can identify $\Xi(\mathcal{A})$ with $\Delta_0 T$, as defined in Section 5.

Now, Γ is hyperbolic relative to the infinite vertex stabilisers and its boundary is precisely $\Delta_0 T$. We have thus shown that the boundary is Γ -equivariantly homeomorphic to M . This proves Theorem 1.4 in the case where $\Pi \neq \emptyset$.

The case where $\Pi = \emptyset$ is precisely Theorem 1.1, which we treat separately. This, in turn can be split into two cases. If every point of M is a conical limit point, then by the result of [Bo1], Γ is hyperbolic with boundary M . It follows that Γ is virtually free, and that its boundary can be identified with the space of ends of Γ .

If not, then let Π' be the Γ -orbit of a non conical limit point. Let Σ be a cofinite Γ -complex with $H_1(\Sigma, \mathbf{Z}_2) = 0$. We map $V(\Sigma)$ Γ -equivariantly to Π' . The argument now proceeds exactly as above to give us a cofinite Γ -tree, T , with finite edge stabilisers whose

boundary, $\Delta_0 T$ is Γ -equivariantly homeomorphic to M . In this case, every vertex of T has finite stabiliser (otherwise it would be parabolic in $\Delta_0 T$, and hence in M). It follows that $\Delta_0 T = \Delta T$ is the same as the space of ends of T . (In retrospect, we deduce that every point of M is a conical limit point, and so this case cannot in fact arise.)

This concludes the proofs of Theorems 1.4 and 1.1. We immediately deduce Theorem 1.3. Theorem 1.2 follows from Theorem 1.4 and the discussion at the end of Section 9.

14. Convergence actions arising from protrees.

Suppose Θ is a Γ -protree such that the stabiliser of each element is finite. We say that Θ is *locally discrete* if each cofinite subprotree is discrete. Note that Γ acts on $\Xi(\Theta)$ by homeomorphism.

Proposition 14.1 : *If Θ is locally discrete, then Γ acts on $\Xi(\Theta)$ as a convergence group.*

Proof : This is easily verified from the criterion described in Section 8. \diamond

If Θ is countable, then we can write it as an increasing union, $\Theta = \bigcup_{n=1}^{\infty} \Theta_n$, of cofinite discrete Γ -protrees, Θ_n . We can identify Θ_n as the directed edge set of a simplicial Γ -tree, T_n . We see that $\Xi(\Theta)$ is an equivariant inverse limit of the spaces $\Xi(\Theta_n) \cong \Delta T_n$, and that $\Xi_0(\Theta)$ is an equivariant inverse limit of the spaces $\Delta_0 T_n$. (This gives another proof of the fact that Γ acts as a convergence group on $\Xi(\Theta)$.) We see that the action of Γ on $\Xi(\Theta)$ is an inverse limit of geometrically finite actions.

Note that an inaccessible group admits a locally discrete action on a non-discrete protree. Dunwoody's example of a finitely generated inaccessible group [Du3] thus gives an example of a non geometrically finite action of such a group on a Cantor set.

We show that examples of this type are typical of convergence actions of (relatively) finitely generated groups on Cantor sets:

Theorem 14.2 : *Suppose that Γ acts as a minimal convergence group on a Cantor set, M , and that \mathcal{G} is a finite collection of parabolic subgroups. Suppose that Γ is finitely generated relative to \mathcal{G} . Then, Γ admits a locally discrete action on a countable protree, Θ , such that M is equivariantly homeomorphic to $\Xi(\Theta)$.*

In particular, the action of Γ on M is an inverse limit of geometrically finite actions.

The proof of Theorem 14.2 proceeds exactly as with that of Theorem 1.4 (and 1.1) as described in Section 13, except that, in this case, we have to make do with a (fine) connected cofinite Γ -graph, K , instead of the 2-complex, Σ . As before, M is equivariantly homeomorphic to $\Xi(\mathcal{A})$ where \mathcal{A} is a Γ -subalgebra of the boolean algebra of K -slices. Proposition 10.3 gives us an invariant nested generating set, \mathcal{E} of \mathcal{A} , which has the structure of a protree, Θ , as described in Section 6. We can canonically identify $\Xi(\Theta)$ as a closed subset of $\Xi(\mathcal{A}) \cong M$, whose complement consists of isolated points and is thus empty in this case. We have equivariantly identified $\Xi(\Theta)$ with M as required.

15. Other applications.

In this section, we sketch two further applications of Theorem 10.1. One concerns constructions of group splittings, and the other relates to the Almost Stability Theorem of [DiD].

Suppose that Γ is a one-ended finitely generated group, and that $G \leq \Gamma$ is any subgroup. Let X be a Cayley graph of Γ (or any cofinite locally finite Γ -graph). As in [Bo5] (cf. [DuS]) we say that G has *codimension-one* in Γ if there is a connected G -invariant subset, $Y \subseteq X$, such that Y/G is compact, and such that $X \setminus Y$ has at least two distinct components neither of which is contained in a uniform neighbourhood of Y . (This is independent of the choice of X .) The following result also follows directly from a result of Niblo [N]:

Proposition 15.1 : *Suppose that Γ is finitely generated and that $G \leq \Gamma$ is a codimension-one subgroup such that Γ is the commensurator of G . Then, Γ splits non-trivially as a graph of groups with G conjugate into one of the vertex groups.*

The “commensurator” condition means that $G \cap gGg^{-1}$ has finite index in G for all $g \in \Gamma$. If we assume that no vertex group is equal to an incident edge group, then it follows that all the vertex and edge groups will be commensurate with G .

To prove Proposition 15.1, let K be a coset graph of G in Γ . In other words, K is a connected cofinite Γ -graph, with $V(K)$ isomorphic as a Γ -set to the set of left translates of G , with Γ acting by left multiplication. Let $x \in V(K)$ be a vertex stabilised by G . Let $Y \subseteq X$ be as in the hypotheses. Thus, we can write $X \setminus Y = P \sqcup Q$ where neither P nor Q is contained in any uniform neighbourhood of Y . Now all but finitely many Γ -images of Y are disjoint from Y , and hence contained in either P or Q . Each Γ -image of Y corresponds to a vertex of K . This therefore assigns all but finitely many elements of $V(K)$ to one of two disjoint infinite subsets, $A, B \subseteq V(K)$, corresponding to P and Q respectively. Assigning the remaining vertices arbitrarily, we can suppose that $B = A^*$.

If $y \in A$ and $z \in A^*$, then any path connecting the corresponding images of Y in X must intersect Y . Since Y/G is finite, it follows that only finitely many Γ -images of any pair $\{y, z\} \subseteq V(K)$ can meet both A and A^* . In other words, A is a slice. We have shown that the algebra of slices, $\mathcal{B}(K)$, of K is non-trivial, Theorem 10.1 now applies to give us a nested set of generators for $\mathcal{B}(K)$. Since the set of Γ -images of Y is locally finite in X , it follows that this generating set is locally discrete. We thus get an action of Γ on a simplicial tree, T . Moreover, there is an equivariant map from $V(K)$ to $V(T)$. This proves Proposition 15.1.

This result is, in some sense, “orthogonal” to the constructions of [Bo5]. Note that we cannot expect the splittings we obtain in this case to be canonical.

Another application of Theorem 10.1 (as observed by Dunwoody) is to give an alternative proof of a version of the Almost Stability Theorem of [DiD]. This can be interpreted as giving a criterion for a Γ -set to be embedded in the vertex set of a simplicial Γ -tree with finite edge stabilisers. Let Γ be a group, and X be a cofinite Γ -set. Let $\mathcal{P}(X)$ be the power set of X , thought of as a Γ -boolean algebra, and let \mathcal{I} be the ideal of finite subsets.

Proposition 15.2 : *Suppose that Γ is a finitely generated group, and that X is a Γ -set with finite point stabilisers. Suppose that $V \subseteq \mathcal{P}(X)$ is a Γ -invariant subset of $\mathcal{P}(X)$ with the property that if $A, B \in V$ then $A + B \in \mathcal{I}$. Then V can be equivariantly embedded in the vertex set of a simplicial Γ -tree.*

Suppose we already know that V embeds in a simplicial Γ -tree, T . We can take X to be the directed edge set of T . To each x in V , we associate the set of directed edges which point towards x . This gives a subset of $\mathcal{P}(X)$ satisfying the hypotheses of Proposition 13.2. Of course, the situation described by the hypotheses may be more general than this.

Proposition 15.2 can be proven as follows. Given any $x \in X$, let $A(x) \in \mathcal{P}(V)$ be the subset of V consisting of those elements of V that contain x . One verifies that $A(x)$ is a slice of V . Let \mathcal{A} be the subalgebra of slices of V generated by $\{A(x) \mid x \in V\}$. (A typical element of \mathcal{A} has the form B or B^* , where B is either finite or is equal to $A(x)$ for some $x \in X$.) We now apply Theorem 10.2. It is easily verified that the resulting generating set is locally discrete.

References.

- [Ber] G.M.Bergman, *On groups acting on locally finite graphs* : Ann. of Math. **88** (1968) 335–340.
- [BesF] M.Bestvina, M.Feighn, *Bounding the complexity of simplicial group actions on trees* : Invent. Math. **103** (1991) 449–469.
- [Bo1] B.H.Bowditch, *A topological characterisation of hyperbolic groups* : J. Amer. Math. Soc. **11** (1998) 643–667.
- [Bo2] B.H.Bowditch, *Connectedness properties of limit sets* : Trans. Amer. Math. Soc. **351** (1999) 3673–3686.
- [Bo3] B.H.Bowditch, *Convergence groups and configuration spaces* : in “Group theory down under” (ed. J.Cossey, C.F.Miller, W.D.Neumann, M.Shapiro), de Gruyter (1999) 23–54.
- [Bo4] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton 1997.
- [Bo5] B.H.Bowditch, *Splittings of finitely generated groups over two-ended subgroups* : preprint, Southampton 1999.
- [CJ] A.Casson, D.Jungreis, *Convergence groups and Seifert fibered 3-manifolds* : Invent. Math. **118** (1994) 441–456.
- [DiD] W.Dicks, M.J.Dunwoody, *Groups acting on graphs* : Cambridge Studies in Advanced Mathematics No. 17, Cambridge University Press (1989).
- [Du1] M.J.Dunwoody, *Cutting up graphs* : Combinatorica **2** (1982) 15–23.
- [Du2] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449–457.

- [Du3] M.J.Dunwoody, *An inaccessible group* : in “Geometric Group Theory, Volume 1” (ed. G.A.Niblo, M.A.Roller) London Math. Soc. Lecture Notes Series No. 181, Cambridge University Press (1993) 75–78.
- [Du4] M.J.Dunwoody, *Groups acting on protrees* : J. London Math. Soc. **56** (1997) 125–136.
- [DuS] M.J.Dunwoody, E.L.Swenson, *The algebraic torus theorem* : Invent. Math. **140** (2000) 605–637.
- [F] E.M.Freden, *Negatively curved groups have the convergence property* : Ann. Acad. Sci. Fenn. Ser. A Math. **20** (1995) 333–348.
- [Ga] D.Gabai, *Convergence groups are fuchsian groups* : Ann. of Math. **136** (1992) 447–510.
- [GehM] F.W.Gehring, G.J.Martin, *Discrete quasiconformal groups I* : Proc. London Math. Soc. **55** (1987) 331–358.
- [Ger] V.Gerasimov, : Manuscript (Russian), Belo Horizonte.
- [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.
- [N] G.A.Niblo, *The singularity obstruction for group splittings* : to appear in Topology Appl.
- [Si] G.F.Simmons, *Introduction to topology and modern analysis* : McGraw-Hill (1963).
- [St] M.H.Stone, *Applications of the theory of Boolean rings to general topology* : Trans. Amer. Math. Soc. **41** (1937) 375–481.
- [T1] P.Tukia, *Homeomorphic conjugates of fuchsian groups* : J. reine angew. Math. **391** (1988) 1–54.
- [T2] P.Tukia, *Convergence groups and Gromov’s metric hyperbolic spaces* : New Zealand J. Math. **23** (1994) 157–187.
- [T3] P.Tukia, *Conical limit points and uniform convergence groups* : J. reine. angew. Math. **501** (1998) 71–98.
- [Y] A.Yaman, *A topological characterisation of relatively hyperbolic groups* : in preparation.